

The Essence of Event-Driven Programming

Jennifer Paykin¹, Neelakantan R. Krishnaswami², and
Steve Zdancewic³

1 University of Pennsylvania, Philadelphia, USA

jpaykin@seas.upenn.edu

2 University of Birmingham, Birmingham, United Kingdom

N.Krishnaswami@cs.bham.ac.uk

3 University of Pennsylvania, Philadelphia, USA

stevez@cis.upenn.edu

Abstract

Event-driven programming is based on a natural abstraction: an event is a computation that can eventually return a value. This paper exploits the intuition relating events and time by drawing a Curry-Howard correspondence between a functional event-driven programming language and a *linear-time temporal logic*. In this logic, the eventually proposition $\Diamond A$ describes the type of events, and Girard’s *linear logic* describes the effectful and concurrent nature of the programs. The correspondence reveals many interesting insights into the nature of event-driven programming, including a generalization of selective choice for synchronizing events, and an implementation in terms of callbacks where $\Diamond A$ is just $\neg \Box \neg A$.

Digital Object Identifier 10.4230/LIPIcs...

1 Introduction

Event-driven programming is a popular approach to functional concurrency in which events, also known as “futures,” “deferred values,” or “lightweight threads,” execute concurrently with each other and eventually produce a result. The abstraction of the event has been extremely successful in producing lightweight, extensible, and efficient concurrent programs. As a result, a wide range of programming languages and libraries use the event-driven paradigm to describe everything from message-passing concurrency [24] and lightweight threads [26] to graphical user interfaces [11] and I/O [13, 21, 25].

Although these systems vary considerably in the details of their implementations and APIs, they share a common, basic structure. They provide:

- the *abstraction* of the event, often given explicitly as a *monad*;
- the ability to *synchronize* across events;
- a continuation-passing style implementation in terms of *callbacks*; and
- a source or sources of *primitive events*.

In this paper we distill event-driven programming to its essence, demonstrating how to derive these components from first principles. Starting from a logical basis and building up the minimal machinery needed to explain the four points above, we proceed as follows:

The logic of events. In Section 2 we identify a Curry-Howard correspondence between events that eventually return a value and the $\Diamond$ (“eventually”) modality from *temporal logic*. We define a core language of pure events where the monad from the event-driven abstraction is identified as $\Diamond$. We formulate the type system in the setting of Girard’s *linear logic* to



© Jennifer Paykin, Neelakantan R. Krishnaswami, and Steve Zdancewic;
licensed under Creative Commons License CC-BY

Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

characterize the fact that events are effectful and execute concurrently. This linear and monadic logic serves as the basic scaffolding for event-driven computations.

Synchronization refers to the ability to execute two events concurrently and record which one happens first. In Section 3 we extend the type system of pure events to include a synchronization operator **choose** in the style of Concurrent ML [24]. We observe that, logically, **choose** corresponds to the linear-time axiom of temporal logic:

$$\Diamond A \wedge \Diamond B \rightarrow \Diamond((A \wedge \Diamond B) \vee (\Diamond A \wedge B)) \quad \text{“eventually, } A \text{ happens before } B \text{ or } B \text{ happens before } A\text{”}$$

This logical interpretation of **choose** suggests a natural generalization in terms of McBride’s derivatives of types [19], which we develop into a new technique for synchronizing events across arbitrary container data structures.

Callbacks, continuations, and the event loop. In the event-driven paradigm, events are implemented using callbacks that interact with an underlying event loop. For the language of pure events, we define in Section 4 a time-aware continuation-passing style (CPS) translation based on the property of temporal logic that $\Diamond A$ is equivalent to $\neg \Box \neg A$, where negation $\neg$ is the type of first-class continuations, and $\Box$ is the “always” operator from temporal logic.

In addition to the logic of pure events, the implementation should take into account the extralogical *sources of concurrency* that interact with the event loop. Throughout this paper we use a range of these concurrency primitives, including nondeterministic events, timeouts, user input, and channels, and argue that the choice of primitive is orthogonal to the logical structure of events.

In Section 5 we extend the CPS translation to account for these axiomatic sources of concurrency. We model the event loop in the *answer type* of the CPS translation [7] and show how to instantiate the answer type for a concrete choice of event primitives.

The essence of events. In this paper we argue that the logical interpretation of events is a unifying idea behind the vast array of real-world event-driven languages and libraries. We complete the story in Section 6 by comparing techniques used in practice with the approaches developed in this paper based on the essence of event-driven programming.

2 The Curry-Howard Connection

The important abstraction of event-driven programming is based on the relationship between events and time: an event is a computation that can eventually return a value. This abstraction takes the form of a *monad*, which means that there is a simple interface for interacting with events:

- **return** e is a trivial event that immediately returns the value e .
- **bind** $x = e_1$ **in** e_2 is an event that first waits for e_1 to return a value, binds that value to x , and then continues as the event e_2 .

In this paper we go even further by identifying the particular monad for events with the “eventually” operator from temporal logic. Consider a simple intuitionistic logic for time in which a proposition is either *true now* or *true later*. A proposition that is true now is also true at every point in the future, but a proposition that is true later may not necessarily be true now.¹ A proposition that is true later is denoted $\Diamond A$, and pronounced “eventually A .”

¹ Other presentations of temporal logic include next ($\circ$), always ($\Box$), and until ($\mathcal{U}$) operators, and do not necessarily assume that just because a proposition is true now, it will always be true.

The “eventually” modality $\Diamond A$ is defined by two rules. The first says that if a proposition is true now, it is also true later. The second rule says that if A is true later, and if A now proves that some B is true later, then B itself is also true later. Through the Curry-Howard correspondence, these proofs correspond to typing rules for **return** and **bind**, respectively.

$$\frac{\Delta \vdash e : A}{\Delta \vdash \text{return } e : \Diamond A} \quad \frac{\Delta_1 \vdash e_1 : \Diamond A \quad \Delta_2, x : A \vdash e_2 : \Diamond B}{\Delta_1, \Delta_2 \vdash \text{bind } x = e_1 \text{ in } e_2 : \Diamond B}$$

2.1 A logic for effects: linear logic

Consider the event **nondet** e that returns e at some nondeterministic point in the future. The program `foo = (let  $x = \text{nondet } e$  in  $\text{in}_1(x, x)$ )` has state $\text{in}_1(\text{undefined}, \text{undefined})$ for some nondeterministic amount of time before simultaneously stepping to $\text{in}_1(\text{return } e, \text{return } e)$. In particular, `foo` is *not* equivalent to the substituted form $\text{in}_1(\text{nondet } e, \text{nondet } e)$, which at some point in the future may have its first component be `return` e and its second component be *undefined*. On the other hand, the computation should not be blocking; the expression `case foo of ( $\text{in}_1 y \rightarrow \text{True} \mid \text{in}_2 z \rightarrow \text{False}$ )`, which tests whether `foo` is a right or a left injection, should resolve immediately to `True`.

Events like `foo` are inherently effectful, because the state of an event changes as computation progresses. In order to describe events in a purely logical way, the Curry-Howard correspondence has to take these effectful relationships into account.

Linear logic [12] is one approach for typing effectful programs that has been successful for concurrency in the settings of session types [5, 27] and uniqueness types [14]. Linear-use (or “one shot”) data structures show up in many event-driven languages, for example in the form of Ivars in Async [21], Futures in Scala [13], and widgets in GUI programming [17]. Linear types are so useful in concurrent programming because they disallow aliasing, meaning that there is no shared state between processes. By using linear variables we can define an *equational operational semantics* that restores the Curry-Howard correspondence with logic.²

2.2 A logic without effects: the unrestricted fragment

The linear and temporal type system will account for effects and events, but event-driven programs also include ordinary computations that don’t require the event or linear-space abstractions. For these we should be able to write terms in a regular programming language—in this case, the (non-linear) simply-typed λ -calculus.

Benton et al. [4] introduced a way to combine linear and non-linear logics that has since been called *adjoint logic* [23] after the categorical structure. In an adjoint logic, we write non-linear types τ to distinguish them from linear types A , and non-linear typing contexts Γ to distinguish them from linear typing contexts Δ . The non-linear typing judgment has the form $\Gamma \vdash t : \tau$, while the linear one has the form $\Gamma; \Delta \vdash e : A$. That is, the linear typing judgment allows unrestricted access to the non-linear variables in Γ , but only linear-space access to the variables in Δ .

Adjoint logic also describes ways in which the linear and non-linear types relate to each other. A linear type A can be embedded into a persistent type $\lceil A \rceil$ when A does not rely

² Other event-driven languages solve this problem in different ways without using linear types. Implementations in strict functional languages like Async [21] ensure that every event normalizes immediately to an asynchronous primitive, which somewhat defeats the purpose of a strict evaluation order. In CML [24], all events evaluate strictly and synchronously unless wrapped in a thunk called `guard`.

$$\begin{aligned}
\tau &::= \text{Unit} \mid \text{Void} \mid \tau \times \tau \mid \tau + \tau \mid \tau \rightarrow \tau \mid [A] \\
A &::= 1 \mid 0 \mid A \otimes A \mid A \oplus A \mid A \multimap A \mid \Diamond A \mid [\tau] \\
t &::= x \mid () \mid \text{case } t \text{ of } () \mid (t_1, t_2) \mid \pi_i t \mid \text{in}_i t \mid \text{case } t \text{ of } (\text{in}_1 x_1 \rightarrow t_1 \mid \text{in}_2 x_2 \rightarrow t_2) \\
&\quad \mid \lambda x. t \mid t_1 t_2 \mid \text{suspend } e \\
e &::= x \mid () \mid \text{let } () = e_1 \text{ in } e_2 \mid \text{case } e \text{ of } () \\
&\quad \mid (e_1, e_2) \mid \text{let } (x_1, x_2) = e_1 \text{ in } e_2 \mid \text{in}_i e \mid \text{case } e \text{ of } (\text{in}_1 x_2 \rightarrow e_1 \mid \text{in}_2 x_2 \rightarrow t_2) \\
&\quad \mid \lambda x. e \mid e_1 e_2 \mid \text{return } e \mid \text{bind } x = e_1 \text{ in } e_2 \mid \text{force } t \mid [t] \mid \text{let } [x] = e_1 \text{ in } e_2
\end{aligned}$$

■ **Figure 1** Syntax of linear and non-linear types, terms, and expressions.

$$\begin{array}{c}
\frac{\Gamma; \cdot \vdash e : A}{\Gamma \vdash \text{suspend } e : [A]} \quad [-]\text{-I} \qquad \frac{\Gamma \vdash t : [A]}{\Gamma; \cdot \vdash \text{force } t : A} \quad [-]\text{-E} \\
\\
\frac{\Gamma \vdash t : \tau}{\Gamma; \cdot \vdash [t] : [\tau]} \quad [-]\text{-I} \qquad \frac{\Gamma; \Delta_1 \vdash e_1 : [\tau] \quad \Gamma, x : \tau; \Delta_2 \vdash e_2 : B}{\Gamma; \Delta_1, \Delta_2 \vdash \text{let } [x] = e_1 \text{ in } e_2 : B} \quad [-]\text{-E}
\end{array}$$

■ **Figure 2** Typing rules for moving between the linear and unrestricted fragments.

on any linear assumptions. On the other hand, a persistent type τ can always be treated as a linear type, written $[\tau]$. Figure 1 shows the syntax of types, (non-linear) terms, and (linear) expressions.

The typing rules for terms and expressions are mostly standard, but Figure 2 shows how to move in between the linear and non-linear fragments via the typing rules for $[A]$ and $[\tau]$. A linear expression e can be suspended to a persistent term $\text{suspend } e$, and can be unsuspended using **force**. On the other hand, a non-linear term can be used linearly in this type system by applying a floor operator, and unpacked in the same way.

The remaining typing rules are shown in Appendix A.

2.3 Completing the Curry-Howard connection

We have shown how the Curry-Howard correspondence relates the propositions of temporal linear logic to types and the proofs of such propositions to event-driven programs. In the remainder of this section we show how the operational semantics of event-based programs relates to the equational theory of proofs.

Events use a call-by-name evaluation strategy, because events themselves are not expected to normalize right away. Consider the nondeterministic event **nondet** e introduced in the beginning of this section. In the expression **let** $x = \text{nondet } e$ **in** $\text{in}_1 x$, the variable x occurs linearly in the rest of the expression and so it is safe to substitute the unresolved event **nondet** e for x .

On the other hand, terms have a call-by-value evaluation strategy, which is safe because a suspended expression **suspend** e is a value in the term language.

A selection of the operational semantics for expressions is shown in Figure 3. Evaluation occurs under contexts, which have the form E^P for contexts with holes for non-linear terms, and E^L for contexts with holes for linear expressions. A full description of the operational

$$\begin{array}{lcl}
\text{bind } x = \text{return } e_1 \text{ in } e_2 \rightsquigarrow e_2\{e_1/x\} & & \frac{t \rightsquigarrow t'}{E^P[t] \rightsquigarrow E^P[t']} \\
\text{force}(\text{suspend } e) \rightsquigarrow e & & \\
\text{let } [x] = [v] \text{ in } e \rightsquigarrow e\{v/x\} & & \\
E^P ::= \dots \mid \text{force } [\] \mid [\] & & \frac{e \rightsquigarrow e'}{E^L[e] \rightsquigarrow E^L[e']} \\
E^L ::= \dots \mid \text{bind } x = [\] \text{ in } e \mid \text{let } [x] = [\] \text{ in } e & &
\end{array}$$

■ **Figure 3** Operational semantics of terms and expressions.

semantics can be found in Appendix B.

► **Theorem 1** (Preservation). *If $\vdash t : \tau$ and $t \rightsquigarrow t'$ then $\vdash t' : \tau$. If $\vdash e : A$ and $e \rightsquigarrow e'$ then $\vdash e' : A$.*

► **Theorem 2** (Progress). *If $\vdash t : \tau$ then either t is a value or t can take a step. If $\vdash e : A$ then either e is in weak head normal form or e can take a step.*

3 Synchronization and Linear Time

In the previous section we described a simple logic for pure events that is not particularly expressive. Although it can express sequential events using **return** and **bind**, it cannot express concurrent events running in parallel. As an example, consider a timeout event **onTimeout** n that returns a unit value after the amount of time specified by n . With this operator we can limit the execution of another event e by running **onTimeout** n and e in parallel and keeping track of which one occurs first. We denote the operation that runs two events in parallel as **choose** and give it the type $\Diamond A \otimes \Diamond B \multimap \Diamond(A \otimes \Diamond B \oplus \Diamond A \otimes B)$. The **choose** operator is a kind of selective choice operator as seen in CML and Async.

```

timeout e |n| = bind z = choose (e, onTimeout |n|) in
  case z of | in1 (a,t) -> let () = drop t in return (Some a)
          | in2 (e,()) -> let () = drop e in return None
  end

```

Here **drop**, of type $\Diamond A \multimap 1$, is an operation that explicitly aborts an event.

The type of **choose** can be interpreted as a property of temporal logic: if both A and B will be true at some point in the future, then either A will come before B , or B will come before A . This axiom distinguishes a *linear-time temporal logic* from a *branching-time temporal logic*, in which two different futures could occur in different timelines.

The logical interpretation of **choose** reveals that it is somehow fundamental to the event-driven interpretation. In fact, we can think of **choose** as part of the operational semantics of events.³ We extend the evaluation contexts so that the two events being synchronized on can evaluate concurrently. Once one of the events resolves to a value, the synchronization operator can take a β -reduction step.

$$\begin{array}{lcl}
E^L ::= \dots \mid \text{choose}([\], e) \mid \text{choose}(e, [\]) & & \text{choose}(\text{return } e_1, e_2) \rightsquigarrow \text{return}(\text{in}_1(e_1, e_2)) \\
& & \text{choose}(e_1, \text{return } e_2) \rightsquigarrow \text{return}(\text{in}_2(e_1, e_2))
\end{array}$$

³ These rules make sense operationally, but not necessarily as part of the Curry-Howard correspondence, because as an equational theory they relate **return**(**in**₁(e_1 , **return** e_2)) and **return**(**in**₂(**return** e_1 , e_2)), which are certainly not equal. This stems from the fact that **choose** is not a pure logical axiom; it relates multiple connectives in a complicated way and is hence neither an introduction nor an elimination rule.

<pre> $\tau ::= \dots \mid \text{Chan } A$ spawn : $\Diamond 1 \multimap 1$ choose (eA, eB) = let win : Chan (1$\oplus$1) = new () in let cA : Chan A = new () in let cB : Chan B = new () in spawn (bind a = eA in spawn (send win (in1 ())); send cA a); spawn (bind b = eB in spawn (send win (in2 ())); send cB b); </pre>	<pre> new : $1 \multimap [\text{Chan } A]$ send : $A \multimap [\text{Chan } A] \multimap \Diamond 1$ receive : $[\text{Chan } A] \multimap \Diamond A$ bind z = receive win in case z of in1 () -> bind a = receive cA in return (in1 (a, receive cB)) in2 () -> bind b = receive cB in return (in2 (receive cA , b)) end </pre>
--	---

■ **Figure 4** Signature of channels and the channel-based implementation of **choose**

3.1 Implementing choose with channels

As we hinted above, the **choose** operator cannot be implemented in the language of pure events. Logically this is because **choose** corresponds to the linear-time *axiom* from temporal logic, so it is not derivable from the other operators. However, as we describe here, we can implement **choose** in terms of primitive sources of concurrency: in this case, linear synchronous channels and the **spawn** operator.

The **spawn** operator takes an event with return type 1 and executes it asynchronously. A synchronous channel is a way to communicate linear information between processes. Although the data transmitted across channels is linear, the reference to the channel itself need not be, so $\text{Chan } A$ is added as a non-linear type τ to our type system. Following Reppy [24], **send** and **receive** are synchronous in that they do not return a value until a send is matched with a receive. The signature of linear message-passing is summarized in Figure 4.

The implementation of **choose**, also shown in Figure 4, uses three channels. The first, called **win**, tracks which event of **eA** or **eB** occurs first. Based on that, **choose** will wait for either **eA** or **eB** explicitly, through the intermediate channels **cA** and **cB**. Then **choose** will return either **in1** or **in2** along with a reference to the intermediate event, the receive event on the unresolved channel.

3.2 Synchronizing more than pairs

For practical programming problems, we often need to synchronize more than two events at a time. For example, the **choose** operators in CML and Async operate over lists instead of pairs. In the linear type system of this paper, the type of such an operator would be

$$\text{chooseList} : \text{List}(\Diamond A) \multimap \text{List}(\Diamond A) \otimes A \otimes \text{List}(\Diamond A)$$

where the input list is partitioned into the prefix and suffix of the first event to return a value. Unfortunately it is not possible to derive **chooseList**, or even a version over triples of events, from the binary version of **choose**. Like **choose** itself we need to implement **chooseList** using channels or some other concurrency primitive. By itself this solution is ad-hoc and unsatisfactory.

In the remainder of this section we describe a way to build up synchronization operators on *arbitrary finite containers of events* by induction on the type of the container. We do this by exploiting a uniform pattern on the structure of these primitives, inspired by Conor

$$\begin{array}{lll}
\partial_\diamond \diamond A = A & \partial_\diamond (A \otimes B) = (\partial_\diamond A \otimes B) \oplus (A \otimes \partial_\diamond B) & \partial_\diamond (A \multimap B) = 0 \\
\partial_\diamond 1 = 0 & \partial_\diamond (A \oplus B) = \partial_\diamond A \oplus \partial_\diamond B & \partial_\diamond [\tau] = 0 \\
\partial_\diamond 0 = 0 & & \\
\\
E^\diamond ::= \dots \mid \text{choose } E^\diamond & \text{fill } [\] e = e & \\
E^\diamond ::= [\] \mid \text{in}_i E^\diamond \mid (E^\diamond, e) \mid (e, E^\diamond) & \text{fill } (\text{in}_i E^\diamond) e = \text{in}_i (\text{fill } E^\diamond e) & \\
& \text{fill } (E^\diamond, e') e = \text{in}_1 (\text{fill } E^\diamond e, e') & \\
\text{choose } E^\diamond [\text{return } e] \rightsquigarrow \text{return} (\text{fill } E^\diamond e) & \text{fill } (e', E^\diamond) e = \text{in}_2 (e', \text{fill } E^\diamond e) &
\end{array}$$

■ **Figure 5** Derivatives with respect to time

McBride’s 2001 observation that the derivative of a regular type is the type of its one-hole contexts [19]. We explore a variation on his idea and show that *the derivative of a type with respect to time is the type of its synchronization operator*.

Derivatives with respect to time. We define the *instant* of an event to be the time at which it returns a value. An event itself can be thought of as a context with a hole for time that is filled in by its instant. For example, the event `return n` consists of the context $[\]n$, where the hole $[\]$ is filled in by its instant “now.”

The semantics of synchronization say that the instant of `choose`(e_1, e_2) is either the instant of e_1 or the instant of e_2 , depending on which occurs first. The context containing that hole has one of two shapes. If e_1 returns a value n before e_2 does, then the context will have the form $([\]n, e_2)$, of type $A \otimes \diamond B$. If e_2 returns a value first, the context will have the type $\diamond A \otimes B$. Thus the return type of `choose`, $(A \otimes \diamond B) \oplus (\diamond A \otimes B)$, describes the possible shapes of its context with a hole for time.

McBride’s partial derivative operation, written $\partial_X A$, records the possible shapes of a one-hole context of A with a hole for the type X .⁴ This intuition extends from finite containers to recursive data types like lists. For example, the one-hole contexts of the type `List A` consist of a one-hole context of A , along with the prefix and suffix lists surrounding the element with the hole. That is, the derivative of `List A` is `List A` $\otimes$ $\partial_X A$ $\otimes$ `List A`, which is reminiscent of the return type of `chooseList`.

In Figure 5 we define the syntactic operation $\partial_\diamond A$ on types that describes the derivative with respect to time.⁵ The derivative of an event type $\diamond A$ is A itself, leaving the time at which the event occurred as the hole.

The general `choose` operator decomposes a type into two parts: its instant (designated by the $\diamond$ prefix) at which synchronization will occur, and a context with a hole for time.

$$\text{choose}_A : A \multimap \diamond(\partial_\diamond A) \quad (\text{when } \partial_\diamond A \text{ is not degenerate, i.e. } \partial_\diamond A \neq 0)$$

This gives us a pattern for synchronizing events across arbitrary finite containers. McBride’s treatment of recursive types provides a way to extend synchronization to arbitrary recursive containers such as lists, but we leave the details to future work.

⁴ The syntax is inspired by the fact that this operation obeys the product and sum rules from calculus. For example, if we write X^2 for $X \times X$ then $\partial_X X^2 \cong 2 \times X = X + X$.

⁵ The one-hole context interpretation of derivatives does not extend to higher-order types, so we make the simplification that the derivative of all higher-order types is 0.

Operational semantics. The operational behavior of choose_A is also shown in Figure 5. The evaluation contexts $E^\perp$ are extended with $\text{choose } E^\diamond$ where $E^\diamond$ is a special kind of context for expressions. We write $\Gamma; \Delta \vdash A :> E^\diamond : B$ to mean that $E^\diamond$ has a hole for expressions of type A .

The β -rule applies when a component of any context returns a value. In this case the synchronization operator must produce an expression of type $\partial_\diamond A$. We define an operation $\text{fill } E^\diamond e$ that constructs the element of the derivative from the context and the component filling in the hole.

► **Lemma 3.** *If $\Gamma; \Delta_1 \vdash \diamond A :> E^\diamond : B$ and $\Gamma; \Delta_2 \vdash e : A$, then $\Gamma; \Delta_1, \Delta_2 \vdash \text{fill } E^\diamond e : \partial_\diamond B$.*

The above lemma is enough to prove preservation in the extended system. The side condition on choose_A that $\partial_\diamond A \not\equiv 0$ ensures that progress also holds by ruling out ill-formed terms like $\text{choose}(\lambda x. e)$.

4 Logic and Callbacks

The interpretation of events as the eventually type from temporal logic has led to some interesting insights about their behavior, including the meaning of synchronization. In the following sections we discuss how the standard implementation of events as callbacks can also be given a logical interpretation.

Callbacks, which are also called continuations and event handlers, are functions that accept some input but do not return a value. Rather, the return type of a callback is invisible, so it can be represented as $A \rightarrow \text{Answer}$ for any type Answer . We could instead write its type as $\neg A$, making the answer type opaque to the programmer.

Consider an implementation of a GUI library using callbacks and an event loop. The library provides a way to register handlers in the event loop that get triggered once the user performs some external action, such as a key press.

$\text{onKeyPress} : (\text{Char} \rightarrow \text{Answer}) \rightarrow \text{Answer}$ or $\text{onKeyPress} : \neg\neg\text{Char}$

What does the type of a callback have to do with temporal logic? The callback being registered in the event loop will not be invoked immediately, but at some point in the future, once the user has pressed a key. The type of the callback itself should then reflect the fact that it will be available in the future. We write $\Box A$ to denote the fact that A will be true at every point in the future, and so the type of onKeyPress should be written

$\text{onKeyPress} : \Box(\text{Char} \rightarrow \text{Answer}) \rightarrow \text{Answer}$ or $\text{onKeyPress} : \neg\Box\neg\text{Char}$ or $\text{onKeyPress} : \Diamond\text{Char}$

This last step follows from the isomorphism $\neg\Box\neg A \cong \Diamond A$, so we conclude: *the act of registering a callback that reacts to a key press is the same as an event that eventually returns the key that was pressed.*

4.1 An alternate temporal logic

While the event-based perspective on programming focused on a temporal logic of the eventually type $\Diamond A$, the callback-based perspective focuses on the “always” or “global” type, written $\Box A$. Under this logic, a proposition that is true now (denoted $A^\mathbf{n}$) is not necessarily true in the future. The modifier $\Box A^\mathbf{n}$ indicates that $A^\mathbf{n}$ is true both now and at every point in the future. However, since the temporal logic is still linear in space, the proposition $\Box A^\mathbf{n}$ is only available until it gets used.

$$\begin{array}{c}
\frac{\Gamma; \Delta, x : A^n \vdash e : \text{Answer}}{\Gamma; \Delta \vdash \lambda x. e : \neg A^n} \quad \frac{\Gamma; \Delta_1 \vdash e_1 : \neg A^n \quad \Gamma; \Delta_2 \vdash e_2 : A^n}{\Gamma; \Delta_1, \Delta_2 \vdash e_1 e_2 : \text{Answer}} \\
\\
\frac{\Gamma; \Box \Delta \vdash e : A}{\Gamma; \Box \Delta \vdash \text{box } e : \Box A} \quad \frac{\Gamma; \Delta \vdash e : \Box A}{\Gamma; \Delta \vdash \text{unbox } e : A}
\end{array}$$

■ **Figure 6** Typing rules for tensor logic

The behavior of $\Box A^n$ can be described using two simple rules. First, if A^n is provable using only hypotheses that are always true, then A^n is always true. On the other hand, a proposition A^n that is always true is also true now.

To describe continuations, we add the negation type $\neg A^n$ and remove all other linear arrows. Because only linear expressions need to undergo a CPS translation, the *non-linear* arrow types are unchanged. We denote the linear propositions in the resulting *adjoint tensor logic* [20] as A^n , and non-linear propositions as τ .

$$\begin{aligned}
\tau &::= \text{Unit} \mid \text{Void} \mid \tau \times \tau \mid \tau + \tau \mid \tau \rightarrow \tau \mid [A]^n \\
A^n &::= 1 \mid 0 \mid A^n \otimes A^n \mid A^n \oplus A^n \mid \neg A \mid \Box A \mid [\tau]
\end{aligned}$$

The syntax of terms and expressions is almost identical to that of the event-based language. The negation type $\neg A^n$ is introduced with a λ -abstraction and is eliminated with an application. The monadic **bind** and **return** operators are replaced by the comonadic **box** and **unbox** operators, as shown in Figure 6. We assume a call-by-value operational semantics for both terms and expressions.

4.2 A temporal CPS translation

The callback-based implementation of events is based on a time-sensitive CPS transformation on types written $\llbracket A \rrbracket$. The CPS translation is time-sensitive in that it keeps track of when hypotheses are available. In particular, while expressions in the event-based language are available at any point in the future, expressions in the callback-based language expire. For example, the translation of the eventually type encodes the fact that the return value can be used at any point in the future: $\llbracket \Diamond A \rrbracket = \neg \Box \neg \Box \llbracket A \rrbracket$. Similarly, a function might use its argument not right away, but at some point in the future, so the type translation of linear functions is $\llbracket A_1 \multimap A_2 \rrbracket = \neg(\Box \llbracket A_1 \rrbracket \otimes \neg \llbracket A_2 \rrbracket)$. The full details of the CPS translation on types are shown in Appendix C.

As expected, the translation on (non-linear) terms is the identity, while the translation on expressions is guided by the type translation in a mostly standard way. We describe a few cases in Appendix C and have the following soundness theorem:

► **Theorem 4.** *If $\Gamma \vdash t : \tau$ then $\llbracket \Gamma \rrbracket \vdash \llbracket t \rrbracket : \llbracket \tau \rrbracket$. If $\Gamma; \Delta \vdash e : A$ then $\llbracket \Gamma \rrbracket; \Box \llbracket \Delta \rrbracket \vdash \llbracket e \rrbracket : \llbracket A \rrbracket$.*

In addition, the translation respects the operational semantics of the source language.

► **Theorem 5.** *If $e \rightsquigarrow e'$ then $\llbracket e \rrbracket \rightarrow^* \llbracket e' \rrbracket$, modulo the administrative redexes introduced by the CPS translation [10].*

5 Concurrency Sources and the Event Loop

The CPS translation in the previous section only describes the language of pure events, and not the additional *sources of concurrency* that are the backbone of event-driven programming. So far in this paper we have considered a number of different sources: nondeterminism

in the event `nondet e`, timeouts of the form `onTimeout e`, synchronous channels that create events `send` and `receive`, and the GUI operation `onKeyPress`.

The following section sketches a way to implement these primitive sources of concurrency as part of the CPS translation. The trick is to integrate the concurrent actions of these primitives with the answer type of the continuation, as described by Claessen [7] in the *poor man's concurrency monad*.

Actions and the answer type. For the pure fragment of the eventually monad (consisting only of `return` and `bind`), the answer type of the continuation is invisible. What we write as $\neg A$ is in fact $A \multimap \text{Answer}$ for some fixed type `Answer`. Claessen observed that this makes the answer type of the continuation the perfect place to hide the presence of effects.

We call these effects *actions* following Claessen. An action can be thought of as a distinct thread of computation [18], the primitive thread operations being **Halt** and **Fork**. These threads are also stateful, operating over an event queue monad, which we write `EventQueue A`. The **Atom** action can execute an arbitrary monadic operation over the event queue. Using Haskell-like notation for the algebraic datatype of actions, we have:

```
data Action =
| Halt : Action
| Fork : Action -> Action -> Action
| Atom : EventQueue Action -> Action
```

Since actions represent threads of computation, they are executed by a scheduler of type `ListAction  $\multimap$  EventQueueAction` that schedules a list of actions inside the event queue monad. For example, the following is a simple round robin scheduler:

```

eventLoop [] = return ()
eventLoop (Halt      :: ls) = eventLoop ls
eventLoop (Fork a1 a2 :: ls) = eventLoop (ls ++ [a1,a2])
eventLoop (Atom mA    :: ls) = bind a = mA in eventLoop (ls++[a])

```

Events and actions interact via a top-level `run` operator that converts a unit-valued event to an action. Its type is $\neg[\![\Diamond 1]\!]$ or $[\![\Diamond 1]\!] \multimap \text{Action}$.

$$\text{run} = \lambda(k : \neg \Box \neg \Box \llbracket 1 \rrbracket).k(\text{box}(\lambda(x : \Box \llbracket 1 \rrbracket).(\text{unbox } x)(\lambda().\text{Halt})))$$

As a sanity check, observe that $\text{run}[\text{return } v]$ evaluates to $\llbracket v \rrbracket(\lambda().\text{Halt})$.

Spawn. Using actions we can encode the event-level concurrency primitives that have been used throughout this paper. For example, **spawn**, of type $\Diamond 1 \multimap 1$, is implemented as a member of its CPS-converted type $\llbracket \Diamond 1 \multimap 1 \rrbracket = \neg(\Box[\Diamond 1] \otimes \neg[\llbracket 1 \rrbracket])$ that takes in an event (of type $\Box[\Diamond 1]$) and a continuation (of type $\neg[\llbracket 1 \rrbracket]$) and produces a **Fork** action that runs the event and calls the continuation in parallel. The definition is found in Appendix D, and we can check that $\text{run}[\text{return}(\text{spawn } e)]$ evaluates to $\text{Fork}(\text{box}[\llbracket e \rrbracket]) \text{Halt}$.

Linear Channels. We conclude this section with a sketch of how to implement synchronous message-passing in the style of CML, which we used to encode `choose` in Section 3. Other sources of concurrency can be implemented in a similar way.

In order to represent channels, the event queue underlying the action type should be stateful over a linear heterogeneous store. Linear references are indexed by some non-linear identifiers, which we write $\text{ld } A$.

$$\begin{array}{ll} \tau ::= \dots \mid \text{Id } A & \text{newId} : A^n \multimap \text{EventQueue}[\text{Id } A^n] \\ A^n ::= \dots \mid \text{EventQueue } A^n & \text{updateId} : [\text{Id } A^n] \multimap (A^n \multimap A^n \otimes B^n) \multimap \text{EventQueue } B^n \end{array}$$

A channel is a reference to a linear cell that consists of either: (a) a list of messages to be sent, along with the event handlers to be triggered after rendezvous, or (b) a list of event handlers waiting for messages. The types of these two possible elements are written `SendElt` and `RecvElt`, respectively.

$$\begin{aligned}\text{SendElt } A^n &= \Box A^n \otimes \Box \neg \Box \llbracket 1 \rrbracket & \text{RecvElt } A^n &= \Box \neg \Box A^n \\ \text{Chan } A^n &= \text{Id}(\text{List}(\text{SendElt } A^n) \oplus \text{List}(\text{RecvElt } A^n))\end{aligned}$$

When a message of type $\Box A^n$ is sent over the channel, we examine the current state of the cell. If the cell contains a list of messages to be sent, the new message will be added to the end of the list. If the cell contains any callbacks waiting for messages, the callback will be applied to the incoming message and stored as an action. This behavior is governed by the function `attachSend` of type $\llbracket \text{Chan } A \rrbracket \multimap \text{SendElt } A \multimap \text{EventQueue Action}$.

```
attachSend c (a,k0) = updateId c (fun s =>
  case s of
  | in1 ls      -> (in1 (ls++[(a,k0)]), Halt)
  | in2 []      -> (in1 [a], Halt)
  | in2 (k::ks) -> (in2 ks, Fork ((unbox k) a) ((unbox k0) [])))
end)
```

The event-level `send` operator has the type $\llbracket \text{Chan } A \rrbracket \multimap A \multimap \Diamond 1$, so its implementation has type $\llbracket \llbracket \text{Chan } A \rrbracket \multimap A \multimap \Diamond 1 \rrbracket$. Then $\llbracket \text{send} \rrbracket$ is a continuation that takes in a channel, a message of type $\Box \llbracket A \rrbracket$, and a continuation of type $\llbracket \Diamond 1 \rrbracket$, and produces an `Atom` performing the monadic computation `attachSend`.

The interpretation of `receive` is governed by a similar protocol `attachReceive` of type $\llbracket \text{Chan } A \rrbracket \multimap \text{RecvElt } A \multimap \text{EventQueue Action}$, the details of which are given in Appendix D.

6 Discussion

In this paper, a linear and temporal logic is the guiding principle in the design of a core language for events. But the connection between logic and programming is only significant if it has a basis in existing event-driven languages, which vary in the ways they embody the event-driven paradigm. To understand these variations and the design decisions of this paper, we revisit the four features of event-driven programming discussed in the introduction: the monadic abstraction of the event, the synchronization operator, the implementation in terms of callbacks, and the primitive sources of events.

Layers of abstraction for the monadic event. The language of pure events presented in Section 2 has two parts: a type of events ($\Diamond A$) and a monadic interface for interacting with them. The monadic structure is explicit in many existing languages, including CML’s `’a event` type [24], Async’s `’a Deferred.t` [21], Lwt’s type for light-weight threads [26], and Scala’s `Future[A]` [13].⁶ Other languages don’t have an explicit event type, but require programmers to work directly in CPS, including Python’s Twisted library [16], JavaScript’s Node.js [25], or Ruby’s EventMachine [6]. Still others have the event abstraction but not in a monadic style, like Racket’s synchronizable events [1] or Go’s goroutines [2].

⁶ Although all of these abstractions are monads, their interfaces are not standard. CML has a return operator `alwaysEvt` but in lieu of `bind` it has a functorial `wrap` along with a `sync` operator of type `’a event -> ’a` that synchronously executes an event. Async and Lwt both use the standard `return` and `bind` terminology but Async has an impure `peek` operation that polls whether or not an event has completed, and Lwt has the ability to explicitly put threads to sleep and wake them up again.

Synchronization. Selective choice as described in Section 3 is less universal than the monadic bind, but has proved useful in CML, Async, Lwt, and Racket, where the default choice operator acts on lists, not pairs, and has type $\text{List}(\Diamond A) \rightarrow \Diamond A$.⁷ In this paper, by considering a linear **choose** operator over pairs instead of lists, we are able to draw a connection with the linear-time axiom of temporal logic, and abstract away from the type of pairs to derive a synchronization operator not only for lists, but for any container data structure.

Implementation: synchronous or asynchronous. In Section 4 we give an implementation of events based on a CPS translation that is independent of the event loop. As a consequence, **return** and **bind** are necessarily synchronous, which means that an event, once running, does not yield control by default. Only once an explicitly asynchronous operation like **choose**, **spawn**, or **receive** is called does control shift back to the scheduler.⁸ In the asynchronous style, used by Async and Scala, the bind operation $\text{bind } x = e_1 \text{ in } e_2$ registers the event e_1 in the event queue so that it executes asynchronously by default. To encode synchrony in an asynchronous system requires a special case, such as Scala’s blocking constructs [13], but the other direction is trivial using **spawn** [22].

Where do events come from? Despite the similarities between event-driven languages, programming in one versus another may feel very different depending on the intended application domain. CML feels most natural for describing message-passing concurrency due to its built-in channels and spawn operator. Async describes shared-state concurrency, where its Ivar data structure is a one-shot kind of shared state. Promises in Scala are more focused on long-running computations like I/O. However, these different concurrency abstractions can be implemented in one another, such as eXene’s implementations of GUIs in CML [9], or Scala’s async libraries [3]. We argue that the choice of concurrency primitive is orthogonal to the design of events themselves, and that the techniques presented in this paper are applicable to a wide range of primitives.

What about FRP? The event-driven paradigm described in this paper is closely connected to functional reactive programming (FRP), which targets many of the same domains. In the FRP model, the input to a program is modeled as a time-varying value, or a stream. FRP programs can be thought of as stream transformers, or as programs that *react* to the current state of the system. Recently, FRP’s connection with linear-time temporal logic [15] was discovered, which in fact prompted us to search for similar connections to event-based programming.

In typical FRP systems, the type $\Box A$ denotes the type of time-varying values, as opposed to our interpretation as an expression that is available now or at any time in the future. FRP programs model events $\Diamond A$ coinductively as $\nu \alpha. A \vee \circ \alpha$, where $\circ$ is the “next” modality. Unfortunately, this forces an implementation based on polling, which means programs continuously check whether an event has resolved yet.⁹ In the event-driven interpretation, the type of events $\Diamond A$ is interpreted as a continuation $\neg \Box \neg A$ and the structure of the event loop avoids polling.

⁷ In CML, **choose** aborts the events that are not chosen by means of its negative acknowledgment mechanism, but Panangaden and Reppy [22] show that this feature is encodable.

⁸ CML and Lwt both use synchronous implementation strategies. Notice that the question of synchronous versus asynchronous events is orthogonal to the choice of synchronous versus asynchronous channels.

⁹ Modern FRP languages work hard to avoid these time and space leaks, either by restricting the expressivity of programs [8] or by mixing ideas from event-driven programming with FRP [9].

Conclusion. The Curry-Howard correspondence reveals many interesting insights into the nature of event-driven programs. Synchronization via selective choice can be thought of as the linear-time axiom from temporal logic, and can be generalized to arbitrary container data structures. The standard implementation using callbacks can be explained in a temporal way by interpreting $\Diamond A$ as $\neg \Box \neg A$, and primitive sources of concurrency are implemented using a clever choice of answer type. The result is a top-to-bottom formulation of the essence of events: computations that eventually return a value.

Acknowledgment. Many thanks to Michael Arntzenius, Antal Spector-Zabusky, and Carter Schonwald for invigorating discussions about this work. This material is based in part upon work supported by the NSF Graduate Research Fellowship under Grant No. DGE-1321851.

References

- 1 Events (Racket documentation). Website. URL: docs.racket-lang.org/reference/sync.html.
- 2 The Go programming language. Website. URL: www.golang.org/.
- 3 scala-async. GitHub repository. URL: github.com/scala/async.
- 4 P.N. Benton. A mixed linear and non-linear logic: Proofs, terms and models. In *Computer Science Logic*. 1995.
- 5 Luís Caires and Frank Pfenning. Session types as intuitionistic linear propositions. In *CONCUR*. 2010.
- 6 Francis Cianfrocca. About EventMachine. Website. URL: www.rubydoc.info/gems/eventmachine.
- 7 Koen Claessen. A poor man’s concurrency monad. *Journal of Functional Programming*, 9:313–323, 1999.
- 8 Antony Courtney and Conal Elliott. Genuinely functional user interfaces. In *Haskell Workshop*, 2001.
- 9 Evan Czaplicki and Stephen Chong. Asynchronous functional reactive programming for GUIs. In *PLDI*, 2013.
- 10 Oliver Danvy and Andrzej Filinski. Representing control: a study of the CPS transformation. *Mathematical Structures in Computer Science*, 2:361–391, 12 1992.
- 11 Emden R. Gansner and John H. Reppy. A multi-threaded higher-order user interface toolkit. In *User Interface Software*, volume 1 of *Software Trends*. 1993.
- 12 Jean-Yves Girard. Linear logic. *Theoretical Computer Science*, 50(1):1–101, 1987.
- 13 Philipp Haller, Aleksandar Prokopec, Heather Miller, Viktor Klang, Roland Kuhn, and Vojin Jovanovic. Futures and Promises (Scala documentation). Website, 2013. URL: <http://docs.scala-lang.org/overviews/core/futures>.
- 14 Dana Harrington. Uniqueness logic. *Theoretical Computer Science*, 354(1):24 – 41, 2006. Algebraic Methods in Language Processing.
- 15 Alan Jeffrey. LTL types FRP: Linear-time temporal logic propositions as types, proofs as functional reactive programs. In *PLPV*, 2012.
- 16 Ken Kinder. Event-driven programming with Twisted and Python. *Linux Journal*, March 2005.
- 17 Neel Krishnaswami and Nick Benton. A semantic model for graphical user interfaces. In *ICFP*, 2011.
- 18 Peng Li and Steve Zdancewic. Combining events and threads for scalable network services: Implementation and evaluation of monadic, application-level concurrency primitives. In *PLDI*, 2007.
- 19 Conor McBride. The derivative of a regular type is its type of one-hole contexts. 2001.
- 20 Paul-André Melliès and Nicolas Tabareau. Resource modalities in tensor logic. *Annals of Pure and Applied Logic*, 161(5):632–653, 2010.
- 21 Yaron Minsky, Anil Madhavapeddy, and Jason Hickey. *Real World OCaml*. O’Reilly Media, 2013.
- 22 Prakash Panangaden and John Reppy. *ML with Concurrency: Design, Analysis, Implementation, and Application*, chapter The Essence of Concurrent ML, pages 5–29. 1997.
- 23 Frank Pfenning and Dennis Griffith. Polarized substructural session types. In *FSSCS*. 2015.
- 24 John H. Reppy. *Concurrent Programming in ML*. Cambridge University Press, 1999.
- 25 S. Tilkov and S. Vinoski. Node.js: Using JavaScript to build high-performance network programs. *IEEE Internet Computing*, 14(6):80–83, Nov 2010.
- 26 Jérôme Vouillon. Lwt: A cooperative thread library. In *ML Workshop*, 2008.
- 27 Philip Wadler. Propositions as sessions. *ICFP*, 2012.

A

 Event-based language

$$\tau ::= \text{Unit} \mid \text{Void} \mid \tau \times \tau \mid \tau + \tau \mid \tau \rightarrow \tau \mid \lceil A \rceil$$

$$A ::= 1 \mid 0 \mid A \otimes A \mid A \oplus A \mid A \multimap A \mid \Diamond A \mid \lfloor \tau \rfloor$$

$$t ::= x \mid () \mid \text{case } t \text{ of } () \mid (t_1, t_2) \mid \pi_i t \mid \text{in}_i t \mid \text{case } t \text{ of } (\text{in}_1 x_1 \rightarrow t_1 \mid \text{in}_2 x_2 \rightarrow t_2) \\ \mid \lambda x. t \mid t_1 t_2 \mid \text{suspend } e$$

$$e ::= x \mid () \mid \text{let } () = e_1 \text{ in } e_2 \mid \text{case } e \text{ of } ()$$

$$\mid (e_1, e_2) \mid \text{let } (x_1, x_2) = e_1 \text{ in } e_2 \mid \text{in}_i e \mid \text{case } e \text{ of } (\text{in}_1 x_2 \rightarrow e_1 \mid \text{in}_2 x_2 \rightarrow e_2)$$

$$\mid \lambda x. e \mid e_1 e_2 \mid \text{return } e \mid \text{bind } x = e_1 \text{ in } e_2 \mid \text{force } t \mid \lfloor t \rfloor \mid \text{let } \lfloor x \rfloor = e_1 \text{ in } e_2$$

$$\begin{array}{c}
\frac{}{\Gamma, x : \tau \vdash x : \tau} \text{VAR} \quad \frac{}{\Gamma \vdash () : \text{Unit}} \text{Unit-I} \quad \frac{\Gamma \vdash t : \text{Void}}{\Gamma \vdash \text{case } t \text{ of } () : \sigma} \text{Void-E} \\
\\
\frac{\Gamma \vdash t_1 : \tau_1 \quad \Gamma \vdash t_2 : \tau_2}{\Gamma \vdash (t_1, t_2) : \tau_1 \times \tau_2} \times\text{-I} \quad \frac{\Gamma \vdash t : \tau_1 \times \tau_2}{\Gamma \vdash \pi_i t : \tau_i} \times\text{-E} \\
\\
\frac{\Gamma \vdash t : \tau_i}{\Gamma \vdash \text{in}_i t : \tau_1 + \tau_2} +\text{-I} \quad \frac{\Gamma \vdash t : \tau_1 + \tau_2 \quad \Gamma, x_1 : \tau_1 \vdash t_1 : \sigma \quad \Gamma, x_2 : \tau_2 \vdash t_2 : \sigma}{\Gamma \vdash \text{case } t \text{ of } (\text{in}_1 x_1 \rightarrow t_1 \mid \text{in}_2 x_2 \rightarrow t_2) : \sigma} +\text{-E} \\
\\
\frac{\Gamma, x : \tau \vdash t : \sigma}{\Gamma \vdash \lambda x. t : \tau \rightarrow \sigma} \rightarrow\text{-I} \quad \frac{\Gamma \vdash t_1 : \tau \rightarrow \sigma \quad \Gamma \vdash t_2 : \tau}{\Gamma \vdash t_1 t_2 : \sigma} \rightarrow\text{-E} \\
\\
\frac{}{\Gamma; x : A \vdash x : A} \text{VAR} \quad \frac{\Gamma; \Delta \vdash e : 0}{\Gamma; \Delta \vdash \text{case } e \text{ of } () : B} 0\text{-E} \\
\\
\frac{}{\Gamma; \cdot \vdash () : 1} 1\text{-I} \quad \frac{\Gamma; \Delta_1 \vdash e_1 : 1 \quad \Gamma; \Delta_2 \vdash e_2 : B}{\Gamma; \Delta_1, \Delta_2 \vdash \text{let } () = e_1 \text{ in } e_2 : B} 1\text{-E} \\
\\
\frac{\Gamma; \Delta_1 \vdash e_1 : A_1 \quad \Gamma; \Delta_2 \vdash e_2 : A_2}{\Gamma; \Delta_1, \Delta_2 \vdash (e_1, e_2) : A_1 \otimes A_2} \otimes\text{-I} \quad \frac{\Gamma; \Delta_1 \vdash e_1 : A_1 \otimes A_2 \quad \Gamma; \Delta_2, x_1 : A_1, x_2 : A_2 \vdash e_2 : B}{\Gamma; \Delta_1, \Delta_2 \vdash \text{let } (x_1, x_2) = e_1 \text{ in } e_2 : B} \otimes\text{-E} \\
\\
\frac{\Gamma; \Delta \vdash e : A_i}{\Gamma; \Delta \vdash \text{in}_i e : A_1 \oplus A_2} \oplus\text{-I} \quad \frac{\Gamma; \Delta_1 \vdash e : A_1 \oplus A_2 \quad \Gamma; \Delta_2, x_1 : A_1 \vdash e_1 : B \quad \Gamma; \Delta_2, x_2 : A_2 \vdash e_2 : B}{\Gamma; \Delta_1, \Delta_2 \vdash \text{case } e \text{ of } (\text{in}_1 x_1 \rightarrow e_1 \mid \text{in}_2 x_2 \rightarrow e_2) : B} \oplus\text{-E} \\
\\
\frac{\Gamma; \Delta, x : A \vdash e : B}{\Gamma; \Delta \vdash \lambda x. e : A \multimap B} \multimap\text{-I} \quad \frac{\Gamma; \Delta_1 \vdash e_1 : A \multimap B \quad \Gamma; \Delta_2 \vdash e_2 : A}{\Gamma; \Delta_1, \Delta_2 \vdash e_1 e_2 : B} \multimap\text{-E} \\
\\
\frac{\Gamma; \Delta \vdash e : A}{\Gamma; \Delta \vdash \text{return } e : \Diamond A} \Diamond\text{-I} \quad \frac{\Gamma; \Delta_1 \vdash e : \Diamond A \quad \Gamma; \Delta, x : A \vdash e' : \Diamond B}{\Gamma; \Delta_1, \Delta \vdash \text{bind } x = e \text{ in } e' : \Diamond B} \Diamond\text{-E} \\
\\
\frac{\Gamma; \cdot \vdash e : A}{\Gamma \vdash \text{suspend } e : \lceil A \rceil} \lceil \cdot \rceil\text{-I} \quad \frac{\Gamma \vdash t : \lceil A \rceil}{\Gamma; \cdot \vdash \text{force } t : A} \lceil \cdot \rceil\text{-E} \\
\\
\frac{\Gamma \vdash t : \tau}{\Gamma; \cdot \vdash \lfloor t \rfloor : \lfloor \tau \rfloor} \lfloor \cdot \rfloor\text{-I} \quad \frac{\Gamma; \Delta_1 \vdash e_1 : \lfloor \tau \rfloor \quad \Gamma, x : \tau; \Delta_2 \vdash e_2 : B}{\Gamma; \Delta_1, \Delta_2 \vdash \text{let } \lfloor x \rfloor = e_1 \text{ in } e_2 : B} \lfloor \cdot \rfloor\text{-E}
\end{array}$$

B Operational Semantics

The normal forms of terms and expressions, respectively, are denoted v and n .

$$\begin{aligned} v &::= () \mid (v_1, v_2) \mid \text{in}_i v \mid \lambda x. t \mid \text{suspend } e \\ n &::= () \mid (e_1, e_2) \mid \text{in}_i e \mid \lambda x. e \mid \text{return } e \mid \lfloor v \rfloor \end{aligned}$$

$$\begin{aligned} \pi_i(v_1, v_2) &\rightsquigarrow v_i & (\lambda x. t_1)v_2 &\rightsquigarrow t_1\{v_2/x\} \\ \text{case in}_i v \text{ of } (\text{in}_1 x_1 \rightarrow t_1 \mid \text{in}_2 x_2 \rightarrow t_2) &\rightsquigarrow t_i\{v/x_i\} \\ \text{let } () = () \text{ in } e &\rightsquigarrow e & \text{let } (x_1, x_2) = (e_1, e_2) \text{ in } e &\rightsquigarrow e\{e_1/x, e_2/x\} \\ (\lambda x. e_1)e_2 &\rightsquigarrow e_1\{e_2/x\} & \text{bind } x = \text{return } e_1 \text{ in } e_2 &\rightsquigarrow e_2\{e_1/x\} \\ \text{force}(\text{suspend } e) &\rightsquigarrow e & \text{let } \lfloor x \rfloor = \lfloor v \rfloor \text{ in } e &\rightsquigarrow e\{v/x\} \\ \text{case in}_i e \text{ of } (\text{in}_1 x_1 \rightarrow e_1 \mid \text{in}_2 x_2 \rightarrow e_2) &\rightsquigarrow e_i\{e/x_i\} \end{aligned}$$

$$\begin{aligned} E^P &::= ([], t) \mid (v, []) \mid \pi_i [] \\ &\mid \text{in}_i [] \mid \text{case } [] \text{ of } (\text{in}_1 x_1 \rightarrow t_1 \mid \text{in}_2 x_2 \rightarrow t_2) \\ &\mid [] t \mid v [] \mid \text{force } [] \mid [[]] \\ E^L &::= \text{let } () = [] \text{ in } e \mid \text{let } (x_1, x_2) = [] \text{ in } e \\ &\mid \text{case } [] \text{ of } (\text{in}_1 x_1 \rightarrow e_1 \mid \text{in}_2 x_2 \rightarrow e_2) \\ &\mid [] e \mid \text{bind } x = [] \text{ in } e \mid \text{let } \lfloor x \rfloor = [] \text{ in } e \end{aligned}$$

$$\frac{t \rightsquigarrow t'}{E^P \llbracket t \rrbracket \rightsquigarrow E^P \llbracket t' \rrbracket}$$

$$\frac{e \rightsquigarrow e'}{E^L \llbracket e \rrbracket \rightsquigarrow E^L \llbracket e' \rrbracket}$$

C CPS translation

$$\begin{aligned} t &::= x \mid () \mid \text{case } t \text{ of } () \\ &\mid (t_1, t_2) \mid \pi_i t \\ &\mid \text{in}_i t \mid \text{case } t \text{ of } (\text{in}_1 x_1 \rightarrow t_1 \mid \text{in}_2 x_2 \rightarrow t_2) \\ &\mid \lambda x. t \mid t_1 t_2 \mid \text{suspend } e \\ e &::= x \mid () \mid \text{let } () = e_1 \text{ in } e_2 \mid \text{case } e \text{ of } () \\ &\mid (e_1, e_2) \mid \text{let } (x_1, x_2) = e_1 \text{ in } e_2 \\ &\mid \text{in}_i e \mid \text{case } e \text{ of } (\text{in}_1 x_1 \rightarrow e_1 \mid \text{in}_2 x_2 \rightarrow e_2) \\ &\mid \lambda x. e \mid e_1 e_2 \\ &\mid \text{box } e \mid \text{unbox } e \\ &\mid \text{force } t \mid \lfloor t \rfloor \mid \text{let } \lfloor x \rfloor = t \text{ in } e \end{aligned}$$

$$\begin{aligned} \llbracket \text{Unit} \rrbracket &= \text{Unit} & \llbracket 1 \rrbracket &= \neg\neg 1 \\ \llbracket \text{Void} \rrbracket &= \text{Void} & \llbracket 0 \rrbracket &= \neg\neg 0 \\ \llbracket \tau_1 \times \tau_2 \rrbracket &= \llbracket \tau_1 \rrbracket \times \llbracket \tau_2 \rrbracket & \llbracket A_1 \otimes A_2 \rrbracket &= \neg\neg(\llbracket A_1 \rrbracket \otimes \llbracket A_2 \rrbracket) \\ \llbracket \tau_1 + \tau_2 \rrbracket &= \llbracket \tau_1 \rrbracket + \llbracket \tau_2 \rrbracket & \llbracket A_1 \oplus A_2 \rrbracket &= \neg\neg(\llbracket A_1 \rrbracket \oplus \llbracket A_2 \rrbracket) \\ \llbracket \tau_1 \rightarrow \tau_2 \rrbracket &= \llbracket \tau_1 \rrbracket \rightarrow \llbracket \tau_2 \rrbracket & \llbracket A_1 \multimap A_2 \rrbracket &= \neg(\Box \llbracket A_1 \rrbracket \otimes \neg \llbracket A_2 \rrbracket) \\ \llbracket \lceil A \rceil \rrbracket &= \lceil \llbracket A \rrbracket \rceil & \llbracket \Diamond A \rrbracket &= \neg \Box \neg \Box \llbracket A \rrbracket \\ & & \llbracket \lfloor \tau \rfloor \rrbracket &= \neg\neg \llbracket \tau \rrbracket \end{aligned}$$

$$\begin{aligned} \llbracket x \rrbracket &= \text{unbox } x & \llbracket \text{return } e \rrbracket &= \lambda k. (\text{unbox } k)(\text{box } \llbracket e \rrbracket) \\ \llbracket \lambda x. e \rrbracket &= \lambda(x, k). k \llbracket e \rrbracket & \llbracket \text{bind } x = e_1 \text{ in } e_2 \rrbracket &= \lambda k. \llbracket e_1 \rrbracket (\text{box } (\lambda x. \llbracket e_2 \rrbracket k)) \\ \llbracket e_1 e_2 \rrbracket &= \lambda k. \llbracket e_1 \rrbracket (\text{box } \llbracket e_2 \rrbracket, \lambda z. zk) \end{aligned}$$

Sketch of Theorem 5. Following Danvy and Filinski [10], we could easily consider a one-pass CPS translation where the administrative reduces—those introduced only by the CPS translation—are treated as meta-operations on terms. These meta-operations are written with an overline, and follow the pattern

$$(\lambda x.e) \overline{\text{@}} v = e\{v/x\} \quad \overline{\text{unbox}}(\text{box } e) = e.$$

The administrative reduces are introduced uniformly in the CPS translation of terms, with the exception of the λ abstraction rule, which requires an extra η -expansion.

$$\begin{aligned} [x] &= \overline{\text{unbox}} x \\ [\lambda x.e] &= \lambda(y, z).(\lambda x.z \overline{\text{@}} [e])y \\ [e_1 e_2] &= \lambda k.[e_1] \overline{\text{@}} (\text{box}[e_2], \lambda z.z \overline{\text{@}} k) \\ [\text{return } e] &= \lambda k.(\overline{\text{unbox}} k)(\text{box}[e]) \\ [\text{bind } x = e_1 \text{ in } e_2] &= \lambda k.[e_1] \overline{\text{@}} (\text{box}(\lambda x.[e_2] \overline{\text{@}} k)) \end{aligned}$$

With this one-pass CPS translation we can prove the theorem directly. ◀

D CPS Implementation of Event Primitives

$$\begin{aligned} \text{spawn} &: \Diamond 1 \multimap 1 \\ \llbracket \text{spawn} \rrbracket : \llbracket \Diamond 1 \multimap 1 \rrbracket &= \neg(\Box \llbracket \Diamond 1 \rrbracket \otimes \neg \llbracket 1 \rrbracket) \\ &= \lambda(x, k).\text{Fork}(\text{run}(\text{unbox } x))(k\llbracket () \rrbracket) \end{aligned}$$

$$\begin{aligned} \text{new} &: 1 \multimap [\text{Chan } A] \\ \llbracket \text{new} \rrbracket : \neg(\Box \llbracket 1 \rrbracket \otimes \neg \llbracket [\text{Chan } A] \rrbracket) \\ &= \lambda(u : \Box \neg \neg 1, k : \neg \llbracket [\text{Chan } A] \rrbracket).(\text{unbox } u)(\lambda().\text{Atom } (\text{bind } i = \text{newId } (\text{inl } []) \text{ in } ki)) \end{aligned}$$

$$\begin{aligned} \text{send} &: [\text{Chan } A] \multimap A \multimap \Diamond 1 \\ \llbracket \text{send} \rrbracket : \neg(\Box \llbracket [\text{Chan } A] \rrbracket \otimes \Box \llbracket A \rrbracket \otimes \neg \llbracket \Diamond 1 \rrbracket) \\ &= \lambda(c : \Box \neg \neg \llbracket [\text{Chan } A] \rrbracket, a : \text{box}[A], k : \llbracket \Diamond 1 \rrbracket). \\ &\quad (\text{unbox } c)(\lambda i.k(\lambda(k' : \Box \neg \Box \llbracket 1 \rrbracket).\text{Atom } (\text{attachSend } i \ k')))) \end{aligned}$$

$$\begin{aligned} \text{receive} &: [\text{Chan } A] \multimap \Diamond A \\ \llbracket \text{receive} \rrbracket : \neg(\Box \llbracket [\text{Chan } A] \rrbracket \otimes \neg \llbracket \Diamond A \rrbracket) \\ &= \lambda(c : \Box \neg \neg \llbracket [\text{Chan } A] \rrbracket, k : \llbracket \Diamond A \rrbracket). \\ &\quad (\text{unbox } c)(\lambda i.k(\lambda(k' : \Box \neg \Box \llbracket A \rrbracket).\text{Atom } (\text{attachReceive } i \ k')))) \end{aligned}$$