

Domain Separation by Construction

William Harrison Mark Tullsen James Hook

Pacific Software Research Center
OGI School of Science & Engineering
Oregon Health & Science University

Abstract. Language-based approaches to security typically use static type systems to control information flow, relying on type inference to distinguish secure programs from insecure ones. This paper advocates a novel approach to language-based security: by structuring software with monads (a form of abstract data type for effects), we are able to maintain separation of effects by construction. The thesis of this work is that well-understood properties of monads and monad transformers aid in the construction and verification of secure software. We introduce a formulation of non-interference based on monads (rather than the typical trace-based formulation). Using this formulation, we prove a non-interference style property for a simple instance of our abstract system model. Because monads may be easily and safely represented within any higher-order, typed functional language, monadic event systems may be directly realized within such a language.

1 Introduction

This paper advocates a novel approach to language-based security: security by construction. Starting from a mathematical model of shared-state concurrency, we develop, through a sequence of refinements, the derivation of an exemplary operating system kernel supporting both standard Unix-like system calls (e.g., fork, sleep, etc.) and a formally verified security policy (domain separation).

The research reported here presents a formal, language-based model of security combining three approaches to system security and language semantics:

- **Security “By Design.”** Some approaches advocate implementation strategies for secure system construction, with the idea that such disciplined strategies are more likely to remain secure. One such strategy used in Java implementations, *sandboxing*, limits the scope of stateful effects by executing threads in disjoint regions of memory as illustrated in Figure 1. Good engineering, however, does not constitute a guarantee of any security policy.
- **Trace-based Formal Security Models.** There are a number of formal security models [10, 40, 20, 41, 21] which characterize permissible interactions between concurrent threads in terms of traces of abstract events. These models make precise the intuition that low-security operations should be “oblivious” to the execution of high-level operations. One drawback of such models is that their precise relationship to actual systems remains unclear.

- **Monadic Language Semantics.** The approach advocated here combines “by design” security with trace-based models into a common framework based on *monads*. Monads provide a mathematically sophisticated theory of effects which has proven useful in denotational semantics [23, 17, 25], functional programming [35], and software verification [15]. Structuring our system specifications with monads yields many benefits, not the least of which are a number of useful properties obtained “by construction” which simplify the verification of our security property.

Language-based Security via Monadic Constructions. We present a formal model of security in which the model itself may be refined to an implementation of a system with secure shared-state concurrency. Essential to our approach is the use of monads and monad transformers to structure our specifications. It is our thesis that systems constructed monadically are more easily verified because of the monadic encapsulation of effects. Monads and monad transformers allow us to reason about our system definitions at the level of denotational semantics. Because monads may be easily realized within any higher-order functional programming language, system specifications are readily executable.

Many formal security models are formulated in terms of sequences of abstract events. For the sake of convenience, we will refer to such models as *event systems*. The intended interpretation of events is that they are imperative operations on a shared state, but this is not made explicit—that is, the events themselves are uninterpreted. Our approach makes this interpretation explicit by considering languages of system behaviors and their denotational semantics. According to our view of language-based security, an arbitrary, interleaved sequence of **Lo** and **Hi** operations in a traditional traced-based model:

$$h_0, l_0, \dots, h_n, l_n$$

is viewed as the imperative *program* describing a particular (partial) system behavior:

$$h_0 ; l_0 ; \dots ; h_n ; l_n$$

We then give such programs a (monadic) denotational semantics:

$$\llbracket - \rrbracket : (Event_{Lo} + Event_{Hi}) \rightarrow R()$$

where R is a monad encapsulating imperative effects and a notion of interleaving concurrency called resumptions [26, 23, 25]. The monad R is constructed especially to isolate **Hi**, **Lo**, and kernel effects from one another. These denotational semantics are developed in Section 5.

Partitioning Effects with Monads. Domain separation is supported by partitioning the state into disjoint pieces, with each piece corresponding to a separate domain. Stateful operations are then given a security level and can only manipulate the storage partition corresponding to its security level. This partitioning

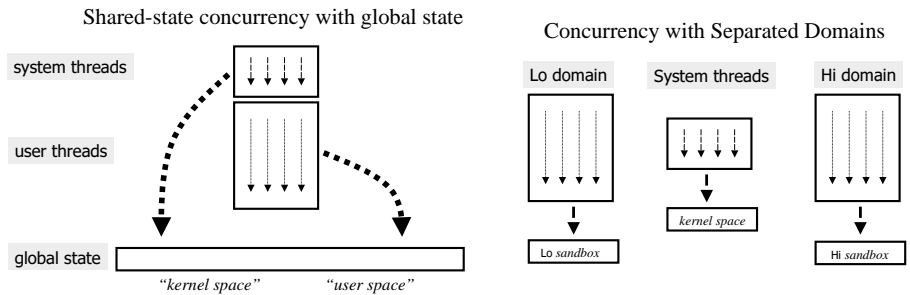


Fig. 1. *Shared-state concurrency with global state (left).* In this system architecture, all threads access the same global state, and so may interfere with one another in ways difficult to control. *Concurrency with separated domains (right).* Secure computation is promoted by partitioning Hi, Lo, and kernel threads into separate domains, each of which has its own storage (or *sandbox*) which only it may affect. Domain separation limits the scope of effects occurring in one domain from others. While this does not guarantee secure computation, it does simplify the implementation, specification, and verification of security properties.

is sometimes referred to as *sandboxing*. With monads, it is a simple matter to partition storage into sandboxes, and this process is particularly straightforward when the monads are constructed with monad transformers.

How this partitioning works is illustrated in Figure 2, where, for the sake of simplicity¹, we assume that there are two security levels Hi and Lo. Corresponding to these levels are separated domains (c) that maintain distinct stores H and L, respectively. Hi and Lo stateful operations are then encapsulated within monads of the same name, created with the monad transformers (StateT H) and (StateT L), respectively. Hi and Lo stateful operations *h* and *l* may be executed by lifting them to the kernel level (b) with liftH and liftL, respectively, and these lift mappings are created by the application of the state monad transformers. Separation of effects is maintained by these lift mappings—precisely how is described in detail in Section 4 below.

Separation By Construction. The approach advocated here achieves secure shared-state concurrency by construction, where “by construction” is used in two different, yet complementary, senses. The process is illustrated in Figure 3, which illuminates the meanings of “by construction”:

- **Stepwise-refinement of System.** The vertical axis of Figure 3 measures the richness of observable system events, and each step along that axis marks an addition to those observables. At point 0, only a single, monolithic process domain exists and all thread scheduling is static. At point 1, threads are

¹ All of our results generalize easily to *n* separate domains.

Monadic Event Systems




Fig. 2. *Monadic Event Systems:* Processes created in isolated security domains Lo and Hi are lifted to the runtime platform K . Laws about monad transformers (StateT and ResT above) aid in proving domain separation.

executable on separate domains. Point 2 allows statically-scheduled multi-tasking, while point 3 allows the scheduling of threads to occur dynamically. Dynamic scheduling is necessary for threads to affect their own execution behavior (as with the Unix system call `sleep` and intradomain synchronisation mechanisms such as semaphores) or to affect the system “wait list” (as with the Unix system call `fork`). Point 4 adds such thread-level control operations, and point 5 allows for secure, interdomain communication (such as message-passing which obeying a “no-write-down” security policy).

- **Properties of Monads & Monad Transformers.** Many of the above enhancements to system functionality are reflected in refinements to the underlying monads and monad transformers of Figure 2. Monads and monad transformers allow the effects of threads of differing security levels to be controlled in a mathematically rigorous manner, and this scoping of effects tames insecure interference between threads. The “by construction” properties of monads and monad transformers are extremely useful in formally demonstrating domain separation.

Security Property. The security property we prove may be intuitively described as the execution of low security events being *oblivious* to the execution of high security events. For any initial sequence of interleaved Hi and Lo events

$$h_0 ; l_0 ; \dots ; h_n ; l_n$$

the effect of its execution on the Lo state should be identical to that of executing the Lo events in isolation:

$$l_0 ; \dots ; l_n$$

Using the aforementioned denotational semantics, we make this notion precise in Section 5.3.

Overview. This paper is structured as follows. Section 2 surveys related work and Section 3 summarizes the background on monads, monad transformers, and resumption-based concurrency necessary for understanding this work. Section 4 describes the useful properties that are obtained *by construction* with monad transformers. Section 5 presents the main results of this research—the stepwise development of an exemplary operating system kernel for secure computation. First, a language of system behaviors is defined and given a resumption-monadic semantics in Section 5.1. Second, a system for shared-state concurrency with global store is transformed into a system with separated domains in Section 5.2. Third, security in this setting—*take separation*—is specified and verified in Section 5.3. Finally, Section 6 summarizes the present work and outlines future directions.

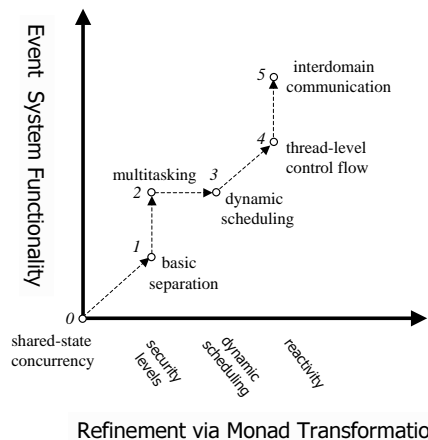


Fig. 3. *Separation by construction:* System functionality is enhanced by refining monad transformers.

2 Related Work

Many formal security models—including non-interference and separability—are formulated in terms of *event systems* [10, 40, 20, 41, 21]. While there are a variety of similar formulations of event systems, they all include the following: an event system S is comprised of a set of abstract *events* E and a set of *traces* T of (potentially infinite) sequences of events in E . Furthermore, each event is assigned a security level, which, for the sake of the present discussion, we will assume to be Lo and Hi ².

Security policies may then be formulated in terms of events systems as relationships between event traces which specify permitted system executions. These

² Event systems also distinguish subsets of E as *input* and *output* events.

typically are of the form: if $\tau_1, \dots, \tau_n \in T$, then $f(\tau_1, \dots, \tau_n) \in T$, where function f combines the traces τ_i in some manner specific to the security policy. For example, *separability* [21, 41] can be defined formally as: for every pair of traces $\tau_1, \tau_2 \in T$, then any trace τ such that $(\tau \downarrow_{\text{Lo}}) = (\tau_1 \downarrow_{\text{Lo}})$ and $(\tau \downarrow_{\text{Hi}}) = (\tau_2 \downarrow_{\text{Hi}})$ is also in T . Here $(\tau \downarrow_{\text{Lo}})$ and $(\tau \downarrow_{\text{Hi}})$ filter out all but the **Lo** and **Hi** events from τ , respectively. (As formulated these techniques are easily generalized to any total order of security levels known statically at compile time.)

Previous work has focused on languages with implicit imperative features, such as Java or ML. This paper instead assumes a language in which all imperative features are captured by a monad. As a result all impure effects (references, exceptions, I/O) are distinguished from pure computations by their types, and thus side-effects are allowed while preserving the semantics of the purely functional subset of the language.

Moggi was the first to observe that the simple structure known as a monad was appropriate for the development of modular semantic theories for programming languages [22]. In his initial development, Moggi showed that most known semantic effects could be naturally expressed monadically. He also showed how a sequential theory of concurrency could be expressed in the Resumption monad [23]. That model of concurrency is used extensively below. Wadler and colleagues at Glasgow University observed that using monads internally in the pure, higher-order, typed language Haskell gave a natural and safe embedding of effectful computation in a pure language.

We are working to develop and formally verify a security kernel as part of the Programatica project at OGI. To formally verify security properties of such a system is a formidable task. One approach to reducing the enormity of this task has been using type-systems for information-flow. There has been a growing emphasis on such *language-based* techniques for information flow security [34, 33, 16, 32, 27] (see [29] for a survey of this work). The chief strength of this type-based approach is that the well-typedness of terms can be checked statically and automatically, yielding high assurance at low cost. Unfortunately, this type-based approach is not as general as one might wish: 1) there will be programs which are secure but which will be rejected by the type system due to lack of precision and 2) there will be programs that have information flow leaks which we want to allow (e.g., a declassification program) but which would be rejected by the type system.

The Programatica team has investigated a more general—but potentially more expensive—approach to building large systems guaranteed to satisfy security properties. In our approach we derive proofs of security properties by hand, and yet proving security properties of large systems is kept tractable. This is due primarily to the following techniques: 1) Using a purely functional, statically-typed, polymorphic language, we are able to reason about most of the system components using only their types. And, 2) We structure the system so as to minimize the amount of code that must be reasoned about explicitly. We do this using monads, a mechanism previously used for modularizing interpreters. Different monadic types are assigned different security levels, and thereby indi-

vidual applications using effects (such as accessing low security memory) may be validated automatically and statically according to a security policy.

There have been a number of attempts to develop secure OS kernels: PSOS [24], KSOS [19], UCLA Secure Unix [37], SAT [5], KIT [2], EROS [31], and MASK [39] among others. There has also been some work using functional languages to develop high confidence system software: the Fox project at CMU [14] is a case in point of how typed functional languages can be used to build reliable systems software (e.g., network protocols, active networks); the Ensemble project at Cornell [4] uses a functional language to build high performance networking software; and the Switchware project [1] at the University of Pennsylvania is developing an active network in which a key part of their system is the use of a typed functional language.

While our approach does not provide full system assurance automatically, as with the type-based approach, many large components of the system are verified automatically (by the type system); this aids verification of the global system security property. This approach to guaranteeing security underlies the construction of the OSKer³ operating system kernel.

Formulating security policies in terms of non-interference goes back to the work of Goguen and Meseguer [10, 11]. Our notion of separation is basically the non-interference of Goguen and Meseguer; although for simplicity of the current presentation our approach only allows for an “interferes” relation that is the identity relation (i.e., no domain can interfere with another). Haigh and Young [13] and Rushby [28] have extended this work to the intransitive case (where the “interferes” relation is not required to be transitive). Non-interference has been extended to concurrent [33, 6] and probabilistic models [12].

Modeling concurrency by resumptions was introduced by Plotkin [26, 30]. Moggi showed how resumptions could be modeled with monads [23] and our formulation of resumptions in terms of monad transformers is that of Papaspyrou [25]. Concurrency can also be handled by continuations [38, 7, 8]. Resumptions can be viewed as a disciplined use of continuations. Using resumptions rather than arbitrary continuations makes reasoning about our system easier.

3 Background

Monads [23] can be understood as abstract data types for defining languages and programs⁴. A monad ADT can encapsulate such language features as state, exceptions, multi-threading, environments, and CPS. Although combinations of such features can be encapsulated in a single monad, a more modular approach, in which each feature can be treated separately, is achieved with *monad transformers* [23, 18, 17]. Monad transformers allow us to easily combine and extend monads. (A monad is extended similarly to how a class is extended using inheritance in object oriented languages: what makes sense in a monad “class” makes sense in the “subclass” created by inheritance.) In this section we will introduce

³ OSKer stands for “Oregon Security Kernel.”

⁴ An introduction to monads can be found in Wadler [36].

monads and monad transformers, we will then describe the resumption monad transformer: a monad for concurrency.

3.1 Monads and Monad Transformers

A monad is a triple $\langle M, \eta, \star \rangle$ consisting of a type constructor M and two operations:

$$\begin{aligned} \eta &: a \rightarrow M a && \text{(unit)} \\ (\star) &: M a \rightarrow (a \rightarrow M b) \rightarrow M b && \text{(bind)} \end{aligned}$$

The η operator is the monadic equivalent of the identity function, it brings an element into the monad. The \star operator gives a form of sequential composition. These operators must satisfy the monad laws:

$$\begin{aligned} (\eta v) \star k &= k v && \text{(left unit)} \\ x \star \eta &= x && \text{(right unit)} \\ x \star (\lambda a. (k a \star h)) &= (x \star k) \star h && \text{(assoc)} \end{aligned}$$

In what follows we often use the \gg operator, defined in terms of \star thus:

$$v \gg k = v \star \lambda_. k$$

The identity monad $\langle \text{Id}, \eta_{\text{Id}}, \star_{\text{Id}} \rangle$ is defined as follows:

$$\begin{aligned} \text{Id } a &= a \\ x \star_{\text{Id}} k &= k x \\ \eta_{\text{Id}} x &= x \end{aligned}$$

Another example is the state monad $\langle \text{St } s, \eta_{\text{St } s}, \star_{\text{St } s} \rangle$ defined as follows:

$$\begin{aligned} \text{St } s a &= s \rightarrow (a \times s) \\ \eta_{(\text{St } s)} x &= \lambda \sigma : s. (x, \sigma) \\ x \star_{(\text{St } s)} f &= \lambda \sigma_0. \text{let } (a, \sigma_1) = x \sigma_0 \text{ in } f a \sigma_1 \\ \text{u}(\Delta) &= \lambda \sigma. ((), \Delta \sigma) \\ \text{g} &= \lambda \sigma. (\sigma, \sigma) \end{aligned}$$

Here, $()$ signifies both the unit type and the single element of that type. The operators u and g , for updating and getting the state respectively, are defined only for the $\text{St } s$ monad. Such monad specific operators are referred to as improper morphisms.

There are various formulations of monad transformers, we follow that given in [18] where a monad transformer consists of a type constructor T , a mapping from a given monad $\langle M, \eta_M, \star_M \rangle$ to a new monad $T M = \langle T M, \eta_{(T M)}, \star_{(T M)} \rangle$, and an associated function lift_T . The function $\text{lift}_T : M a \rightarrow T M a$ performs a “lifting” of computations in M to computations in $(T M)$; and will generally satisfy the *Lifting Laws* [17]:

$$\begin{aligned} \text{lift} \circ \eta_M &= \eta_{(T M)} \\ \text{lift}(x \star_M f) &= (\text{lift } x) \star_{(T M)} (\text{lift} \circ f) \end{aligned}$$

These laws ensure that a monad transformer adds features without changing features of the base monad M .

As an example, the state monad $\text{St } s$ can be written more generally as the monad transformer $\text{StateT } s$ which transforms M . The definition is in Figure 4. Note that the monad transformer $\text{StateT } s$ applied to the identity monad Id gives the original $\text{St } s$ monad.

$\text{StateT } s \text{ M } a = s \rightarrow M(a \times s)$ $\eta x = \lambda \sigma. \eta_M(x, \sigma)$ $x \star f = \lambda \sigma_0. (x \sigma_0) \star_M (\lambda(a, \sigma_1). f a \sigma_1)$ $\text{lift } x = \lambda \sigma. x \star_M \lambda y. \eta_M(y, \sigma)$ $u(\Delta : s \rightarrow s) = \lambda \sigma. \eta_M((), \Delta \sigma)$ $g = \lambda \sigma. \eta_M(\sigma, \sigma)$	$\text{ResT } M a = \mu R. \text{Done } a + \text{Pause}(M(R a))$ $\eta x = \text{Done } x$ $(\text{Done } v) \star f = f v$ $(\text{Pause } m) \star f = \text{Pause}(m \star_M \lambda r. \eta(r \star_M f))$ $\text{step } \phi = \text{Pause}(\phi \star_M \lambda v. \eta_M(\text{Done } v))$
---	--

Fig. 4. The State and Resumption Monad Transformers ($\text{StateT } s$) and ResT . The unit and bind of the new monads are η and \star , respectively, while those of the transformed monad M are subscripted.

3.2 Resumption-based Concurrency and ResT

This section introduces concurrency based on the resumption monad transformer, the definition of which can be found in Figure 4. How resumption based concurrency works is best explained by an example. We define *thread* to be a (possibly infinite) sequence of “atomic operations.” We make this notion precise below, but for the moment, assume that an atomic operation is a single machine instruction and that a thread is a stream of such instructions characterizing a program execution. Consider first that we have two simple threads $a = [a_0; a_1]$ and $b = [b_0]$. According to the “concurrency as interleaving” model of concurrency, concurrent execution of threads a and b means the set of all their possible interleavings: $\{[a_0; a_1; b_0], [a_0; b_0; a_1], [b_0; a_0; a_1]\}$.

But how do computations in a resumption monad correspond to threads? If the atomic operations of a and b are computations of type $M()$, then the computations of type $\text{ResT } M()$ are the set of possible interleavings:

$$\begin{aligned} & \text{Pause}(a_0 \gg_M \eta_M(\text{Pause}(a_1 \gg_M \eta_M(\text{Pause } b_0 \gg_M \eta_M(\text{Done } ()))))) \\ & \text{Pause}(a_0 \gg_M \eta_M(\text{Pause}(b_0 \gg_M \eta_M(\text{Pause } a_1 \gg_M \eta_M(\text{Done } ()))))) \\ & \text{Pause}(b_0 \gg_M \eta_M(\text{Pause}(a_0 \gg_M \eta_M(\text{Pause } a_1 \gg_M \eta_M(\text{Done } ()))))) \end{aligned}$$

Closer comparison of these two reveals that, instead of using the lazy “cons” operation ($h : t$) as in the stream definition of concurrency above, the monadic version uses something similar: $\text{Pause}(h \gg_M \eta_M t)$. This is important because threads may be infinite, and the laziness of Pause allows infinite *computations* to be constructed in ($\text{ResT } M$) just as the laziness of ($h : t$) allows infinite *streams* to be constructed.

Finally, we note that the resumption semantics of concurrency involves the elaboration of all possible thread interleavings, and that, while such a semantics may be expressed monadically via the non-determinism monad [23, 18], it is not computationally tractable. We choose instead to pick out a single particular interleaving via a scheduler.

4 “By Construction” Properties of Monad Transformers

We get a number of useful properties by construction through the use of monad transformers. The state monad transformer’s lift mappings have two principal uses here. First, lifting preserves stateful behavior. In Figure 2 for example, this means that `Hi` and `Lo` operations behave the same when lifted to the kernel monad `K` as at their respective base monads. This is formally captured in Section 4.1 below. Furthermore, liftings also delimit the effects of stateful operations on separate domains, and this phenomenon—which we call *atomic non-interference*—is described in detail in Section 4.2.

4.1 State Monads and Their Axiomatization

This section presents an algebraic characterization of state monads. Intuitively, a state monad is a monad with non-proper morphisms to manipulate state. The behavior of these non-proper morphisms is captured by axioms below. Not surprisingly, it is then demonstrated that the state monad transformer creates new state monads, and preserves existing state monads.

State Monad Structure. The quintuple $\langle M, \eta, \star, u, g, s \rangle$ is a *state monad structure* when:

1. $\langle M, \eta, \star \rangle$ is a monad with operations: unit $\eta : \alpha \rightarrow M \alpha$ and bind $\star : M \alpha \rightarrow (\alpha \rightarrow M \beta) \rightarrow M \beta$
2. update operation on s is $u : (s \rightarrow s) \rightarrow M()$
3. get operation on s is $g : M s$

We will refer to a state monad structure $\langle M, \eta, \star, u, g, s \rangle$ simply as M if the associated operations and state type are clear from context.

State Monad Axioms. The following axiomatization of the state monad is not meant to be complete. Rather, it reflects the properties of state monads required later in the proofs.

Let $M = \langle M, \eta, \star, u, g, s \rangle$ be a state monad structure. M is a *state monad* if the following equations hold for any $f, g : s \rightarrow s$,

$$\begin{aligned}
 u f \star \lambda _ . u g &= u (g \circ f) && \text{(sequencing)} \\
 g \star \lambda \sigma_0 . u (\lambda _ . \sigma_0) &= \eta() && \text{(cancellation)}
 \end{aligned}$$

The (sequencing) axiom shows how updating by f and then updating by g is the same as just updating by their composition ($g \circ f$). The (cancellation) axiom requires that overwriting the state with the result of an immediately-preceding g operation has no effect.

For state monad $\langle M, \eta, \star, u, g, \sigma \rangle$, a consequence of (sequencing) we use later is:

$$u \ f \gg \text{init} = \text{init} \quad (\text{clobber})$$

where init is defined as: $u \ (\lambda_.\sigma_0)$ for some constant state σ_0 .

Theorem 1 (StateT creates a state monad). *For any monad M , let monad $M' = \text{StateT } sto \ M$, $u : (sto \rightarrow sto) \rightarrow M'()$ and $g : M' \ sto$ be the non-proper morphisms added by $(\text{StateT } sto)$. Then $\langle M', \eta_{M'}, \star_{M'}, u, g, sto \rangle$ is a state monad.*

Theorem 2 (StateT preserves stateful behavior). *For any state monad $M = \langle M, \eta, \star, u, g, sto \rangle$, the following state monad structure is a state monad:*

$$\langle \text{StateT } s \ M, \eta', \star', \text{lift} \circ u, \text{lift}g, s \rangle$$

where η' , \star' , and lift are the monadic unit, bind, and lifting operations, respectively, defined by $(\text{StateT } s)$.

4.2 Formalizing Atomic Non-interference

The second “by construction” property of monad transformers relates to how their associated lift mappings delimit stateful effects in monads created from multiple applications of the state monad transformers (e.g., the kernel monad K in Figure 2). This property—*atomic non-interference*—is useful for proving the security property below in Section 5.3.

Atomic Non-interference Relation $\#$. For any monad M with bind operation \star , define the non-interference relation $\# \subseteq M() \times M()$ so that, for $\varphi, \gamma : M()$, $\varphi \# \gamma$ holds if and only if:

$$\varphi \gg \gamma = \gamma \gg \varphi$$

Theorem 3 (Atomic Non-interference of updates). *Let M be the state monad $\langle M, \eta_M, \star_M, u_A, g_A, A \rangle$. By Theorem 1, the following structure is also a state monad: $M' = \langle \text{StateT } B \ M, \eta_{M'}, \star_{M'}, u_B, g_B, B \rangle$. Then, for all $\delta_A : A \rightarrow A$ and $\delta_B : B \rightarrow B$,*

$$(u_B \ \delta_B) \#_{M'} \text{lift}(u_A \ \delta_A)$$

Theorem 4 (Non-interference Preserved by Monad Transformers). *Let M be a monad with two operations, $a : M()$ and $b : M()$ such that $a \#_M b$. Then,*

$$(\text{lift } a) \#_{(\text{T } M)} (\text{lift } b)$$

where T is a monad transformer and $\text{lift} : M \ a \rightarrow (\text{T } M) \ a$.

5 Stepwise Development of a Secure OS Kernel

In monadic event systems, events are programs in the *Event* language, and traces of events are the denotations of these programs according to a resumption-monadic semantics. The abstract syntax and resumption semantics of the language of events are presented below. The *Event* language contains sufficient expressiveness to allow for potentially infinite streams of operations because of the inclusion of loops. More importantly, it allows for expression of potentially interfering programs as well.

Abstract Syntax for the Event Language. Below is an abstract syntax for *Event*:

$$\begin{aligned} \textit{Event} ::= & \textit{Var} := \textit{Exp} \mid \textit{skip} \mid \textit{Event}; \textit{Event} \mid \\ & \textit{if } \textit{Exp} \textit{ then } \textit{Event} \textit{ else } \textit{Event} \mid \textit{while } \textit{Exp} \textit{ do } \textit{Event} \\ \textit{Exp} ::= & \textit{Var} \mid \textit{Integer} \end{aligned}$$

5.1 Shared-State Concurrency with Global State (system 0)

In this section, we introduce the most basic model for shared-state concurrency is constructed. The monadic event system described in this section corresponds to the point labelled with “0” in Figure 3.

Monad Hierarchy Using the monad transformers `ResT` and `StateT` defined in Section 3, we define the monad hierarchy for shared-state concurrency with global state. This monadic event system is degenerate in the sense that the For any monad M , this hierarchy is:

$$\text{Gl} = \text{StateT } G \text{ } M \quad K = \text{Gl} \quad R = \text{ResT } K$$

Here, state type G is $Name \rightarrow Integer$.

Associated with these monad constructions are a number of non-proper morphisms (i.e., monadic operators other than η and \star). These are:

$$\begin{aligned} u_G & : (G \rightarrow G) \rightarrow \text{Gl}() & \textit{stepG} & : K a \rightarrow R a \\ g_G & : \text{Gl } G & \textit{stepG } \varphi & = \textit{PauseLo} (\varphi \star_{\kappa} \lambda v. \eta_{\kappa}(\textit{Done } v)) \\ \textit{liftG} & = id \end{aligned}$$

The morphism u_G applies a state-to-state map to the current G state in the Gl monad, while the morphism g_G reads and returns the current G state. The lifting \textit{liftG} reinterprets Gl computations in the kernel monad K . The aforementioned morphisms are all defined by applications of the `StateT` monad transformer. The morphisms \textit{stepH} and \textit{stepL} create a “paused” K computation in either the Hi or Lo security levels. The “step” functions result from the application of the `ResT` transformer.

Semantics for the Event Language. The resumption-monadic denotational semantics for system behaviors is listed below. It is just what one would expect given recent work on the resumption-monadic semantics of concurrency [25].

$$\begin{aligned}
ev : Event &\rightarrow \mathbb{R} () \\
ev(x := e) &= (exp\ e) \star_{\mathbb{R}} \lambda v. (load \circ u_G)[x \mapsto v] \\
ev(e_1; e_2) &= (ev\ e_1) \gg_{\mathbb{R}} (ev\ e_2) \\
ev\ skip &= (load \circ u_G) (\lambda i. i) \\
ev(\text{if } b \text{ then } e_1 \text{ else } e_2) &= exp\ b \star_{\mathbb{R}} \lambda v. \text{if } v = 0 \text{ then } (ev\ e_1) \text{ else } (ev\ e_2) \\
ev(\text{while } b \text{ do } c) &= mwhile\ (exp\ b)\ (ev\ c) \\
mwhile\ b\ \varphi &= b \star_{\mathbb{R}} \lambda v. \text{if } v = 0 \text{ then } \varphi \gg_{\mathbb{R}} (mwhile\ b\ \varphi) \\
&\quad \text{else } (load \circ u_G) (\lambda i. i)
\end{aligned}$$

$$\begin{aligned}
exp : Exp &\rightarrow \mathbb{R}\ Integer & load : G\!l\ a &\rightarrow \mathbb{R}\ a \\
exp\ x &= (load\ g_G) \star_{\mathbb{R}} \lambda \sigma. \eta_{\mathbb{R}}(\sigma\ x) & load &= step \circ lift \\
exp\ i &= \eta_{\mathbb{R}}\ i
\end{aligned}$$

Please observe that because each observable event (namely u or g) is wrapped by a *load*, each such event is *Pause*'d.

5.2 Basic Separation (system 1)

Event systems contain trace projections based on security level—we write these as “ \downarrow_{Hi} ” and “ \downarrow_{Lo} .” Monadic event systems need a similar capability, which is achieved by refining the resumption monad transformer of Papaspyrou [25] to reflect the Hi and Lo security levels. The refined resumption monad transformer is:

$$ResT\ M\ a = \mu R. Done\ a + PauseLo\ (M(R\ a)) + PauseHi\ (M(R\ a))$$

The unit (η) is *Done*, and the bind (\star) of the transformed monad is defined just as one would expect:

$$\begin{aligned}
(Done\ v) \star f &= f\ v \\
(PauseLo\ \varphi) \star f &= PauseLo\ (\varphi \star_{\mathbb{M}} \lambda r. \eta_{\mathbb{M}}(r \star f)) \\
(PauseHi\ \varphi) \star f &= PauseHi\ (\varphi \star_{\mathbb{M}} \lambda r. \eta_{\mathbb{M}}(r \star f))
\end{aligned}$$

Using this resumption transformer and `StateT`, we define the monad hierarchy from Figure 3 for any monad M as:

$$\begin{aligned}
Hi &= StateT\ H\ M & K &= StateT\ H\ (StateT\ L\ M) \\
Lo &= StateT\ L\ M & R &= ResT\ K
\end{aligned}$$

Here, L and H are state types equal to $Name \rightarrow Integer$.

Associated with these monad constructions are a number of non-proper morphisms (i.e., monadic operators other than η and \star). These are:

$$\begin{array}{ll}
\mathbf{u}_L & : (L \rightarrow L) \rightarrow \mathbf{Lo}() \\
\mathbf{g}_L & : \mathbf{Lo} L \\
\mathit{lift}L & : \mathbf{Lo} a \rightarrow \mathbf{K} a \\
\mathit{step}L & : \mathbf{K} a \rightarrow \mathbf{R} a \\
\mathit{step}L \varphi = \mathit{PauseLo} (\varphi \star_{\mathbf{K}} \lambda v. \eta_{\mathbf{K}}(\mathit{Done} v)) &
\end{array}
\qquad
\begin{array}{ll}
\mathbf{u}_H & : (H \rightarrow H) \rightarrow \mathbf{Hi}() \\
\mathbf{g}_H & : \mathbf{Hi} H \\
\mathit{lift}H & : \mathbf{Hi} a \rightarrow \mathbf{K} a \\
\mathit{step}H & : \mathbf{K} a \rightarrow \mathbf{R} a \\
\mathit{step}H \varphi = \mathit{PauseHi} (\varphi \star_{\mathbf{K}} \lambda v. \eta_{\mathbf{K}}(\mathit{Done} v)) &
\end{array}$$

The morphisms \mathbf{u}_L and \mathbf{u}_H apply a state-to-state map to the current state in their respective monads, while the morphisms \mathbf{g}_L and \mathbf{g}_H read and return the current state. The liftings $\mathit{lift}L$ and $\mathit{lift}H$ reinterpret \mathbf{Lo} and \mathbf{Hi} computations, resp., in the kernel monad \mathbf{K} . The aforementioned morphisms are all defined by applications of the \mathbf{StateT} monad transformer. The morphisms $\mathit{step}H$ and $\mathit{step}L$ create a “paused” \mathbf{K} computation in either the \mathbf{Hi} or \mathbf{Lo} security levels. The “step” functions result from the application of the \mathbf{ResT} transformer.

By Theorems 1 and 2, we know that the monad \mathbf{K} is a state monad with the operations lifted from the \mathbf{Hi} and \mathbf{Lo} monads. That is, the following are state monads:

$$\langle \mathbf{K}, \eta_{\mathbf{K}}, \star_{\mathbf{K}}, (\mathit{lift}L \circ \mathbf{u}_L), (\mathit{lift}L \mathbf{g}_L), L \rangle \quad \langle \mathbf{K}, \eta_{\mathbf{K}}, \star_{\mathbf{K}}, (\mathit{lift}H \circ \mathbf{u}_H), (\mathit{lift}H \mathbf{g}_H), H \rangle$$

Event Semantics (system 1) There are two semantics for *Event* corresponding to the \mathbf{Hi} and \mathbf{Lo} security levels—these are $\mathit{ev}H$ and $\mathit{ev}L$, respectively. The low-security resumption semantics $\mathit{ev}L$ and $\mathit{exp}L$. The high-security semantics (not shown), $\mathit{ev}H$ and $\mathit{exp}H$, is analogous⁵.

$$\begin{array}{ll}
\mathit{ev}L :: \mathit{Event} \rightarrow \mathbf{R} () & \\
\mathit{ev}L (x := e) & = (\mathit{exp}L e) \star_{\mathbf{R}} \lambda v. (\mathit{load}L \circ \mathbf{u}_L)[x \mapsto v] \\
\mathit{ev}L (e_1; e_2) & = (\mathit{ev}L e_1) \gg_{\mathbf{R}} (\mathit{ev}L e_2) \\
\mathit{ev}L \mathit{skip} & = (\mathit{load}L \circ \mathbf{u}_L) (\lambda i. i) \\
\mathit{ev}L (\mathit{if} b \mathit{then} e_1 \mathit{else} e_2) & = \mathit{exp}L b \star_{\mathbf{R}} \lambda v. \mathit{if} v = 0 \mathit{then} (\mathit{ev}L e_1) \mathit{else} (\mathit{ev}L e_2) \\
\mathit{ev}L (\mathit{while} b \mathit{do} c) & = \mathit{mwhile} (\mathit{exp}L b) (\mathit{ev}L c) \\
\mathit{mwhile} b \varphi & = b \star_{\mathbf{R}} \lambda v. \mathit{if} v = 0 \mathit{then} \varphi \gg_{\mathbf{R}} (\mathit{mwhile} b \varphi) \\
& \qquad \qquad \qquad \mathit{else} (\mathit{load}L \circ \mathbf{u}_L) (\lambda i. i) \\
\\
\mathit{exp}L : \mathit{Exp} \rightarrow \mathbf{R} \mathit{Integer} & \qquad \mathit{load}L : \mathbf{Lo} a \rightarrow \mathbf{R} a \\
\mathit{exp}L x = (\mathit{load}L \mathbf{g}_L) \star_{\mathbf{R}} \lambda \sigma. \eta_{\mathbf{R}}(\sigma x) & \qquad \mathit{load}L = \mathit{step}L \circ \mathit{lift}L \\
\mathit{exp}L i = \eta_{\mathbf{R}} i &
\end{array}$$

The semantics $\mathit{ev}L$ ($\mathit{ev}H$) creates traces by injecting the \mathbf{Lo} (\mathbf{Hi}) operations into the $\mathit{PauseLo}$ ($\mathit{PauseHi}$) side of \mathbf{R} . Let e be $(x := 1)$ in the following:

$$(\mathit{ev}L e) = \mathit{PauseLo} (\mathbf{u}_L[x \mapsto 1]) \gg_{\mathbf{K}} \eta_{\mathbf{K}}(\mathit{Done} ()) \quad (1)$$

$$(\mathit{ev}H e) = \mathit{PauseHi} (\mathbf{u}_H[x \mapsto 1]) \gg_{\mathbf{K}} \eta_{\mathbf{K}}(\mathit{Done} ()) \quad (2)$$

Note that the assignment in e is mapped by $\mathit{ev}L$ and $\mathit{ev}H$ into lifted operations on *different* monads (i.e., \mathbf{Lo} and \mathbf{Hi} in (1) and (2), resp.). Then, this security

⁵ The $\mathit{ev}H$ semantics is obtained from $\mathit{ev}L$ by replacing all “L”-suffixed operations (e.g., “ \mathbf{u}_L ”) by their corresponding “H”-suffixed operations (e.g., “ \mathbf{u}_H ”).

assignment is maintained in the resumption-trace by the *PauseLo* and *PauseHi* tags. Because u_L and u_H operate on different states, these assignments can not interfere (in the sense of Section 4.2). The monadic structure is the key to making this approach work.

Below are two schedulers for the basic separation model, *withHi* and *withoutHi*, for a simple monadic event system. The schedule (*withHi lo hi*) creates a Lo and Hi threads from event behaviors *lo* and *hi*, resp., in a round-robin fashion. Schedule (*withoutHi lo*) creates a single Lo thread.

$$\begin{aligned}
\textit{withHi} & : \textit{Event} \rightarrow \textit{Event} \rightarrow \mathbf{R}() & \textit{withoutHi} & : \textit{Event} \rightarrow \mathbf{R}() \\
\textit{withHi lo hi} & = \textit{weave} (\textit{evL lo}) (\textit{evH hi}) & \textit{withoutHi lo} & = \textit{evL lo} \\
\textit{weave} & : \mathbf{R}() \rightarrow \mathbf{R}() \rightarrow \mathbf{R}() \\
\textit{weave} (\textit{PauseLo } \varphi) (\textit{PauseHi } \gamma) & = \\
& \textit{PauseLo} (\varphi \star_{\mathbf{K}} \lambda r_{lo}. \eta_{\mathbf{K}}(\textit{PauseHi} (v \star_{\mathbf{K}} \lambda r_{hi}. \eta_{\mathbf{K}}(\textit{weave } r_{lo} r_{hi}))))
\end{aligned}$$

5.3 Security for this setting: *take separation*

This section develops a non-interference style specification of process separation for monadic event systems. The question we answer is: given a resumption computation representing a schedule of threads on separated domains, how do we specify that those processes do not interfere? The answer we provide in this section adapts a technique for proving properties of streams to the resumption-monadic setting.

A common technique for proving a property of infinite lists is to show that the property holds of all finite approximations (i.e., finite initial prefixes) of the list. A well-known version of this technique is the *take lemma* [3, 9]:

$$s_1 = s_2 \Leftrightarrow \forall (n < \omega). \textit{take } n \ s_1 = \textit{take } n \ s_2$$

where (*take n [v₁, ..., v_n, ...]*) = [v₁, ..., v_n]. To show two streams s_1 and s_2 are equal using the take lemma, one shows that, for any non-negative integer n , each length n prefix of s_1 and s_2 are equal.

The security property proved here is analogous to the take lemma—hence its name. If, for any initial sequence of interleaved Hi and Lo events with n Lo events obtained from a system execution:

$$h_0 ; l_0 ; \dots ; h_n ; l_n$$

the effect of its execution on the Lo state should be identical to that of executing the Lo events in isolation:

$$l_0 ; \dots ; l_n$$

Using the denotational semantics, we make this notion precise:

$$\llbracket h_0 ; l_0 ; \dots ; h_n ; l_n \rrbracket \gg_{\mathbf{K}} \textit{maskNonLo} = \llbracket l_0 ; \dots ; l_n \rrbracket \gg_{\mathbf{K}} \textit{maskNonLo}$$

Here, *maskNonLo* is a stateful operation on the K monad which overwrites all non-Lo states in the K monad. The particular definition of *maskNonLo* differs as the monadic event system is refined, but for our case, it is merely

$liftHi(u_H (\lambda_.h_0))$ The operator $maskNonLo$ plays the rôle of \downarrow_{Lo} in a trace-based event system. We define two additional helper functions, $takeLo$ and run . The helper function $takeLo$ picks out the initial sequences containing n Lo events, and the run function projects schedulings in the resumption monad R back into the kernel monad K for execution:

$$\begin{array}{ll}
takeLo : Integer \rightarrow R() \rightarrow R() & run :: Ra \rightarrow Ka \\
takeLo 0 x & = Done() & run (Done v) & = \eta_K v \\
takeLo n (PauseLo \varphi) = & & run (PauseLo \varphi) & = \varphi \star_K run \\
\quad PauseLo (\varphi \star_K (\eta_K \circ (takeLo (n - 1)))) & & run (PauseHi \varphi) & = \varphi \star_K run \\
takeLo n (PauseHi \varphi) = & & & \\
\quad PauseHi (\varphi \star_K (\eta_K \circ (takeLo n))) & & &
\end{array}$$

We may now formulate and prove take separation for the Basic Separation system described in Section 5.2:

Theorem 5 (Take Separation). *Let $lo, hi \in Event$, then for all natural numbers n ,*

$$\begin{aligned}
run (takeLo n (withoutHi (evL lo))) &\gg; initH \\
&= run (takeLo n (withHi (evL lo) (evH hi))) \gg; initH
\end{aligned}$$

where $initH = liftHi(u_H (\lambda_.h_0))$ for any arbitrary fixed $h_0 \in H$.

Proof of Theorem 5 by induction on n . All binds (\star) and units (η) are in the K monad.

Case $n = 0$.

$$\begin{aligned}
run (takeLo 0 (withoutHi (evL lo))) &\gg; initH \\
\{\text{defn takeLo}\} &= run (Done ()) \gg; initH \\
&= run (takeLo 0 (withHi (evL lo) (evH hi))) \gg; initH
\end{aligned}$$

Case $n = k + 1$.

$$\begin{aligned}
run (takeLo (k + 1) (withHi (evL lo) (evH hi))) &\gg; initH \\
&= (l_1 \gg; h_1 \gg; \dots \gg; l_{(k+1)} \gg; h_{(k+1)} \gg; \eta()) \gg; initH \\
\{\text{right unit}\} &= l_1 \gg; h_1 \gg; \dots \gg; l_{(k+1)} \gg; h_{(k+1)} \gg; initH \\
\{\text{clobber}\} &= l_1 \gg; h_1 \gg; \dots \gg; l_{(k+1)} \gg; initH \\
\{l_{(k+1)} \# initH\} &= l_1 \gg; h_1 \gg; \dots \gg; initH \gg; l_{(k+1)} \\
\{\text{ind. hyp.}\} &= \underbrace{l_1 \gg; \dots \gg; h_1}_{h_i \text{ excised}} \gg; initH \gg; l_{(k+1)} \\
\{l_i \# initH\} &= l_1 \gg; \dots \gg; l_{(k+1)} \gg; initH \\
\{\text{right unit}\} &= l_1 \gg; \dots \gg; l_{(k+1)} \gg; \eta() \\
&= run (takeLo (k + 1) (withoutHi (evL lo)))
\end{aligned}$$

6 Conclusion

Monads and monad transformers provide a powerful framework for the structuring of high-assurance software systems because they permit an explicit yet flexible partitioning of the domains-of-effect within a computation. Pure, higher-order, typed languages that support this style of programming, such as Haskell, must be considered as modeling and implementation languages for complex, security-critical software systems.

Monads were originally introduced into pure, higher-order functional programming languages to allow principled, effectful programming. Monads permitted functional programmers to program imperatively without imperative languages. This work demonstrates that the monadic approach to effects has benefits far beyond the mere imitation of imperative-style; the precise scoping of effects provided by monads allows information-flow to be managed in a straightforward yet mathematically rigorous manner. How such scoping of effects within imperative languages might be achieved remains an open question.

This paper presents an overview of an approach for constructing and formally verifying domain separation for systems developed in pure, higher-order, typed languages such as Haskell. The approach relies fundamentally on the ability of the language to express all internal effects using monads and to express monad transformers as a functional construction. Though we use the pure functional language Haskell as our implementation language, we do not use laziness or type classes in any essential way; thus our techniques could be applied using the *pure* subset of ML. Note that the purity of the language *is* essential to our approach.

In this paper we have taken a basic model of concurrency, reflected it inside a programming language type system, and exploited it to build a rich model of provably non-interfering concurrent execution. We have refined this to the point where basic mechanisms for process communication are provided.

This work is being done as part of the Programatica project at OGI, where we are developing tools and techniques to support high-assurance software development in Haskell. These techniques are being applied in the validation of the OSKer security kernel—a novel high-assurance operating system.

The models presented in this paper have been developed in the context of the development of a multi-level secure operating system written in Haskell called OSKer. Although the verification of OSKer’s security properties are not complete, the challenges of that verification have inspired this formalism. In particular, this formalism is rich enough to be extended to realistic system architectures in which trusted processes (such as cryptographic servers) are allowed to reclassify information. When such extensions are introduced the verification of system assurance becomes more complex, and involves some local application of traditional methods, but the overall system validation is still modular and scalable.

Because monad transformers give a notion of composable formal specification [17, 15], monad structuring of the event system allows modification without complete re-verification. Scalable techniques for the development of complex systems with high confidence in their security properties remains a grand challenge

of computer science. We believe that the ultimate solution to this challenge will draw heavily from the theory and practice of pure, higher-order, typed languages.

References

1. D. Alexander, W. Arbaugh, M. Hicks, P. Kakkar, A. Keromytis, J. Moore, C. Gun-der, S. Nettles, and J. Smith. The switchware active network architecture, 1998.
2. W. R. Bevier. Kit: A study in operating system verification. *IEEE Transactions on Software Engineering*, 15(11):1382–1396, 1989.
3. Richard J. Bird. *Introduction to Functional Programming using Haskell*. Prentice-Hall Series in Computer Science. Prentice-Hall Europe, London, UK, second edition, 1998.
4. K. Birman, R. Constable, M. Hayden, C. Kreitz, O. Rodeh, R. van Renesse, and W. Vogels. The Horus and Ensemble projects: Accomplishments and limitations. In *Proceedings of the DARPA Information Survivability Conference & Exposition (DISCEX '00)*, Hilton Head, South Carolina USA, January 2000.
5. W.E. Boebert, W.D. Young, R.Y. Kain, and S.A. Hansohn. Secure ada target: Issues, system design, and verification. In *Proc. IEEE Symposium on Security and Privacy*, pages 176–183, 1985.
6. Gérard Boudol and Ilaria Castellani. Noninterference for concurrent programs. *Lecture Notes in Computer Science*, 2076:382+, 2001.
7. Koehn Claessen. A poor man’s concurrency monad. *Journal of Functional Programming*, 9(3):313–323, 1999.
8. Matthew Flatt, Robert Bruce Findler, Shriram Krishnamurthi, and Matthias Felleisen. Programming languages as operating systems (or revenge of the son of the lisp machine). In *International Conference on Functional Programming*, pages 138–147, 1999.
9. Jeremy Gibbons and Graham Hutton. Proof methods for structured corecursive programs. In *1st Scottish Functional Programming Workshop, Stirling, Scotland, August 1999*.
10. J. A. Goguen and J. Meseguer. Security policies and security models. In *Proceedings of the 1982 Symposium on Security and Privacy (SSP '82)*, pages 11–20, Los Alamitos, Ca., USA, April 1990. IEEE Computer Society Press.
11. J.A. Goguen and J. Meseguer. Unwinding and inference control. In *Proc. IEEE Symposium on Security and Privacy*, pages 75–86, 1984.
12. J.W. Gray III. Probabilistic interference. In *Proc. IEEE Symposium on Research in Security and Privacy*, pages 170–179, 1990.
13. J.T. Haigh and W.D. Young. Extending the non-interference version of MLS for SAT. In *Proc. IEEE Symposium on Security and Privacy*, pages 232–239, 1986.
14. Robert Harper, Peter Lee, and Frank Pfenning. The Fox project: Advanced language technology for extensible systems. Technical Report CMU-CS-98-107, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, January 1998. (Also published as Fox Memorandum CMU-CS-FOX-98-02).
15. William Harrison. *Modular Compilers and Their Correctness Proofs*. PhD thesis, University of Illinois at Urbana-Champaign, 2001.
16. Nevin Heintze and Jon G. Riecke. The SLam calculus: programming with secrecy and integrity. In ACM, editor, *Conference record of POPL '98: the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Diego, California, 19–21 January 1998*, pages 365–377, New York, NY, USA, 1998. ACM Press.

17. Sheng Liang. *Modular Monadic Semantics and Compilation*. PhD thesis, Yale University, 1998.
18. Sheng Liang, Paul Hudak, and Mark Jones. Monad transformers and modular interpreters. In *22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, 1995*, pages 333–343, New York, NY, USA, 1995. ACM Press.
19. E. J. McCauley and P. J. Drongowski. KSOS—the design of a secure operating system. In *Proc. AFIPS National Computer Conference*, volume 48, pages 345–353, 1979.
20. D. McCullough. Noninterference and the composability of security properties. In *Proc. IEEE Symposium on Security and Privacy*, pages 177–187, 1988.
21. J. McLean. A general theory of composition for trace sets closed under selective interleaving functions. In *Proc. IEEE Symposium on Research in Security and Privacy*, pages 79–93, 1994.
22. Eugenio Moggi. Computational lambda-calculus and monads. In *Proceedings 4th Annual IEEE Symp. on Logic in Computer Science, LICS'89, Pacific Grove, CA, USA, 5–8 June 1989*, pages 14–23. IEEE Computer Society Press, Washington, DC, 1989.
23. Eugenio Moggi. An abstract view of programming languages. Technical Report ECS-LFCS-90-113, Dept. of Computer Science, Edinburgh Univ., 1990.
24. P. G. Neumann, R. S. Boyer, R. J. Feiertag, K. N. Levitt, and L. Robinson. A provably secure operating system: The system, its applications, and proof. Technical Report CSL-116, SRI, May 1980.
25. Nikos S. Papaspyrou. A resumption monad transformer and its applications in the semantics of concurrency. In *Proceedings of the 3rd Panhellenic Logic Symposium*, Anogia, Crete, 2001.
26. G. D. Plotkin. A powerdomain construction. *SIAM Journal of Computing*, 5(3), 1976.
27. François Pottier and Vincent Simonet. Information flow inference for ML. In *Proceedings of the 29th ACM Symposium on Principles of Programming Languages (POPL'02)*, pages 319–330, Portland, Oregon, January 2002.
28. John Rushby. Noninterference, transitivity, and channel-control security policies. Technical report, SRI, dec 1992.
29. A. Sabelfeld and A. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 21(1), January 2003.
30. David A. Schmidt. *Denotational Semantics*. Allyn and Bacon, Boston, 1986.
31. Jonathan S. Shapiro, Jonathan M. Smith, and David J. Farber. EROS: a fast capability system. In *Symposium on Operating Systems Principles*, pages 170–185, 1999.
32. G. Smith. A new type system for secure information flow, 2001.
33. Geoffrey Smith and Dennis Volpano. Secure information flow in a multi-threaded imperative language. In *Conference Record of POPL 98: The 25TH ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Diego, California*, pages 355–364, New York, January 1998. ACM.
34. Dennis M. Volpano and Geoffrey Smith. A type-based approach to program security. In *TAPSOFT*, pages 607–621, 1997.
35. Philip Wadler. Theorems for free. In *Functional Programming Languages and Computer Architecture*, pages 347–359. ACM, 1989.
36. Philip Wadler. The Essence of Functional Programming. In *Proceedings of the 19th Symposium on Principles of Programming Languages*, pages 1–14, Albuquerque, New Mexico, January 19 – 22, 1992. ACM Press.

37. Bruce J. Walker, Richard A. Kemmerer, and Gerald J. Popek. Specification and verification of the ucla unix security kernel. *Communications of the ACM*, 23(2):118–131, 1980.
38. Mitchell Wand. Continuation-based multiprocessing. In J. Allen, editor, *Conference Record of the 1980 LISP Conference*, pages 19–28, Palo Alto, CA, 1980. The Lisp Company.
39. P. White, W. Martin, A. Goldberg, and F.S. Taylor. Formal construction of the mathematically analyzed separation kernel. In *The Fifteenth IEEE International Conference on Automated Software Engineering (ASE'00)*, pages 133–141, September 2000.
40. Aris Zakinthinos and E. S. Lee. The composability of non-interference. In *Proceedings of the Eighth Computer Security Foundations Workshop (CSFW '95)*, pages 2–8, Washington - Brussels - Tokyo, June 1995. IEEE.
41. Aris Zakinthinos and E. Stewart Lee. A general theory of security properties. In *Proceedings of the 18th IEEE Computer Society Symposium on Research in Security and Privacy*, 1997.

A Theorems and Proofs

Theorem 1 (StateT creates a state monad) *For any monad M , let monad $M' = \text{StateT } sto M$ and also $u : (sto \rightarrow sto) \rightarrow M'()$ and $g : M' sto$ be the non-proper morphisms added by $(\text{StateT } sto)$. Then $\langle M', \eta_{M'}, \star_{M'}, u, g, sto \rangle$ is a state monad.*

Proof.

Case: Sequencing.

$$\begin{aligned}
& \mathbf{u}f \star_{M'} \lambda_{\cdot} \mathbf{u}g \\
&= \lambda \sigma_0. (\mathbf{u}f) \sigma_0 \star_M (\lambda \langle v, \sigma_1 \rangle. (\lambda_{\cdot} \mathbf{u}g) v \sigma_1) && \{\text{def. } \star_{M'}\} \\
&= \lambda \sigma_0. (\mathbf{u}f) \sigma_0 \star_M (\lambda \langle v, \sigma_1 \rangle. (\mathbf{u}g) \sigma_1) && \{\beta\} \\
&= \lambda \sigma_0. (\lambda \sigma. \eta_M \langle () , f \sigma \rangle) \sigma_0 \star_M (\lambda \langle v, \sigma_1 \rangle. (\lambda \sigma. \eta_M \langle () , g \sigma \rangle) \sigma_1) && \{\text{def. } \mathbf{u}(\times 2)\} \\
&= \lambda \sigma_0. (\eta_M \langle () , f \sigma_0 \rangle) \star_M (\lambda \langle v, \sigma_1 \rangle. (\eta_M \langle () , g \sigma_1 \rangle)) && \{\beta(\times 2)\} \\
&= \lambda \sigma_0. (\eta_M \langle () , g(f \sigma_0) \rangle) && \{\text{left unit}\} \\
&= \mathbf{u}(f; g) && \{\text{def. } \mathbf{u}(\times 2)\}
\end{aligned}$$

Case: Cancellation.

$$\begin{aligned}
& \mathbf{g} \star_{M'} \lambda s. \mathbf{u}(\lambda_{\cdot} s_0) \\
&= \lambda \sigma. \eta_M \langle \sigma, \sigma \rangle \star_M \lambda \langle v, \sigma_1 \rangle. (\lambda s. \lambda \sigma'. \eta_M \langle () , s_0 \rangle) v \sigma_1 && \{\text{def. } \mathbf{g}, \mathbf{u}, \star_M\} \\
&= \lambda \sigma. \eta_M \langle \sigma, \sigma \rangle \star_M \lambda \langle v, \sigma_1 \rangle. (\eta_M \langle () , s_0 \rangle) && \{\beta\} \\
&= \lambda \sigma. (\eta_M \langle () , s_0 \rangle) && \{\text{left unit}\} \\
&= \mathbf{u}(\lambda_{\cdot} s_0) && \{\text{def. } \mathbf{u}\}
\end{aligned}$$

□

Theorem 3 (StateT preserves stateful behavior) *For any state monad $M = \langle M, \eta, \star, \mathbf{u}, \mathbf{g}, sto \rangle$, the following state monad structure is a state monad:*

$$\langle \text{StateT } s M, \eta', \star', (\mathbf{u}; \text{lift}), \text{lift}(\mathbf{g}) \rangle$$

where η' , \star' , and lift are the monadic unit, bind, and lifting operations, respectively, defined by $(\text{StateT } s)$ in Section 3.1.

Proof. Let $M' = \top M$ below. Each step follows from the *Lifting Laws* listed in Section 3.1.

Case: Sequencing.

$$\text{lift}(uf) \star_{M'} \lambda_{\cdot} \cdot \text{lift}(ug) = \text{lift}(uf \star_M \lambda_{\cdot} \cdot ug) = \text{lift}(u(f; g))$$

Case: Cancellation.

$$\text{lift}(g) \star_{M'} \lambda s \cdot \text{lift}(u(\lambda_{\cdot} \cdot s_0)) = \text{lift}(g \star_M \lambda s \cdot u(\lambda_{\cdot} \cdot s_0)) = \text{lift}(u(\lambda_{\cdot} \cdot s_0))$$

□

Theorem 4 (Non-interference of updates) *Define state monads M and M' as the following:*

$$\begin{aligned} M &= \langle M, \eta_M, \star_M, u_A, g_A, A \rangle \\ M' &= \langle \text{StateT } B M, \eta_{M'}, \star_{M'}, u_B, g_B, B \rangle \end{aligned}$$

Then, for all $\delta_A : A \rightarrow A$ and $\delta_B : B \rightarrow B$,

$$(u_B \delta_B) \#_{M'()} \text{lift}(u_A \delta_A)$$

Proof.

$$\begin{aligned} &(u_B \delta_B) \gg_{M'} \text{lift}(u_A \delta_A) \\ &= \lambda \sigma_0. ((u_B \delta_B) \sigma_0) \star_M \lambda \langle \cdot, \sigma_1 \rangle. (\lambda_{\cdot} \cdot \text{lift}(u_A \delta_A)) \langle \cdot \rangle \sigma_1 \\ &= \lambda \sigma_0. ((u_B \delta_B) \sigma_0) \star_M \lambda \langle \cdot, \sigma_1 \rangle. (\text{lift}(u_A \delta_A)) \sigma_1 \\ &= \lambda \sigma_0. ((\lambda \sigma. \eta_M \langle \cdot, \delta_B \sigma \rangle) \sigma_0) \star_M \lambda \langle \cdot, \sigma_1 \rangle. (\text{lift}(u_A \delta_A)) \sigma_1 \\ &= \lambda \sigma_0. (\eta_M \langle \cdot, \delta_B \sigma_0 \rangle) \star_M \lambda \langle \cdot, \sigma_1 \rangle. (\text{lift}(u_A \delta_A)) \sigma_1 \\ &= \lambda \sigma_0. (\text{lift}(u_A \delta_A)) (\delta_B \sigma_0) \\ &= \lambda \sigma_0. (\lambda \sigma. (u_A \delta_A) \star_M \lambda v. \eta_M \langle v, \sigma \rangle) (\delta_B \sigma_0) \\ &= \lambda \sigma_0. (u_A \delta_A) \star_M \lambda v. \eta_M \langle v, \delta_B \sigma_0 \rangle \\ &= \lambda \sigma_0. (u_A \delta_A) \star_M \lambda v. (\lambda \sigma. \eta_M \langle v, \delta_B \sigma \rangle) \sigma_0 \\ &= \lambda \sigma_0. (u_A \delta_A) \star_M \lambda v. (u_B \delta_B) \sigma_0 \\ &= \lambda \sigma_0. (\lambda \sigma. (u_A \delta_A \star_M \lambda_{\cdot} \cdot \eta_M \langle \cdot, \sigma \rangle) \sigma_0) \star_M \lambda \langle v, \sigma \rangle. (u_B \delta_B) \sigma \\ &= \lambda \sigma_0. (\text{lift}(u_A \delta_A)) \sigma_0 \star_M \lambda \langle v, \sigma \rangle. (u_B \delta_B) \sigma \\ &= (\text{lift}(u_A \delta_A)) \star_M \lambda_{\cdot} \cdot (u_B \delta_B) \end{aligned}$$

□

Theorem 5 (Non-interference Preserved by Monad Transformers) *Let M be a monad with two operations, $a : M()$ and $b : M()$ such that $a \#_M b$. Then,*

$$(\text{lift } a) \#_{(\top M)} (\text{lift } b)$$

where \top is a monad transformer and $\text{lift} : M a \rightarrow (\top M) a$.

Proof.

$$\begin{aligned} & \text{lift } a \star_{(\text{TM})} \lambda. \text{lift } b \\ &= \text{lift}(a \star_{\text{M}} \lambda. b) \\ &= \text{lift}(b \star_{\text{M}} \lambda. a) \\ &= \text{lift } b \star_{(\text{TM})} \lambda. \text{lift } a \end{aligned}$$

□