Non-Interference for a Practical DIFC-Based Operating System

Preliminary version 2008-11-12

Maxwell Krohn[1] and Eran Tromer[2]

[1] CyLab,
Carnegie Mellon University,
Pittsbrugh, PA, USA
krohn@mit.edu

[2] Computer Science and Artificial Intelligence Laboratory,
Massachusetts Institute of Technology,
Cambridge, MA 02139,
tromer@csail.mit.edu

Abstract. The *Flume* system is an implementation of decentralized information flow control (DIFC) at the operating system level. Prior work has shown Flume can be implemented as a practical extension to the Linux Operating System, allowing real Web applications to achieve useful security guarantees. However, the question remains if the Flume system is actually secure. This paper compares Flume with other recent DIFC system like Asbestos, arguing that the latter is susceptible to wide-bandwidth covert channels, and proving Flume is not by means of a formal non-interference proof.

1 Introduction

Recent work in operating systems [1,2,3] makes the case that Distributed Information Flow Control (DIFC) [4] solves important application-level security problems for real systems. For example, modern dynamic web sites are trusted to safeguard private data for millions of users, but they often fail in their task (e.g., [5,**?**,6,7,8,9]). Web applications built atop DIFC operating systems can achieve better security properties, by factoring security-critical code into small, isolated, trustworthy processes, while allowing the rest of the application to balloon without affecting the TCB.

To achieve such a split between trustworthy and untrustworthy components, DIFC must monitor and regulate the information flows among them. Enforcement today exists in two forms: static, as an extension of a programming language's type system; and dynamic, as a feature of an OS's system call interface. These two styles have their strengths and weaknesses: DIFC for programming languages gives fine-grained guarantees about which parts of the program have been influenced by which types of data, but it requires rewrites of existing applications using one of a few compiled languages. DIFC at the OS level gives coarser-grained information flow tracking (each process is its own security domain) but supports existing applications and languages. In particular, popular Web applications written in popular interpreted languages (e.g. Python, PHP, Perl and Ruby) can achieve improved security on DIFC OSes. A case can be made for both techniques, but to date, only DIFC at the language level has enjoyed formal security guarantees.

Consider the Flume system: a DIFC OS implemented as a 30,000-line extension to a standard Linux kernel [3]. Flume allows legacy processes to run as before, while confining those that need strong security guarantees (like web servers and network applications) to a tightly-controlled sandbox, from which all of their communication must conform to DIFC-based rules. This technique produces real security improvements in real web applications, like the popular Python-based MoinMoin Wiki package. But all claims

of application-level security presuppose a correct OS kernel, appealing to intuition alone to justify the OS's security. Intuition can mislead: other seemingly-secure OSes [1,10] have inadvertently included high-bandwidth covert channels in their very interface, allowing information to leak against intended security policies (see Section 3 for more details).

This paper presents the first formal non-interference security argument for a real DIFC operating system (in this case Flume). A DIFC OS with provable security guarantees is an important foundation for provable application-level security, both in Web applications like MoinMoin, and other security-sensitive applications yet to be explored.

The roadmap is follows. Section 2 reviews the Flume system and its intended policies at a high level: first and foremost, that untrustworthy applications can compute with private data without being able to reveal (i.e. leak) it. Section 3 describes potential pitfalls, motivating a formal approach. Section 4 describes the important parts of the Flume System using a formal process algebra, namely Communicating Sequential Processes (CSP). This model captures both a trustworthy kernel and an untrustworthy collection of user-space applications. Next, Section 5 proves that this model fulfills non-interference, a strong definition of security that means untrustworthy user-space processes cannot leak secrets they were not authorized to reveal. Though the arguments focus on secrecy, the same model and proof also applies to integrity.

In sum, this paper contributes the following new results:

1. A formal model for a real DIFC Linux-based operating system, that captures both a trustworthy kernel and an untrustworthy user-space; and
2. A formal proof of non-interference.

These are important limitations. First, the actual Flume implementation is not guaranteed to follow the model described in the process-algebra. Second, there are no guarantees that covert channels do not exist in parts of the system that the model abstracts. In particular, the Flume model does not capture physical hardware, so covert channels might of course exist in Flume's use of the processor, the disk, memory, etc. What our result does imply is that those operating systems the follow the given interface (like Flume) have a chance of achieving good security properties; i.e., wide leaks are not "baked" into their specifications. We leave to future work the mapping of the given model to a verified implementation.

2 Review of Flume

This section recapitulates Flume's security primitives, previously reported elsewhere [3]. Flume uses *tags* and *labels* to track data as it flows through a system. Let $\mathcal{T}$ be a very large set of opaque tokens called *tags*. A tag carries no inherent meaning, but processes generally associate each tag with some category of secrecy or integrity. For example, a tag $b \in \mathcal{T}$ may label Bob's private data.

Labels are subsets of $\mathcal{T}$. Labels form a lattice under the partial order of the subset relation [11]. Each Flume process p has two labels, S_p for secrecy and I_p for integrity. Both labels serve to (1) summarize which types of data have influenced p in the past and (2) regulate where p can read and write in the future. Consider a process p and a tag t. If $t \in S_p$, then the system conservatively assumes that p has seen some private data tagged with t. In the future, p can read more private data tagged with t but requires consent from an authority who controls t before it can reveal any data publicly. If there are multiple tags in S_p, then p requires independent consent for each tag before writing publicly. Process p's integrity label I_p serves as a lower bound on the purity of its influences: if $t \in I_p$, then every input to p has been endorsed as having integrity for t. To maintain this property, the system only allows p to read from other sources that also have t in their integrity labels. Files (and other objects) also have secrecy and integrity labels; they can be thought of as passive processes.

Distributed Information Flow Control (DIFC) is a generalization of centralized IFC. In centralized IFC, only a trusted "security officer" can create new tags, subtract tags from secrecy labels (*declassify*

information), or add tags to integrity labels (*endorse* information). In Flume DIFC, any process can create new tags, which gives that process the privilege to declassify and/or endorse information for those tags.

2.1 Capabilities

Flume represents privilege using two *capabilities* per tag. Capabilities are objects from the set $\mathcal{O} = \mathcal{T} \times \{-, +\}$. For tag t, the capabilities are denoted t^+ and t^-. Each process p *owns* a set of capabilities $O_p \subseteq \mathcal{O}$. A process with $t^+ \in O_p$ *owns* the t^+ capability, giving it the privilege to add t to its labels; and a process with $t^- \in O_p$ can remove t from its labels. In terms of secrecy, t^+ lets a process add t to its secrecy label, granting itself the privilege to receive secret t data, while t^- lets it remove t from its secrecy label, effectively declassifying any secret t data it has seen. In terms of integrity, t^- lets a process remove t from its integrity label, allowing it to receive low-t-integrity data, while t^+ lets it add t to its integrity label, endorsing the process's current state as high-t-integrity. A process that owns both t^+ and t^- has *dual privilege* for t and can completely control how t appears in its labels. The set D_p where

$$D_p \triangleq \{t \mid t^+ \in O_p \ \wedge \ t^- \in O_p\}$$

represents all tags for which p has dual privilege.

Any process p can allocate a tag. Tag allocation yields a fresh tag $t \in \mathcal{T}$ and sets $O_p \leftarrow O_p \cup \{t^+, t^-\}$, granting p dual privilege for t. Tag allocation should not expose any information about system state.

For a set of capabilities $O \subseteq \mathcal{O}$, we define the notation:

$$O^+ \triangleq \{t \mid t^+ \in O\} \ , \quad O^- \triangleq \{t \mid t^- \in O\} \ .$$

2.2 Global Capabilities

Flume also supports a *global* capability set $\hat{O}$, Every process has access to every capability in $\hat{O}$, useful for implementing key security policies (see Section 2.4). A process p's effective set of capabilities is given by:

$$\bar{O}_p \triangleq O_p \cup \hat{O}$$

Similarly, its effective set of dual privileges is given by:

$$\bar{D}_p \triangleq \{t \mid t^+ \in \bar{O}_p \ \wedge \ t^- \in \bar{O}_p\}$$

Tag allocation can update $\hat{O}$; an allocation parameter determines whether the new tag's t^+, t^-, or neither is added to $\hat{O}$ (and thus to every current and future process's $\bar{O}_p$).

Flume restricts access to the shared set $\hat{O}$, lest processes manipulate it to leak data. A first restriction is that processes can only add a tag to $\hat{O}$ during its allocation (otherwise a process p could leak information to a process q by either adding or refraining from adding a pre-specified tag to $\hat{O}$). A second restriction is that no process p can enumerate $\hat{O}$ or $\bar{O}_p$ (otherwise, p could poll $\|\hat{O}\|$ while q allocated new tags, allowing q to communicate bits to p). Processes can, however, enumerate their non-global capabilities (those in O_p), since they do not share this resource with other processes.

A process p can grant capabilities in O_p to process q so long as p can send a message to q. p can also subtract capabilities from O_p as it sees fit.

2.3 Security

The Flume model assumes many processes running on the same machine and communicating via messages, or "flows". The model's goal is to track data flow by regulating both process communication and process label changes.

Safe Label Changes In the Flume model (as in HiStar), the labels S_p and I_p of a process p can be changed only by an explicit request from p itself. Other models allow a process's label to change as the result of receiving a message [1,12,10], but implicit label changes turn the labels themselves into covert channels [11,2] (see Section 3). Only those label changes permitted by a process's capabilities are safe:

Definition 1 (Safe label change).
For a process p, let the label L be S_p or I_p, and let L' be the requested new value of the label. The change from L to L' is safe if and only if:

$$L' - L \subseteq (\bar{O}_p)^+ \quad \textit{and} \quad L - L' \subseteq (\bar{O}_p)^- \ .$$

For example, say process p wishes to subtract tag t from S_p, to achieve a new secrecy label S'_p. In set notation, $t \in S_p - S'_p$, and such a transition is safe only if p owns the subtraction capability for t (i.e. $t^- \in \bar{O}_p$). Likewise, t can be added only if $t^+ \in \bar{O}_p$.

Safe Messages Information flow control restricts process communication to prevent data leaks. Essentially, the Flume model restricts communication among unprivileged processes as in traditional IFC: p can send a message to q only if $S_p \subseteq S_q$ ("no read up, no write down" [13]) and $I_q \subseteq I_p$ ("no read down, no write up" [14]).

Flume relaxes these rules to better accommodate declassifiers (see Appendix .2 for an example). Specifically, if two processes *could* communicate by changing their labels, sending a message using the centralized IFC rules, and then restoring their original labels, then the model can safely allow the processes to communicate without actually performing label changes. A process can make such a temporary label change only for tags in $\bar{D}_p$, i.e., those for which it has dual privilege. A process p with labels S_p, I_p would get maximum latitude in sending messages if it were to lower its secrecy to $S_p - \bar{D}_p$ and raise its integrity to $I_p \cup \bar{D}_p$. It could receive the most messages if it were to raise secrecy to $S_p \cup \bar{D}_p$ and lower integrity to $I_p - \bar{D}_p$.

The following definition captures these *hypothetical* label changes to determine what messages are safe:

Definition 2 (Safe message).
A message from p to q is safe iff

$$S_p - \bar{D}_p \subseteq S_q \cup \bar{D}_q \quad \textit{and} \quad I_q - \bar{D}_q \subseteq I_p \cup \bar{D}_p \ .$$

External Sinks and Sources Any data sink or source outside of Flume's control, such as a remote host, the user's terminal, a printer, and so forth, is modeled as an unprivileged process x with permanently empty secrecy and integrity labels: $S_x = I_x = \{\}$ and also $O_x = \{\}$. As a result, a process p can only write to the network or console if it could reduce its secrecy label to $\{\}$ (the only label with $S_p \subseteq S_x$), and a process can only read from the network or keyboard if it could reduce its integrity label to $\{\}$ (the only label with $I_x \subseteq I_p$).

2.4 Security Policies

The most important security policy in Flume is *export protection*, wherein untrustworthy processes can compute with secret data without the ability to revealing it. An export protection tag is a tag t such that $t^+ \in \hat{O}$ and $t^- \notin \hat{O}$. For a process p to achieve such a result, it creates a new tag t and grants t^+ to the global set $\hat{O}$, while closely guarding t^-. To protect a file f, p creates the file with secrecy label $\{t\}$. If a process q wishes to read f, it must first add t to S_q, which it can do since $t^+ \in \bar{O}_q$. Now, q can only send messages to other processes with t in their labels. It requires p's authorization to remove t from its label or send data to the network. Additional Flume policies exist, like those that protect integrity.

3 Covert Channels in Dynamic Label Systems

As described in Section 2.3, processes in Flume change their labels explicitly; labels do not change implicitly upon message receipt, as they do in Asbestos [1] or IX [10]. We show by example why implicit label changes (also known as "floating" labels) enable high-throughput information leaks (as predicted by Denning [11]).

Consider a process p with secrecy label $S_p = \{t\}$ and a process q with $S_q = \{\}$, both with empty ownership sets $O_p = O_q = \{\}$. In a floating label system like Asbestos, p can send a message to q, and q will successfully receive it, upon which the kernel automatically raises $S_q = \{t\}$. Thus, the kernel can track which processes have seen secrets tagged with t, even if those processes are uncooperative. Such a scheme introduces new problems: what if a process q doesn't want its label to change from $S_q = \{\}$? For this reason, Asbestos also introduces "receive labels," which serve to filter out incoming traffic to a process, letting it avoid unwanted label changes.

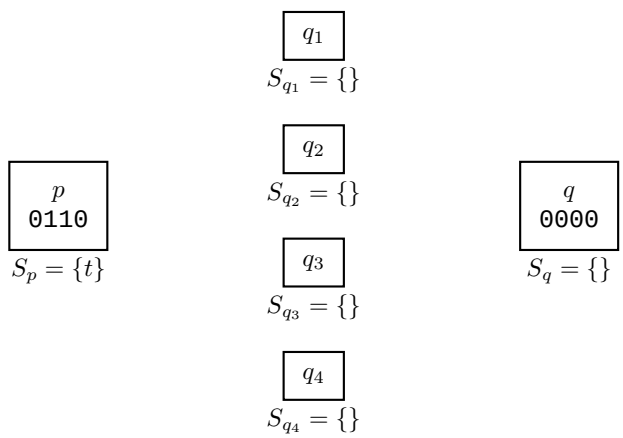

Fig. 1. The "leaking" system initializes.

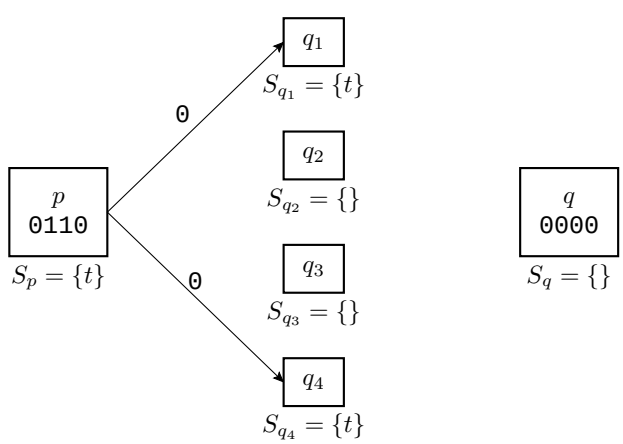

Fig. 2. p sends a "0" to q_i if the ith bit of the message is 0.

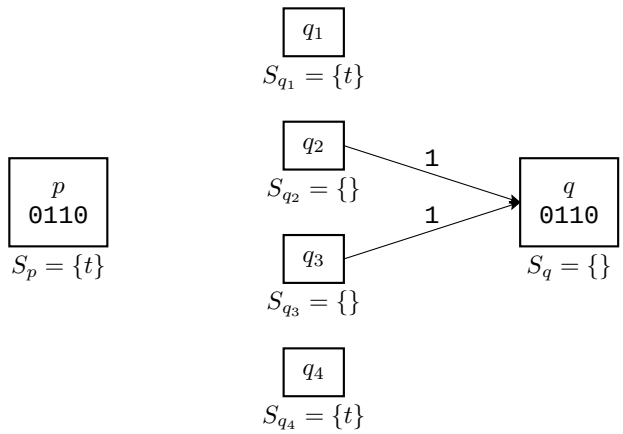

Fig. 3. If q_i did not receive a "0" before the timeout, it assumes an implicit "1" and writes "1" to q at position i.

The problem with floating is best seen through example (see Figures 1–3). Imagine processes p and q as above, with a sender process p wanting to leak the 4-bit secret "0110" to a receiver process q. Their goal is to transmit these bits without q's label changing. Figure 1 shows the initialization: q launches 4 helper processes (q_1 through q_4), each with a label initialized to $S_{q_i} = \{\}$. q's version of the secret starts out initialized to all 0s, but it will overwrite some of those bits during the attack.

Next, p communicates selected bits of the secret to its helpers. If the ith bit of the message is equal to 0, then p sends the message "0" to the process q_i. If the ith bit of the message is 1, p does nothing. Figure 2 shows this step. Note that as a result of receiving these 0 bits, q_1 and q_4 have changed labels! Their labels floated up from $\{\}$ to $\{t\}$, as the kernel accounts for how information flowed in the system.

In the last step (Figure 3), the q_i processes wait for a predefined time limit before giving up. At the timeout, each q_i which did not receive a message (here, q_2 and q_3) sends a message "1" to q, and upon receipt of this message q updates its bit position i to 1. The remaining processes (q_1 and q_4) do not write to q, nor could they without affecting q's label. Now, q has the exact secret, copied bit-for-bit from p. This example shows 4 bits of data leak, but by forking n processes, p and q can leak n bits per timeout period. Because Asbestos's *event process* abstraction makes forking very fast, this channel on Asbestos can leak kilobits of data per second.

This attack fails against the Flume system. In Figure 2, each q_i must each make a decision: should it raise its secrecy label to $S_{q_i} = \{t\}$, or leave its as is? If q_i raises S_{q_i} then it will receive messages from p, but it won't be able to write to q. Otherwise, q_i will never receive a message from p. In either case, q_i cannot alter its messages to q in response to messages from p. And crucially, q_i must decide whether to upgrade S_{q_i} *before* receiving messages from p.

4 The Formal Flume Model

Though the above attack fails against the Flume model, no formal reasoning proves that other attacks could not succeed. This section and the next seek a formal separation between the Asbestos style of "floating" labels and the Flume style of "explicitly specified" labels. The ultimate goal is to prove that Flume exhibits *non-interference* properties: for example, that processes with empty ownership and whose secrecy label contains t cannot in any way alter the execution of those processes with empty labels. Such

a non-interference result requires a formal model of Flume, which we build up here. Section 5 provides the proof that the Flume Model meets a standard definition of non-interference with high probability.

We present a formal model for the Flume System in the Communicating Sequential Processes (CSP) process algebra [15]. (The CSP algebra is reviewed in Appendix .1.) The model captures a kernel, a system call interface, and arbitrary user processes that can interact via the system call interface. Processes can communicate with one another over IPC, changing labels, allocating tags, and forking new processes. In model dictates which kernel details are safe to expose to user-level applications, where I/O can safely happen, which return codes from system calls to provide, etc. It does not capture lower-level hardware details, like CPU, cache, memory, network or disk usage. Therefore, it is powerless to disprove the existence of covert channels that modulate CPU, cache, memory, network or disk usage to communicate data from one process to another.

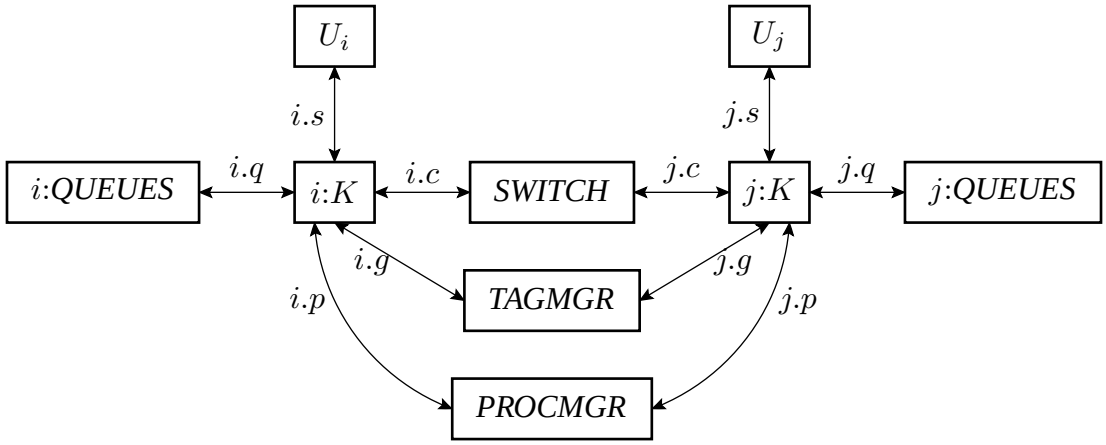

Fig. 4. Two user processes, U_i and U_j, in the CSP model for Flume. i:K and j:K are the kernel halves of these two processes (respectively), *TAGMGR* is the process that manages the global set of tags and associated privileges, *PROCMGR* manages the process ID space, and *SWITCH* enables all user-visible interprocess communication. Arrows denote CSP communication channels.

Figure 4 depicts the Flume model organization. At a high level, the model splits each Unix-like process i running on a system (e.g., a web server or text editor) into two logical components: a *user half* process U_i which can take almost any form, and a *kernel half* process i:K which behaves according to a strict state machine.[3] The user half of a process can communicate with its kernel half (and thus, indirectly, with user processes) via a system call interface. This interface takes the form of a CSP channel $i.s$ between the U_i and i:K. Inside the kernel, the kernel-halves of processes can communicate with one another to deliver IPCs initiated by user processes. Also inside the kernel, a global process (*TAGMGR*) manages the circulation of tags and globally-shared privileges; another global process (*PROCMGR*) manages the process ID space. The process *SWITCH* is involved with communication between user-level processes. The remainder of this section fills out the details of the Flume model.

[3] The CSP notation i:K means the i-th copy of a template process K.

4.1 System Call Interface

At a high level, user-level processes communicate with the kernel (and thus, indirectly, with each other) through a "system-call" interface consisting of events on the channel $i.s$ between U_i and $i{:}K$. Each user-level process has access to the following system calls:

- $t \leftarrow \mathsf{create\_tag}(\mathit{which})$
  Allocate a new tag t, and depending on the parameter *which*, make the associated capabilities for t globally accessible. Here, *which* can be one of None, Remove or Add. For Remove, add t^- to $\hat{O}$, essentially granting it to all other processes; likewise, for Add, add t^+ to $\hat{O}$.
- $\mathrm{rc} \leftarrow \mathsf{change\_label}(\mathit{which}, L)$
  Change the process's *which* label to L. Return Ok on success and Error on failure. Here, *which* can be either Secrecy or Integrity.
- $L \leftarrow \mathsf{get\_label}(\mathit{which})$
  Read this process's own label out of the kernel's data structures. Here, *which* can be either Secrecy or Integrity, controlling which label is read.
- $O \leftarrow \mathsf{get\_caps}()$
  Read this process's ownership set out of the kernel's data structures.
- $\mathsf{send}(j, \mathit{msg}, X)$
  Send message *msg* and capabilities X to process j. (Crucially, the sender gets no indication whether the transmission failed due to label checks.)
- $(\mathrm{msg}, X) \leftarrow \mathsf{recv}(j)$
  Receive message *msg* and capabilities X from process j. Block until a message is ready.
- $j \leftarrow \mathsf{fork}()$
  Fork the current process; yield a process j. fork returns j in the parent process and 0 in the child process.
- $i \leftarrow \mathsf{getpid}()$
  Return i, the ID of the current process.
- $\mathsf{drop\_caps}(X)$
  Set $O_i \leftarrow O_i - X$.

The Flume model places no restrictions on what the user portions of processes can do other than: (1) such processes cannot communicate with each other; and (2) they can only communicate with the kernel via the prescribed system call interface. Formally, let $C_i = \{c \mid (c.v) \in \alpha U_i\}$ be the set of channels that U_i has. For instance, $i.s \in C_i$ where $i.s$ is the channel that process i uses to make system calls into the kernel. The first requirement on U_i is that for all $j \neq i$, $C_i \cap C_j = \{\}$. That is, no process U_i can communicate directly with another process U_j. Also, process i cannot tamper with the system call interface of any other process j, meaning $j.s \notin C_i$ for all j, $j \neq i$. Finally, each kernel process $j{:}K$ has four other channels, $\{j.b, j.g, j.c, j.p\}$, all discussed below. No user process can access any of these channels directly. That is, for all i, j: $\{j.b., j.g., j.c, j.p\} \cap C_i = \{\}$. Of course, C_i is not empty. For all i, $i.s \in C_i$, where $i.s$ is U_i's dedicated channel for sending system calls to the kernel and receiving replies. C_i can also contain channels from the process U_i to itself.

4.2 Kernel Processes

Each process i has an instantiation of the kernel process K that obeys a strict state machine. We apply CSP's standard technique for "relabeling" the interior states of a process, giving $i{:}K$. Because by definition, $i{:}K$ and $j{:}K$ have different alphabets for $i \neq j$, their operations cannot interfere. Each process $i{:}K$ takes on a state configuration based on process i's labels. That is, $i{:}K_{S,I,O}$ denotes the kernel half of process i, with secrecy label $S_i \subseteq \mathcal{T}$, integrity label $I_i \subseteq \mathcal{T}$, and ownership of capabilities given by $O_i \subseteq \mathcal{O}$.

At a high level, a kernel process K starts idle, then springs to life once receiving an activation message. Once active, it receives either system calls from its user half, or internal messages from other kernel processes on the system. It eventually dies when the user process exits. In CSP notation:

$$K \;=\; b?(S, I, O) \rightarrow K_{S,I,O}$$

where b is the channel that K listens on for its "birth" message. It expects arguments of the form (S, I, O), to instruct it which labels and capabilities to start its execution with. Subsequently, $K_{S,I,O}$ handles the bulk of the kernel process's duties:

$$K_{S,I,O} \;=\; SYSCALL_{S,I,O} \mid INTRECV_{S,I,O}$$

where $SYSCALL$ is a subprocess tasked with handling all system calls, and $INTRECV$ is the *internal receiving* sub-process, tasked with receiving internal messages from other kernel processes.

For any process i, the subprocess $i{:}K_{S,I,O}$ handles system calls by listening for incoming messages from U_i along a shared channel $i.s$. In the definition of $K_{S,I,O}$, each system call gets its own dedicated subprocess:

$$\begin{aligned} SYSCALL_{S,I,O} \;=\; & NEWTAG_{S,I,O} \mid \\ & CHANGELABEL_{S,I,O} \mid \\ & READMYLABEL_{S,I,O} \mid \\ & READMYCAPS_{S,I,O} \mid \\ & DROPCAPS_{S,I,O} \mid \\ & SEND_{S,I,O} \mid \\ & RECV_{S,I,O} \mid \\ & FORK_{S,I,O} \mid \\ & GETPID_{S,I,O} \mid \\ & EXIT_{S,I,O} \end{aligned}$$

Section 4.4 presents all of these subprocesses in more detail.

4.3 Process Alphabets

In the next section, we will prove properties about the system, in particular, that messages between "high processes" (those that have a specified tag in their secrecy label) do not influence the activity of "low processes." The standard formula for such proofs is to split the system's alphabet into two disjoint sets: "high" symbols, those that the secret influences; and "low" symbols, those that should not be affected by the secret. We must provide the appropriate alphabets for these processes so that any symbol in the model unambiguously belongs to one set or the other.

Consider some examples. Take process i with secrecy label $S_i = \{t\}$ and integrity label $I_i = \{\}$. When U_i issues a system call (say $\mathsf{create\_tag}(\mathsf{Add})$) to its kernel half $i{:}K$, the trace for U_i is of the form

$$\langle \ldots, i.s!(\mathsf{create\_tag}, \mathsf{Add}), \ldots \rangle$$

and the trace for the kernel half $i{:}K$ is of the form

$$\langle \ldots, i.s?(\mathsf{create\_tag}, \mathsf{Add}), \ldots \rangle \;.$$

That is, U_i is sending a message $(\mathsf{create\_tag}, \mathsf{Add})$ on the channel $i.s$, and $i{:}K$ is receiving it. The problem, however, is that looking at these traces does not capture the fact that i's secrecy label contains t and therefore that U_i is in a "high" state in which it should not affect low processes. Such a shortcoming does not inhibit the accuracy of the model, but it does inhibit the proof of non-interference in Section 5.

A solution to the problem is to include a process's labels in the messages it sends. That is, once i has a secrecy label of $S = \{t\}$, its kernel process should be in a state such as $K_{\{t\},\{\},\{\}}$. When a kernel process is in this state, it will only receive system calls of the form $i.s?(\{t\}, \{\}, \mathsf{create\_tag}, \mathsf{Add})$. Thus, U_i must now send system calls in the form:

$$i.s!(\{t\}, \{\}, \mathsf{create\_tag}, \mathsf{Add})$$

In this regime, U_i must do its own accounting of S_i and I_i. Failure to do so prevents U_i from making successful system calls but does *not* affect security guarantees.

Messages of the form $c!(S, I, \dots)$ and $c?(S, I, \dots)$, for various channels c, occur often in our model; for clarity we use the following notation. For any α:

$$c\underset{S,I}{!}(\alpha) \triangleq c!(S, I, \alpha)\ , \qquad c\underset{S,I}{?}(\alpha) \triangleq c?(S, I, \alpha)$$

In the context of a kernel process $K_{S,I,O}$, we need not specify S and I explicitly; they are inferred from the kernel's state. That is, when appearing inside a process $K_{S,I,O}$, we define:

$$c\underset{\kappa}{!}(\alpha) \triangleq c!(S, I, \alpha)\ , \qquad c\underset{\kappa}{?}(\alpha) \triangleq c?(S, I, \alpha)$$

4.4 System Calls

We now define the sub-processes of $K_{S,I,O}$ that correspond to kernel's implementation of each system call.[4] The first system call subprocess handles a user process's request for new tags. Much of this system call is handled by a subroutine call to the global tag manager *TAGMGR*. Note that after tag allocation, the kernel process transitions to a different state, reflecting the new privilege(s) it acquired for tag t.

$$\begin{aligned} NEWTAG_{S,I,O} = \big(& s\underset{\kappa}{?}(\mathsf{create\_tag}, w) \rightarrow \\ & g\underset{\kappa}{!}(\mathsf{create\_tag}, w)?(t, O_{\text{new}}) \rightarrow \\ & s\underset{\kappa}{!}t \ \rightarrow \ K_{S,I,O\cup O_{\text{new}}}\big) \end{aligned}$$

The *CHANGELABEL* subprocess is split into two cases, for changes to secrecy labels vs. changes to integrity labels:

$$\begin{aligned} CHANGELABEL_{S,I,O} = {} & S\text{-}CHANGELABEL_{S,I,O} \ \mid \\ & I\text{-}CHANGELABEL_{S,I,O} \end{aligned}$$

$$\begin{aligned} S\text{-}CHANGELABEL_{S,I,O} = \big(& chk : CHECK_{S,I,O} /\!/ \\ & \big(s\underset{\kappa}{?}(\mathsf{change\_label}, \mathsf{Secrecy}, S') \rightarrow \\ & chk!(S, S')?r \rightarrow \\ & \textbf{if } r \textbf{ then } s\underset{\kappa}{!}\mathsf{Ok} \rightarrow K_{S',I,O} \\ & \quad \textbf{else } s\underset{\kappa}{!}\mathsf{Error} \rightarrow K_{S,I,O}\ \big)\big) \end{aligned}$$

[4] A note on notation: the channel between U_i and $i{:}K$ is named $i.s$ as stated earlier. However, the "$i.$" prefix is induced by the CSP renaming operator on the "template" kernel process K. In this section we define the subprocesses of K, so the channel is named merely s.

$$
\begin{aligned}
\textit{I-CHANGELABEL}_{S,I,O} = \big(&chk : \textit{CHECK}_{S,I,O} /\!\!/ \\
&\big(s \underset{K}{?} (\mathsf{change\_label}, \mathsf{Integrity}, I') \rightarrow \\
&chk!(I, I')?r \rightarrow \\
&\textbf{if } r \textbf{ then } s \underset{K}{!} \mathsf{Ok} \rightarrow K_{S,I',O} \\
&\quad \textbf{else } s \underset{K}{!} \mathsf{Error} \rightarrow K_{S,I,O} \big)\big)
\end{aligned}
$$

In both cases, the user process specifies a new label, and the *CHECK* subroutine determines if that label change is valid. In the success case, the kernel process transitions to a new state, reflecting the new labels. In the failure case, the kernel process remains in the same state. The *CHECK* process computes the validity of the label change based on the process's current capabilities, as well as the global capabilities $\hat{O}$ (represented by a global tag register process listening on the g channel):

$$
\begin{aligned}
\textit{CHECK}_{S,I,O} = &?(L, L') \rightarrow \\
&g!(\mathsf{check\text{-}}, L - L' - O^-) \;\rightarrow\; g?r \;\rightarrow \\
&g!(\mathsf{check+}, L' - L - O^+) \;\rightarrow\; g?a \;\rightarrow \\
&!(r \wedge a) \;\rightarrow\; \textit{CHECK}_{S,I,O}
\end{aligned}
$$

As we will see below, the global tag register replies True to $(\mathsf{check\text{-}}, L)$ iff $L \subseteq \hat{O}^-$, and replies True to $(\mathsf{check+}, L)$ iff $L \subseteq \hat{O}^+$. Thus, we have that the user process can only change from label L to L' if it can subtract all tags in $L - L'$ and add all tags in $L' - L$, either by its own capabilities or those globally owned (see Definition 1 in Section 2.3).

The user half of a process can call the kernel half to determine its own S or I labels and its capabilities O. These calls are handled simply by the following subprocesses:

$$
\begin{aligned}
&\textit{READMYLABEL}_{S,I,O} = \\
&\quad \big(s \underset{K}{?} (\mathsf{get\_label}, \mathsf{Secrecy}) \rightarrow s \underset{K}{!} S \rightarrow K_{S,I,O} \mid \\
&\quad s \underset{K}{?} (\mathsf{get\_label}, \mathsf{Integrity}) \rightarrow s \underset{K}{!} I \;\rightarrow K_{S,I,O}\big)
\end{aligned}
$$

$$
\textit{READMYCAPS}_{S,I,O} = \big(s \underset{K}{?} (\mathsf{get\_caps}) \rightarrow s \underset{K}{!} O \;\rightarrow K_{S,I,O}\big)
$$

A process also can discard capabilities using *DROPCAPS*:

$$
\textit{DROPCAPS}_{S,I,O} = \big(s \underset{K}{?} (\mathsf{drop\_caps}, X) \;\rightarrow\; K_{S,I,O-X}\big)
$$

On a successful drop of capabilities, the process transitions to a new kernel state, reflecting the reduced ownership set.

The next process to cover is forking. Recall that each active task i on the system has two components: a user component U_i and a kernel component $i{:}K$. The Flume model does not capture what happens to U_i when it calls fork,[5] but the kernel-side behavior of fork is specified as follows:

$$
\begin{aligned}
\textit{FORK}_{S,I,O} = \big(&s \underset{K}{?} (\mathsf{fork}) \;\rightarrow \\
&p \underset{K}{!} (\mathsf{fork}, O) \rightarrow p?j \;\rightarrow \\
&s \underset{K}{!} j \;\rightarrow\; K_{S,I,O}\big)
\end{aligned}
$$

[5] In Flume's concrete implementation of this model, forking has the familiar semantics of copying the address space and configuring the execution environment of the child process.

Recall that $i.p$ is a channel from the i-th kernel process to the process manager in the kernel, *PROCMGR*. (The latter allocates the child's process ID j and gives birth to $j{:}K$.)

The process handling getpid is straightforward:

$$\begin{aligned} GETPID_{S,I,O} = \big(s \underset{K}{?} (\mathsf{getpid}) \ \rightarrow & \\ & p!(\mathsf{getpid}) \rightarrow p?i \ \rightarrow \\ & s \underset{K}{!} i \ \rightarrow \ K_{S,I,O}\big) \end{aligned}$$

And user processes issue an exit system call as they terminate:

$$\begin{aligned} EXIT_{S,I,O} = \big(s \underset{K}{?} (\mathsf{exit}) \ \rightarrow & \\ & q!(\mathsf{clear}) \ \rightarrow \ p!(\mathsf{exit}) \ \rightarrow \ SKIP\big) \end{aligned}$$

Once a process with a given ID has run and exited, its ID is retired, never to be used again.

4.5 Communication

The communication subprocesses are the crux of the Flume CSP model. They require care to ensure that subtle state transitions in high processes do not result in observable behavior by low processes. At the same time, they must make a concerted effort deliver messages, so that the system is useful.

The beginning of a message delivery sequence is the process $i{:}SEND_{S,I,O}$, invoked when U_i wishes to send a message to U_j. To this succeed as often as possible, the kernel attempts to shrink the process's S label and to grow its integrity label I as much as allowed by the process's privileges (by querying the tag manager on channel $i.g$). The message is passed through the switchboard process *SWITCH* via channel $i.c$, and the switchboard forwards the message onto the destination j.

$$\begin{aligned} SEND_{S,I,O} = \big(s \underset{K}{?} (\mathsf{send}, j, X, m) \ \rightarrow & \\ & g!(\mathsf{dual\_privs}, O) \rightarrow g?D \rightarrow \\ & c!(S - D, I \cup D, j, X \cap O, m) \ \rightarrow \ K_{S,I,O}\big) \end{aligned}$$

The process *SWITCH* listens on the other side of the receive channel $i.c$. It accepts messages of the form $i.c?(S, I, j, X, m)$ and forwards them to the process $j{:}K$ as $j.c!(S, I, j, X, m)$:

$$\begin{aligned} SWITCH \ = \ |_{\forall i}\big(i.c?(S, I, j, X, m) \ \rightarrow & \\ & (j.c!(S, I, i, X, m) \ \rightarrow SKIP \ \| \ SWITCH)\big) \end{aligned}$$

The *SWITCH* process sends messages in parallel with the next receive operation. This parallelism avoids deadlocking the system if the receiving process has exited, not yet started, or is waiting to send a message. In other words, the *SWITCH* process is always willing to receive a new message, delegating potentially-blocking send operations to an asynchronous child process.

Once the message leaves the switch, the receiver process handles it with its *INTRECV* subprocess. After performing the label checks given by Definition 2 in Section 2.3, this process enqueues the incoming message for later retrieval:

$$\begin{aligned} INTRECV_{S,I,O} \ = c?(S_{\text{in}}, I_{\text{in}}, j, X, m) \ \rightarrow & \\ & g!(\mathsf{dual\_privs}, O) \rightarrow g?D \rightarrow \\ & \textbf{if } (S_{\text{in}} \subseteq S \cup D) \ \wedge \ (I - D \subseteq I_{\text{in}}) \\ & \textbf{then } q!(\mathsf{enqueue}, (X, m)) \rightarrow K_{S,I,O} \\ & \textbf{else } K_{S,I,O} \end{aligned}$$

The final link in the chain is the actual message delivery in user space. For a user process to receive a message, it calls into the kernel, asking it to dequeue and deliver any waiting messages. Receiving also updates the process's ownership, to reflect new capabilities it gained.

$$\begin{aligned} RECV_{S,I,O} = \big(& s \underset{\kappa}{?} (\mathsf{recv}, j) \ \rightarrow \\ & q!(\mathsf{dequeue}, j) \rightarrow q?(X, m) \ \rightarrow \\ & s \underset{\kappa}{!} m \ \ \rightarrow \ \ K_{S,I,O \cup X} \big) \end{aligned}$$

4.6 Helper Processes

It now remains to fill in the details for the helper processes that the various $K_{S,I,O}$ processes call upon. They are: *TAGMGR*, which manages all global tag allocation and global capabilities; *QUEUES*, which manages receive message queues, one per process; and finally *PROGMGR*, which manages process creation, deletion, etc.

The Tag Manager (*TAGMGR*) The tag manager keeps track of a global universe of tags ($\mathcal{T}$), and the global set of privileges available to all processes ($\hat{O}$). It also tabulates which tags have already been allocated, so as never to reissue the same tag. The set $\hat{T}$ refers to those tags that were allocated in the past. Thus, the task manager's state is parameterized as $TAGMGR_{\hat{O},\hat{T}}$. Initially, $\hat{O}$ and $\hat{T}$ are empty:

$$TAGMGR = TAGMGR_{\{\},\{\}}$$

Once active, the tag manager services several calls:

$$\begin{aligned} TAGMGR_{\hat{O},\hat{T}} \ = \ & NEWTAG{+}_{\hat{O},\hat{T}} \ | \\ & NEWTAG{-}_{\hat{O},\hat{T}} \ | \\ & NEWTAG0_{\hat{O},\hat{T}} \ | \\ & DUALPRIVS_{\hat{O},\hat{T}} \ | \\ & CHECK{+}_{\hat{O},\hat{T}} \ | \\ & CHECK\text{-}_{\hat{O},\hat{T}} \end{aligned}$$

Many of these subprocesses will call upon a subroutine that randomly chooses an element from a given set. We define that subroutine here. Given a set Y:

$$CHOOSE_Y = ?(S, I) \ \rightarrow \ \sqcap_{y \in Y} (!y) \ \rightarrow \ STOP$$

That is, the subprocess *CHOOSE* nondeterministically picks an element y from Y and returns it to the caller. As we will see in Section 5, *CHOOSE*'s refinement (i.e., its instantiation) has an important impact on security. It can, and in some cases should, take into account the labels on the kernel process on whose behalf it operates.

The first set of calls involve allocating new tags, such as:

$$\begin{aligned} NEWTAG{+}_{\hat{O},\hat{T}} = \ & choose{:}CHOOSE_{\mathcal{T}-\hat{T}} /\!\!/ \\ & \big|_{\forall i} \big(i.g \underset{\kappa}{?} (\mathsf{create\_tag}, \mathsf{Add}) \ \rightarrow \\ & \quad choose!(S, I)?t \ \rightarrow \\ & \quad i.g!(t, \{t^-\}) \ \rightarrow \\ & \quad TAGMGR_{\hat{O} \cup \{t^+\}, \hat{T} \cup \{t\}} \big) \end{aligned}$$

That is, the subprocess *NEWTAG+* picks a channel i such that $i.g$ has input available. Then, it chooses an unallocated tag t via *CHOOSE* and returns that tag to the calling kernel process. It services the next request in a different state, reflecting that a new capability is available to all processes (t^+) and the tag t is now allocated.

We next define *NEWTAG–* and *NEWTAG0* similarly:

$$\begin{aligned}
NEWTAG\text{–}_{\hat{O},\hat{T}} = \mathit{choose} : & CHOOSE_{\mathcal{T}-\hat{T}} /\!\!/ \\
& \Big|_{\forall i} \big(i.g \underset{\kappa}{?} (\mathsf{create_tag}, \mathsf{Remove}) \ \rightarrow \\
& \quad \mathit{choose}!(S, I)?t \ \rightarrow \\
& \quad i.g!(t, \{t^+\}) \ \rightarrow \\
& \quad TAGMGR_{\hat{O} \cup \{t^-\}, \hat{T} \cup \{t\}} \big) \\
NEWTAG0_{\hat{O},\hat{T}} = \mathit{choose} : & CHOOSE_{\mathcal{T}-\hat{T}} /\!\!/ \\
& \Big|_{\forall i} \big(i.g \underset{\kappa}{?} (\mathsf{create_tag}, \mathsf{None}) \ \rightarrow \\
& \quad \mathit{choose}!(S, I)?t \ \rightarrow \\
& \quad i.g!(t, \{t^-, t^+\}) \ \rightarrow \\
& \quad TAGMGR_{\hat{O}, \hat{T} \cup \{t\}} \big)
\end{aligned}$$

The purpose of the *DUALPRIVS* subprocess is to augment a user process's ownership set O with all of the globally-held privileges available in $\hat{O}$. That is, to return $\bar{O}_i = O_i \cup \hat{O}$ for a given process i. However, it must be careful to do so without allowing a process to enumerate the contents of $\hat{O}$. To achieve both ends, we specialize the interface to *TAGMGR*. Given a set O_i, the tag manager process will return the set of tags that U_i has dual privilege for. Since there are no tags t such that $\{t^-, t^+\} \subseteq \hat{O}$, it follows that the process must own at least one privilege for t to get dual privilege for it. Thus, the *DUALPRIVS* call will not alert any process to the existence of any tags it did not already know of:

$$\begin{aligned}
DUALPRIVS_{\hat{O},\hat{T}} = \Big|_{\forall i} \big(& i.g?(\mathsf{dual_privs}, O_i) \ \rightarrow \\
& i.g!((O_i^+ \cup \hat{O}^+) \cap (O_i^- \cup \hat{O}^-)) \ \rightarrow \\
& TAGMGR_{\hat{O},\hat{T}} \big)
\end{aligned}$$

Finally, the behavior of *CHECK+* has already been hinted at. Recall this subprocess checks to see if the supplied set of tags is globally addable:

$$\begin{aligned}
CHECK+_{\hat{O},\hat{T}} = \Big|_{\forall i} \big(& i.g?(\mathsf{check+}, L) \ \rightarrow \\
& (\textbf{if } L \subseteq \hat{O}^+ \\
& \ \textbf{then } i.g!\mathsf{True} \\
& \ \textbf{else } i.g!\mathsf{False}) \ \rightarrow \\
& TAGMGR_{\hat{O},\hat{T}} \big)
\end{aligned}$$

And similarly:

$$\begin{aligned}
CHECK\text{-}_{\hat{O},\hat{T}} = \Big|_{\forall i} \big(& i.g?(\mathsf{check\text{-}}, L) \ \rightarrow \\
& (\textbf{if } L \subseteq \hat{O}^- \\
& \ \textbf{then } i.g!\mathsf{True} \\
& \ \textbf{else } i.g!\mathsf{False}) \ \rightarrow \\
& TAGMGR_{\hat{O},\hat{T}} \big)
\end{aligned}$$

The Process Manager (*PROCMGR*) The main job of the process manager is to allocate process identifiers when kernel processes call fork. We assume a large space of process identifiers, $\mathcal{P}$. The process manager keeps track of subset $\hat{P} \subseteq \mathcal{P}$ to account for which of those processes identifiers have already been used. In then allocates from $\mathcal{P} - \hat{P}$.

$$PROCMGR_{\hat{P}} = PM\text{-}FORK_{\hat{P}} \mid PM\text{-}GETPID_{\hat{P}} \mid PM\text{-}EXIT_{\hat{P}}$$

To answer the fork operation, the process manager picks an unused process ID (j) for the child, gives birth to the child ($j{:}K$) with the message $j.b!(S, I, O)$, and returns child's process ID to the caller (parent):

$$\begin{aligned} PM\text{-}FORK_{\hat{P}} = & \mathit{choose} : CHOOSE_{\mathcal{P}-\hat{P}} /\!\!/ \\ & \big|_{\forall i} \big(i.p?(S, I, \mathsf{fork}, O) \ \rightarrow \\ & \quad \mathit{choose}!(S, I)?j \ \rightarrow \\ & \quad j.b!(S, I, O) \ \rightarrow \ i.p!(j) \ \rightarrow \\ & \quad PROGMGR_{\hat{P} \cup \{j\}} \big) \end{aligned}$$

Trivially:

$$PM\text{-}GETPID = \big|_{\forall i} \big(i.p?(\mathsf{getpid})!i \ \rightarrow \ PROCMGR \big)$$

Kernel processes notify the process manager of their exits. Of course, such notification would give it opportunity to update its accounting and to return the exiting process identifier into circulation. But for simplicy, we model this as a no-op:

$$PM\text{-}EXIT = \big|_{\forall i} \big(i.p?(\mathsf{exit}) \ \rightarrow \ PROCMGR \big)$$

A final task for the process manager is to initialize the system, launching the first kernel process. This process runs with special process ID init, off-limits to other processes. Thus:

$$PROCMGR0 = \mathsf{init}.b!(\{\}, \mathcal{T}, \{\}, \{\}) \rightarrow PROCMGR_{\mathcal{P}-\{\mathsf{init}\}}$$

See Appendix .3 for a description of the *QUEUES* process, which is mostly an implementation detail.

4.7 High Level System Definition

The overall system *SYS* is an interleaving of all the processes specified. Consider some subset $J \subseteq \mathcal{P}$ of all possible process IDs. The user-half of the system, restricted to those processes in J, is:

$$UPROCS_J = |||_{j \in J} U_j$$

The kernel processes are:

$$KS_J = |||_{j \in J} j{:}\big(\big(K \underset{\{j.q\}}{\parallel} QUEUES \big) \setminus \alpha QUEUES \big) \tag{1}$$

Adding in the helper process gives the complete kernel:

$$\begin{aligned} KERNEL1_J &= \big(KS_J \underset{\{j.c \mid j \in J\}}{\parallel} SWITCH \big) \setminus \alpha SWITCH \\ KERNEL2_J &= \big(KERNEL1_J \underset{\{j.g \mid j \in J\}}{\parallel} TAGMGR \big) \\ &\quad \setminus \alpha TAGMGR \\ KERNEL_J &= \big(KERNEL2_J \underset{\{j.p \mid j \in J\}}{\parallel} PROCMGR0 \big) \\ &\quad \setminus \alpha PROCMGR0 \end{aligned}$$

Finally:

$$SYS_J = UPROCS_J \underset{\{j.s|j \in J\}}{\|} KERNEL_J$$

Of course, the whole system is captured simply by $SYS_{\mathcal{P}}$.

This assembly of kernel process makes extensive use of the CSP hiding operator ("\"). That is, the combined process SYS does not show direct evidence of internal state transitions such as: communications between any i:K and the switch; communications with the tag manager; communications with the process manager; etc. In fact the only events that remain visible are the workings of the user processes U_i and their system calls given by $i.s\underset{\kappa}{?}$ and $i.s\underset{\kappa}{!}$. By implication, kernels that implement the Flume model should hide the system's inner workings from unprivileged users (which is indeed the case for the Flume implementation). In practical terms, the CSP model for SYS shows what a non-root Unix user might see if examining his processes with the `strace` utility.

5 Non-Interference

A mature definition in the literature for models like Flume's is *non-interference*. Informally [16]:

> One group of users, using a certain set of commands, is *noninterfering* with another group of users if what the first group does with those commands has no effect on what the second group of users can see.

In Flume, this means that for an export-protection tag t and a process p running with $S_p = \{t\}$, a process q running with $S_q = \{\}$ should have an execution path that is entirely independent of p's. If p could somehow influence q, then it could reveal to q information tagged with t, violating the export-protection policy.

This chapter explores the non-interference properties of Flume's CSP model. Previous work by Ryan and Schneider [17] informs which definition of non-interference to apply (see Section 5.1). A proof that Flume fits the definition follows (see Section 5.3).

5.1 Definition

We use Ryan and Schneider's definition of non-interference [17], where process equivalence follows the *stable failures model* [18,19]. See Appendix .1 for more details. This definition considers all possible pairs of traces for S that vary only by elements in the high alphabet (i.e., they are equal when projected to low). For each pair of traces, two experiments are considered: advancing S over the elements in left trace, and advancing S over the elements in the right trace. The two resulting processes must look equivalent from a "low" perspective. Formally:

Definition 3 (Non-Interference for System S). *For a CSP process S, and an alphabet of low symbols $LO \subseteq \alpha S$, the predicate $\mathsf{NI}_{LO}(S)$ is true iff*

$$\forall tr, tr' \in traces(S) \ : \ tr \upharpoonright LO = tr' \upharpoonright LO \ \Rightarrow$$
$$\mathcal{SF}[\![S/tr]\!] \upharpoonright LO \ = \mathcal{SF}[\![S/tr']\!] \upharpoonright LO \ .$$

We say that the process S exhibits non-interference with respect to the low alphabet LO iff $\mathsf{NI}_{LO}(S)$ is true.

That is, after being advanced by tr and tr', the two processes must accept all of the same traces (projected to low) and refuse all of the same refusal sets (projected to low).

Stability and Divergence There are several complications. The first is the issue of whether or not the stable failures model is adequate. For instance, if a high process caused the kernel to diverge (i.e., *hang*), a low process could record such an occurrence on reboot, thereby leaking a bit (very slowly!) to low. By construction, the Flume kernel never diverges. One can check this property by examining each system call and verifying that only a finite number of internal events can occur before the process is ready to receive the next call. User-space process (e.g., U_i) can diverge, but they cannot observe each other's divergence, and so their divergence is inconsequential in our security analysis.

If divergence attacks were a practical concern, we could capture divergent behavior with the more general Failures, Divergences, Infinite Traces (FDI) model [19]. We conjecture that Flume's non-interference results under the stable failures model also hold in the FDI model, but the proof mechanics are yet more complicated.

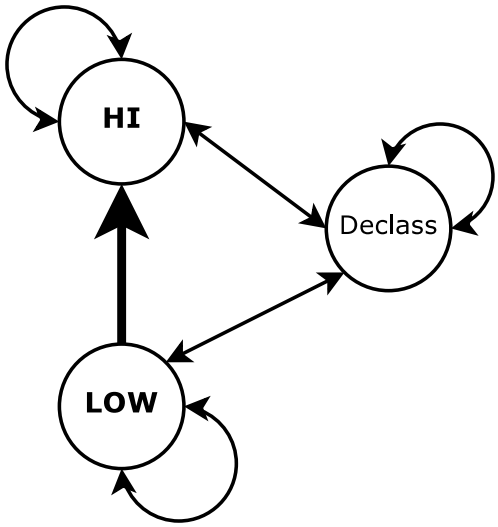

Fig. 5. Intransitive Non-interference. Arrows depict allowed influence. All influences are allowed *except* high to low.

Declassification The second complication is *declassification*, also known as *intransitive non-interference*. That is, the system should allow certain flows of information from "high" processes to "low" processes, if that flow traverses the appropriate declassifier. Figure 5 provides a pictorial representation: the system allows low processes and the declassifier to influence all other processes, and the high processes to influence other high processes and declassifiers but *not* to influence low processes. However, in the transitive closure, all processes can influence all other processes, which negates any desired security properties. Previous work assumed a global security policy, and modifies existing non-interference definitions to rule out flows not in the given policy [20].

In this paper, we simplify the problem. Consider an export-protection tag t, for which $t^+ \in \hat{O}$ and $t^- \notin \hat{O}$. All results then cover only those processes that cannot declassify t. Let N_t be the list of the these processes: $N_t = \{j \mid t^- \notin O_j\}$. Given N_t, we can define the high and low alphabets. The high symbols

emanate from a process $i \in N_t$ with $t \in S_i$:

$$\begin{aligned} HI_t \triangleq & \{i.b.(S, I, \dots) \mid i \in N_t \wedge t \in S\} \cup \\ & \{i.s.(S, I, \dots) \mid i \in N_t \wedge t \in S\} \end{aligned}$$

The low symbols are the complement set:

$$\begin{aligned} LO_t \triangleq & \{i.b.(S, I, \dots) \mid i \in N_t \wedge t \notin S\} \cup \\ & \{i.s.(S, I, \dots) \mid i \in N_t \wedge t \notin S\} \end{aligned}$$

5.2 Allocation of Global Identifiers

The model presented in Chapter 4 is almost fully-specified, with an important exception: the process *CHOOSE*:

$$CHOOSE_Y = ?(S, I) \rightarrow \sqcap_{y \in Y}(!y) \rightarrow STOP$$

The "internal (nondeterministic) choice" operator ($\sqcap$) implies that the model requires further *refinement*. The question becomes: how to allocate tags and process identifiers?

An idea that does not work is sequential allocation, yielding the tag (or process ID) sequence $\langle 1, 2, 3, \dots \rangle$. To attack this scheme, a low process forks, retrieving a child ID i. Then the high process forks k times, to communicate the value k to t. The next time low process forks, it gets process ID $i + k$, and by subtracting i recovers the high message. There are two problems: (1) low and high processes share the same process ID space; and (2) they can manipulate it in a predictable way.

The second weakness is exploitable even without the first. In a different attack, a high process communicates a "1" by allocating a tag via create_tag(Add), and communicates a "0" by refraining from allocating. If a low process could guess which tag was allocated (call it t), it could then attempt to change its label to $S = \{t\}$. If the change succeeds, then the low process had access to t^+ through $\hat{O}$, meaning the high process allocated the tag. If the change fails, the high process must not have allocated. The weakness here is that the low process "guessed" the tag t without the high process needing to communicate it. If such guesses were impossible (or very unlikely), the attack would fail.

Another idea—common to all DIFC kernels (c.f., Asbestos [1], HiStar [2] and the Flume implementation)—is random allocation from a large pool. The random allocation scheme addresses the second weakness—predictability—but not the first. That is, operations like process forking and tag creation always have globally observable side affects: a previously unallocated resource becomes claimed. Consider, as an example, this trace for the Flume system:

$$\begin{aligned} tr = \langle & i.b.(\{t\}, \{\}, \{\}, \{\}), \\ & i.s.(\{t\}, \{\}, \mathsf{fork}), \\ & j.b.(\{t\}, \{\}, \{\}, \{\}), \\ & i.s.(\{t\}, \{\}, j), \dots \rangle \end{aligned}$$

A new process i is born, with secrecy label $S_i = \{t\}$, and empty integrity and ownership. Thus, i's actions fall into the HI_t alphabet. Once i starts, it forks a new process, which the kernel randomly picks as j. The child j runs with secrecy $S_j = \{t\}$, inheriting its parent's secrecy label.

Projecting this trace onto the low alphabet yields the empty sequence ($tr \upharpoonright LO_t = \langle\rangle$). Thus, this trace should have no impact on the system from a low process k's perspective. Unfortunately, this is not the

case. Before tr occurred, l could have forked off process j, meaning:

$$\begin{aligned} tr' = \langle & k.b.(\{\},\{\},\{\},\{\}), \\ & k.s.(\{\},\{\},\mathsf{fork}), \\ & j.b.(\{\},\{\},\{\},\{\}), \\ & k.s.(\{\},\{\},j),\ldots\rangle \end{aligned}$$

was also a valid trace for the system. But after tr occurs, tr' is no longer possible, since the process j can only be born once. In other words, $tr \frown tr'$ is not a valid trace for the system but tr' is by itself. This contradicts the definition of non-interference in the stable failures model of process equivalence.

To summarize, we have argued that allocation of elements from $\hat{O}$ and $\hat{P}$, must obey two properties: (1) unpredictability and (2) partitioning. Our approach is to design a randomized allocation scheme that achieves both. Define parameters:

$$\begin{aligned} \alpha &\triangleq \log_2(\text{the number of tags}) \\ \beta &\triangleq \log_2(\text{maximum number of operations}) \\ \epsilon &\triangleq -\log_2(\text{acceptable failure probability}) \end{aligned}$$

A reasonable value for β is 80, meaning that no instance of the Flume system will attempt more that 2^{80} operations. Since tag allocation, forking and constructing labels count as operations, it the system expresses fewer than 2^{β} tags, process IDs, or labels in its lifetime. A reasonable value for ϵ is 100, meaning the system fails catastrophically at any moment with probability at most 2^{-100}.

Define a lookup table $s(\cdot)$, that given any label $L \subseteq \mathcal{T}$, $s(L)$ outputs a integer in $[0, 2^{\beta})$ that uniquely identifies L. The serialization can be predictable. Next consider the family of all *injective* functions:

$$G : (\{0,1\}^{\beta}, \{0,1\}^{\beta}, \{0,1\}^{\beta}) \rightarrow \{0,1\}^{\alpha}$$

The Flume system, upon startup, picks an element $g \in G$ at random. When called upon by a process with labels S, I to allocate a new tag or process ID, it returns $g(s(S), s(I), x)$, for some heretofore unused $x \in \{0,1\}^{\beta}$. The output is a tag in $\{0,1\}^{\alpha}$. Appendix .4 derives $\alpha \geq \max(\epsilon + 1, 3)$, meaning $\alpha = 240$ for our example parameters.

Thus, we let $\mathcal{T} = \mathcal{P} = \{0,1\}^{\alpha}$, for a sufficiently large α. The kernel picks $g \in G$ at random upon startup. Then *CHOOSE* is refined as:

$$CHOOSE_Y = ?(S,I) \ \rightarrow \ \sqcap_{y \in \mathcal{G}(S,I,Y)} (!y) \ \rightarrow \ STOP$$

where

$$\mathcal{G}(S,I,Y) = \left\{ t \mid x \in \mathcal{T} \wedge t = g(s(S), s(I), x) \wedge t \in Y \right\} \ .$$

Note that $\mathcal{G}(S,I,Y) \subseteq Y$, so the nature of the refinement is just to restrict the set of IDs that $CHOOSE_Y$ will ever output, based on the secrecy and integrity labels of the calling process.

5.3 Theorem and Proof

The main theorem is as follows:

Theorem 1 (Non-Interference in Flume). *For any security parameter ϵ, there exists an instantiation of CHOOSE such that: for any export-protection tag t, for any Flume instance SYS_{N_t}, $\Pr[\mathsf{NI}_{LO_t}(SYS_{N_t})] \geq 1 - \epsilon$.*

Note than many instances of SYS_{N_t} exist; they differ from one another based on their user portions ($UPROCS_{N_t}$). The theorem must hold for any such instance. Also, the theorem is probabilistic. There is a small chance that an instance of SYS_{N_t} can guess the output of *CHOOSE*, and in that case, the non-interference property fails. We argue such a guess succeeds with arbitrarily small probability.

The proof is by induction over the number of low symbols in the two traces, tr and tr'. (Recall that tr and tr' are equivalent when projected to the low alphabet). For the base case, tr and tr' have no low symbols, and therefore have no high symbols, since the system accepts only low symbols in its initial state. Since $tr = tr' = \{\}$, the theorem follows trivially.

We prove the inductive step casewise, considering each system call and whether the kernel is looking to accept a new system call or reply to an outstanding call. Most cases reason about the causal relationships among events in the trace. A more involved case is change_label, which must consider the unlikely case that a low process guessed which tags a high process received from the kernel when calling create_tag. See Appendix .5 for details.

6 Discussion

To review, we have described the Flume kernel both informally and with CSP formalism and proven that the CSP model upholds a definition of non-interference. In this section, we discuss the implications of these results, and how they can be translated to a practical system.

6.1 Refinement

Due to the well-known *refinement paradox*, the Flume model might satisfy non-interference, but implementations (i.e. *refinements*[6]) of it might not [21]. To circumvent this paradox, Lowe has recently strengthened notions of non-interference, requiring that a system like *SYS* and all of its refinements exhibit non-interference [22]. Other work has suggested restricting refinement to a set of operators known to preserve security guarantees [23]. We follow Lowe's approach as best as possible, arguing that non-interference holds for most refinements.

The parts of the Flume model that need refinement are those that display non-determinism via the $\sqcap$ or $\Box$ operators: (1) the *CHOOSE* process; (2) the user processes U_i; (3) "scheduling"; and (4) the *tock* events in timed CSP. As for (1), the proof in Section .5 holds for some refinements of *CHOOSE*, such as the random function specified in Section 5.2 and a more practical hash-based approach considered below. An Flume implementation should refine *CHOOSE* as specified, or with a method known to preserve non-interference. As for (2), the proof in Section .5 holds for *arbitrary* refinements of user processes U_is, as long as they communicate only through the designated system calls (see Section 4.1). In practice, we cannot hope to isolate the U_is completely from one another: they can communicate by manipulating shared hardware resources (e.g., disks, CPU cache, CPU cycles, network bandwidth, etc.)

As for (3), the Flume model hides scheduling for simplicity: any interleaving of processes is admissible, by Equation 1 in Section 4.7. However, a practical refinement of Flume would implement "fairness" restrictions on scheduling, disallowing one process from consuming more than its "fair" share of resources. Scheduling refinements, in and of themselves, do not affect the proof of security: non-interference holds in any reordering of high and low processes, as long as all processes get to run eventually, and do not have any comprehension of time. As for (4), Appendix .6 explores extending the Flume model to show an explicit passage of time: the event *tock* denotes one clock tick, and all parts of the model must stand aside and yield to *tock*. We conjecture that if $tock \in HI_t$, then the proof holds for all refinements of the scheduler,

[6] Q is a refinement of P iff $\mathcal{SF}[\![Q]\!] \subseteq \mathcal{SF}[\![P]\!]$.

and that if $tock \in LO_t$, that the proof holds for only some refinements. Further exploration is deferred to future work.

In sum, we believe the Flume model to maintain non-interference under important refinements—tag allocation, arbitrary user processes, and scheduling—if timing channels are excluded. Of course, this is a far cry from proving non-interference in a working implementation, but a significant improvement over the status quo. Systems such as Asbestos and IX have gaping covert channels baked into their very specifications, so *any* refinements of those systems are insecure. By contrast, an OS developer has a fighting chance to realize a secure Flume implementation.

6.2 Kernel Organization

The Flume DIFC model is a "monolithic" kernel design, in which the kernel is a hidden black box, and user-level processes have a well-specified system call interface. Some modern approaches to kernel design (e.g. the Exokernel [24] and the Infokernel [25]) expose more of the inner workings of the kernel to give application developers more flexibility. However, such transparency in an information-flow control setting can leak information: imagine a high process issuing create_tag, and a low process observing *TAGMGR*'s transitions. The simplest way to work around this problem is to conceal the inner workings of the kernel (as we have done). Another, more complicated solution, is to model more parallelism inside the kernel, so that the tag manager can serve both high and low concurrently.

The Flume model captures most of the kernel processes—like the $i{:}K$, the tag manager, and the process manager—as single-threaded processes. For instance, if the tag manager is responding to a request for $i.g.(\mathsf{create\_tag}, w)$, it cannot service $j.g.(\mathsf{create\_tag}, w)$ until it replies to $i.g.(\mathsf{create\_tag}, w)$. In practical implementations of this CSP model, such serialization might be a bottleneck for performance. More parallelism internal to the kernel is possible, but would require explicit synchronization through locks, and more complexity overall.

6.3 Tag Allocations

The use of a truly random function for *CHOOSE* is impractical, as are tag sizes of 240 bits. In practice, a weaker cryptographic primitive suffices, such as a Message Authentication Code (MAC) [26] with a collision-resistant hash function [27]. Let M be such a MAC of suitable input length. The kernel picks a random secret key k on startup, and then computes new tags and process IDs using $M_k(S, I, x)$ for a counter variable x. This construction approximates both important properties. The unforgability property of the MAC implies that an adversary cannot find $(S, I, x) \neq (S', I', x')$ such that $\mathrm{HMAC}_k(S, I, x) = \mathrm{HMAC}_k(S', I', x')$, so a high process with secrecy $\{t\}$ and a low process with secrecy $\{\}$ will not get the same tag. Similarly, user processes cannot predict the output of $\mathrm{HMAC}_k(S, I, x)$ without knowing k.

6.4 Integrity

Though we have focused on secrecy, the same arguments hold for integrity. Analogously, one would pick an *integrity-protection* tag t, one for which $t^- \in \hat{O}$ and $t^+ \notin \hat{O}$. The low symbols are those whose integrity tags contain t, and the high symbols are those that do not. The same proof shows that the high events do not interfere with the low.

7 Related Work

Information flow control (IFC) at the operating system level dates back to the centralized military systems of the 70s, and 80s [28,14,29]. Several systems like IX [10] and SELinux [30] integrated information-flow ideas with Unix-like operating systems in 90s. Denning first pointed out that dynamically-adjusted

security labels could leak data [11] and suggested instead static checking, which later found fruition as type-analysis [31]. *Decentralized* declassification and endorsement proved a key relaxation, making IFC practical for language-based systems [4], and eventual spurring a revitalization of the idea in operating systems and web-serving settings with the Asbestos [1], HiStar [2] and Flume [3] systems. HiStar introduced the idea of "self-tainting," solving the wide covert channel described in Section 3. Flume later adopted a similar strategy, but within a streamlined label system.

Taint-tracking is another technique for tracking information flow through legacy software written in arbitrary languages [32,33]. Such systems run a target application as rewritten binary, without the cooperation or recognition of the application in question, meaning they must infer label changes. Therefore taint-tracking systems are susceptible to the covert channel attacks described in Section 3, and cannot uphold non-interference.

Goguen and Meseguer introduced the idea of non-interference for security protocols [16], while Volpano et al. first showed that type systems could provably uphold the idea [34]. More recently, Zheng and Myers [35] and also Tse and Zdancewic [36] proved that statically-typed systems with runtime principles could still obey non-interference. Relative to Flume, the advantage of these approaches are first that information flow is monitored at a finer granularity, and second that the compiler ensures the model and refinement are one in the same (and hence the refinement paradox does not apply). On the other hand, the Flume model provides more flexibility as to how the various user processes U_i behave: it restricts these processes from accessing all but one communication channel, but otherwise they can act arbitrarily and need not be type-checked.

The proofs offered here are manual. In future work, we hope to investigate shaping the Flume model used in the proof to a form that is amenable to automated analysis. Lowe first used an automated checker to break protocol previous assumed secure [37]. Ryan and Schneider also describe how the FDR automated checker can verify standard security protocols [38],

8 Conclusion

This paper presented the first formal security argument for a DIFC-based operating system. It modelled Flume using the CSP formalism, and proved that the model fulfills non-interference—an end-to-end property that protects secrecy and integrity even against subtle covert channel attacks. The model and proof are not substantially weakened by the refinement paradox, since the proof holds for many refinements of the model. Future work calls for further investigation of timing-based covert and side channels, and automation of the proof techniques, for both the model and its implementation.

References

1. P. Efstathopoulos, M. Krohn, S. VanDeBogart, C. Frey, D. Ziegler, E. Kohler, D. Mazières, F. Kaashoek, and R. Morris, "Labels and event processes in the Asbestos operating system," in *Proceedings of 20th ACM Symposium on Operating Systems Principles*, Brighton, UK, October 2005.
2. N. B. Zeldovich, S. Boyd-Wickizer, E. Kohler, and D. Mazières, "Making information flow explicit in HiStar," in *Proceedings of 5th Symposium on Operating Systems Design and Implementation (OSDI)*, Seattle, WA, Nov. 2006.
3. M. Krohn, A. Yip, M. Brodsky, N. Cliffer, M. F. Kaashoek, E. Kohler, and R. Morris, "Information flow control for standard OS abstractions," in *Proceedings of the 21st Symposium on Operating Systems Principles (SOSP)*, Stevenson, WA, October 2007.
4. A. C. Myers and B. Liskov, "A decentralized model for information flow control," in *Proceedings of16th ACM Symposium on Operating Systems Principles*, Saint-Malô, France, Oct. 1997, pp. 129–142.
5. J. Fielding, "UN website is defaced via SQL injection," *Tech Republic*, Aug. 2007, `http://blogs.techrepublic.com.com/networking/?p=312`.

6. News10, "Hacker accesses thousands of personal data files at CSU Chico," Mar. 2005, `http://www.news10.net/display_story.aspx?storyid=9784`.
7. K. Poulsen, "Car shoppers' credit details exposed in bulk," *SecurityFocus*, September 2003, `http://www.securityfocus.com/news/7067`.
8. ——, "FTC investigates petco.com security hole," *SecurityFocus*, December 2003, `http://www.securityfocus.com/news/7581`.
9. R. Trounson, "Major breach of UCLA's computer files," *Los Angeles Times*, Dec. 12 2006.
10. M. D. McIlroy and J. A. Reeds, "Multilevel security in the UNIX tradition," *Software—Practice and Experience*, vol. 22, no. 8, pp. 673–694, 1992.
11. D. E. Denning, "A lattice model of secure information flow," *Communications of the ACM*, vol. 19, no. 5, pp. 236–243, May 1976.
12. T. Fraser, "LOMAC: Low water-mark integrity protection for COTS environments," in *Proceedings of2000 IEEE Symposium on Security and Privacy*, Oakland, CA, May 2000, pp. 230–245.
13. D. E. Bell and L. J. LaPadula, "Secure computer system: Unified exposition and multics interpretation," MITRE Corparation, Bedford, MA, Tech. Rep. MTR-2997, Rev. 1, March 1976.
14. K. J. Biba, "Integrity considerations for secure computer systems," MITRE Corp., Bedford, MA, Tech. Rep. MTR-3153, Rev. 1, 1976.
15. C. A. R. Hoare, *Communicating Sequential Processes*. Englewood Cliffs, New Jersey: Prentice/Hall International, 1985.
16. J. A. Goguen and J. Meseguer, "Security policies and security models," in *IEEE Symposium on Research in Security and Privacy*, 1982.
17. P. A. Ryan and S. A. Schneider, "Process algebra and non-interference," *Journal of Computer Security*, no. 9, pp. 75–103, 2001.
18. A. W. Roscoe, *A Theory and Practice of Concurrency*. London, UK: Prentice Hall, 1998.
19. S. Schneider, *Concurrent and Real-Time Systems: The CSP Approach*. Chichester, UK: John Wiley & Sons, LTD, 2000.
20. A. Roscoe and M. H. Goldsmith, "What is intransitive noninterference?" in *PCSFW: Proceedings of The 12th Computer Security Foundations Workshop*. IEEE Computer Society Press, 1999.
21. J. Jacob, "On the derivation of secure components," in *Proceedings of the IEEE Symposium on Security and Privacy*, Oakland, CA, 1989.
22. G. Lowe, "On information flow and refinement-closure," in *Proceedings of the Workshop on Issues in the Theory of Security (WITS07)*, 2007.
23. H. Mantel, "Preserving Information Flow Properties under Refinement," in *Proceedings of 2001 IEEE Symposium on Security and Privacy*, May 2001.
24. D. R. Engler, M. F. Kaashoek, and J. O'Toole, "Exokernel: An operating system architecture for application-level resource management," in *Proceedings of15th ACM Symposium on Operating Systems Principles*, Copper Mountain Resort, Colorado, Dec. 1995, pp. 251–266.
25. A. C. Arpaci-Dusseau, R. H. Arpaci-Dusseau, N. C. Burnett, T. E. Denehy, T. J. Engle, H. S. Gunawi, J. A. Nugent, and F. I. Popovici, "Transforming policies into mechanisms with infokernel," in *Proceedings of19th ACM Symposium on Operating Systems Principles*, Bolton Landing, Lake George, New York, Oct. 2003, pp. 90–105.
26. M. Bellare, R. Canetti, and H. Krawczyk, "Keying hash functions for message authentication," in *Proceedings of the 16th Annual International Cryptology Conference on Advances in Cryptology (CRYPTO)*, Aug. 1996, pp. 1–15.
27. FIPS 180-2, *Secure Hash Standard*, U.S. Department of Commerce/N.I.S.T., National Technical Information Service, Springfield, VA, Aug. 2002.
28. D. E. Bell and L. J. LaPadula, "Secure computer systems: Mathematical foundations," MITRE Corparation, Bedford, MA, Tech. Rep. Technical Report 2547, Volume I, Mar. 1973.
29. *Trusted Computer System Evaluation Criteria (Orange Book)*, Dod 5200.28-std ed., Department of Defense, December 1985.
30. P. Loscocco and S. Smalley, "Integrating flexible support for security policies into the Linux operating system," in *Proceedings of 2001 USENIX Annual Technical Conference*, San Diego, CA, June 2001, fREENIX track.
31. J. Palsberg and P. Ørbæk, "Trust in the λ-Calculus," in *Proceedings of2nd International Symposium on Static Analysis*, Sept. 1995.

32. J. Newsome and D. Song, "Dynamic taint analysis for automatic detection, analysis, and signature generation of exploits on commodity software," in *Network and Distributed System Security Symposium (NDSS)*, Feb. 2005.
33. S. McCamant and M. Ernst, "Quantitative Information Flow as Network Flow Capacity," in *Proceedings ofACM SIGPLAN Conference on Programming Language Design and Implementation*, 2008.
34. D. Volpano, C. Irvine, and G. Smith, "A sound type system for secure flow analysis," vol. 4, no. 2-3, 1996, pp. 167–187.
35. L. Zheng and A. C. Myers, "Dynamic security labels and noninterference," in *Proceedings of 2nd International Workshop on Formal Aspects in Security and Trust (FAST)*, Aug. 2004.
36. S. Tse and S. Zdancewic, "Run-time principals in information-flow type systems," *ACM Transactions on Programming Language Systems*, vol. 30, no. 1, 2007.
37. G. Lowe, "Some new attacks upon security protocols," in *Proceedings of9th IEEE Computer Security Foundations Workshop*, 1996.
38. P. Ryan and S. Schneider, *The Modelling and Analysis of Security Protocols: The Csp Approach.* Addison-Wesley Professional, 2001.
39. J. Ouaknine, "A framework for model-checking timed CSP," Oxford University, Tech. Rep., 1999.

.1 Review of CSP

Communicating Sequential Processes (CSP) is a process algebra useful in specifying systems as a set of parallel state machines that sometimes synchronize on events. We offer a brief review of it here, borrowing heavily from Hoare's book [15].

CSP Processes Among the most basic CSP examples is Hoare's vending machine:

$$VMS = in25 \rightarrow choc \rightarrow VMS$$

This vending machine waits for the event *in25*, which corresponds to the input of a quarter into the machine. Next, it accepts the event *choc*, which corresponds to a chocolate falling out of the machine. Then it returns to the original state, with a recursive call to itself. The basic operator at use here is the *prefix* operator. If x is an event, and P is a process, then $(x \rightarrow P)$, pronounced "x then P," represents a process that engages in event x and then behaves like process P. For a process P, the notation αP describes the "alphabet" of P. It is a set of all of the events that P is ever willing to engage in. For example, $\alpha VMS = \{in25, choc\}$.

For any CSP process P, we can discuss a *trace* of events that P may accept. For the *VMS* example, various traces include:

$$\langle\rangle$$
$$\langle in25\rangle$$
$$\langle in25, choc\rangle$$
$$\langle in25, choc, in25, choc, in25\rangle$$

The next important operator is "choice," denoted by "|". If x and y are distinct events, then:

$$(x \rightarrow P \mid y \rightarrow Q)$$

denotes a process that accepts x and then behaves like P or accepts y and then behaves like Q. For example, a new vending machine can accept either a coin and output a chocolate, or accept a bill and output an ice cream cone:

$$VMS2 = (bill \rightarrow cone \rightarrow VMS2 \mid in25 \rightarrow choc \rightarrow VMS2)$$

CSP offers a more general choice function (for choosing between many inputs succinctly), but the Flume model only requires simple choice.

A related operator is *internal (nondeterministic) choice*, is denoted "$\sqcap$". In simple choice, the machine reacts exactly to events it fields from the machine's user. In nondeterministic choice, the machine behaves unpredictably from the perspective of the user, maybe because the machine's description is underspecified, or maybe because the machine is picking from a random number generator. For instance, a change machine might return coins in any order, depending on how the machine was last serviced:

$$\begin{aligned} CHNG = (in25 \rightarrow (&out10 \rightarrow out10 \rightarrow out5 \rightarrow CHNG \sqcap \\ &out10 \rightarrow out5 \rightarrow out10 \rightarrow CHNG)) \end{aligned}$$

That is, the machine takes as input a quarter, and returns two dimes and a nickel in one of two orderings. Another standard operator, "external choice" denoted "$\Box$", has different semantics but does not appear in Flume's model.

CSP provides useful predefined processes like *STOP*, the process that accepts no events, and *SKIP*, the process that shows a successful termination and then behaves like *STOP*. Other processes like *DIV*, *RUN* and *CHAOS* are standard in the literature, but are not required here.

The next class of operators relate to parallelism. The notation:

$$P \underset{A}{\parallel} Q$$

denotes P running in parallel with Q, synchronizing on events in A.[7] This means a stream of incoming events can be arbitrarily assigned to either P or Q, assuming those events are not in A. However, for events in A, both P and Q must accept them in synchrony. As an example, consider the vending machine and the change machine running in parallel, synchronizing on the event *in25*:

$$FREELUNCH = VMS \underset{\{in25\}}{\parallel} CHNG$$

Possible traces for this new process are the various interleavings of the traces for the two component machines that agree on the event *in25*. For instance:

$$\begin{aligned} &\langle in25, choc, out10, out10, out5, \ldots\rangle \\ &\langle in25, out10, choc, out10, out5, \ldots\rangle \\ &\langle in25, out10, out10, choc, out5, \ldots\rangle \\ &\langle in25, out10, out10, out5, choc, \ldots\rangle \\ &\langle in25, choc, out10, out5, out10, \ldots\rangle \\ &\langle in25, out10, choc, out5, out10, \ldots\rangle \\ &\langle in25, out10, out5, choc, out10, \ldots\rangle \\ &\langle in25, out10, out5, out10, choc, \ldots\rangle \end{aligned}$$

are possible execution paths for *FREELUNCH*.

Another variation on parallel composition is arbitrary interleaving, denoted: $P \,|||\, Q$. In interleaving, P and Q never synchronize, operating independently of one another. $P \,|||\, Q$ is therefore equivalent to $P \parallel_{\{\}} Q$, which means P and Q run in parallel and synchronize on the empty set.

Processes that run in parallel can communicate with one another over *channels*. A typical channel c can carry various values v, denoted $c.v$.[8]. This is represented as the sending process accepting the

[7] Parallelism differs between Hoare's original CSP formulation and more modern formulations, like Schneider's. We use Schneider's "interface parallelism" in this model.

[8] If the channel has a compound name like $i.c$, its values are respectively denoted $i.c.v$. Channel names are prefix-free so this is never ambigous

event $c!v$ while the receiving process accepts the event $c?x$ (where x is thus far unbound) and sets x to v. Communication on a channel is possible only when the sender and receiver processes are in the respective states simultanoeusly. If one process is at the suitable state and the other is not, the ready process waits until its partner becomes ready. In a slight deviation from Hoare's semantics, channels here are bidirectional: messages can travel independently in either direction across a channel. The Flume model uses channels extensively.

The next important CSP feature is *concealment* or *hiding*. For a process P and a set of symbols C, the process $P\backslash C$ is P with symbols in C hidden or concealed. The events in C become internal transitions, that cannot be observed by other processes through synchronization or channel communication. Concealment can induce *divergence*—an infinite sequence of internal transitions. For instance, the process $P = (c \rightarrow P)\backslash\{c\}$ diverges immediately, never to be useful again. The use of concealment in the Flume model is careful never to induce divergence in this manner.

Concealment enables *subroutines* (or *subordination*, in Hoare's terminology). For two process P and Q whose alphabets fulfill $\alpha P \subseteq \alpha Q$, the new process $P \mathbin{/\!\!/} Q$ is defined as $(P \parallel Q)\backslash\alpha P$. This means that the subroutine P is available within Q, but not visible to the outside world. The notation $p{:}P \mathbin{/\!\!/} Q$ means a particular instance p of the subroutine P is available in Q. Then an event such as $p!x?y$ within Q means that Q is calling subroutine p with argument x, and that the return value is placed into y. Within P, the event $?x$ means receive the argument x from the caller, and the event $!y$ means return the result y to the caller.

A final important language feature is *renaming*. Given a "template" process P, the notation $i{:}P$ means a "renaming" of P with all events prefixed by i. That is, if the event $c!v$ appears in P, then the event $i.c!v$ appears in $i{:}P$, where $i.c$ is the channel c that has been renamed to $i.c$. Thus, for any $i \neq j$, the alphabets of $i{:}P$ and $j{:}P$ are disjoint: $\alpha(i{:}P) \cap \alpha(j{:}P) = \{\}$. This concludes our whirlwind tour of CSP features. We refer the reader to Hoare's [15], Schneider's [19] and Roscoe's [18] books for many more details.

The Stable Failures Model We now expand upon the idea of traces to develop an idea of process equivalence in CSP. The traces of P (denoted $\mathit{traces}(P)$) is the set of all traces accepted by the process P. The refusals of P (denoted $\mathit{refusals}(P)$) is a set of sets. A set X is in $\mathit{refusals}(P)$ if and only if P deadlocks when offered any event from X. For instance, consider the process P_0:

$$P_0 = (a \rightarrow \mathit{STOP} \sqcap b \rightarrow \mathit{STOP}) \quad .$$

We write that $\mathit{refusals}(P_0) = \{\{a\}, \{b\}\}$. That is, P_0 can nondeterministically choose the left branch, in which case it will only accept $\{a\}$ and will refuse $\{b\}$. On the other hand, if it nondeterministically chooses the left branch, it will accept $\{b\}$ and refuse $\{a\}$. Thus, due to nondeterminism, we write $\mathit{refusals}(P)$ as above, and *not* as the flattened union $\{a, b\}$.

The notation $Q \downarrow$ is a predicate that denotes the process Q is "stable." Unstable states are those that transition internally, or those that diverge. For example, consider the process P_0 begins at an unstable state, since it can make progress in either the left or right direction without accepting any input. However, once it makes its first internal transition, arriving at either $a \rightarrow \mathit{STOP}$ or $b \rightarrow \mathit{STOP}$, it becomes stable. A process that diverges, such as $(c \rightarrow P)\backslash\{c\}$, has no stable states. Conversely, stable states are those that can make no internal progress.

Though the CSP literature explores many notions of process equivalence, this paper uses the "stable failures" model, given in Hoare's book [15] and rephrased by Schneider [19] and Roscoe [18]. For a process P, the *stable failures* of P, written $\mathcal{SF}[\![P]\!]$, are defined as:

$$\mathcal{SF}[\![P]\!] = \{(s, X) \mid s \in \mathit{traces}(P) \ \wedge \ P/s \downarrow \ \wedge \\ X \in \mathit{refusals}(P/s)\}.$$

In other words, the failures of P captures which traces P accepts, and which sets it refuses after accepting those traces. For example:

$$\mathcal{SF}[\![P_0]\!] = \{(\langle\rangle, \{a\}), (\langle\rangle, \{b\}), (\langle a\rangle, \{a, b\}), (\langle b\rangle, \{a, b\})\} \quad .$$

In the stable failures model, two processes P and Q are deemed equivalent if and only if $\mathcal{SF}[\![P]\!] = \mathcal{SF}[\![Q]\!]$.

Lastly, CSP offers a way to identify processes in states other than their initial states: the process P/tr is P advanced to the state after the trace tr has occurred. Next, we often talk about the effects of "purging" certain events from traces and process states. The operator "$\restriction$" denotes *projection*. The trace $tr \restriction A$ is the trace tr projected onto the set A, meaning all events not in A are removed. For instance, if $A = \{a\}$, and $tr = \langle a, b, c, d, a, b, c\rangle$, then $tr \restriction A = \langle b, c, d, b, c\rangle$. For a set C, the set $C \restriction A$ is simply the intersection of the two. Define two projected processes $P \restriction A$ and $Q \restriction A$ equivalent if and only if $\mathcal{SF}[\![P]\!] \restriction A = \mathcal{SF}[\![Q]\!] \restriction A$, where:

$$\mathcal{SF}[\![P]\!] \restriction A = \{(tr \restriction A, X \cap A) \mid (tr, X) \in \mathcal{SF}[\![P]\!]\}$$

and similarly for Q.

.2 Declassifiers in Flume

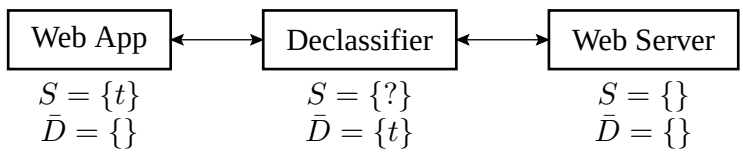

Fig. 6. An example web system; what should the declassifier's label be?

For example, consider Figure 6. Here, a web application runs with secrecy $S = \{t\}$, meaning it can read and compute on data tagged with secrecy t. Since the application is unprivileged ($\bar{D} = \{\}$), it cannot export this data on its own; it relies on a privileged declassifier ($\bar{D} = \{t\}$) to do so. If the declassifier decides to declassify, it sends the data out to the network via the web server, which runs with an empty secrecy label. Thus, data can flow in this example from high ($S = \{t\}$) to low ($S = \{\}$) with the help of the declassifier in the middle.

The question becomes, what should the declassifier's secrecy label be? One idea is for the declassifier to explicitly switch between $S = \{\}$ and $S = \{t\}$, as it communicates with its two partners. Though this solution sometimes works for *sending*, it is impractical for asynchronous *receives*: the declassifier has no way of knowing when each of its partners will send, and therefore cannot make the necessary label change ahead of time. Another idea is that the declassifier run with $S = \{t\}$ and only lower its label to $S = \{\}$ when it sends to its right. But this approach does not generalize—imagine a similar scenario in which the web server runs with secrecy $S = \{u\}$ and the declassifier has dual privileges for both t and u. This second approach is also ungainly for multithreaded declassifiers with blocking I/O operations.

.3 Per-process Queues (*QUEUES*)

Each kernel process $i{:}K$ needs it own set of queues, to handle messages received asynchronously from other processes. For convenience, we package up all of the queues in a single process $i{:}QUEUES$, which $i{:}K$ can access in all of its various states. The channel q serves communication between the queues and the kernel

process. The building block of this process is a single $QUEUE$ process, similar to that defined in Hoare's book. This process is parameterized by the value stored in the queue, and of course the queue starts out empty:

$$QUEUE = QUEUE_{<>}$$

From here, we define state transitions:

$$\begin{aligned}
QUEUE_{\langle\rangle} \;=\; & \big(?(\mathsf{enqueue}, x) \;\rightarrow\; QUEUE_{\langle x\rangle} \mid \\
& \quad ?(\mathsf{select}, j)!\{\} \;\rightarrow\; QUEUE_{\langle\rangle}\big) \\
QUEUE_{\langle x\rangle \frown s} \;=\; & \big(?(\mathsf{enqueue}, y) \;\rightarrow\; \mathsf{if}\ \#s + 1 < N_Q \\
& \qquad\qquad \mathsf{then}\ QUEUE_{\langle x\rangle \frown s \frown \langle y\rangle} \\
& \qquad\qquad \mathsf{else}\ QUEUE_{\langle x\rangle \frown s} \mid \\
& \quad ?(\mathsf{dequeue})!x \;\rightarrow\; QUEUE_s \mid \\
& \quad ?(\mathsf{select}, j)!\{j\} \;\rightarrow\; QUEUE_{\langle x\rangle \frown s}\big)
\end{aligned}$$

Note that these queues are bounded beneath N_Q elements. Attempts to enqueue messages on filled queues result in dropped messages. The model combines many $QUEUE$ subprocesses into a collection processes called $QUEUESET$:

$$QUEUESET = \|_{i \in \mathcal{P}}\ i{:}QUEUE$$

The process called $QUEUES$ communicates with kernel processes. Recall that $i.q$ is the channel shared between $i{:}K$ and $i{:}QUEUES$:

$$\begin{aligned}
QUEUES = s\ :\ & QUEUESET \ /\!/\ sel\ :\ QSELECT_s \ /\!/\ \mu X \bullet \\
& \big(q?(\mathsf{enqueue}, j, m) \;\rightarrow\; s.j!(\mathsf{enqueue}, m) \\
& \quad \rightarrow\; X \mid \\
& \quad q?(\mathsf{dequeue}, j) \;\rightarrow\; s.j!(\mathsf{dequeue})?m \\
& \quad \rightarrow\; q!m \;\rightarrow\; X \mid \\
& \quad q?(\mathsf{select}, Y) \;\rightarrow\; sel!Y?Z \;\rightarrow\; q!Z \;\rightarrow\; X \mid \\
& \quad q?(\mathsf{clear}) \;\rightarrow\; QUEUES\big)
\end{aligned}$$

Finally, the point of $QSELECT$ is to determine which of the supplied queues have pending messages. This process uses tail recursion to add to the variable Z as readied queues are found.

$$\begin{aligned}
QSELECT_s = & Z\ :\ VAR \ /\!/\ ?Y \;\rightarrow \\
& \quad Z := \{\}\,; \\
& \quad \big(\mu X \bullet (\mathbf{if}\ Y = \{\} \\
& \qquad\qquad \mathbf{then}\ (!Z \;\rightarrow\; QSELECT_s) \\
& \qquad\qquad \mathbf{else}\ \mathbf{pick}\ j \in Y\,; \\
& \qquad\qquad\quad Y := Y - \{j\}\,; \\
& \qquad\qquad\quad (s.j!(\mathsf{select}, j) \rightarrow s.j?A \;\rightarrow \\
& \qquad\qquad\quad (Z := Z \cup A\,;\ X))\big)
\end{aligned}$$

.4 Determining Tag Sizes

We can solve for how big α must be in terms of β and ϵ. *Partitioning* requires that functions from G must be injective, giving $2^{3\beta} \leq 2^{\alpha}$, or equivalently, $\alpha \geq 3\beta$. As for *unpredictability*, g is chosen randomly from G, so it will output elements in $\{0,1\}^{\alpha}$ in random order. After 2^{β} calls, g outputs elements from a set sized $2^{\alpha} - 2^{\beta}$ at random. Since $\alpha \geq 3\beta$, this "restricted" range for g still has well in excess of $2^{\alpha-1}$ elements. Failure occurs when a process can predict the output of g, which happens with probability no greater than $2^{\alpha-1}$. Thus, $\alpha - 1 \geq \epsilon$. Combining these two restrictions, $\alpha \geq \max(\epsilon + 1, 3\beta)$. For $\beta = 80$ and $\epsilon = 100$ we get $\alpha = 240$.

.5 Proof

Alphabets The relevant high and low alphabets were defined in Section 5.1. The rest of the events in the Flume model (like communication through the switch, to the process or tag manager, etc.) are all hidden by the CSP-hiding operators, as given in Section 4.7. For convenience, we define the set of events that correspond to kernel process i's incoming system calls, and a set of event that correspond to process i's responses:

$$\begin{aligned} C_i \triangleq & \big\{ i.s.(S, I, \mathsf{create\_tag}, w) \mid \\ & \quad S, I \subset \mathcal{T} \wedge w \in \{\mathsf{Add}, \mathsf{Remove}, \mathsf{None}\} \big\} \cup \\ & \big\{ i.s.(S, I, \mathsf{change\_label}, w, L) \mid \\ & \quad S, I, L \subset \mathcal{T} \wedge w \in \{\mathsf{Integrity}, \mathsf{Secrecy}\} \big\} \cup \\ & \big\{ i.s.(S, I, \mathsf{get\_label}, w) \mid \\ & \quad w \in \{\mathsf{Integrity}, \mathsf{Secrecy}\} \big\} \cup \cdots \end{aligned}$$

and so on for all system calls. Similarly for return values from system calls:

$$\begin{aligned} R_i \triangleq & \big\{ i.s.(S, I, t) \mid S, I \subset \mathcal{T} \wedge t \in \mathcal{T} \big\} \cup \\ & \big\{ i.s.(S, I, r) \mid S, I \subset \mathcal{T} \wedge r \in \{\mathsf{Ok}, \mathsf{Error}\} \big\} \cup \\ & \big\{ i.s.(S, I, L) \mid S, I, L \subset \mathcal{T} \big\} \cup \\ & \big\{ i.s.(S, I, O) \mid S, I \subset \mathcal{T} \wedge O \subset \mathcal{O} \big\} \cup \\ & \big\{ i.s.(S, I, p) \mid S, I \subset \mathcal{T} \wedge p \in \mathcal{P} \big\} \end{aligned}$$

The only visible events for process $i{:}K$ are system calls, and system call replies:

$$\alpha(i{:}K) = C_i \cup R_i$$

A final notational convenience: we often describe the failures of a process P projected onto the low alphabet LO_t and abbreviate it:

$$\mathcal{L}_t[\![P]\!] \triangleq \mathcal{SF}[\![P]\!] \upharpoonright LO_t$$

proof Consider any two traces tr and tr' such that $tr \approx_{LO_t} tr'$. The proof technique is induction over the length of the traces tr and tr'. We invent a new function $\lambda(\cdot)$

$$\lambda(tr) \triangleq \#(tr \upharpoonright LO_t)$$

that outputs the number of low events in a trace. Because $tr \approx tr'$, it follows that $\lambda(tr) = \lambda(tr')$. We first show the theorem holds for all traces tr and tr' such that $\lambda(tr) = \lambda(tr') = 0$. We then assume it holds for all traces with $\lambda(tr) = \lambda(tr') = k - 1$ and prove it holds for all traces with $\lambda(tr) = \lambda(tr') = k$.

Base Case For the base case, consider all $tr, tr' \in traces(SYS_{N_t})$ such that $\lambda(tr) = \lambda(tr') = 0$. In other words, $tr, tr' \in HI_t^*$.

At the system startup (SYS_{N_t} after no transitions), all of the kernel process $i{:}K$ are waiting on a message of the form $i.b$ before they spring to life. Until such a message arrives, $i{:}K$ will refuse all events C_i and R_i. The one exception is the process init, which is already waiting to accept incoming system calls when the system starts. By construction $S_{\mathsf{init}} = \{\}$ and $I_{\mathsf{init}} = \mathcal{T}$. Since $t \notin S_{\mathsf{init}}$, $C_{\mathsf{init}} \cup R_{\mathsf{init}} \subseteq LO_t$. Therefore, the system refuses all high events at startup, and $tr = \langle\rangle$ is the only trace of SYS_{N_t} without low symbols (and for which $\lambda(tr) = 0$). For $tr = tr' = \langle\rangle$, the lemma trivially holds.

Inductive Step For the inductive step, assume the lemma holds for all traces tr, tr' of SYS_{N_t} such that $tr \approx_{LO_t} tr'$ and also $\lambda(tr) = \lambda(tr') = k - 1$. Now, we seek to show the lemma holds for all equivalent traces with one more low symbol.

Given an arbitrary trace $tr \in traces(SYS_{N_t})$ such that $\lambda(tr) = k$, write tr in the form $tr = p \frown l \frown h$, where p is prefix of tr, $l \in LO_t$ is a single low event, and $h \in HI_t^*$ are traces of high events. Similarly for $tr' \in traces(SYS_{N_t})$ where $tr \approx_{LO_t} tr'$: write $tr' = p' \frown l \frown h'$. It suffices to show that $\mathcal{L}_t[\![S/tr]\!] = \mathcal{L}_t[\![S/(p \frown l)]\!]$. If we have shown this equality for arbitrary tr, then the same applies for S/tr', meaning $\mathcal{L}_t[\![S/tr']\!] = \mathcal{L}_t[\![S/(p' \frown l)]\!]$. By inductive hypothesis, $\mathcal{L}_t[\![S/p]\!] = \mathcal{L}_t[\![S/p']\!]$, and therefore $\mathcal{L}_t[\![S/(p \frown l)]\!] = \mathcal{L}_t[\![S/(p' \frown l)]\!]$. By transitivity, we have that $\mathcal{L}_t[\![S/tr]\!] = \mathcal{L}_t[\![S/tr']\!]$, which is what needs to be proven. Thus, the crux of the argument is to show that the high events of tr do not affect low's view of the system; the second trace tr' is immaterial.

We consider the event l case-by-case over the different events in SYS_{N_t}:

- $l \in R_i$ for some i
  That is, l is a return from a system call into user space. Because l is a low event, l is of the form $i.s.(S, I, \dots)$ where $t \notin S$. After this event, $i{:}K$ is in a state ready to receive a new system call ($i : K_{S,I,O}$). Because all events in h are high events, none are system calls of the form $i.s.(S, I, \dots)$ with $t \notin S$, and therefore, none can force $i{:}K$ into a different state. In other words, the events h can happen either before or after l; SYS_{N_t} will accept (and refuse) the same events after either ordering. That is:
  $$\mathcal{L}_t[\![SYS_{N_t}/(p \frown l \frown h)]\!] = \mathcal{L}_t[\![SYS_{N_t}/(p \frown h \frown l)]\!].$$
  We can apply the inductive hypothesis to deduce that:
  $$\mathcal{L}_t[\![SYS_{N_t}/(p \frown h)]\!] = \mathcal{L}_t[\![SYS_{N_t}/p]\!]$$
  Appending the same event l to the tail of each trace gives:
  $$\mathcal{L}_t[\![SYS_{N_t}/(p \frown h \frown l)]\!] = \mathcal{L}_t[\![SYS_{N_t}/(p \frown l)]\!]$$
  and by transitivity:
  $$\mathcal{L}_t[\![SYS_{N_t}/(p \frown l \frown h)]\!] = \mathcal{L}_t[\![SYS_{N_t}/(p \frown l)]\!]$$
  which proves the claim for this case.
- $l = i.s.(S, I, \mathsf{create\_tag}, w)$ for some $i \in \mathcal{P}$, and some $w \in \{\mathsf{Add}, \mathsf{Remove}, \mathsf{None}\}$
  After accepting this event, the process $i{:}K$ can no longer accept system calls; it can only accept a response in the form $i.s.(S, I, t')$ for some tag t'. Since $l \in LO_t$, it follows that $t \notin S$ for both the system call and its eventual reply. The high events in h could affect the return value to this system call (and therefore $\mathcal{SF}[\![S/tr]\!]$) if the space of t's returned somehow depends on h, because h changed the state of the shared tag manager. An inspection of the tag manager shows that its state only changes as a result of a call to $e = j.g.(S', I', \mathsf{create\_tag}, w)$ for some process j, and labels S' and I'. Such a

call would result in a tag such as $t' = g(S', I', x)$ being allocated, for some arbitrary x. Because $e \in h$ is a high event, $t \in S'$. Because l is a low event, $t \notin S$. Thus, $S' \neq S$, and assuming g is injective, it follows that $t' \neq t$, for all x. Therefore, events in h cannot influence which tags t' might be allocated as a result of a call to create_tag. We apply the same argument as above, that h and l can happen either before or after one another without changing the failures of the system. Hence, the claim holds in this case.

- $l = i.s.(S, I, \mathsf{change\_label}, w, L)$ for some $i \in \mathcal{P}$, $w \in \{\mathsf{Add}, \mathsf{Remove}, \mathsf{None}\}$ and $L \subseteq \mathcal{T}$.
  After accepting l, the process $i{:}K$ is expecting an event of the form $i.s.(S, I, r)$ for $r \in \{\mathsf{Ok}, \mathsf{Error}\}$, to indicate whether the label change succeeded or failed. It will transition to another internal state (and will behave differently in the future) on success. The only way an event in h can influence this outcome is to alter the composition of $\hat{O}$, which the tag manager checks on i's behalf by answering $i.g.(\mathsf{check+})$ and $i.g.(\mathsf{check\text{-}})$ within the *CHECK* subprocess.
  Consider the case in which h contains an event e such that $e = j.s.(S', I', \mathsf{create\_tag}, w)$, and j is the high process that issued e (that is, $t \in S'$). After e, the kernel might have performed the internal events necessary to serving this system call, meaning a new tag t' was allocated, and the tag manager switched to a new state reflecting ${t'}^{+} \in \hat{O}$ or ${t'}^{-} \in \hat{O}$. If $t' \in L$, then h's occurrence allows l to succeed, and h's absence causes l to fail. $t' \in L$ if and only if $\mathcal{L}_t[\![SYS_{N_t}/(p \frown l)]\!] \neq \mathcal{L}_t[\![SYS_{N_t}(p \frown l \frown h)]\!]$. However, we claim that t' is a member of L only if U_i "predicted" the output of g, which it can do with negligible probability ($2^{-\epsilon}$). With extremely high probability, L could only contain t' if the event e happened before the event l. But our inductive hypothesis has already ruled out this possibility.
- $l = i.s.(S, I, \mathsf{get\_label}, w)$ for some w.
  This call only outputs information about what state a kernel process is in; this state only updates as a result of low events $i.s.(S, I, \mathsf{change\_label})$. All events $e \in h$ do not fit this template since they are high events. Therefore, h does not impact the result of system call l.
- $l = i.s.(S, I, \mathsf{get\_caps})$.
  There are three state transitions that can alter the reply to the get_caps system call: $i.s.(S, I, \mathsf{create\_tag}, w)$, $i.s.(S, I, \mathsf{drop\_caps}, L)$ or $i.s.(S, I, \mathsf{recv}, j)$. None of these calls are equal to an event in h, since they are low events and h contains only high events.
- $l = i.s.(S, I, \mathsf{drop\_caps}, X)$ for some X.
  By definition of the *DROPCAPS* sub-process, a transition to a new $K_{S,I,O'}$ can follow a reply to l. If $e \in h$ can influence O', then it can change $i{:}K$'s failures. However, e cannot influence O' since O' is set to $O - X$ on a successful operation. If the event e is to allocate a new tag t', we can apply the same argument as above to see that ${t'}^{+} \notin O$ and ${t'}^{-} \notin O$, and therefore e cannot affect O'.
- $l = i.s.(S, I, \mathsf{fork})$
  The only event $i{:}K$ will accept after l is $i.s.(S, I, k)$ where k is the process ID of the newly-forked child. By definition of *CHOOSE* above, there exists some x such that $k = g(S, I, x)$. If an event $e \in h$ causes a process ID to be chosen, it would be of the form $p = g(S', I', y)$, for some y, and some S' such that $t \in S'$. That l is a low symbol implies that $t \notin S$ and $S \neq S'$. If g is injective then $k \neq p$. Therefore, event e will never change the value k that this kernel process might output next as its reply to the system call l.
  The other result of the fork system call is that now, a new process k is running. That is, $k : K$ has moved out of the "birth state" and is now willing to accept incoming system calls in state $(k : K_{S,I,O})$. The same arguments as above apply here. Because k was forked by a low process, it too is a low process, expecting only low symbols before it transitions to a new state. Therefore, the events in h cannot affect its state machine.
- $l = i.s.(S, I, \mathsf{getpid})$
  After accepting l, the process $i{:}K$ will only accept $i.s.(S, I, i)$ in this state, so h obviously has no effect.
- $l = i.s.(S, I, \mathsf{exit})$
  Regardless of h, a kernel process will accept no more events after exiting.

– $l = i.s.(S, I, \mathsf{send}, j, X, m)$ for some j, X, m.
The outcome of the send operation depends only on whether $X \subseteq O$ or not. It therefore does not depend on h.
– $l = i.s.(S, I, \mathsf{recv}, j)$
The event after l that $i{:}K$ accepts is $i.s.(S, I, m)$ for some message m. It might also change to a different state if the process j sent capabilities. The relevant possibility for $e \in h$ to consider is $e = j.s.(S', I', \mathsf{send}, i, \{t'^{+}\}, \langle\rangle)$, for some high process j with $t \in S'$. The claim is that this message will never be enqueued at i and therefore will not affect i's next visible event. Say that process j has ownership O' and dual privileges D'. Because we assumed that $t^{-} \notin O \cup O' \cup \hat{O}$, t cannot appear in either D or D'. Also, because i is a low process $t \notin S$. Therefore, $t \in S' - D'$ and $t \notin S \cup D$, which implies that $S' - D' \not\subseteq S \cup D$, and the kernel will not enqueue or deliver j's message to i. Again, we have that h does not affect the i's possibilities for the next message it receives.

We have covered all of the relevant cases, and the theorem follows by induction. ∎

.6 select and Timing

Consider a new system call, select, that involves an explicit timeout:

– $Y \leftarrow \mathsf{select}(t, X)$
Given a set of process indices X, return a set $Y \subseteq X$ such that all $j \in Y$, calling $\mathsf{recv}(j)$ will yield immediate results. This call will block until Y is non-empty, or until t clock ticks expire.

A new kernel subprocess *SELECT* handles the new system call; it allows a user program to wait for the first available receive channel to become readable:

$$
\begin{aligned}
SELECT_{S,I,O} = \big(& s \underset{\kappa}{?} (\mathsf{select}, t, A) \ \rightarrow \\
& \big(\mu X \bullet (q!(\mathsf{select}, A) \rightarrow q?B \ \rightarrow \\
& \qquad \textbf{if } B = \{\} \\
& \qquad \textbf{then } INTRECV^{*}{}_{S,I,O} \,;\, X \\
& \qquad \textbf{else } s \underset{\kappa}{!} B \ \rightarrow \ K_{S,I,O}) \\
& \triangle_t \ \big(s \underset{\kappa}{!} \{\} \ \rightarrow \ K_{S,I,O}\big)\big)
\end{aligned}
$$

There are three new CSP operators here. The first is $\mu X \bullet F(X)$, which according to Hoare's original syntax is the recursion operator: the process X such that $X = F(X)$. The syntax $P; Q$ denotes the process P followed by Q upon P's successful termination. Successful termination is denoted by the special CSP process *SKIP*. Lastly, the "timed interrupt operator" $\triangle_t$ [19] interrupts the selection process after t clicks of the clock and outputs an empty result set.

In the above formulation, the process *SELECT* calls subprocess *INTRECV**, which behaves mostly like *INTRECV*, except it keeps receiving until an admissible message arrives:

$$
\begin{aligned}
INTRECV^{*}{}_{S,I,O} \ = \ & c?(S_{\text{in}}, I_{\text{in}}, j, X, m) \ \rightarrow \\
& g!(\mathsf{dual\_privs}, O) \rightarrow g?D \rightarrow \\
& \textbf{if } (S_{\text{in}} \subseteq S \cup D) \ \wedge \ (I - D \subseteq I_{\text{in}}) \\
& \textbf{then } q!(\mathsf{enqueue}, (X, m)) \rightarrow SKIP \\
& \textbf{else } INTRECV^{*}{}_{S,I,O}
\end{aligned}
$$

With the inclusion of the select operation, the Flume CSP model now explicitly models time. We must update our definitions and proof accordingly. Schneider develops a full notion of process equivalence in timed CSP [19], but the mechanics are complex. Instead, we suggest a technique introduced by Ouaknine [39] and also covered by Schneider [19]: convert a timed model into an untimed model with the introduction of the event *tock*, which represents a discrete unit of time's passage. In particular, Schneider provides the Ψ function for mapping processes from timed CSP to discrete-event CSP with a *tock* event representing the passage of time. For example:

$$\begin{aligned} \Psi(a \rightarrow Q) = \quad & P_0 = a \rightarrow \Psi(Q) \\ & \Box \ \mathit{tock} \ \rightarrow P_0 \\ \Psi(\mathit{WAIT}\, n+1) \ = \ & \mathit{tock} \ \rightarrow \ \Psi(\mathit{WAIT}\, n) \end{aligned}$$

In the proof of non-interference, we consider *tock* a low event, that is not hidden by any concealment operator. We then can assume that the Ψ translation is applied to all states of the Flume model.