






Adjoint Reactive GUI Programming

Christian Uldal Graulund¹  , Dmitrij Szamozvancev², and Neel Krishnaswami² 

¹ IT University of Copenhagen, 2300 Copenhagen, DK cgra@itu.dk

² University of Cambridge, Cambridge CB3 0FD, UK
nk480@c1.cam.ac.uk, ds709@c1.cam.ac.uk

Abstract. Most interaction with a computer is via graphical user interfaces. These are traditionally implemented imperatively, using shared mutable state and callbacks. This is efficient, but is also difficult to reason about and error prone. Functional Reactive Programming (FRP) provides an elegant alternative which allows GUIs to be designed in a declarative fashion. However, most FRP languages are synchronous and continually check for new data. This means that an FRP-style GUI will “wake up” on each program cycle. This is problematic for applications like text editors and browsers, where often nothing happens for extended periods of time, and we want the implementation to sleep until new data arrives. In this paper, we present an *asynchronous* FRP language for designing GUIs called λ_{Widget} . Our language provides a novel semantics for widgets, the building block of GUIs, which offers both a natural Curry–Howard logical interpretation and an efficient implementation strategy.

Keywords: Linear Types · FRP · Asynchrony · GUIs

Introduction

Many programs, like compilers, can be thought of as functions – they take a single input (a source file) and then produce an output (such as a type error message). Other programs, like embedded controllers, video games, and integrated development environments (IDEs), engage in a dialogue with their environment: they receive an input, produce an output, and then wait for a new input that depends on the prior input, and produce a new output which is in turn potentially based on the whole history of prior inputs.

The usual techniques for programming interactive applications are often confusing, since different parts of the program are not written to interact via structured control flow (e.g., by passing and return values from functions). Instead, they communicate indirectly, via state-manipulating callbacks which are implicitly invoked by an event loop. This makes program reasoning very challenging, since each of aliased mutable state, higher-order functions, and concurrency is tricky on its own, and interactive programs rely upon their *combination*.

This challenge has led to a great deal of work on better abstractions for programming reactive systems. Two of the main lines of work on this problem are *synchronous dataflow* and *functional reactive programming*. The synchronous

dataflow languages, like Esterel [5], Lustre [9], and Lucid Sychrone [28], feature a programming model inspired by Kahn networks. Programs are networks of stream-processing nodes which communicate with each other, each node consuming and producing a fixed number of primitive values at each clock tick. The first-order nature of these languages makes them strongly analysable, which lets them offer powerful guarantees on space and time usage. This means they see substantial use in embedded and safety-critical contexts.

Functional reactive programming, introduced by Elliott and Hudak [13], also uses time-indexed values, dubbed signals, rather than mutable state as its basic primitive. However, FRP differs from synchronous dataflow by sacrificing static analysability in favour of a much richer programming model. Signals are true first-class values, and can be used freely, including in higher-order functions and signal-valued signals. This permits writing programs with a dynamically-varying dataflow network, which simplifies writing programs (such as GUIs) in which the available signals can change as the program executes. Over the past decade, a long line of work has refined FRP via the Curry–Howard correspondence [21,18,17,19,20,10,1]. This approach views functional reactive programs as the programming counterpart for proofs of formulas in linear temporal logic [27], and has enabled the design of calculi which can rule out spacetime leaks [20] or can enforce temporal safety and liveness properties [10].

However, both synchronous dataflow and FRP (in both original and modal flavours) have a *synchronous* (or “pull”) model of time – time passes in ticks, and the program wakes up on every tick to do a little bit more computation. This is suitable for applications in which something new happens at every time step (e.g., video games), but many GUI programs like text editors and spreadsheets spend most of their time doing nothing. That is, even at each event, most of the program will continue doing nothing, and we only want to wake up a component when an event directly relevant to it occurs. This is important both from a performance point of view, as well as for saving energy (and extending battery life). Because of this need, most GUI programs continue to be written in the traditional callbacks-on-mutable-state style.

In this paper, we give a reactive programming language whose type system both has a very straightforward logical reading, and which can give natural types to stateful widgets and the event-based programming model they encourage. We also derive a denotational semantics of the language, by first working out a semantics of widgets in terms of the operations that can be performed upon them and the behaviour they should exhibit. Then, we find the categorical setting in which the widget semantics should live, and by studying the structure this setup has, we are able to interpret all of the other types of the programming language.

Contributions The contributions of this paper are:

- We give a descriptive semantics for widgets in GUI programming, and show that this semantics correctly models a variety of expected behaviours. For example, our semantics shows that a widget which is periodically re-set to the colour red is different from a widget that was only persistently set to

the colour red at the first timestep. Our semantic model can show that as long as neither one is updated, they look the same, but that they differ if they are ever set to blue – the first will return to red at reset time, and the second will remain blue.

- From this semantics, we find a categorical model within which the widget semantics naturally fits. This model is a Kripke–Joyal presheaf semantics, which is morally a “proof-relevant” Kripke model of temporal logic.
- We give a concrete calculus for event-based reactive programming, which can be implemented in terms of the standard primitives for modern GUI programming, scene graphs (or DOM) which are updated via callbacks invoked upon events. We then show that our model can soundly interpret the types of our calculus in an entirely standard way, showing that the types of our reactive programming language can be interpreted as time-varying sets.
- Furthermore, this calculus has an entirely standard logical reading in terms of the Curry–Howard correspondence. It is a “linear temporal linear logic”, with the linear part of the language corresponding to the Benton–Wadler [3] LNL calculus for linear logic, and the temporal part of the language corresponding to S4.3 linear temporal logic. We also give a proof term for the S_t 4.3 axiom enforcing the linearity of time, and show that it corresponds to the `select` primitive of concurrent programming.

The Language

We now present λ_{Widget} through the API of the `Widget` type. This API mirrors how one would work with a GUI at the browser level. An important feature of a well-designed GUI is that it should not do anything when not in use. In particular, it should not check for new inputs in each program cycle (*pull*-based reactive programming), but rather sleep until new data arrives (*push*-based reactive programming). Many FRP languages are *synchronous* languages and have some internal notion of a timestep. These languages are mostly pull-based, whereas more traditional imperative reactive languages are push-based. The former have clear semantics and are easy to reason about, the latter have efficient implementations. In λ_{Widget} we would like to combine these aspects and get a language that is easy to reason about with an efficient implementation.

In general, we think of a widget as a *state through time*, i.e., at each timestep, the widget is in some state which is presented to the user. The widget is modified by *commands*, which can update the state. To program with widgets, the programmer applies commands at various times.

The proper type system for a language of widgets should thus be a system with both state and time. If we consider what a *logic* for widgets should be, there are two obvious choices. A logic for state is linear logic [14], and a logic for time is linear temporal logic [27]. The combination of these two is the correct setting for a language of widgets, and, going through Curry–Howard, the corresponding type theory is a linear, linear temporal type theory.

Widget API To work with widgets, we define a API which mirrors how one would work with a browser level GUI:

$$\begin{aligned}
\text{newWidget} &: \mathbb{I} \multimap \exists (i : \text{Id}), \text{Widget } i \\
\text{dropWidget} &: \forall (i : \text{Id}), \text{Widget } i \multimap \mathbb{I} \\
\text{setColor} &: \forall (i : \text{Id}), \text{F Color} \otimes \text{Widget } i \multimap \text{Widget } i \\
\text{onClick} &: \forall (i : \text{Id}), \text{Widget } i \multimap \text{Widget } i \otimes \diamond \mathbb{I} \\
\text{onKeypress} &: \forall (i : \text{Id}), \text{Widget } i \multimap \text{Widget } i \otimes \diamond (\text{F Char}) \\
\text{out} &: \diamond A \multimap \exists (n : \text{Time}), A @ n \\
\text{into} &: \exists (n : \text{Time}), A @ n \multimap \diamond A \\
\text{split} &: \forall (i : \text{Id}) (t : \text{Time}), \text{Widget } i \multimap \text{Prefix } i t \otimes (\text{Widget } i) @ t \\
\text{join} &: \forall (i : \text{Id}) (t : \text{Time}), \text{Prefix } i t \otimes (\text{Widget } i) @ t \multimap \text{Widget } i
\end{aligned}$$

The first two commands creates and deletes widgets, respectively. The \multimap should be understood as *state passing*. We read the type of `newWidget` as “consuming no state, produce a new identifier index and a widget with that identifier index”. The identifier indices are used to ensure the correct behavior when using the `split` and `join` commands explained below. The existential quantification describes the *non-deterministic* creation of an identifier index. The use of non-determinism is crucial in our language and will be explaining in further detail in [section 1](#). Since λ_{Widget} has a linear type system, we need an explicit construction to delete state. For widgets, this is `dropWidget`. The type is read as “for any identifier index, consume a widget with that identifier index and produce nothing”.

The first command that modifies the state of a widget is `setColor`. Here we see the adjoint nature of the calculus with `F Color`. A color is itself *not* a linear thing, and as such, to use it in the linear setting, we apply `F`, which moves from the non-linear (Cartesian) fragment and into the linear fragment. The second new thing is the linear product \otimes . This differs from the regular non-linear product in that we do not have projection maps. Again, because of the linearity of our language, we cannot just discard state. We can now read the type of `setColor` as “Given a color and a identified widget, consume both and produce a new widget”. The produced widget is the same as the consumed widget, but with the color attribute updated.

The next two commands, `onClick` and `onKeypress`, are roughly similar. Both register a handle on the widget, for a mouse click and a key press, respectively. Here we see the first use of the \diamond modality, which represents an *event*. The type $\diamond A$ represents that *at some point in the future* we will receive something of type A . Importantly, because of the asynchronous nature of λ_{Widget} , we do not know *when* it happens. We can then read the type of `onClick` as “Consuming an identified widget, produce an updated widget together with a mouse click event”. The same holds for `onKeypress` except a key press event is produced.

The two commands `out` and `into` allows us to work with events in a more precise way. Given an event, we can use `out` to “unfold” it into an existential. The $@$ connective describes a type that is only available at a certain timestep, i.e., $A @ n$ means “at the timestep n , a term of type A will be available”. The `into` commands is the reverse of `out` and turns an existential and an $@$ into an event.

Note besides the above ways of constructing events, we can also turn any value into an event using the `evt` construction which is part of the core calculus. Given some element $a : A$, we get `evt a : $\diamond A$` which represents the event that returns immediately.

So far, we have only applied commands to a widget in the current timestep, but to program appropriately with widgets, we should be able to react to events and apply commands “in the future”. This is exactly what the `split` and `join` commands allows us to do. The type of `split` is read as “Given any time step and any identified widget, split the widget into all the states *before* that time and the widget *at* that time”. We denote the collection of states before a given time a *prefix* and give it the type `Prefix`. Given the state of the widget at a given timestep, we can now apply commands *at that timestep*. Note that both the prefix and the widget is indexed by the same identifier index. This is to ensure that when we use `join`, we combined the correct prefix and future.

Widget Programming To see the API in action, we now proceed with several examples of widget programming. For each example, we will add a comment on each line with the type of variables, and then explain the example in text afterwards.

One of the simplest things we can do with a widget is to perform some action when the widget is clicked. In the following example, we register a handler for mouse clicks, and then we use the click event to change the color of the widget to red at the time of the click. To do this, we use the `out` map to get the time of the event, then we split the widget and apply `setColor` at that point in the future.

```

1  turnRedOnClick :  $\forall (i : \text{Id}), \text{Widget } i \multimap \text{Widget } i$ 
2  turnRedOnClick i w0 =
3  let (w1, c0)      = onClick i w0 in   -- w1 : Widget i, c0 :  $\diamond I$ 
4  let unpack (x, c1) = out c0 in       -- x : Time, c1 : I @ x
5  let c2 @ x         = c1 in           -- c2 : I at x
6  let ⟨ ⟩ @ x         = c2 in
7  let (p, w2)       = split i x w1 in  -- p : Prefix i x, w2 : Widget i @ x
8  let w3 @ x         = w2 in           -- w3 : Widget i at time x
9  let w4             =                  -- w4 : Widget i @ x
10   (setColor (F Red) w3) @ x in
11  join i x (p, w4)

```

To see why this type checks, we go through the example line by line. In line 3, we register a handle for a mouse click on the widget. In line 4, we turn the click event into an existential. In line 5, we get c_2 which is a binding that is only available at the timestep x . Since we only need the *time* of the click, we discharge the click itself in line 6. In line 7 and 8, we split the widget using the timestep x and bind w_3 to the state of the widget at that timestep. In line 9-10, we change the color of the widget to red at x and in line 11 we recompose the widget.

In general, we will allow pattern matching in eliminations and since widget identity indices can always be inferred, we will omit them. In this style, the above example become:

```

1 turnRedOnClick :  $\forall (i : \text{Id}), \text{Widget } i \multimap \text{Widget } i$ 
2 turnRedOnClick  $w_0 =$ 
3   let  $(w_1, c_0) = \text{onClick } w_1 \text{ in}$  --  $w_1 : \text{Widget } i, c_0 : \diamond 1$ 
4   let  $\text{unpack } (x, \langle \rangle @ x) = \text{out } c_0 \text{ in}$  --  $x : \text{Time}$ 
5   let  $(p, w_2 @ x) = \text{split } x \ w_1 \text{ in}$  --  $p : \text{Prefix } i \ x, w_2 : \text{Widget } i \text{ at time } x$ 
6   join  $x \ (p, (\text{setColor } (\text{F Red}) \ w_2) @ x)$ 

```

We will use the same sugared style throughout the rest of the examples.

The above example turns a widget red exactly at the time of the mouse click, but will not do anything with successive clicks. To also handle further mouse clicks, we must register an event handler *recursively*. This is a simple modification of the previous code:

```

1 keepTurningRed :  $\forall (i : \text{Id}), \text{Widget } i \multimap \text{Widget } i$ 
2 keepTurningRed  $w_0 =$ 
3   let  $(w_1, c_0) = \text{onClick } w_1 \text{ in}$  --  $w_1 : \text{Widget } i, c_0 : \diamond 1$ 
4   let  $\text{unpack } (x, \langle \rangle @ x) = \text{out } c_0 \text{ in}$  --  $x : \text{Time}$ 
5   let  $(p, w_2 @ x) = \text{split } x \ w_1 \text{ in}$  --  $p : \text{Prefix } i \ x, w_2 : \text{Widget } i \text{ at time } x$ 
6   join  $(p, (\text{setColor } (\text{F Red}) \ (\text{keepTurningRed } w_2) @ x))$ 

```

By calling itself recursively, this function will make sure a widget will always turn red on a mouse click.

To understand the difference between two above examples, consider the code *turnBlueOnClick(keepTurningRed w)*, where w is some widget. On the first click, the widget will turn blue, on the second click it will turn red and on any subsequent click, it will keep turning red, i.e., stay red unless further modified.

When working with widgets, we will often register multiple handlers on a single widget. For example, a widget should have one behavior for a click and another behavior for a key press. To choose between two events, we use the *select* construction. This construction is central to our language and how to think about a push-based reactive language.

Given two events, $t_1 : \diamond A, t_2 : \diamond B$, there are three possible behaviors: Either t_1 returns first, and we wait for t_2 or t_2 returns first and we wait for t_1 or they return at the same time. In general, we want to select between n events, but if we need to handle all possible cases, this will give 2^n cases, so to keep the syntax linear in size, we will omit the last case. In the case events *actually* return at the same time, we do a non-deterministic choice between them. The syntax for *select* is

$$\text{select } (t_1 \text{ as } x \mapsto t'_1 \mid t_2 \text{ as } y \mapsto t'_2)$$

where $x : A, y : B, t'_1 : A \multimap \diamond B \multimap \diamond C$ and $t'_2 : B \multimap \diamond A \multimap \diamond C$. The second important thing to understand when working with *select* is that given we are working with events, we do not actually know at which timestep the events will trigger, and hence, we do not know what the (linear) context contains. Thus,

when using `select`, we will *only* know either $a : A, t_2 : \diamond B$ or $t_1 : \diamond A, b : B$. We can think of the `select` rule a *case-expression* that must respect time.

In the following example, we register two handlers, one for clicks and one for key presses, and change the color of the widget based on which returns first. We will only annotate the new parts.

```

1 widgetSelect :  $\forall (i : \text{Id}), \text{Widget } i \multimap \text{Widget } i$ 
2 widgetSelect w0 =
3   let (w1, c)           = onClick w0 in      -- c :  $\diamond \text{I}$ .
4   let (w2, k)           = onKeypress w1 in   -- k :  $\diamond(\text{F char})$ .
5   let col                =                  -- col :  $\diamond(\text{F Color})$ 
6     select
7       ( c as x  $\rightarrow$  let  $\langle \rangle = x$  in          -- x : I, k :  $\diamond(\text{F Color})$ .
8         let unpack (t,  $\langle \rangle$ ) @ t)
9         = out (mapE (fun F (-)  $\rightarrow$   $\langle \rangle$ ) k) in
10        evt (F Red)
11     | k as y  $\rightarrow$  let F k' = y in          -- y : F char, c :  $\diamond \text{I}$ 
12        let unpack (t,  $\langle \rangle$ ) @ t = c in
13        evt (F Blue))
14   let unpack (x, col' @ x) = out col in      -- col' : F Color at time x.
15   let (p, w3 @ x)         = split x w2 in
16   join (p, (setColor col' w3) @ x)

```

In line 3 and 4, we register the two handlers. In line 5-13, we use the `select` construction. In the first case, the click happens first and we return the color red. In the second case, the key press happens first and we return the color blue. In both cases, because of the linear nature of the language, we need to discharge the unit and char, respectively, and the event that does not return first. In line 14, we turn the color event into an existential. In line 15, we use the timestep of the color event to split the widget, and in line 16, we change the color of the widget at that time and recompose it.

To see how λ_{Widget} differs from more traditional synchronous FRP languages, we will examine how to encode a kind of streams. Since our language is *asynchronous*, the stream type must be encoded as

$$\text{Str } A := \nu \alpha. \diamond(A \otimes \alpha)$$

This asynchronous stream will *at some point in the future* give a head and a tail. We do not know when the first element of the stream will arrive, and after each element of the stream is produced, we will wait an indeterminate amount of time for the next element. The reason why the stream type in λ_{Widget} must be like this is essentially that we want a *push-based* language, i.e., we do not want to wake up and check for new data in each program cycle. Instead, the program should sleep until new data arrives.

To show the difference between the asynchronous stream and the more traditional synchronous stream, we will look at some examples. With a traditional stream, a standard operation is zipping two streams: that is, given $\text{Str } A$ and $\text{Str } B$, we can produce $\text{Str } A \times B$, which should be the element-wise pairing of the two streams. It should be clear that this is not possible for our asynchronous

streams. Given two streams, we can wait until the first stream produces an element, but the second stream may only produce an element after a long period of time. Hence, we would need to buffer the first element, which is not supported in general. Remember, when using `select`, we can not use any already defined linear variables, since we do not know if they will be available in the future.

Rather than zipping stream, we can instead do a kind of *interleaving* as shown below. We use `fold` and `unfold` to denote the folding and unfolding of the fixpoint.

```

1 interleave : Str A  $\multimap$  Str B  $\multimap$  Str (A  $\oplus$  B)
2 interleave xs ys = fold (
3   select
4     (unfold xs as xs'  $\rightarrow$ 
5       let (x, xs'') = xs' in -- xs' : A  $\otimes$  Str A, x : A, xs'' : Str A
6         evt (inl x, interleave xs'' ys)
7     | unfold ys as ys'  $\rightarrow$ 
8       let (y, ys'') = ys' in -- ys' : B  $\otimes$  Str B, y : B, ys'' : Str A
9         evt (inr y, interleave xs ys''))

```

Here, we use `select` to choose between which stream returns first, and then we let that element be the first element of the new stream.

On the other hand, some of the traditional FRP functions on streams can be translated. For instance, we can map of function over a stream, given that *it is available at each step in time*:

```

1 map : F (G (A  $\multimap$  B))  $\multimap$  Str A  $\multimap$  Str B
2 map f0 xs =
3   let F f1 = f0 in -- f1 : G(A  $\multimap$  B)
4   let (y, (x, xs') @ y) = -- y : Time, x : A, xs' :  $\diamond$ Str A at time y
5     out (unfold xs) in
6     fold (evt ((runG f1) x, map f0 xs'))

```

The type $F(G(A \multimap B))$ is read as a linear function with no free variables that can be used in a non-linear fashion, i.e., duplicated. This restriction to such “globally available functions” is reminiscent of the “box” modality in Bahr et al. [1] and Krishnaswami [20], and the F and G construction can be understood as decomposing the box modality into two separate steps. This relationship will be made precise in the logical interpretation of λ_{Widget} in section 1

As a final example, we will show how to dynamically update the GUI, i.e., how to add new widgets on the fly. Before we can give the example, we need to extend our widget API, to allow composition of widgets. To that end, we add the `vAttach` command to our API.

$$\text{vAttach} : \forall(i, j : \text{Id}), \text{Widget } i \multimap \text{Widget } j \multimap \text{Widget } i$$

This command should be understood as an abstract version the `div` tag in HTML. In the following example, we think of the widget as a simple button that when clicked, will create a new button. When *any* of the buttons gets clicked, a new button gets attached.


```

1 buttonStack :  $\forall i, \text{Widget } i \multimap \text{Widget } i$ 
2 buttonStack  $w_0 =$ 
3   let ( $w_1, c$ ) = onClick  $w_0$  in
4   let ( $x, \langle \rangle @ x$ ) = out  $e$  in
5   let ( $p, w_2 @ x$ ) = split  $x w_1$  in
6   let  $w_3$  = (let ( $y, w$ ) = newWidget  $\langle \rangle$  in
7             vAttach  $w_2$  (buttonStack  $w$ )) @  $x$  in
8   join ( $p, w_3$ )
    
```

The important step here is in line 6 and 7. Here the new button is attached at the time of the mouse click, and *buttonStack* is called recursively on the newly created button.

Formal Calculus

This sections gives the rules, meta-theory and logical interpretation of λ_{Widget} . Briefly, the language is a mixed linear-non-linear adjoint calculus in the style of Benton–Wadler [4,3]. The non-linear fragment, also called Cartesian in the following, is a minimal simply typed lambda calculus whereas the linear fragment contains several non-standard judgments used for widget programming.

Contexts and Typing Judgments We have three typing judgments: one for indices, one for Cartesian (non-linear) terms, and one for linear terms. These are distinguished by a subscript on the turnstile, i for indices, c for Cartesian terms and l for linear terms. These depend on different contexts. The index judgment depends only on a index context, whereas the Cartesian and linear judgments depends on both an index and a linear and/or a Cartesian context. The rules for context formation is given in Figure 1. These are mostly standard except for the dependence on a previously defined context and the fact that the linear context contains variables of the form $a :_{\tau} A$, i.e., temporal variables. The judgment $a :_{\tau} A$ is read as “ a has the type A at the timestep τ ”. In the linear setting we will write $a : A$ instead of $a :_0 A$, i.e., a judgment in the current timestep.

Indices:	$\frac{}{\vdash_i \cdot}$	$\frac{\vdash_i \Theta \quad s \notin \text{dom}(\Theta) \quad \sigma \in \{\text{Id}, \text{Time}\}}{\vdash_i \Theta, s : \sigma}$
Cartesian:	$\frac{}{\vdash_c \cdot}$	$\frac{\Theta \vdash_c \Gamma \quad x \notin \text{dom}(\Gamma) \quad \Theta \vdash_c X}{\Theta \vdash_c \Gamma, x : X}$
Linear:	$\frac{}{\vdash_l \cdot}$	$\frac{\Theta \vdash_l \Delta \quad x \notin \text{dom}(\Delta) \quad \Theta \vdash_l A \quad \Theta \vdash_l \tau : \text{Time}}{\Theta \vdash_l \Delta, a :_{\tau} A}$

Fig. 1. Context Formation

The index judgment describes how to introduce indices. The typing rules are given in Figure 2. The judgment $\Theta \vdash_i \tau : \sigma$ contains a single context, Θ , for index variables. There are only two sorts of indices, identifiers and timesteps.

Index Judgments:

$$\frac{\tau \in \text{Time}}{\Theta \vdash_i \tau : \text{Time}} \text{TIME} \qquad \frac{\iota \in \text{Id}}{\Theta \vdash_i \iota : \text{Id}} \text{ID} \qquad \frac{i : \sigma \in \Theta}{\Theta \vdash_i i : \sigma} \text{VAR}$$

Fig. 2. Index Typing rules

The Cartesian judgment describes the Cartesian, or non-linear, fragment. This is a minimal simply typed lambda calculus with the addition of the G type, used for moving between the linear and Cartesian fragment, and explained further below. The judgment $\Theta; \Gamma \vdash_c t : A$ has two contexts; Θ for indices and Γ for Cartesian variables.

The linear fragment is most of the language, and a selection of typing rules is given in [Figure 3](#). The judgment is done w.r.t three contexts, Θ for index variables, Γ for Cartesian variables and Δ for linear variables. Many of the rules are standard for a linear calculus, except for the presence of the additional contexts. We will not describe the standard rules any further.

The first non-standard rule is for \diamond . The introduction and elimination rules follow from the fact that \diamond is a non-strong monad. More interesting is the *select* rule. Here we see the formal rule corresponding to the informal explanation in [section 1](#). The important thing here is that we can not use any previously defined linear variable when typing t'_1 and t'_2 , since we do not actually know *when* the typing happens. Note, we can see the *select* rule as a binary version of the \diamond let-binding. This could be extended to an n -ary version, but we do not do this in our core calculus. The rules for $A @ \tau$ shows how to move between the judgment $t : A @ \tau$ and $t :_{\tau} A$. That is, moving from knowing in the current timestep that t will have the type A at time τ and knowing at time τ that t has type A . The (F-I), (F-E), (G-I) and (G-E) rules show the adjoint structure of the language. The (G-I) rule takes a closed linear term of type A and gives it the Cartesian type $\mathsf{G} A$. Note, because it has no free linear variables, it is safe to duplicate. The (G-E) rule lets us get an A without needing any linear resources. Conversely, the (F-I) rule embeds a intuitionistic term into the linear fragment and the (F-E) rule binds an intuitionistic variable to let us freely use the value. The (Delay) rule shows what happens when we actually *know* the timestep. The important part is $\Delta' = \Delta \downarrow^{\tau}$ which means two things. One, all the variables in Δ are on the form $a :_{\tau} A$, i.e., judgments at time τ and two, we shift Δ into the future such that all the variables of Δ' is of the form $a : A$. The way to understand this is, if all the variables in Δ are typed at time τ and the conclusion is at time τ , it is enough to “move to” time τ and then type w.r.t that timestep. Finally, we have $(!_{\tau}\text{-E})$ and $(\otimes_{\tau}\text{-E})$. These allow us to work with linear unit and products at time τ . These are added explicitly since they can not be derived by the other rules, and are needed for typing certain kinds of programs.

Unfolding Events to Exists The type system as given above contains both $\diamond A$ and $A @ k$, as two distinct ways to handle time. The former means that

$$\begin{array}{c}
\frac{\Theta \vdash_i \tau : \text{Time} \quad \Theta; \Gamma; \Delta \vdash_l t :_{\tau} A}{\Theta; \Gamma; \Delta \vdash_l t @ \tau : A @ \tau} \text{(@-I)} \\
\\
\frac{\Theta \vdash_i t : \text{Time} \quad \Theta; \Gamma; \Delta_1 \vdash_l t_1 : A @ \tau \quad \Theta; \Gamma; \Delta_2, a :_{\tau} A \vdash_l t_2 : B}{\Theta; \Gamma; \Delta_1, \Delta_2 \vdash_l \text{let } a @ \tau = t_1 \text{ in } t_2 : B} \text{(@-E)} \\
\\
\frac{\Theta; \Gamma \vdash_c e : \mathbf{G} A}{\Theta; \Gamma; \cdot \vdash_l \text{runG } e : A} \text{(G-E)} \qquad \frac{\Theta; \Gamma \vdash_c e : X}{\Theta; \Gamma; \cdot \vdash_l \mathbf{F} e : \mathbf{F} x} \text{(F-I)} \\
\\
\frac{\Theta; \Gamma; \Delta_1 \vdash_l t_1 : \mathbf{F} X \quad \Theta; \Gamma, x : X; \Delta_2 \vdash_l t_2 : B}{\Theta; \Gamma; \Delta_1, \Delta_2 \vdash_l \text{let } \mathbf{F} x = t_1 \text{ in } t_2 : B} \text{(F-E)} \\
\\
\frac{\Theta, i : \sigma; \Gamma; \Delta \vdash_l t : A}{\Theta; \Gamma; \Delta \vdash_l \Lambda(i : \sigma).t : \forall(i : \sigma).A} \text{(\forall-I)} \quad \frac{\Theta \vdash_i s : \sigma \quad \Theta; \Gamma; \Delta \vdash_l t : \forall(i : \sigma).A}{\Theta; \Gamma; \Delta \vdash_l t_s : \{s/i\}A} \text{(\forall-E)} \\
\\
\frac{\Theta \vdash_i s : \sigma \quad \Theta; \Gamma; \Delta \vdash_l t : \{s/i\}A}{\Theta; \Gamma; \Delta \vdash_l \{s, t\} : \exists(i : \sigma).A} \text{(\exists-I)} \\
\\
\frac{\Theta; \Gamma; \Delta_1 \vdash_l t_1 : \exists(i : \sigma).A \quad \Theta, s : \sigma; \Gamma; \Delta_2, a : \{s/i\}A \vdash_l t_2 : B}{\Theta; \Gamma; \Delta_1, \Delta_2 \vdash_l \text{let unpack } \{s, a\} = t_1 \text{ in } t_2 : B} \text{(\exists-E)} \\
\\
\frac{\Theta; \Gamma; \Delta_1 \vdash_l t_1 : \diamond A \quad \Theta; \Gamma; \Delta_2 \vdash_l t_2 : \diamond B \quad \Theta; \Gamma; a : A, t_2 : \diamond B \vdash_l t'_1 : \diamond C \quad \Theta; \Gamma; b : B, t_1 : \diamond A \vdash_l t'_2 : \diamond C}{\Theta; \Gamma; \Delta_1, \Delta_2 \vdash_l \text{select } (t_1 \text{ as } a \mapsto t'_1 \mid t_2 \text{ as } b \mapsto t'_2) : \diamond C} \text{(SELECT)} \\
\\
\frac{\Theta \vdash_i \tau : \text{Time} \quad \Delta' = \Delta \downarrow^{\tau} \quad \Theta; \Gamma; \Delta' \vdash_l t : A}{\Theta; \Gamma; \Delta \vdash_l t :_{\tau} A} \text{(DELAY)} \\
\\
\frac{\Theta \vdash_i \tau : \text{Time} \quad \Theta; \Gamma; \Delta_1 \vdash_l t_1 :_{\tau} \mathbf{I} \quad \Theta; \Gamma; \Delta_2 \vdash_l t_2 : B}{\Theta; \Gamma; \Delta_1, \Delta_2 \vdash_l \text{let } \langle \rangle @ \tau = t_1 \text{ in } t_2 : B} \text{(I}_{\tau}\text{-E)} \\
\\
\frac{\Theta \vdash_i \tau : \text{Time} \quad \Theta; \Gamma; \Delta_1 \vdash_l t_1 :_{\tau} A @ B \quad \Theta; \Gamma; \Delta_2, a :_{\tau} A, b :_{\tau} B \vdash_l t_2 : C}{\Theta; \Gamma; \Delta_1, \Delta_2 \vdash_l \text{let } (a, b) @ \tau = t_1 \text{ in } t_2 : C} \text{(@}_{\tau}\text{-E)}
\end{array}$$

Fig. 3. Selected Linear Typing rules

something of type A will arrive at *some* point in the future, whereas the latter means an A arrives at a *specific* point in the future. The strength of \diamond is that it gives easy and concise typing rules, whereas the strength of $A@k$ is that it allows for a more precise usage of time. To connect these two, we add the linear isomorphism $\diamond A \cong \exists k. A@k$ to our language, which is witnessed by `out` and `into`, as part of the widget API. This isomorphism is true semantically, but can not be derived in the type system. In particular, this isomorphism allows the `select` rule to be given with \diamond , while still allowing the use timesteps when working with the resulting event. If we were to give the equivalent definition using timesteps, one would need to have some sort of *constraint system* for deciding which events happens first. Avoiding such constraints also allows for a simpler implementation, as everything in our type system can be inferred.

Meta-theory of Substitution The meta-theory of λ_{Widget} is given in the form of a series of substitution lemmas. Since we have three different contexts, we will end up with six different substitutions into terms. The Cartesian to Cartesian, Cartesian to linear and linear to linear are the usual notion of mutual recursive substitution. More interesting is the substitution of indices into Cartesian and linear terms and types. We prove the following lemma, showing that typing is preserved under index substitution:

Lemma 1 (Preservation of Typing under Index Substitution).

$$\frac{\zeta : \Theta' \rightarrow \Theta \quad \Theta; \Gamma \vdash_c e : X}{\Theta'; \zeta(\Gamma) \vdash_c \zeta(e) : \zeta(X)} \quad \frac{\zeta : \Theta' \rightarrow \Theta \quad \Theta; \Gamma; \Delta \vdash_l t :_{\tau} A}{\Theta'; \zeta(\Gamma); \zeta(\Delta) \vdash_l \zeta(t) :_{\tau} \zeta(A)}$$

Both are these (and all other cases for substitution) are proved by a lengthy but standard induction over the typing tree. See the technical appendix for full proofs of all six substitution lemmas.

Logical Interpretation Our language has a straightforward logical interpretation. The logic corresponding to the Cartesian fragment is a propositional intuitionistic logic, following the usual Curry–Howard interpretation. The logic corresponding to the substructural part of the language is a linear, linear temporal logic. The single-use condition on variables means that the syntax and typing rules correspond to the rules of intuitionistic linear logic (i.e., the first occurrence of linear in “linear, linear temporal”). However, we do not have a comonadic exponential modality $!A$ as a primitive. Instead, we follow the Benton–Wadler approach [4,3] and decompose the exponential into the composition of a pair of adjoint functors mediating between the Cartesian and linear logic.

In addition to the Benton–Wadler rules, we have a temporal modality $\diamond A$, which corresponds to the eventually modality of linear temporal logic (i.e., the second occurrence of “linear” in “linear, linear temporal logic”). This connective is usually written $F A$ in temporal logic, but that collides with the F modality of the Benton–Wadler calculus. Therefore we write it as $\diamond A$ to reflect its nature as a possibility modality (or monad). In our calculus, the axioms of S4.3 are

derivable:

$$\begin{aligned} (T) &: A \multimap \diamond A \\ (4) &: \diamond \diamond A \multimap \diamond A \\ (.3) &: \diamond(A \otimes B) \multimap \diamond((\diamond A \otimes B) \oplus \diamond(A \otimes \diamond B) \oplus \diamond(A \otimes B)) \end{aligned}$$

Since the ambient logic is linear, intuitionistic implication $X \rightarrow Y$ is replaced with the linear implication $A \multimap B$, and intuitionistic conjunction $X \wedge Y$ is replaced with the linear tensor product $A \otimes B$. It is easy to see that the first two axiom corresponds to the monadic structure of \diamond , and the .3 axiom corresponds to the *select* rule (with our syntax for *select* corresponding to immediately waiting for and then pattern-matching on the sum type). In the literature, the .3 axiom is often written in terms of the box modality $\Box A$ [8], but we present it here in a (classically) equivalent formulation mentioning the eventually modality $\diamond A$. We do not need to an explicit box modality $\Box A$, since the decomposition of the exponential $F(GA)$ from the linear-non-linear calculus serves that role.

In our system, *we do not offer* the next-step operator $\triangleright A$. Since we model asynchronous programs, we do not let programmers write programs which wake up in a specified amount of time. We only offer an iterated version of this connective, $A @ n$, which can be interpreted as $\triangleright^n A$, and our term syntax has no numeric constants which can be used to demand a specific delay.


Finally, the universal and existential quantifiers (in both the intuitionistic and linear fragments) are the usual quantifier rules for first-order logic.

Semantics

In this section we give a denotational model for λ_{Widget} . It is a linear-non-linear (LNL) hyperdoctrine [24,16] with the non-linear part being **Set** and the linear part being the category of internal relations over a suitable “reactive” category. The hyperdoctrine structure is used to interpret the quantification over indices. This model is nearly entirely standard: the most interesting thing is the reactive base category and the interpretation of widgets. It is well known that any symmetric monoidal closed category (SMCC) models multiplicative intuitionistic linear logic (MILL), and it is similarly well known that the category of relations over **Set** can be give the structure of a SMCC by using the Cartesian product as both the monoidal product and monoidal exponential. This construction lift directly to any category of internal relations over a category that is suitably “Set-like”, i.e., a topos. Our base category is a simple presheaf category, and hence, we use this construction to model the linear fragment of λ_{Widget} .

The Base Reactive Category The base reactive category is where the notion of time will arise and is it this notion that will be lifted all the way up to the LNL hyperdoctrine. The simplest model of “time” is $\mathbf{Set}^{\mathbb{N}}$, which can be understood as “sets through time” [23]. This can indeed be used as a model for a reactive setting, but for our purposes it is too simple, and further, depending on which ordering is considered for \mathbb{N} , may have undesirable properties for the reactive

setting. Instead, we use the only slightly more complicated $\text{Set}^{\mathbb{N}+1}$, henceforth denoted \mathcal{R} , where the ordering on $\mathbb{N} + 1$ is the discrete ordering on \mathbb{N} and 1 is related to everything else. Adding this “point at infinity” allows global reasoning about objects, an intuition that is further supported by the definition of the sub-object classifier below. Further, this model is known to be able to differentiate between least and greatest fixpoints [15], and even though we do not use this for λ_{Widget} , we consider it a useful property for further work (see section 1). Objects in \mathcal{R} can be visualized as



We can think of A_∞ as the global view of the object and A_n as the local view of the object at each timestep. Morphisms are natural transformations between such diagrams and the naturality condition means that having a map from A_∞ to B_∞ must also come with coherent maps at each timestep.

In \mathcal{R} we define two endofunctors, which can be seen as describing the passage of time:

Definition 1. We define the later and previous endofunctors on \mathcal{R} , denoted \triangleright and \triangleleft , respectively:

$$(\triangleright A)_n := \begin{cases} 1 & n = 0 \\ A_{n'} & n = n' + 1 \\ A_\infty & n = \infty \end{cases} \quad (\triangleleft A)_n := \begin{cases} A_{n+1} & n \neq \infty \\ A_\infty & n = \infty \end{cases}$$

Note that when we apply the later functor, the global view does not change, but the local views are shifted forward in time.

Theorem 1. The later and previous endofunctors form an adjunction.

Definition 2. The sub-object classifier, denoted Ω , in \mathcal{R} is the object

$$\Omega_\infty = \mathcal{P}(\mathbb{N}) + 1 \quad \Omega_n = \{0, 1\}$$

For each $n \in \mathbb{N}$, Ω_n denotes whether a given proposition is true at the n th timestep. Ω_∞ gives the “global truth” of a given proposition. The left injection is some subset of \mathbb{N} that denotes at which points in time something is true. The right injection denotes that something is true “at the limit”, and in particular, also at all timesteps. Note, a proposition can be true at all timesteps but not at the limit. This extra point at infinity is precisely what allows us the differentiate between least and greatest fixpoints.

The Category of Internal Relations To interpret the linear fragment of the language, we will use the category of internal relations on \mathcal{R} . Given two objects A and B in \mathcal{R} , an *internal relation* is a sub-object of the product $A \times B$. This can equivalently be understood as a map $A \times B \rightarrow \Omega$. The category of internal relations in the category where the objects are the objects of \mathcal{R} and the morphisms $A \rightarrow B$ are internal relations $A \times B \rightarrow \Omega$ in \mathcal{R} . We denote the category of internal relations as $\text{Rel}_{\mathcal{R}}$.

Theorem 2. Using $A \otimes B = A \times B$ and $A \multimap B = A \times B$ as monoidal product and exponential, respectively, $\text{Rel}_{\mathcal{R}}$ is a symmetric monoidal closed category.

Theorem 3. There is an adjunction $\triangleleft \vdash \triangleright$ in $\text{Rel}_{\mathcal{R}}$ where \triangleleft and \triangleright are the lifting of the previous and later functors from \mathcal{R} to $\text{Rel}_{\mathcal{R}}$.

Definition 3. We define the iterated later modality or the “at” connective as a successive application of the later modality.

$$\begin{aligned}\triangleright^0 A &= A \\ \triangleright^{(k+1)} A &= \triangleright(\triangleright^k A)\end{aligned}$$

and we will alternatively write $A @ k$ to mean $\triangleright^k A$.

Definition 4. We define the event functor on $\text{Rel}_{\mathcal{R}}$ as an iterated later.

$$\begin{aligned}\diamond A &: \text{Rel}_{\mathcal{R}} \rightarrow \text{Rel}_{\mathcal{R}} \\ (\diamond A)_{\infty} &= A_{\infty} \\ (\diamond A)_n &= \Sigma(k : \mathbb{N}).(\triangleright^k A)_n\end{aligned}$$

The event functor additionally carries a monadic structure (see [29] and the technical appendix).

Theorem 4. We have the isomorphism $\diamond A \cong \Sigma(n : \mathbb{N}).A @ n$ for any A

Theorem 5. We have the following adjunctions between Set , \mathcal{R} and $\text{Rel}_{\mathcal{R}}$:

$$\begin{array}{ccc} \text{Set} & \begin{array}{c} \xrightarrow{\Delta} \\ \perp \\ \xleftarrow{\text{lim}} \end{array} & \mathcal{R} & \begin{array}{c} \xrightarrow{I} \\ \perp \\ \xleftarrow{P} \end{array} & \text{Rel}_{\mathcal{R}} \end{array}$$

where Δ is the constant functor, lim is the limit functor, I is the inclusion functor and P is the image functor. This induces an adjunction between Set and $\text{Rel}_{\mathcal{R}}$.

The Widget Object One of the most important objects in $\text{Rel}_{\mathcal{R}}$ is the *widget* object. This object is used to interpret widgets and prefixes. The widget object will be defined with respect to an ambient notion of identifiers, which we will denote Id . These will be part of the hyperdoctrine structure define below, and for now, we will just assume such an object to exists. We will also use a notion of timesteps internal to the widget object. Note that this timestep is different from the abstract timestep used for defining $\text{Rel}_{\mathcal{R}}$, but are related as defined below. We denote the abstract timesteps with Time .

Before we can define the widget object, we need to define an appropriate object of commands. In our minimal Widget API, the only *semantic* commands will be `setColor`, `onClick` and `onKeyPress`. The rest of the API is defined as morphisms on the widget object itself. To work with the semantics commands, we additionally need a *compatibility* relation. This relation describes what commands can be applied at the same time. In our setting this relation is minimal, but can in principle be used to encode whatever restrictions is needed for a given API.

Definition 5. We define the command object as

$$\text{Cmd} = \{\text{setColor}, \text{color}\}, \text{onClick}, \text{onKeypress}\}$$

where color is an element of a “color” object. The compatibility relations are:

$$(\text{op}, \text{arg}) \bowtie (\text{op}', \text{arg}') \text{ iff } (\text{op} = \text{op}' \Rightarrow \text{arg} = \text{arg}')$$

The only non-compatible combination of commands is two application of the `setColor` command, the idea being that you can not set the color twice in the same timestep.

We can now define the widget and prefix objects

Definition 6. The widget object, denoted Widget , is indexed by $i \in \text{Id}$ and is defined as

$$\text{Widget}_\infty i = \{(w, i) \mid w \in \mathcal{P}(\text{Time} \times \text{Cmd}), (t, c) \in w \wedge (t, c') \in w \rightarrow c \bowtie c'\}$$

$$\text{Widget}_n i = \{(w, i) \subset \text{Widget}_\infty i \mid \forall (t, c) \in w, t \leq n\}$$

The prefix object, denoted Prefix , is indexed by $i \in \text{Id}$ and $t \in \text{Time}$ and is:

$$\text{Prefix}_\infty i t = \{(P, i) \subset \text{Widget}_\infty i \mid \forall (t', c) \in P, t' \leq t\}$$

$$\text{Prefix}_n i t = \begin{cases} \{(P, i) \subset \text{Prefix}_\infty i t \mid \forall (t', c) \in P, t' \leq n\} & n < t \\ \mathbb{1} & \text{otherwise} \end{cases}$$

The widget object is a collection of times and commands keeping track of what has happened to it at various times – imagine a *logbook* with entries for each time step. At the point at infinity, the “global” behavior of the widget is defined, i.e., the full logbook of the widget. For each n , Widget_n is simply what has happened to the widget so far, i.e., a truncated logbook. The prefix object is a widget object that is only defined up to some timestep, and is the unit after that. This yields a semantic difference between the widget where the color is set only once, and the widget where the color is set at every timestep. This reflects a real difference in actual widget behavior: if *turnRedOnClick* w later set to be blue, it will remain blue, but *keepTurningRed* w will turn back to being red.

To manipulate widgets we define two “restriction” maps.

Definition 7. We define the following on widgets and prefixes

$$\text{shift } t : \text{Widget } i \rightarrow_{\text{Rel}_{\mathcal{R}}} \text{Widget } i$$

$$(\text{shift } t W)_n = \{(t' - t, c) \mid (t', c) \in W \wedge t \leq t'\}$$

$$\text{prefix } t i : \text{Widget } i \rightarrow_{\text{Rel}_{\mathcal{R}}} \text{Prefix } i t$$

$$(\text{prefix } t i W)_n = \begin{cases} \{(t', c) \in W \mid t' < t\} & n < t \\ \mathbb{1} & n \geq t \end{cases}$$

The intuition behind these is that $\text{prefix } t i$ “cuts off” the widget after t , giving a prefix, whereas $\text{shift } t$ shifts forward all entries in the widget by t .

Using the above, we can now define the `split` and `join` morphisms. These are again given w.r.t ambient Id and Time objects, which will be part of the full hyperdoctrine structure:

Definition 8. We define the following morphisms on the widget object

$$\begin{aligned} \text{split } i t &: \text{Widget } i \rightarrow_{\text{Rel}_{\mathcal{R}}} \text{Prefix } i t \otimes \text{Widget } i @ t \\ (\text{split } i t w)_n &= (\text{prefix } t i w, \text{shift } t w)_n \\ \\ \text{join } i t &: \text{Prefix } i t \otimes \text{Widget } i @ t \rightarrow_{\text{Rel}_{\mathcal{R}}} \text{Widget } i \\ (\text{join } i t (p, w))_n &= \begin{cases} p_n & n < t \\ w_{n-t} & n \geq t \end{cases} \end{aligned}$$

Linear-non-linear Hyperdoctrine So far we have not explained in details how to model the quantifiers in our system. To do this, we use the notion of a *hyperdoctrine* [22]. For first-order logic, this is a functor from a category of contexts and substitutions to the category of Cartesian closed categories, with the idea that we have one CCC for each valuation of the free first-order variables.

As our category of contexts, we use a Cartesian category to interpret our index objects, `Time` and `Id`. The former is interpreted as $\mathbb{N} + 1$ and the latter as \mathbb{N} . In our case, both `Set` and $\text{Rel}_{\mathcal{R}}$ are themselves hyperdoctrines w.r.t to this category of contexts, the former a first-order hyperdoctrine and the latter a multiplicative intuitionistic linear logic (MILL) hyperdoctrine. Together these form a linear-non-linear hyperdoctrine through the adjunction given in [Theorem 5](#).

Definition 9. A linear-non-linear hyperdoctrine is a MILL hyperdoctrine L together with a first-order hyperdoctrine C and a fiber-wise monoidal adjunction $F : L \rightleftarrows C : G$.

Theorem 6. The categories `Set` and $\text{Rel}_{\mathcal{R}}$ form a linear-non-linear hyperdoctrine w.r.t the interpretation of the indices objects, with the adjunction given as in [Theorem 5](#).

We refer the reader to the accompanying technical appendix for the full details.

Denotational Semantics We the above, we have enough structure to give an interpretation of λ_{Widget} . Again, most of this interpretation is standard in the use of the hyperdoctrine structure, and we interpret \diamond in the obvious way using the linear hyperdoctrine structure on $\text{Rel}_{\mathcal{R}}$. As an example, we sketch the interpretation of the widget object and the `setColor` command below.

Definition 10. We interpret the `Widget` i and `Prefix` i types using the widget and prefix objects:

$$\begin{aligned} \llbracket \Theta \vdash \text{Widget } i \rrbracket &= \text{Widget } \llbracket \Theta \vdash_s i : \text{Id} \rrbracket \\ \llbracket \Theta \vdash \text{Prefix } i t \rrbracket &= \text{Prefix } \llbracket \Theta \vdash_s i : \text{Id} \rrbracket \llbracket \Theta \vdash_s t : \text{Time} \rrbracket \end{aligned}$$

and we interpret the `setColor` commands as:

$$\begin{aligned} \llbracket \text{setColor} : \forall (i : \text{Id}), \text{Widget } i \otimes \text{F Color} \multimap \text{Widget } i \rrbracket = \\ \{w \cup_W \{(0, (\text{setColor}, col))\} \mid w \in \llbracket \text{Widget } i \rrbracket, col \in \llbracket \text{Color} \rrbracket\} \end{aligned}$$

where \cup_W is a “widget union”, which is a union of sets such that identifiers indices and compatibility of commands are respected

This interpretation shows that a widget is indeed a logbook of events. Using the `setColor` command simply adds an entry to the logbook of the widget. Note we only set the color in the current timestep. To set the color in the future, we combine the above with appropriate uses of splits and joins. The interpretation of `split` and `join` are done using their semantic counterparts, and the interpretation of `onClick` and `onKeyPress` are done, using our non-deterministic semantics, by associating a widget with *all possible occurrences* of the corresponds event.

Soundness of Substitution Finally, we prove that semantic substitution is sound w.r.t syntactic substitution. As with the proofs of type preservation for syntactic substitution, there are several cases for the different kinds of substitution, but the main results is again concerned with substitution of indices:

Theorem 7. *Given $\zeta : \Theta' \rightarrow \Theta$, $\Theta; \Gamma \vdash_c e : X$ and $\Theta; \Gamma; \Delta \vdash_l t : A$ then*

$$\begin{aligned} \llbracket \zeta \rrbracket \llbracket \Theta; \Gamma \vdash_c e : X \rrbracket &= \llbracket \Theta'; \zeta(\Gamma) \vdash_c \zeta(e) : \zeta(X) \rrbracket \\ \llbracket \zeta \rrbracket \llbracket \Theta; \Gamma; \Delta \vdash_l t : A \rrbracket &= \llbracket \Theta'; \zeta(\Gamma); \zeta(\Delta) \vdash_l \zeta(t) : \zeta(A) \rrbracket \end{aligned}$$

Proofs for all six substitutions lemmas can be found in the technical appendix.

Related and Future Work

Much work has aimed at a logical perspective on FRP via the Curry–Howard correspondence [21,18,17,19,20,10,1]. As mentioned earlier, most of this work has focused on calculi that have a Nakano-style later modality [25], but this has the consequence that it makes it easy to write programs which wake up on every clock tick. In this paper, we remove the explicit next-step modality from the calculus, which opens the door to a more efficient implementation style based on the so-called “push” (or event/notification-based) implementation style. Elliott [12] also looked at implementing a push-based model, but viewed it as an optimization rather than a first-class feature in its own right. In future work, we plan on implementing a language based upon this calculus, with the idea that we can compile to Javascript, and represent widgets with DOM nodes, and represent the $\diamond A$ and $A @ n$ temporal connectives using doubly-negated callback types (in Haskell notation, `Event A = (A -> IO ()) -> IO ()`). This should let us write GUI programs in functional style, while generating imperative, callback-based code in the same style that a handwritten GUI program would use.

Our model, in terms of $\text{Set}^{\mathbb{N}+1}$, enriches LTL’s semantics from time-indexed truth-values to time-indexed sets. The addition of the global view or point at infinity enables our model to distinguish between least and greatest fixed points [15] (i.e., inductive and coinductive types), unlike in models of guarded recursion where guarded types are bilimit-compact [6]. This lets us encode temporal liveness and safety properties using inductive and coinductive types [10,2].

A recent development for comonadic modalities is the introduction of the so-called ‘Fitch-style’ calculi [7,11] as an alternative to the Pfenning–Davies pattern-style elimination [26]. These calculi have been used successfully for FRP [1], and one interesting question is whether they extend to adjoint calculi as well – i.e., can the $F(X)$ modality support a direct-style eliminator?

References

1. Bahr, P., Graulund, C.U., Møgelberg, R.E.: Simply RaTT: a Fitch-style modal calculus for reactive programming without space leaks. *Proceedings of the ACM on Programming Languages* **3**(ICFP), 1–27 (2019)
2. Bahr, P., Graulund, C.U., Møgelberg, R.: Diamonds are not forever: Liveness in reactive programming with guarded recursion (2020)
3. Benton, N., Wadler, P.: Linear logic, monads and the lambda calculus. *Proceedings 11th Annual IEEE Symposium on Logic in Computer Science* pp. 420–431 (1996)
4. Benton, P.N.: A mixed linear and non-linear logic: Proofs, terms and models. In: Pacholski, L., Tiuryn, J. (eds.) *Computer Science Logic*. pp. 121–135. Springer Berlin Heidelberg, Berlin, Heidelberg (1995)
5. Berry, G., Cosserat, L.: The ESTEREL synchronous programming language and its mathematical semantics. In: Brookes, S.D., Roscoe, A.W., Winskel, G. (eds.) *Seminar on Concurrency*. pp. 389–448. Springer Berlin Heidelberg, Berlin, Heidelberg (1985)
6. Birkedal, L., Møgelberg, R.E., Schwinghammer, J., Støvring, K.: First steps in synthetic guarded domain theory: Step-indexing in the topos of trees. In: *In Proc. of LICS* (2011)
7. Bizjak, A., Grathwohl, H.B., Clouston, R., Møgelberg, R.E., Birkedal, L.: Guarded dependent type theory with coinductive types. In: *International Conference on Foundations of Software Science and Computation Structures*. pp. 20–35. Springer (2016)
8. Blackburn, P., de Rijke, M., Venema, Y.: *Modal Logic*. Cambridge Tracts in Theoretical Computer Science, Cambridge University Press (2002)
9. Caspi, P., Pilaud, D., Halbwegs, N., Plaice, J.A.: LUSTRE: A Declarative Language for Real-time Programming. In: *Proceedings of the 14th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*. pp. 178–188. POPL '87, ACM, New York, NY, USA (1987). <https://doi.org/10.1145/41625.41641>
10. Cave, A., Ferreira, F., Panangaden, P., Pientka, B.: Fair Reactive Programming. In: *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. pp. 361–372. POPL '14, ACM, San Diego, California, USA (2014). <https://doi.org/10.1145/2535838.2535881>
11. Clouston, R.: Fitch-style modal lambda calculi. In: Baier, C., Dal Lago, U. (eds.) *Foundations of Software Science and Computation Structures*. pp. 258–275. Springer International Publishing, Cham (2018)
12. Elliott, C.: Push-pull functional reactive programming. In: *Haskell Symposium* (2009), <http://conal.net/papers/push-pull-frp>
13. Elliott, C., Hudak, P.: Functional reactive animation. In: *Proceedings of the Second ACM SIGPLAN International Conference on Functional Programming*. pp. 263–273. ICFP '97, ACM, New York, NY, USA (1997). <https://doi.org/10.1145/258948.258973>
14. Girard, J.Y.: Linear logic. *Theoretical Computer Science* **50**(1), 1 – 101 (1987). [https://doi.org/10.1016/0304-3975\(87\)90045-4](https://doi.org/10.1016/0304-3975(87)90045-4)
15. Graulund, C.: *Lambda Calculus for Reactive Programming*. Master's thesis, IT University of Copenhagen (2018)
16. Haim, M., Malherbe, O.: Linear Hyperdoctrines and Comodules. arXiv e-prints arXiv:1612.06602 (Dec 2016)

17. Jeffrey, A.: LTL types FRP: linear-time temporal logic propositions as types, proofs as functional reactive programs. In: Proceedings of the sixth workshop on Programming Languages meets Program Verification, PLPV 2012, Philadelphia, PA, USA, January 24, 2012. pp. 49–60. Philadelphia, PA, USA (2012). <https://doi.org/10.1145/2103776.2103783>
18. Jeltsch, W.: Towards a common categorical semantics for linear-time temporal logic and functional reactive programming. *Electronic Notes in Theoretical Computer Science* **286**, 229–242 (2012)
19. Jeltsch, W.: Temporal Logic with "Until", Functional Reactive Programming with Processes, and Concrete Process Categories. In: Proceedings of the 7th Workshop on Programming Languages Meets Program Verification. pp. 69–78. PLPV '13, ACM, New York, NY, USA (2013). <https://doi.org/10.1145/2428116.2428128>
20. Krishnaswami, N.R.: Higher-order Functional Reactive Programming Without Spacetime Leaks. In: Proceedings of the 18th ACM SIGPLAN International Conference on Functional Programming. pp. 221–232. ICFP '13, ACM, Boston, Massachusetts, USA (2013). <https://doi.org/10.1145/2500365.2500588>
21. Krishnaswami, N.R., Benton, N.: Ultrametric semantics of reactive programs. In: 2011 IEEE 26th Annual Symposium on Logic in Computer Science. pp. 257–266. IEEE Computer Society, Washington, DC, USA (June 2011). <https://doi.org/10.1109/LICS.2011.38>
22. Lawvere, F.W.: Adjointness in foundations. *Dialectica* **23**(3-4), 281–296 (1969). <https://doi.org/10.1111/j.1746-8361.1969.tb01194.x>
23. MacLane, S., Moerdijk, I.: *Sheaves in Geometry and Logic: A First Introduction to Topos Theory*. Universitext, Springer New York (1994). <https://doi.org/10.1007/978-1-4612-0927-0>
24. Maietti, M.E., de Paiva, V., Ritter, E.: Categorical models for intuitionistic and linear type theory. In: Tiuryn, J. (ed.) *Foundations of Software Science and Computation Structures*. pp. 223–237. Springer Berlin Heidelberg, Berlin, Heidelberg (2000)
25. Nakano, H.: A modality for recursion. In: Proceedings Fifteenth Annual IEEE Symposium on Logic in Computer Science (Cat. No.99CB36332). pp. 255–266. IEEE Computer Society, Washington, DC, USA (June 2000). <https://doi.org/10.1109/LICS.2000.855774>
26. Pfenning, F., Davies, R.: A judgmental reconstruction of modal logic. *Mathematical Structures in Computer Science* **11**(4), 511–540 (2001). <https://doi.org/10.1017/S0960129501003322>
27. Pnueli, A.: The temporal logic of programs. In: 18th Annual Symposium on Foundations of Computer Science (sfcs 1977). pp. 46–57. IEEE (1977)
28. Pouzet, M.: *Lucid Synchrone, version 3*. Tech. rep., Laboratoire d'Informatique de Paris (2006)
29. Szamozvancev, D.: *Semantics of temporal type systems*. Master's thesis, University of Cambridge (2018)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

