

A Core Calculus for Secure Hardware: Its Formal Semantics and Proof System

Thomas N. Reynolds^{*}, Adam Procter[†], William L. Harrison^{*}, Gerard Allwein[‡]

^{*}Department of Computer Science, University of Missouri

[†]Intel Corporation, Hillsboro, Oregon, USA

[‡]US Naval Research Laboratory, Washington, DC

Abstract—Constructing high assurance, secure hardware remains a challenge, because to do so relies on both a verifiable means of hardware description and implementation. However, production hardware description languages (HDL) lack the formal underpinnings required by formal methods in security. Still, there is no such thing as high assurance systems without high assurance hardware. We present a core calculus of secure hardware description with its formal semantics, security type system and mechanization in Coq. This calculus is the core of the functional HDL, ReWire, shown in previous work to have useful applications in reconfigurable computing. This work supports a full-fledged, formal methodology for producing high assurance hardware.

I. INTRODUCTION

ReWire is a functional hardware description language (HDL): it is a functional language—a subset of Haskell—from which circuits are synthesized automatically. Previous work has introduced ReWire’s language design and implementation as well as its application to the construction of high assurance hardware [1]–[4]. This paper describes the Coq formalization of ReWire intended to support the verification of hardware designs and, in particular, the sort of information flow properties described in our previous work [1], [2], [5].

The aforementioned previous work, in panoramic view, used “by-construction” properties of layered monads to verify properties by hand. For the moment, we rely on the reader’s intuition to explain the contributions of the present work at a high level (Section II presents an overview of these concepts in more detail). Assume, for example, that ReWire devices h and l are written respectively in terms of state monad layers, StT Hi and StT Lo . Then, device h (resp., l) only accesses internal storage of type Hi (resp., Lo). In a composite device written in terms of monad $\text{M} = \text{StT Hi} (\text{StT Lo Id})$, it is guaranteed by semantic properties of the layers StT Hi and StT Lo to disallow covert channels between the Hi and Lo storage.

The challenge is, then, the formalization of ReWire and, in particular, ReWire’s underlying layered monad language and its semantic properties within an automated proof system. The contributions of this work are as follows. (1) A static effect-type system (extending Wadler [6]) that disallows covert storage channels in ReWire. This type system extends state layers with effect labels, so that, continuing the example above, h (resp., l) is written in monad $\text{StT RW Hi} (\text{StT } \langle \rangle \text{ Lo Id})$ (resp., $\text{StT } \langle \rangle \text{ Hi} (\text{StT RW Lo Id})$). The effect label “RW” means h can both read and write on the Hi layer and while “ $\langle \rangle$ ” means it

can do neither on the Lo layer (and, *vice versa*, for l). The soundness of our type system (Theorems 6 and 7) guarantees freedom from covert storage channels. (2) A small-step semantics for ReWire formalized in Coq that justifies (3) a typed equational logic (Figure 11) capturing the semantic properties of monads and state layers used in by-hand proofs in our previous work. Finally, (4) a number of related metatheorems (e.g., progress, preservation, strong normalization, etc.) have been proved in Coq. All of the definitions and theorems in this paper have been checked with the Coq proof checker (Coq scripts are available [here](#)).

The most direct approach to formalizing ReWire in Coq would seem to be the transliteration of monad transformer declarations from Haskell into Coq, but this quickly runs afoul of Coq’s strict positivity requirement. ReWire relies on reactive resumption monad transformers (see Section II) as a model of synchronous parallelism and this transformer is a coinductive construction, which can be tricky to formalize, even with Coq’s coinduction library. Another approach considers formalizing ReWire’s denotational semantics [7], building on existing work by Huffman [8] or Schröder and Mossakowski [9] in Isabelle/HOLCF. Instead, we chose to formalize a small-step, operational semantics for ReWire in Coq, in part, because the authors have more experience with Coq than with HOLCF, but also because developing and formalizing a small-step operational semantics seemed more straightforward than mechanizing denotational semantics. The semantic properties of ReWire’s underlying monads on which the by-hand verifications of our previous work rely are then captured as a typed equational logic whose rules are derived from the formalized operational semantics.

The remainder of this section discusses related work. Section II presents an overview of ReWire to motivate the syntax and semantics of the formal calculus, RWC. Section III defines the syntax and small-step operational semantics of RWC. Section IV describes RWC’s metatheory and a number of related metatheorems (e.g., progress, preservation, strong normalization, etc.) are demonstrated. A type-directed equational logic for RWC is defined in Section V and Section VI discusses conclusions and future work.

Related Work

Effect systems are a static semantics of effects while monads [10] are a dynamic semantics of effects. Effect sys-

tems [11] are commonly associated with impure, strongly-typed functional languages in which the effect annotations make explicit the side effects already present implicitly in the language itself. Monads are used to mimic side-effecting computations within pure, strongly-typed functional languages (e.g., Haskell) in which there are no implicit side effects.

Layered monads—i.e., monads constructed by monad transformers [12]—provide modularity to the semantics of computational effects and functional programs alike by integrating multiple effects within a single monad. This modularity-via-integration, however, has consequences for formal verification: because its effects are all encapsulated within the single monad, they are not distinguished syntactically within the type system of a specification language itself. Wadler [6] “married” effect types to monads, and previous work by the authors [13] seems to be the first marriage of effect types to layered monads. This latter marriage seems to be important for exploiting monadic semantics in formal methods: layered monads provide a modular semantics of effects including by-construction properties and effect types allow the expression of these properties in a formal proof system like Coq (e.g., Figure 11).

As a concept for formal (i.e., machine-checked) verification, monads are less common, although not unheard of [9], [14]–[16] and the use of both effect types and layered monads distinguishes the current work from these.

Formal methods for secure hardware are generally spread across two categories: (1) type-based approaches [17]–[19]; and (2) logic-based approaches (including theorem-proving [20], and BDDs and model-checking [21]), in which a hardware design and desired properties are formulated in a logic and scrutinized in a (semi-)automatic manner. Types-based approaches have support for security concerns integrated into a domain-specific language for hardware description. As with any security type system, however, the question of its permissiveness arises—i.e., does it reject secure designs? The types-based approach offers no recourse to the rejection of a secure design—you simply can’t argue with a type checker. A logic-based approach avoids this pitfall, but comes with overhead—e.g., your own theory of security—and, furthermore, it is not connected directly to any implementation path.

One language-based approach to hardware security is to extend an existing HDL with security types. Caisson [17], Sapper [18], and SecVerilog [19] each extend a subset of Verilog with security types and annotations. The type systems of Caisson and SecVerilog reject programs that violate information flow policies, while Sapper uses static analysis to automatically insert dynamic checks to enforce information flow policies at runtime. SecVerilog has an operational semantics, albeit not one formalized in a theorem prover with a proof system [22]. ReWire (or, RWC, rather) differs fundamentally from these language- and type-based approaches in three respects: (1) it is a pure functional language; (2) it possesses a formal semantics mechanized in Coq; and (3) its type system is based on effect types. We discuss the significance of item (3) in Section VI.

The SAFE project focuses on the clean slate design of a provably secure computer system stack (e.g., hardware, operating system, etc.). In a recent publication [23], the SAFE team describes an operational semantics of the SAFE hardware’s instruction set and its role in the end-to-end verification in Coq of a non-interference security property. The ReWire project has complimentary, but orthogonal, goals to SAFE: developing a verifiable toolchain for producing high assurance, secure hardware. Interesting follow-on research would explore implementations of the SAFE hardware in the ReWire language.

One traditional approach to hardware verification starts from a design expressed in a production HDL, creates an abstract specification “by hand” as it were, encodes this specification in the logic of an automated theorem prover, and proceeds towards formal verification [20]. This approach relies heavily on the faithfulness of the abstraction step. One reason that this approach must be accomplished “by hand” is that production HDLs do not possess rigorous semantics. Although attempts have been made in the past to define them semantically, none of these projects were evidently completed [24], [25]. By contrast with production HDLs like Verilog or VHDL, ReWire possesses a rigorous semantics for which the present work provides a Coq mechanization. ReWire becomes a vehicle for expressing and implementing hardware designs and for verifying them as well. In previous work [1], [2], we presented several case studies in hardware verification based in ReWire, but there the verifications were not machine-checked.

Goncharov and Schröder [26] extend Moggi’s computational λ -calculus with constructs for concurrency and shared state; RWC’s design is inspired, in part, by their treatment of corecursion. Crary et al. [27] consider a logical characterization of information flow security that incorporates Moggi’s computational λ -calculus at its core. With their approach, monads are, in effect, logical modalities signifying the potential presence of effects at a security level. In contrast, Harrison and Hook’s treatment of information flow security [5] is more semantic and model-theoretic than Crary’s logical and type-theoretic approach, relying as it does on structural properties of monads and monad transformers to construct secure systems. Security verifications of ReWire designs [1] are based on Harrison and Hook’s approach, and the present work formally supports that approach in Coq.

Ghica and Jung [28] provide a categorical semantics for a class of digital circuits in terms of monoidal categories and are motivated by the need for supporting syntactic, equational reasoning. ReWire specifications may be reasoned about equationally in the usual manner of functional languages; this was the approach taken in our previous ReWire verification work [1], [2]. By contrast with Ghica and Jung’s work, ReWire specifications are, more or less, ordinary functional programs that are compiled into circuits.

II. BACKGROUND: REWIRE’S PROGRAMMING MODEL


The purpose of this section is twofold: (1) to make this article as self-contained as possible by providing sufficient background on ReWire and (2) to motivate RWC’s type

system and operational semantics. Throughout this section, we explicitly link this background material to subsequent sections on RWC. ReWire is a subset of Haskell and uses ideas from monadic semantics as an organizing principle of the language. It is, therefore, assumed of necessity that the reader is, at least, somewhat familiar with functional programming and monads. For those requiring more information about ReWire and its programming model and language design, please consult the references [1], [7].

ReWire is a subset of the Haskell functional programming language [29]—i.e., ReWire programs are Haskell programs, but not necessarily *vice versa*. All ReWire programs can be compiled to synthesizable VHDL with the ReWire compiler. The principal difference between Haskell and ReWire is that recursion in ReWire is restricted to tail recursion so that every ReWire program requires only a finite, bounded memory footprint. Unbounded recursion requires an unbounded stack or heap for compilation and such dynamic control structures are anathema to hardware’s fixed storage.

ReWire has type constructors for devices where a *device* represents a clocked computation that, for each clock cycle, takes an input of type i , produces an output of type o , and may possess internal storage of type s (see inset figure).

The type of d as shown would be $d :: \text{ReT } i \circ (\text{StT } s \text{ Id}) ()$, where ReT and StT are the reactive resumption and state monad transformers and Id is the identity monad (about all of which we say more below in the next sections). Device d is clocked, as illustrated in the inset figure, although the clock is represented by the underlying structure of reactive resumptions rather than as an explicit parameter. A device is created in ReWire by either iterating a function or through composition of existing devices. Previous work [4] introduced operators for constructing devices and composing them into larger, interconnected devices; Section II-C presents a simple device specification template in ReWire.



Device d

A. Background: Monads

A monad is a triple $\langle M, \text{return}, \gg= \rangle$ consisting of a type constructor M and two operations:

$$\begin{aligned} \text{return} &: a \rightarrow M a && \text{--- “unit”} \\ (\gg=) &: M a \rightarrow (a \rightarrow M b) \rightarrow M b && \text{--- “bind”} \end{aligned}$$

These operations must obey the well-known *monad laws* [10], [12] (these are (LEFT-UNIT), (RIGHT-UNIT), and (ASSOCIATIVITY) in Figure 11). The return operator is the monadic analogue of the identity function, injecting a value into the monad. The $\gg=$ operator is a form of sequential application. The “null bind” operator, $\gg : M a \rightarrow M b \rightarrow M b$, is defined as: $x \gg k = x \gg= \lambda_.k$. The binding (i.e., “ $\lambda_$ ”) acts as a dummy variable, ignoring the value produced by x .

B. Background: Monad Transformers

The organizing principle underlying ReWire are *reactive resumption monads with state* [30] (RRMS), which encapsulate

a notion of computation appropriate to hardware—namely, synchronous parallelism. RRMS support the expression of structural hardware designs in a functional style [4]. RWC is a computational λ -calculus whose syntax and semantics formalizes RRMS in Coq. In particular, RWC’s type system includes constructors that correspond to the state and reactive resumption monad transformers. For the sake of being self-contained, we provide the reader with Haskell definitions of the StT and ReT monad transformers. This code is meant only to aid the reader in comprehending the intended semantics of RWC. If more background is required on RRMS, please consult the references [30].

1) *State Monad Transformer*: The state monad transformer is a well-documented structure in functional programming and semantics [12]. The Haskell code for the state monad transformer, StT , along with its lifting functions is below:

```
data StT s m a = StT (s -> m (a,s))
liftstT :: m a -> StT s m a
liftstT m
  = StT (\ s -> m \>>= \ v -> returnm (v,s))
get :: StT s m s
get = StT (\ s -> returnm (s,s))
put :: s -> StT s m ()
put s = StT (\ _ -> returnm ((),s))
```

The lift converts or “lifts” an $m a$ computation into an $\text{StT } s \text{ m } a$ computation. The get operation returns the current value of the s -store while the $\text{put } s$ operation replaces the current store with store s . In the definitions above, the binds and returns for the m monad are affixed with a subscript to disambiguate them from the operations being defined.

2) *Reactive Resumption Monad Transformer*: Computations in $\text{ReT } i \circ m a$ may be viewed intuitively as (potentially infinite) sequences of m computations. If that sequence terminates, it produces an a -value, otherwise it produces an o -output value and a continuation. Both lift operations convert an m computation into respective enriched computations. Computations in ReT over layered state monads correspond closely to synchronous hardware as discussed in previous work [1].

The Haskell code for the reactive resumption monad transformer, ReT , along with its associated functions is below:

```
data ReT i o m a
  = Pause (m (Either a (o,i -> ReT i o m a)))
liftreT :: m a -> ReT i o m a
liftreT m = Pause (m \>>= returnm . Left)
signal :: o -> ReT i o m i
signal o = Pause (returnm
  (Right (o,returnm . returnreT)))
data Either a b = Left a | Right b
```

Recall that function composition (i.e., “ $.$ ”) and sum types (i.e., Either) are built-in to Haskell. In terms of the device d example above, the operation $\text{signal } o$ represents the end of a clock cycle and sets the output signal of d to o . RWC includes a pause primitive in the term syntax (Fig. 2) as a means of representing signal .

3) *By-construction Properties of Layered Monads*: Layered state monads—monads with multiple StT applications (e.g., $M = \text{StT } s_1 (\text{StT } s_2 \text{ Id})$)—have a number of useful properties by construction [5], including:

```

put s' >> put s    = put s
put s >> liftStT φ = liftStT φ >> put s

```

The first rule is an inter-layer property (a.k.a., “clobber”) while the second is an inter-layer property (a.k.a., “atomic non-interference”). Clobber states that the `put s` cancels earlier effects on the same layer. By convention for a fixed state s_0 , we define $\text{mask} = \text{put } s_0$; the mask included in the term syntax of RWC generalizes this idea. Atomic non-interference states that effects from different state layers commute. The equational logic derived in Coq for RWC presented in Section V gives generalizations of both properties.

C. Defining Devices in ReWire

Simple ReWire devices are generally defined as tail recursive functions whose codomain is written in terms of the ReT layer. Assume that we have functions defined which specify the internal and external behaviors of device d that have function types: `internal :: i -> StT s Id v` and `external :: i -> v -> o`. Function `internal` takes the input i , performs some computation with the current internal storage s , and produces an intermediate result v . Function `external` takes the input i and the result v and produces the next output signal for d .

Given an initial input i_0 , $d = \text{dev } i_0$ where corecursive function `dev` is defined as:

```

dev :: i -> ReT i o (StT s Id) ()
dev i = liftReT (internal i) >>= \ v ->
  signal (external i v) >>= \ i' ->
  dev i'

```

At the beginning of a clock cycle, `dev` first consumes input, i , then performs `internal i` computation on the internal storage s , and then outputs the `external i v` signal at the end of the clock cycle.

Device definitions are expressed with an explicit corecursion operator, `unfold`; for example, device d would be written:

```

unfold i_0
  (\ i -> internal i >>= \ v ->
    return (Right (external i v, id)))

```

For this reason, Figure 2 includes syntax for an `unfold` primitive and its semantics are defined in subsequent sections.

D. Background: Goguen-Meseguer Non-interference

The essence of the Goguen-Meseguer noninterference information flow model [31] and its many descendants is that systems, broadly construed, are state machines whose inputs and outputs are partitioned by security level. The definition of information flow is formulated in terms of sequences of stateful operations of mixed security levels and stipulates that high-level operations must not affect low-level outputs. More concretely, for any mixed-level sequence, $s = (l_1; h_1; \dots; l_n; h_n)$, the low-level outputs of s must be identical to those produced by $(l_1; \dots; l_n)$, which is the result of filtering out from s all high-level operations.

E. The Marriage of Effects and Layered State Monads

“By construction” properties of layered state monads [5] tell us that high- and low-security operations commute (a.k.a., atomic non-interference) and that mask_H cancels high-level operations (i.e., $\varphi_H \gg \text{mask}_H = \text{mask}_H$). This cancelling property is known as the “clobber rule” [5]. The atomic non-interference and clobber rules are helpful in demonstrating that monadic noninterference equations (like that of the previous section) hold for particular software and hardware applications [1], [5].

The Goguen-Meseguer model was recast in monadic terms previously [5], so that high-level effects must be cancellable without affecting the low-level effects. Here, the utility of the RWC effect type system becomes evident, because it can statically distinguish computations occurring on distinct layers. For the sake of concreteness, consider the case of a monad, M , with a high- and low-security stores types, H and L . High and low operations may be distinguished by the RWC effect type system by annotating the layers with effect labels:

$$\begin{aligned} \varphi_H &: \text{StT RW } H (\text{StT } \langle \rangle L \text{Id}) () \\ \varphi_L &: \text{StT } \langle \rangle H (\text{StT RW } L \text{Id}) () \end{aligned}$$

Note that φ_H (resp., φ_L) only has read-write effects (RW) on the outer (resp., inner) state layer of M . Furthermore, we assume the existence of an operation, mask_H which initializes the H state layer. The mask_H operation can be assumed to be `put s0` on the H -layer, where s_0 is an arbitrary, fixed value in H . Then, the monadic formulation of non-interference boils down to demonstrating that equations like the following hold:

$$\varphi_H \gg \varphi_L \gg \text{mask}_H = \varphi_L \gg \text{mask}_H$$

Put simply, this means that reinitializing the H layer cancels the effects of high-security operations like φ_H . This is the monadic analogy of Goguen and Meseguer’s filtering out of high-security operations.

III. RWC: THE REWIRE CORE CALCULUS

This section introduces the syntax (Section III-A), type system (Section III-B) and operational semantics (Section III-C) of the ReWire Calculus (RWC). RWC is a computational λ -calculus that extends the functional features of a typed lambda calculus with support for stateful effects and reactive parallelism. These effects are encapsulated through the use of *monads* [10], enabling us to provide a useful equational theory in the presence of effects. The addition of effects to a computational λ -calculus was examined in [6].

A. Syntax

This section introduces the syntax of RWC, which is a variety of computational λ -calculus extended with operations for synchronous, stateful parallelism. Here, the stateful component is organized as *layered* state monads—i.e., monads created by multiple applications of the state monad transformer. Layered state monads have by-construction properties that support information flow security verification [1], [5]; we defer presenting the general formalization of these by-construction

$$\begin{aligned}
\ell \in \text{EffectLabel} &::= \langle \rangle \mid R \mid W \mid RW \\
S \in \text{StateMonad} &::= \text{Id} \mid \text{StT } \ell \tau S \\
M \in \text{Monad} &::= S \mid \text{ReT } \tau \tau' S \\
\tau, \tau' \in \text{Type} &::= \tau \rightarrow \tau' \mid \tau \times \tau' \mid \tau + \tau' \mid () \mid M \tau
\end{aligned}$$

Figure 1: Syntax of RWC types

$$\begin{aligned}
\text{Identifier} &::= x \mid y \mid z \mid w \mid \text{etc.} \\
t \in \text{Term} &::= x \mid t t' \mid \lambda x : \tau. t \mid () \mid \langle t, t' \rangle \mid \text{proj } t t' \mid \\
&\quad \mid \text{inl}_\tau t \mid \text{inr}_\tau t \mid \text{case } t t' t'' \mid \text{return}_M t \mid t \gg= t' \\
&\quad \mid \text{lift}_M t \mid \text{elevate}_S t \mid \text{get}_S \mid \text{put } t \mid \text{pause}_{M, \tau} t \\
&\quad \mid \text{runSt } t t' \mid \text{runld } t \mid \text{unfold}_{M, \tau, \tau'} t t' \mid \text{runRe}_\tau t \\
v, s \in \text{Value} &::= \lambda x : \tau. t \mid () \mid \langle v, v' \rangle \mid \text{inl}_\tau v \mid \text{inr}_\tau v \mid \text{return}_M v \\
&\quad \mid v \gg= v' \mid \text{lift}_M v \mid \text{elevate}_M v \mid \text{get}_S \mid \text{put } v \\
&\quad \mid \text{pause}_{M, \tau} v \mid \text{runSt } v v' \mid \text{runRe}_\tau v \\
&\quad \mid \text{unfold}_{M, \tau, \tau'} v v' \\
\Sigma \in \text{Store} &::= \text{nil} \mid s :: \Sigma \\
c \in \text{Config} &::= \langle t, \Sigma \rangle \\
D \in \text{DoneConfig} &::= \langle \text{return}_M v, \Sigma \rangle \mid \langle \text{pause}_{M, \tau} v, \Sigma \rangle
\end{aligned}$$

Figure 2: Syntax of terms, stores, and configurations

properties until Section V. Section II provides the reader with some background on monad transformers, although readers requiring more should consult the references.

1) *Types*: Figure 1 shows the syntax of types.

As a computational λ -calculus, RWC extends the simply-typed λ -calculus with unit, sum, and product types along with a notion of *computational* types: if M is a monad and τ is a type, then $M \tau$ is the type of computations in the monad M with a result value of type τ . Exactly which monad stands in for M will determine what sort of computational effects are possible. RWC permits the use of monads built in terms of the Id (identity) monad and the ReT (reactive resumption), and StT (state) monad transformers, where ReT must be the outermost monad transformer application (if it is present). RWC’s monads encompass the combination of resumption and layered state monads found in [30] with the addition of *effect labels* attached to each StT . The presence of an effect label ℓ at a given layer certifies that the computation has at *most* the effects ℓ at that layer. For example, the effect label W reflects the *possibility* that a computation will write, not the necessity, and certifies that the computation will *not* read.

We note in passing that the *denotational* semantics of these monads corresponds exactly to the semantics of their Haskell equivalents, up to the erasure of the effect labels and with the considerable simplification that lifted domains are not necessary due to the absence of general recursion; see [7] for further details.

2) *Terms*: Figure 2 shows the syntax of terms. Note the widespread use of type and monad subscripts. These are necessary to ensure that every term has a unique type, and to handle overloading of monadic operations. We will sometimes

omit these subscripts, as long as doing so does not introduce ambiguity.

We will not remark on the standard λ -calculus machinery, other than to note that the constructs used for destructing pairs and elements of sum type are slightly nonstandard. The term constructor proj , used for destructing pairs, takes two subterms: the first corresponding to the pair being deconstructed—suppose it has type $\tau \times \tau'$ —and the second corresponding to a *function* of type $\tau \rightarrow \tau' \rightarrow \tau''$ that produces a value from the pair’s elements. (Note that the conventional left- and right-projection operators can be constructed in terms of the proj operator.) The term constructor case , used for destructing elements of sum type, takes three subterms: the first is the scrutinee of type $\tau + \tau'$, the second to a function f_1 of type $\tau \rightarrow \tau''$, and the third to a function f_2 of type $\tau' \rightarrow \tau''$. If the scrutinee evaluates to $\text{inl } v$ (resp., $\text{inr } v$), then v will be passed to f_1 (resp., f_2).

Computations are defined in terms of certain primitives. The (overloaded) term constructors return and \star correspond respectively to the unit and bind operations of the monads, and lift to the lift operation of each monad transformer. Terms typed in a state monad may read and write to the store using the get and put operations. The term constructor elevate adds effect labels – e.g., W or R – to the effect labels, if any, on a state monad computation; thereby, converting state monad computations with a less permissive types to a more permissive type (where “permissiveness” is understood as in Figure 4). For example, a term t of type $\text{StT } R \tau \text{Id } \tau'$ can be typecast into the more permissive type $\text{StT } RW \tau \text{Id } \tau'$ via elevate , essentially de-certifying that t does not write. (A cast in the “other direction”, to $\text{StT } \langle \rangle \tau \text{Id } \tau'$, is not permitted by the type system.) Reactive computations are defined in terms of the primitives pause and unfold . The term $\text{pause } t$ is essentially a suspended computation that is waiting for an input value, and unfold can be used to produce “looping” computations; we postpone a discussion of their exact semantics until we have discussed the type system in greater detail. Finally, the term constructors runRe , runSt , and runld allow the effects of a given monad transformer to be reflected into the base monad. It may be helpful to view runRe as executing a single step of a resumption-monadic computation, runSt as supplying the initial state for the uppermost state layer, and runld as moving from the effect-free Id monad into the universe of non-monadic terms.

3) *Stores and Configurations*.: Figure 2 (bottom) shows the syntax of *stores* and *configurations*, which will be used to specify the semantics of computations. A store is a list of terms, each of which corresponds semantically to a state monad transformer, and a configuration $\langle t, \Sigma \rangle$ pairs a term t with a state Σ . Generally, we use the metavariables s, s', s'' to refer store values.

B. Type System

Typing rules for terms are given in Figure 3. Typing judgments take the form $\Gamma \vdash t : \tau$, where Γ is a set of assumptions (i.e., a mapping of variables to types). For the

$$\begin{array}{c}
\frac{\Gamma \vdash t : \tau}{\Gamma \vdash \text{return}_M t : M \tau} \text{(RETURN)} \quad \frac{\Gamma \vdash t : M \tau \quad \Gamma \vdash t' : \tau \rightarrow M \tau'}{\Gamma \vdash t \gg t' : M \tau'} \text{(BIND)} \\
\frac{\Gamma \vdash t : S \tau}{\Gamma \vdash \text{lift}_{(\text{StT} \ell \tau' S)} t : \text{StT} \ell \tau' S \tau} \text{(LIFTST)} \\
\frac{\Gamma \vdash t : S \tau}{\Gamma \vdash \text{lift}_{(\text{ReT} \tau' \tau'' S)} t : \text{ReT} \tau' \tau'' S \tau} \text{(LIFTRE)} \\
\frac{R \leq \ell}{\Gamma \vdash \text{get}_{(\text{StT} \ell \tau S)} : \text{StT} \ell \tau S \tau} \text{(GET)} \quad \frac{\Gamma \vdash t : \tau \quad W \leq \ell}{\Gamma \vdash \text{put} t : \text{StT} \ell \tau S ()} \text{(PUT)} \\
\frac{\Gamma \vdash t : \text{StT} \ell \tau' S \tau \quad \Gamma \vdash t' : \tau'}{\Gamma \vdash \text{runSt} t t' : S (\tau \times \tau')} \text{(RUNST)} \quad \frac{\Gamma \vdash t : \text{Id} \tau}{\Gamma \vdash \text{runId} t : \tau} \text{(RUNID)} \\
\frac{\Gamma \vdash t : S (\tau' \times (\tau \rightarrow \text{ReT} \tau \tau' S \tau''))}{\Gamma \vdash \text{pause}_{(\text{ReT} \tau \tau' S, \tau'')} t : \text{ReT} \tau \tau' S \tau''} \text{(PAUSE)} \\
\frac{\Gamma \vdash t : \tau''' \quad \Gamma \vdash t' : \tau'''' \rightarrow S (\tau'' + (\tau' \times (\tau \rightarrow \tau''')))}{\Gamma \vdash \text{unfold}_{(\text{ReT} \tau \tau' S, \tau'', \tau''')} t t' : \text{ReT} \tau \tau' S \tau''} \text{(UNFOLD)} \\
\frac{\Gamma \vdash t : \text{ReT} \tau \tau' S \tau''}{\Gamma \vdash \text{runRe}_\tau t : S (\tau'' + (\tau' \times (\tau \rightarrow \text{ReT} \tau \tau' S \tau'')))} \text{(RUNRE)} \\
\frac{\Gamma \vdash t : S \tau \quad S \leq S'}{\Gamma \vdash \text{elevate}_{S'} t : S' \tau} \text{(ELEVATE)}
\end{array}$$

Figure 3: Typing judgments for terms. Rules for λ -calculus are omitted.




Figure 4: Ordering on effect labels (given by the diagram) and on state monads.

empty context, we write $\{\}$. Many of the rules are standard, reflecting the rules of computational λ -calculus. The rules for get, put, and elevate require special attention, as they directly involve effect labels. Rule T-GET restricts the effect label on the top monad transformer to include a read label, and T-PUT restricts it to include a write label. These restrictions are expressed in terms of an ordering on effect labels (which is really nothing more than the subset relation) given in Figure 4 at left. For rule T-ELEVATE, we require that the target monad S' has (non-strictly) more effect labels than the source monad S ; the precise meaning of this is expressed in Figure 4 at right. The intuition is that elevate permits us to *decertify* that a computation does *not* read or write at any given state layers, but not to remove existing effect labels.

Stores and configurations also have a notion of type, defined by the rules of Figure 5. A store Σ is said to *match* a monad M if the types of its elements correspond, in order, to the state types of the state monad transformers in M . For this, we simply write that Σ matches M . A configuration $\langle t, \Sigma \rangle$, then, has type $M \tau$ if and only if Σ matches M and $\{\} \vdash t : M \tau$. We write this $\langle t, \Sigma \rangle \triangleright M \tau$.

Theorem 1 (Uniqueness of Types): If $\Gamma \vdash t : \tau$ and $\Gamma \vdash t : \tau'$, then $\tau = \tau'$. Also, if $\langle t, \Sigma \rangle \triangleright \tau$ and $\langle t, \Sigma \rangle \triangleright \tau'$, then $\tau = \tau'$.

$$\frac{\frac{\Sigma \text{ matches } S}{\Sigma \text{ matches } \text{ReT} \tau \tau' S} \text{(M-RET)} \quad \frac{\{\} \vdash s : \tau \quad \Sigma \text{ matches } S}{s :: \Sigma \text{ matches } \text{StT} \ell \tau S} \text{(M-STT)}}{\{\} \vdash t : M \tau \quad \Sigma \text{ matches } M} \text{(T-CONFIG)} \quad \langle t, \Sigma \rangle \triangleright M \tau$$

Figure 5: Typing judgments for stores (top) and configurations (bottom).

C. Small-Step Operational Semantics

In this section we describe the semantics of RWC in a small-step operational style. As a computational λ -calculus, RWC contains both functional features (functional abstraction and application) as well as effectful ones (mutable state and reactive parallelism). The operational semantics is structured around this dichotomy, with two interdefined notions of reduction: *pure* and *effectful* reduction. Pure reduction reflects the notion of effect-free evaluation. A pure reduction judgment takes the form $t \rightsquigarrow t'$; note that this makes no mention of any store. Effectful reduction provides semantics to computational terms which may have effects. Thus an effectful reduction judgment takes the form $\langle t, \Sigma \rangle \rightsquigarrow \langle t', \Sigma' \rangle$.

The rules for pure and effectful reduction are given in Figures 6 and 7, respectively. We adopt a call-by-value evaluation strategy, as this is (we feel) simpler to work with metatheoretically than call-by-name or -need. This may seem strange in light of ReWire’s antecedents in Haskell (which is a non-strict language), but since ReWire is a strongly normalizing subset of Haskell, it does not matter whether we choose an eager or lazy evaluation strategy from a “backwards compatibility” point of view: since there is no “bottom” value, strictness is not a concern.

A few of the rules require close inspection. To begin with, we note that pure and effectful reduction are interdefined. Rule STM-ST of Figure 7 allows pure reduction to be “lifted” into the universe of effectful reduction: if the term component t of a configuration $\langle t, \Sigma \rangle$ still has not been evaluated to a value, we will continue to evaluate it without changing the store. Dually, if less obviously, the rule ST-RUNIDMO in Figure 6 allows monadic evaluation in the identity monad (and *only* in the identity monad) to be reified in a pure setting. If we wish to run a computation in a more complex monad, we may use runRe and runSt to “peel off” one monad transformer at a time, until we reach the Id monad at the core. In the runSt case, we must supply an initial value for the corresponding state layer, producing a computation in the base monad which will return the post-value for that layer. The runRe operator will produce a computation in the base monad that either returns a final result value, or an output value paired with a continuation waiting on more input.

Note also the interaction between the rule STM-LIFTST, STM-GET, and STM-PUT. The get and put operations always operate on the ‘head’ (leftmost) item in the store. Applying lift_{StT} to these operations allows us to access items deeper in the store, by executing the underlying computation against the “tail” of the store and leaving the “head” item unchanged.

$$\begin{array}{c}
\frac{}{(\lambda x : \tau. t)v \rightsquigarrow t[x := v]} \text{ (ST-APPABS)} \quad \frac{t \rightsquigarrow t''}{t t' \rightsquigarrow t'' t'} \text{ (ST-APP1)} \\
\frac{t \rightsquigarrow t'}{v t \rightsquigarrow v t'} \text{ (ST-APP2)} \quad \frac{t \rightsquigarrow t'}{\text{inl}_\tau t \rightsquigarrow \text{inl}_\tau t'} \text{ (ST-INL)} \quad \frac{\langle v, \text{nil} \rangle \rightsquigarrow \langle t, \text{nil} \rangle}{\text{runld } v \rightsquigarrow \text{runld } t} \text{ (ST-RUNIDMO)}
\end{array}$$

Figure 6: Reduction Rules for Lambda Calculus Reduction. These rules, mostly omitted, specify a call-by-value evaluation strategy on RWC.

The rule STM-UNFOLD may be justified directly by the Haskell definition of unfold. Rule STM-PAUSE is more subtle. The basic idea, however, is that if a pause arises to the left of a \star , we should “absorb” what comes to the right of the \star into the pause’s continuation, guaranteeing that we make progress towards a “done” configuration.

As stated in Theorem 2, the reduction relation the results from the rules for pure and effectful reduction is deterministic.

Theorem 2 (Evaluation is Deterministic): If $t \rightsquigarrow t'$ and $t \rightsquigarrow t''$, then $t' = t''$. Also, if $\langle t, \Sigma \rangle \rightsquigarrow \langle t', \Sigma' \rangle$ and $\langle t, \Sigma \rangle \rightsquigarrow \langle t'', \Sigma'' \rangle$, then $\langle t', \Sigma' \rangle = \langle t'', \Sigma'' \rangle$.

IV. METATHEORY

In this section we discuss the metatheory of RWC, in particular *type safety* (Section IV-A), *strong normalization* (Section IV-B), and *soundness* of effect labels (Section IV-C).

A. Type Safety

As is standard in operational semantics, we take type safety to be the conjunction of *progress*, meaning that any well-typed term (resp. configuration) that is not a value (resp. is not “done”) always reduces to something, and *preservation*, meaning that reduction preserves the types of terms (and configurations). Together, these properties imply that well-typed programs can’t go wrong—i.e., evaluation of well-typed programs never “gets stuck”.

Theorem 3 (Progress): If $\{\} \vdash t : \tau$, then either t is a value or there exists t' such that $t \rightsquigarrow t'$. Also, if $\langle t, \Sigma \rangle \triangleright M \tau$, then either $\langle t, \Sigma \rangle$ is done, or there exist t' and Σ' such that $\langle t, \Sigma \rangle \rightsquigarrow \langle t', \Sigma' \rangle$.

Theorem 4 (Preservation): If $\{\} \vdash t : \tau$ and $t \rightsquigarrow t'$, then $\{\} \vdash t' : \tau$. Also, if $\langle t, \Sigma \rangle \triangleright M \tau$ and $\langle t, \Sigma \rangle \rightsquigarrow \langle t', \Sigma' \rangle$, then $\langle t', \Sigma' \rangle \triangleright M \tau$.

Perhaps surprisingly, the addition of computational features does not substantially complicate the proof of type safety relative to a pure λ -calculus.

B. Normalization

Unlike Haskell, RWC enjoys the property of *strong normalization*, which means that all well-typed terms (resp. configurations) will eventually reduce to some value (resp. done configuration). This property is especially important in hardware applications for the reason that hardware cannot be allowed to “loop forever” between clock ticks. The computation time between clock ticks must have a static, finite upper bound—this issue is discussed in detail in the references [1, 7]. Strong normalization also makes defined equality easier

$$\begin{array}{c}
\frac{}{\text{nil} \text{ld}(\underline{\mathbb{W}}) \text{nil}} \text{ (SWNW-ID)} \quad \frac{\Sigma \underline{S(\underline{\mathbb{W}})} \Sigma'}{\Sigma \text{ReT } \tau \underline{S(\underline{\mathbb{W}})} \Sigma'} \text{ (SWMW-RE)} \\
\frac{\Sigma \underline{S(\underline{\mathbb{W}})} \Sigma'}{(s :: \Sigma) \text{StT } \tau \underline{S(\underline{\mathbb{W}})} (s' :: \Sigma')} \text{ (SWNW-N)} \quad \frac{\Sigma \underline{S(\underline{\mathbb{W}})} \Sigma'}{(s :: \Sigma) \text{StT } \tau \underline{S(\underline{\mathbb{W}})} (s' :: \Sigma')} \text{ (SWNW-R)} \\
\frac{\Sigma \underline{S(\underline{\mathbb{W}})} \Sigma'}{(s :: \Sigma) \text{StT } \tau \underline{S(\underline{\mathbb{W}})} (s' :: \Sigma')} \text{ (SWNW-W)} \\
\frac{\Sigma \underline{S(\underline{\mathbb{W}})} \Sigma'}{(s :: \Sigma) \text{StT } \tau \underline{S(\underline{\mathbb{W}})} (s' :: \Sigma')} \text{ (SWNW-RW)}
\end{array}$$

Figure 8: The ‘same where no write’ relation.

$$\begin{array}{c}
\frac{}{\text{nil} \text{ld}(\underline{\mathbb{W}}) \text{nil}} \text{ (SWR-ID)} \quad \frac{\Sigma \underline{S(\underline{\mathbb{W}})} \Sigma'}{\Sigma \text{ReT } \tau \underline{S(\underline{\mathbb{W}})} \Sigma'} \text{ (SWR-RE)} \\
\frac{\Sigma \underline{S(\underline{\mathbb{W}})} \Sigma'}{(s :: \Sigma) \text{StT } \tau \underline{S(\underline{\mathbb{W}})} (s' :: \Sigma')} \text{ (SWR-N)} \quad \frac{\Sigma \underline{S(\underline{\mathbb{W}})} \Sigma'}{(s :: \Sigma) \text{StT } \tau \underline{S(\underline{\mathbb{W}})} (s' :: \Sigma')} \text{ (SWR-R)} \\
\frac{\Sigma \underline{S(\underline{\mathbb{W}})} \Sigma'}{(s :: \Sigma) \text{StT } \tau \underline{S(\underline{\mathbb{W}})} (s' :: \Sigma')} \text{ (SWR-W)} \\
\frac{\Sigma \underline{S(\underline{\mathbb{W}})} \Sigma'}{(s :: \Sigma) \text{StT } \tau \underline{S(\underline{\mathbb{W}})} (s' :: \Sigma')} \text{ (SWR-RW)}
\end{array}$$

Figure 9: The ‘same where read’ relation.

to work with, as it eliminates the need to account for equality of diverging computations.

We shall write \rightsquigarrow^* for the multistep reduction relation, i.e. the reflexive, transitive closure of \rightsquigarrow . We say that a term t *halts* if and only if there exists a (not necessarily distinct) value v , such that $t \rightsquigarrow^* v$. In a similar fashion, a configuration $\langle t, \Sigma \rangle$ *halts* if and only if there exists a done configuration D such that $\langle t, \Sigma \rangle \rightsquigarrow^* D$.

Theorem 5 (Normalization): If $\{\} \vdash t : \tau$, then t halts. Also, if $\langle t, \Sigma \rangle \triangleright M \tau$, then $\langle t, \Sigma \rangle$ halts.

The proof of Theorem 5 uses an adaptation of the standard *logical relations* technique [32]. Given a property \mathbf{P} , a *logical relation*, $\mathcal{R}_{\{\mathbf{T} \in \mathcal{T}\}}$ (with respect to \mathbf{P}), is a collection of type-indexed relations such that for every $\mathbf{R}_{\mathbf{T}} \in \mathcal{R}_{\{\mathbf{T} \in \mathcal{T}\}}$, every element $t \in \mathbf{R}_{\mathbf{T}}$, either has, or preserves \mathbf{P} . In the case of strong normalization, halting is the property of interest.

Proving Theorem 5 in Coq required developing novel techniques. Because resumptions involve potentially infinite computations, proving that strong normalization holds for configurations requires the use of proofs by coinduction. The use of coinduction allows the \mathbf{R} property to be appropriately applied over potentially infinite computations. The use of coinduction and the corresponding notion of bisimulation have been attributed to David Park [33].

C. Soundness of Effect Labels

Since effect labels are meant to track effects and their potential propagation, soundness of effect labels (roughly) corresponds to preservation of security levels indicated by the label, and that stores track such features accordingly. Thus, given well-typed configurations, establishing soundness of effect labels amounts to verifying that monadic-reduction

$$\begin{array}{c}
\frac{t \rightsquigarrow t'}{\langle t, \Sigma \rangle \rightsquigarrow \langle t', \Sigma \rangle} \text{ (STM-ST)} \quad \frac{\langle v, \Sigma \rangle \rightsquigarrow \langle t, \Sigma' \rangle}{\langle v \gg= v', \Sigma \rangle \rightsquigarrow \langle t \gg= v', \Sigma' \rangle} \text{ (STM-BIND)} \quad \frac{\langle \text{return}_M v \gg= v', \Sigma \rangle \rightsquigarrow \langle v', \Sigma, \Sigma \rangle}{\langle v, \Sigma \rangle \rightsquigarrow \langle t, \Sigma' \rangle} \text{ (STM-BINDRET)} \\
\frac{\langle \text{lift}_{(\text{StT } \ell \tau \text{ S})} v, s :: \Sigma \rangle \rightsquigarrow \langle \text{lift}_{(\text{StT } \ell \tau \text{ S})} t, s :: \Sigma' \rangle}{\langle \text{lift}_M (\text{return}_M v), \Sigma \rangle \rightsquigarrow \langle \text{return}_M v, \Sigma \rangle} \text{ (STM-LIFTST)} \quad \frac{\langle \text{lift}_{(\text{ReT } \tau \tau' \text{ S})} v, \Sigma \rangle \rightsquigarrow \langle \text{lift}_{(\text{ReT } \tau \tau' \text{ S})} t, \Sigma' \rangle}{\langle \text{lift}_M (\text{return}_M v), \Sigma \rangle \rightsquigarrow \langle \text{return}_M v, \Sigma \rangle} \text{ (STM-LIFTRE)} \\
\frac{\langle \text{get}_S, s :: \Sigma \rangle \rightsquigarrow \langle \text{returns } s, s :: \Sigma \rangle}{\langle \text{lift}_M (\text{return}_M v), \Sigma \rangle \rightsquigarrow \langle \text{return}_M v, \Sigma \rangle} \text{ (STM-PUT)} \quad \frac{\langle \text{put } v, s :: \Sigma \rangle \rightsquigarrow \langle \text{returns } (), v :: \Sigma \rangle}{\langle \text{lift}_M (\text{return}_M v), \Sigma \rangle \rightsquigarrow \langle \text{return}_M v, \Sigma \rangle} \text{ (STM-PUT)} \\
\frac{\langle t, \Sigma \rangle \rightsquigarrow \langle t', \Sigma' \rangle}{\langle \text{elevate}_S t, \Sigma \rangle \rightsquigarrow \langle \text{elevate}_S t', \Sigma' \rangle} \text{ (STM-ELEVATE)} \quad \frac{\langle \text{elevate}_{S'} (\text{returns } v), \Sigma \rangle \rightsquigarrow \langle \text{returns}_{S'} v, \Sigma \rangle}{\langle \text{elevate}_S t, \Sigma \rangle \rightsquigarrow \langle \text{elevate}_S t', \Sigma' \rangle} \text{ (STM-ELEVATERET)} \\
\frac{\langle v, s :: \Sigma \rangle \rightsquigarrow \langle t, s' :: \Sigma' \rangle}{\langle \text{runSt } v, s, \Sigma \rangle \rightsquigarrow \langle \text{runSt } t, s', \Sigma' \rangle} \text{ (STM-RUNST)} \quad \frac{\langle \text{runSt } (\text{return}_{(\text{StT } \ell \tau \text{ S})} v) v', \Sigma \rangle \rightsquigarrow \langle \text{returns } \langle v, v' \rangle, \Sigma \rangle}{\langle \text{runSt } v, s, \Sigma \rangle \rightsquigarrow \langle \text{runSt } t, s', \Sigma' \rangle} \text{ (STM-RUNSTRET)} \\
\frac{\langle v, \Sigma \rangle \rightsquigarrow \langle t, \Sigma' \rangle}{\langle \text{runRe}_\tau v, \Sigma \rangle \rightsquigarrow \langle \text{runRe}_\tau t, \Sigma' \rangle} \text{ (STM-RUNRE)} \quad \frac{\langle \text{runRe}_{\tau''} (\text{pause}_{(\text{ReT } \tau \tau' \text{ S})} v), \Sigma \rangle \rightsquigarrow \langle v * \lambda x. \text{return } (\text{inr}_{\tau''} x), \Sigma \rangle}{\langle \text{runRe}_\tau v, \Sigma \rangle \rightsquigarrow \langle \text{runRe}_\tau t, \Sigma' \rangle} \text{ (STM-RUNREPAUSE)} \\
\frac{\langle \text{runRe}_{\tau''} (\text{return}_{(\text{ReT } \tau \tau' \text{ S})} v), \Sigma \rangle \rightsquigarrow \langle \text{return}_S (\text{inl}_{(\tau \rightarrow (\tau' \times (\text{ReT } \tau \tau' \text{ S } \tau''))})} v), \Sigma \rangle}{\langle \text{runRe}_{\tau''} (\text{pause}_{(\text{ReT } \tau \tau' \text{ S})} v), \Sigma \rangle \rightsquigarrow \langle \text{runRe}_{\tau''} (\text{return}_{(\text{ReT } \tau \tau' \text{ S})} v), \Sigma \rangle} \text{ (STM-RUNRERET)} \\
\frac{\langle \text{unfold } v v', \Sigma \rangle \rightsquigarrow \langle \text{lift } (v' v) \gg= \lambda u. \left(\text{case } u \left(\begin{array}{l} \lambda w. \text{return } w \\ \lambda w. \text{proj } w (\lambda x. \lambda y. \text{pause} (\text{return } \langle x, \lambda z. \text{unfold } (y z) v' \rangle)) \end{array} \right), \Sigma \right) \rangle}{\langle \text{runRe}_{\tau''} (\text{pause}_{(\text{ReT } \tau \tau' \text{ S})} v), \Sigma \rangle \rightsquigarrow \langle \text{return}_S (\text{inl}_{(\tau \rightarrow (\tau' \times (\text{ReT } \tau \tau' \text{ S } \tau''))})} v), \Sigma \rangle} \text{ (STM-UNFOLD)} \\
\frac{\langle (\text{pause } v) * v', \Sigma \rangle \rightsquigarrow \langle \text{pause } (v \gg= \lambda w. (\text{proj } w (\lambda x. \lambda y. \text{return } \langle x, \lambda z. (y z) \gg= v' \rangle)) \rangle, \Sigma \rangle}{\langle \text{runRe}_{\tau''} (\text{pause}_{(\text{ReT } \tau \tau' \text{ S})} v), \Sigma \rangle \rightsquigarrow \langle \text{return}_S (\text{inl}_{(\tau \rightarrow (\tau' \times (\text{ReT } \tau \tau' \text{ S } \tau''))})} v), \Sigma \rangle} \text{ (STM-PAUSEBIND)}
\end{array}$$

Figure 7: Evaluation rules for monadic reduction. For the sake of readability, type annotations in STM-UNFOLD and STM-PAUSEBIND are elided.

$$\begin{array}{c}
\frac{\langle \text{nil}, \text{nil} \rangle \text{ wc } \langle \text{nil}, \text{nil} \rangle}{\langle \Sigma_1, \Sigma_2 \rangle \text{ wc } \langle \Sigma'_1, \Sigma'_2 \rangle} \text{ (WC-ID)} \\
\frac{\langle s_1 :: \Sigma_1, s_2 :: \Sigma_2 \rangle \text{ wc } \langle s_1 :: \Sigma'_1, s_2 :: \Sigma'_2 \rangle}{\langle \Sigma_1, \Sigma_2 \rangle \text{ wc } \langle \Sigma'_1, \Sigma'_2 \rangle} \text{ (WC-UNCHANGED)} \\
\frac{\langle s_1 :: \Sigma_1, s_2 :: \Sigma_2 \rangle \text{ wc } \langle s :: \Sigma'_1, s :: \Sigma'_2 \rangle}{\langle \Sigma_1, \Sigma_2 \rangle \text{ wc } \langle \Sigma'_1, \Sigma'_2 \rangle} \text{ (WC-CHANGED)}
\end{array}$$

Figure 10: The write consistency relation.

does not alter stores where no writes are allowed (Theorem 6); and moreover, that monadic-reduction does not reveal any changes to stores relative to monads with effect labels where only reads are allowed (Theorem 7). To that end, we make use of three relations: “same where no writes”, “same where read”, and “write consistency”, written $M \stackrel{(\text{W})}{\equiv}$, $M \stackrel{(\text{R})}{\equiv}$, and wc – and defined in Figure 8, Figure 9, and Figure 10 – respectively.

Stores (semantically) correspond to state monad transformers. Given a well-typed configuration, the associated store will contain appropriate elements relative to each layer in the state monad transformer stack. In order to update a store, its state monad must contain a write label.

Theorem 6 (No Forbidden Updates): If $\langle t, \Sigma \rangle \triangleright M \tau$, then $\langle t, \Sigma \rangle \rightsquigarrow \langle t', \Sigma' \rangle$ implies $\Sigma \stackrel{M(\text{W})}{\equiv} \Sigma'$.

Similarly, reading from a store takes place only relative to state monads that have a read label. This is reflected in the type judgments for put (resp., get) that require a write (resp., read) label in order to be well-typed.

Theorem 7 (No Forbidden Reads): Suppose $\Sigma_1 \stackrel{M(\text{R})}{\equiv} \Sigma_2$ and that $\langle t, \Sigma_1 \rangle \triangleright M \tau$ and $\langle t, \Sigma_2 \rangle \triangleright M \tau$. Then if $\langle t, \Sigma_1 \rangle \rightsquigarrow \langle t'_1, \Sigma'_1 \rangle$ and $\langle t, \Sigma_2 \rangle \rightsquigarrow \langle t'_2, \Sigma'_2 \rangle$, it follows that $t'_1 = t'_2$ and $\langle \Sigma_1, \Sigma_2 \rangle$ is write consistent with $\langle \Sigma'_1, \Sigma'_2 \rangle$.

The intuition underlying write consistency is that when considering a pair of stores Σ_1 and Σ_2 , prior to a reduction and a pair of matching stores Σ'_1 and Σ'_2 , after a reduction it is either the case that the pre-reduction stores do not differ from

their corresponding post-reduction stores (i.e. because no write takes place) or are equal to each other (i.e., because the same value was written to both Σ_1 and Σ_2).

V. TYPE-DIRECTED EQUATIONAL LOGIC FOR REWIRE CALCULUS

The rules provided in Figure 11 represent the properties of monads present in RWC. Rules (LEFT-UNIT), (RIGHT-UNIT), and (ASSOCIATIVITY) are the well-known “monad laws” and Rules (LIFT-RETURN) and (LIFT-*) are the “lifting laws” of Liang [12]. Rules (PUT-PUT), (PUT-GET), and (GET-GET), specify the interaction of stateful operations and are drawn from previous work [5]. The \leq relation on state monads is defined in Figure 4. We defer discussion of the remaining rules until the next section.

The equational logic of RWC has both atomic noninterference and clobber formalized as consequences of the RWC semantics in Coq; here, we refer to the last three rules of Figure 11. These are particular instances for a two layer state monad of the more general rules found in the Coq script [repository](#). Note that, in its Coq formalization, mask computes the appropriate definition from a monad type term taken as an argument. The exact details of this definition need not concern us here, and the interested reader may consult the repository.

VI. CONCLUSIONS AND FUTURE WORK

The ReWire methodology differs fundamentally from the type-based approach to secure hardware (e.g., that of Caisson [17], Sapper [18], and SecVerilog [19]) in three important respects. Firstly, ReWire is a functional language (a subset of Haskell) and has the benefit, we would argue, of the expressiveness of functional languages. Secondly, ReWire possesses a formal semantics and equational theory mechanized in the Coq theorem proving system, allowing security verification

$$\begin{array}{c}
\frac{t = t' : \tau \in \Gamma}{\Gamma \vdash t = t' : \tau} \text{ (AXIOM)} \quad \frac{\Gamma \vdash t = t' : \tau}{\gamma, \Gamma \vdash t = t' : \tau} \text{ (WEAKENING)} \quad \frac{\Gamma \vdash M : \tau}{\Gamma \vdash t = t : \tau} \text{ (REFL)} \quad \frac{\Gamma \vdash t' = t : \tau}{\Gamma \vdash t = t' : \tau} \text{ (SYM)} \quad \frac{\Gamma \vdash t = t' : \tau \quad \Gamma \vdash t' = t'' : \tau}{\Gamma \vdash t = t'' : \tau} \text{ (TRANS)} \\
\frac{y \notin \text{FV}(t)}{\Gamma \vdash \lambda x : \tau. t = \lambda y : \tau. t[x := y] : \tau \rightarrow \tau'} \text{ (\alpha)} \quad \frac{\Gamma \vdash \lambda x : \tau. t : \tau \rightarrow \tau' \quad \Gamma \vdash t' : \tau}{\Gamma \vdash (\lambda x : \tau. t)t' = t[t' := x] : \tau'} \text{ (\beta)} \\
\frac{\Gamma \vdash t' : \tau \quad \Gamma \vdash t : \tau \rightarrow M \tau'}{\Gamma \vdash (\text{return}_M t') \gg t = t t' : M \tau'} \text{ (LEFT-UNIT)} \quad \frac{\Gamma \vdash t : M \tau}{\Gamma \vdash t \gg \lambda x : \tau. (\text{return}_M x) = t : M \tau} \text{ (RIGHT-UNIT)} \\
\frac{\Gamma \vdash t : M \tau \quad \Gamma \vdash t' : \tau \rightarrow M \tau' \quad \Gamma \vdash t'' : \tau' \rightarrow M \tau'' \quad x \notin \text{FV}(t')}{\Gamma \vdash (t \gg t') \gg t'' = t \gg (\lambda x : \tau. t' x \gg t'') : M \tau''} \text{ (ASSOCIATIVITY-*)} \\
\frac{\Gamma \vdash \text{return}_M t : M \tau}{\Gamma \vdash \text{lift}_{M'}(\text{return}_M t) = \text{return}_{M'} t : M' \tau} \text{ (LIFT-RETURN)} \quad \frac{\Gamma \vdash t : M \tau \quad \Gamma \vdash t' : \tau \rightarrow M \tau' \quad x \notin \text{FV}(t')}{\Gamma \vdash \text{lift}_M(t \gg t') = (\text{lift}_M t) \gg (\lambda(x : \tau). \text{lift}_{M'}(t' x))} \text{ (LIFT-*)} \\
\frac{\Gamma \vdash \text{put } t : \text{StT } \ell \tau S() \quad \Gamma \vdash \text{put } t' : \text{StT } \ell \tau S()}{\Gamma \vdash (\text{put } t \gg \text{put } t') = \text{put } t' : \text{StT } \ell \tau S()} \text{ (PUT-PUT)} \\
\frac{\Gamma \vdash t : \tau}{\Gamma \vdash (\text{put } t \gg \text{get}_{(\text{StT RW } \tau S)}) = \text{put } t \gg \text{return}_{(\text{StT RW } \tau S)} t : \text{StT RW } \tau S} \text{ (PUT-GET)} \\
\frac{\langle R \rangle \leq \ell, \text{ such that } M = \text{StT } \ell \tau S}{\Gamma \vdash \text{get}_M \gg \lambda x : \tau. \text{get}_M \gg \lambda y : \tau. \text{return}_M(x, y) = \text{get}_M \gg \lambda z : \tau. \text{return}_M(z, z) : M(\tau \times \tau)} \text{ (GET-GET)} \\
\frac{\Gamma \vdash t : S \tau \quad S \text{ is StT RW } \tau(\text{StT } \langle \rangle \tau' \text{Id})}{\Gamma \vdash t \gg (\text{mask } S) = \text{mask } S : S()} \text{ (CLOBBER-LO)} \quad \frac{\Gamma \vdash \text{lift}_S t : S() \quad S \text{ is StT } \langle \rangle \tau(\text{StT RW } \tau' \text{Id})}{\Gamma \vdash \text{lift}_S t \gg (\text{mask } S) = (\text{mask } S) \gg \text{lift}_S t : S()} \text{ (ATOMIC NONINTERFERENCE)}
\end{array}$$

Figure 11: Type-directed Equational Logic for RWC

to be automatically checked with the attendant increased assurance. Thirdly, and most importantly, ReWire’s type system is not a security type system in the usual sense [34]. Security verification in ReWire is not fully automatic via a security type system, but, rather, the equational style of security verification of our previous work [1], [5] is supported by an effects type system based on the marriage of effects and monads [6]. However, we believe that ReWire’s being a pure functional language will support the adaptation of ideas from language-based security to the construction of high assurance, secure hardware via extensions to the ReWire type system.

The ReWire methodology, therefore, occupies a middle ground between the security-via-typechecking approach of Caisson and SecVerilog and traditional hardware verification with theorem provers [20]. It combines the advantages of both—static checking on the one hand and deductive reasoning on the other—with the expressive power of functional languages. Our previous work [1]–[3] demonstrates that ReWire possesses what the creators of the Delite framework refer to as “the three P’s” [35]: productivity, performance and portability. The current work shows ReWire also possesses a fourth “P”: provability. Follow-on articles will present the formalizations of previously published verifications of ReWire devices [1], [2].

The CompCert [36] project mechanizes both a source language’s semantics and compiler in Coq, thereby providing the foundation for (1) verifying properties of C source programs and (2) compiling those programs to efficient implementations in a verifiably property-preserving manner. One particular strength of the CompCert approach is that other tools may be mechanized in Coq as well (e.g., static analysis tools, etc., from the Verified Software Toolchain [37]) to provide increased automation and trust to the whole workflow. The current work is motivated by the goal of producing trusted hardware in the same manner as CompCert supports trusted

C implementations. This is, admittedly, a very ambitious goal, but the current work is an early, yet important, step in this program. The current work also provides an important first step towards the formal verification of the ReWire compiler.

REFERENCES

- [1] A. Procter, W. Harrison, I. Graves, M. Becchi, and G. Allwein, “A principled approach to secure multi-core processor design with ReWire,” *ACM TECS*, vol. 16, no. 2, pp. 33:1–33:25, Feb. 2017.
- [2] I. Graves, W. Harrison, A. Procter, and G. Allwein, “Provably correct development of reconfigurable hardware designs via equational reasoning,” in *IEEE Inter. Conf. on Field-Programmable Technology (ICFPT)*, 2015, pp. 160–171.
- [3] I. Graves, A. Procter, W. Harrison, M. Becchi, and G. Allwein, “Hardware synthesis from functional embedded domain-specific languages: A case study in regular expression compilation,” in *Applied Reconfigurable Computing*, ser. LNCS, vol. 9040, 2015, pp. 41–52.
- [4] W. Harrison, A. Procter, I. Graves, M. Becchi, and G. Allwein, “A programming model for reconfigurable computing based in functional concurrency,” in *11th Inter. Symp. on Reconfigurable Communication-Centric Systems-on-Chip*, 2016.
- [5] W. Harrison and J. Hook, “Achieving information flow security through monadic control of effects,” *JCS*, vol. 17, pp. 599–653, Oct 2009.
- [6] P. Wadler, “The marriage of effects and monads,” in *ICFP*, 1998, pp. 63–74.
- [7] A. Procter, “Semantics-driven design and implementation of high-assurance hardware,” Ph.D. dissertation, University of Missouri, 2014.
- [8] B. Huffman, “HOLCF ’11: A definitional domain theory for verifying functional programs,” Ph.D. dissertation, Portland State University, 2012.
- [9] L. Schröder and T. Mossakowski, “Hascal: Integrated higher-order specification and program development,” *Theoretical Computer Science*, vol. 410, no. 12, pp. 1217 – 1260, 2009.
- [10] E. Moggi, “Notions of computation and monads,” *Information and Computation*, vol. 93, no. 1, pp. 55–92, July 1991.
- [11] F. Nielson, H. Nielson, and C. Hankin, *Principles of Program Analysis*, 1999.
- [12] S. Liang, P. Hudak, and M. Jones, “Monad transformers and modular interpreters,” in *POPL*, 1995, pp. 333–343.
- [13] A. P. W. Harrison and G. Allwein, “The confinement problem in the presence of faults,” in *ICFEM*, 2012, pp. 182–197.
- [14] D. Cock, G. Klein, and T. Sewell, “Secure microkernels, state monads and scalable refinement,” in *TPHOLs*, 2008, pp. 167–182.

- [15] A. Nanevski, G. Morrisett, A. Shinnar, P. Govereau, and L. Birkedal, “Ynot: Dependent types for imperative programs,” in *ICFP*, 2008, pp. 229–240.
- [16] W. Swierstra, “A hoare logic for the state monad,” in *TPHOLs*, 2009, pp. 440–451.
- [17] X. Li, M. Tiwari, J. Oberg, V. Kashyap, F. Chong, T. Sherwood, and B. Hardekopf, “Caius: a hardware description language for secure information flow,” in *PLDI*, 2011, pp. 109–120.
- [18] X. Li, V. Kashyap, J. Oberg, M. Tiwari, V. R. Rajarathinam, R. Kastner, T. Sherwood, B. Hardekopf, and F. Chong, “Sapper: A language for hardware-level security policy enforcement,” in *ASPLOS*, 2014.
- [19] D. Zhang, Y. Wang, G. E. Suh, and A. Myers, “A hardware design language for efficient control of timing channels,” Dept. of Computer Science, Cornell University, Tech. Rep. 2014-04-10, 2014, extended version of the authors’ ASPLOS15 paper.
- [20] T. Melham, *Higher Order Logic and Hardware Verification*, ser. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 1993, vol. 31.
- [21] G. Cabodi and M. Murciano, “BDD-based hardware verification,” in *6th Inter. Conf. on Formal Methods for the Design of Computer, Communication, and Software Systems*, ser. SFM’06, 2006, pp. 78–107.
- [22] A. Myers, personal communication, Mar. 2017.
- [23] A. Azevedo de Amorim, N. Collins, A. DeHon, D. Demange, C. Hrițcu, D. Pichardie, B. Pierce, R. Pollack, and A. Tolmach, “A verified information-flow architecture,” in *POPL*, 2014, pp. 165–178.
- [24] M. Gordon, “The semantic challenge of Verilog HDL,” in *Logic in Computer Science, 1995. LICS ’95. Proceedings., Tenth Annual IEEE Symposium on*, Jun 1995, pp. 136–145.
- [25] C. Kloos and P. Breuer, Eds., *Formal Semantics for VHDL*. Kluwer Academic Publishers, 1995.
- [26] S. Goncharov and L. Schröder, “A coinductive calculus for asynchronous side-effecting processes,” in *Proc. of the 18th International Conf. on Fundamentals of Computation Theory*, 2011, pp. 276–287.
- [27] K. Crary, A. Kligler, and F. Pfenning, “A monadic analysis of information flow security with mutable state,” *JFP*, vol. 15, no. 2, pp. 249–291, Mar. 2005.
- [28] D. Ghica and A. Jung, “Categorical semantics of digital circuits,” in *FMCAD*, 2016.
- [29] S. Peyton Jones, Ed., *Haskell 98 Language and Libraries, the Revised Report*. Cambridge University Press, 2003.
- [30] W. Harrison, “The essence of multitasking,” in *Algebraic Methodology and Software Technology*, 2006, pp. 158–172.
- [31] J. Goguen and J. Meseguer, “Unwinding and inference control,” in *IEEE Symp. on Security and Privacy*, 1984, pp. 75–86.
- [32] J. Mitchell, *Foundations for Programming Languages*. MIT Press Cambridge, 1996.
- [33] D. Sangiorgi, “On the origins of bisimulation and coinduction,” *ACM Trans. Program. Lang. Syst.*, vol. 31, no. 4, pp. 15:1–15:41, May 2009.
- [34] A. Sabelfeld and A. Myers, “Language-based information-flow security,” *IEEE Journ. on Sel. Areas in Commun.*, vol. 21, no. 1, Jan. 2003.
- [35] H. Lee, K. Brown, A. Sujeeeth, H. Chafi, T. Rompf, M. Odersky, and K. Olukotun, “Implementing domain-specific languages for heterogeneous parallel computing,” *IEEE Micro*, vol. 31, no. 5, pp. 42–53, Sep. 2011.
- [36] X. Leroy, “Formal verification of a realistic compiler,” *Commun. ACM*, vol. 52, no. 7, pp. 107–115, Jul. 2009.
- [37] “Verified Software Toolchain,” <http://vst.cs.princeton.edu>.