Higher-Order Functional Reactive Programming in Bounded Space

Neelakantan R. Krishnaswami
MPI-SWS
neelk@mpi-sws.org

Nick Benton
Microsoft Research
nick@microsoft.com

Jan Hoffmann
Yale University
jan.hoffmann@yale.edu

Abstract

Functional reactive programming (FRP) is an elegant and successful approach to programming reactive systems declaratively. The high levels of abstraction and expressivity that make FRP attractive as a programming model do, however, often lead to programs whose resource usage is excessive and hard to predict.

In this paper, we address the problem of space leaks in discrete-time functional reactive programs. We present a functional reactive programming language that statically bounds the size of the dataflow graph a reactive program creates, while still permitting use of higher-order functions and higher-type streams such as streams of streams. We achieve this with a novel linear type theory that both controls allocation and ensures that all recursive definitions are well-founded.

We also give a denotational semantics for our language by combining recent work on metric spaces for the interpretation of higher-order causal functions with length-space models of space-bounded computation. The resulting category is doubly closed and hence forms a model of the logic of bunched implications.

Categories and Subject Descriptors D.3.2 [*Dataflow Languages*]

General Terms languages, design, theory

Keywords functional reactive programming, dataflow, space-bounded computation, linear logic, bunched implications

1. Introduction

Reactive systems engage in an ongoing interaction with their environment, consuming input events and producing corresponding output events. Examples of such systems range from embedded controllers and sensor networks up to complex graphical user interfaces, web applications, games and simulations. Programming reactive systems in a general-purpose imperative language can be unpleasant, as different parts of the program interact not by structured control flow, but by dynamically registering state-manipulating callback functions with one another. The complexity of writing and reasoning about programs written in such a higher-order imperative style, as well as the critical nature and resource requirements of many reactive systems, has inspired extensive research into domain-specific languages (DSL), libraries and analysis techniques for reactive programming.

Synchronous dataflow languages, such as Esterel [3], Lustre [4], and Lucid Synchrone [21], implement a domain-specific computational model deriving from Kahn networks. A program corresponds to a fixed network of stream-processing nodes that communicate with one another, each consuming and producing a statically-known number of primitive values at every clock tick. Synchronous languages have precise, analysable semantics, provide strong guarantees about bounded usage of space and time, and are widely used in applications such as hardware synthesis and embedded control software.

Functional reactive programming (FRP), as introduced by Elliott and Hudak [8], also works with time-varying values (rather than mutable state) as a primitive abstraction, but provides a much richer model than the synchronous languages: signals (behaviours) can vary continuously as well as discretely, values can be higher-order (including both first-class functions and signal-valued signals), and the overall structure of the system can change dynamically. FRP has been applied in problem domains including robotics, animation, games, web applications and GUIs. However, the expressivity and apparently simple semantics of the classic FRP model come at a price. Firstly, the intuitively appealing idea of modelling A-valued signals as elements of the stream type A^ω (or, in the continuous case, $A^\mathbb{R}$) and reactive systems as stream functions $\mathrm{Input}^\omega \to \mathrm{Output}^\omega$ does not rule out systems that violate causality (the output today can depend upon the input tomorrow) or reactivity (ill-founded feedback can lead to undefined behaviour). Secondly, as the model is highly expressive and abstracts entirely from resource usage, the time and space behaviour of FRP programs is hard to predict and, even with sophisticated implementation techniques, can often be poor. It is all too easy to write FRP programs with significant space leaks, caused by, for example, inadvertently accumulating the entire history of a signal.[1]

Subsequent research has attempted to reduce 'junk' and alleviate performance problems by imposing restrictions on the classic FRP model. The Yale Haskell Group's Yampa [11, 19], for example, is an embedded DSL for FRP that constructs signal processing networks using Hughes's arrow abstraction [12]. Signals are no longer first-class, and signal-processing functions must be built from well-behaved casual primitives by causality-preserving combinators. Arrowized FRP allows signals to carry complex values but is essentially first-order (there is no exponential at the level of signal functions), though certain forms of dynamism are allowed via built in 'switching' combinators. Yampa does not enforce reactiv-

[1] Closely related are 'time leaks', which occur when sampling a time-dependent value can invoke an arbitrarily lengthy computation to 'catch up' with the current time.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
POPL '12 January 25–27, 2012, Philadelphia, PA, USA.
Copyright © 2012 ACM 978-1-4503-1083-3/12/01. . . $10.00

ity or provide resource guarantees but, empirically at least, makes certain kinds of leaks less likely.

Krishnaswami and Benton [13] recently described a semantic model for higher-order, discrete-time functional reactive programs based on ultrametric spaces, identifying causal functions with non-expansive maps and interpreting well-founded feedback via Banach's fixpoint theorem. They gave an associated language, featuring a Nakano-style [18] temporal modality for well-founded recursion, and showed the correctness of an implementation using an imperatively-updated dataflow graph. This implementation is much more efficient than directly running the functional semantics, but nothing prevents the dataflow graph from growing unboundedly as a program executes, leading to undesirable space leaks. In this paper, we solve the problem of such leaks by extending the ultrametric approach to FRP with linearly-typed resources that represent the permission to perform heap-allocation, following the pattern of Hofmann's work [9, 10] on non-size-increasing computation.

We give a denotational model for bounded higher-order reactive programming in terms of 'complete ultrametric length spaces', which carry both an ultrametric distance measure and a size function. Maps between such spaces must be non-expansive and non-size-increasing. Intuitively, the metric is used to enforce good temporal behaviour (causality and productivity of recursive definitions), whilst the size measure enforces good spatial behaviour, bounding the number of cells in the dataflow graph. The category of complete ultrametric length spaces is doubly-closed, forming a model of the logic of bunched implications [20] and exposing a (perhaps) surprising connection between the type theory of stream programming and separation logic.

We define a term language, with a rather novel type theory, that corresponds to our model and allows us to write bounded reactive programs. Judgements are all time-indexed, with the successor operation on times internalized in a modality $\bullet A$. Terms are typed in three contexts, one carrying linear (actually affine) resources of type $\Diamond$, giving permission to allocate; one binding pure, resource-free variables; and one binding potentially resourceful variables. Resource-freedom is internalized via a $!A$ modality, in the style of linear logic. We give the language an interesting staged operational semantics, which separates the normalizing reduction that takes place within each time step from the transitions that take place when the clock advances, and show that this is soundly modelled in the denotational semantics. The operational semantics uses terms of the language itself to encode the heap context within which evaluation takes place.

We also give a number of examples in our language that illustrate properties of the model and the applicability of our approach, showing that one can work with recursively-defined higher-order functions and streams in a natural way while still ensuring causality, productivity and bounded space usage.

To improve readability, we invert the 'logical' order of presentation: Section 2 gives an informal account of the language that suffices for presenting some of the motivating examples. Section 3 formally defines the language and the type system. We then, in Section 4, define the operational semantics and, in Section 5 and 6, present the details of the denotational model. Finally, in Section 7, we discuss our work and relate it to existing research.

2. Programming Language

The language is essentially a simply-typed λ-calculus with a type constructor $\mathsf{S}(-)$ for infinite streams, extended with three non-standard notions: delay types, resource types, and pure types. We treat streams as much as possible as mathematical sequences, but want to ensure that we can also interpret them as the successive values of signals generated by implementable clocked systems. Thus a definition (in a generic functional syntax for now)

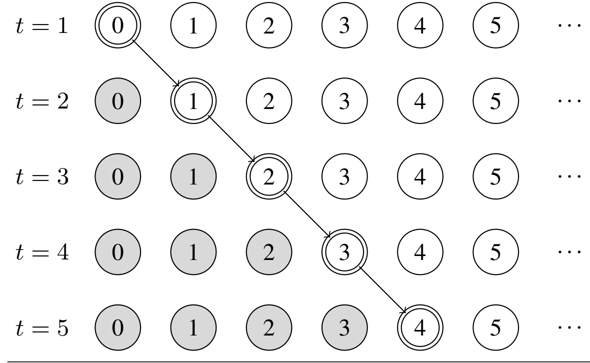

Figure 1. Memory Usage of Streams

$$\begin{array}{l} \mathit{nats} \;:\; \mathbb{N} \to \mathsf{S}(\mathbb{N}) \\ \mathit{nats}\; n = \mathsf{cons}(n,\; \mathit{nats}(n{+}1)) \end{array}$$

denotes a parameterised infinite stream but can also be understood as a stateful process ticking out successive natural numbers as time advances. It is necessary to restrict recursive definitions to ensure that signals are well-defined at all times. The recursion above is clearly *guarded*: the recursive call to nats only occurs underneath a cons constructor, so successively unfolding the stream at each clock tick is productive. That is, we always discover at least one new cons constructor which we can examine to find the head and the tail of the stream at the current time.

However, simple syntactic guardedness checks (used by languages as varied as Lucid Synchrone, Agda and Coq) do not integrate well with higher-order. For example, one might want a stream functional that abstracts over the constructor:

$$\mathit{higher_order}\;\; f\;\; v = f(v,\;\; \mathit{higher_order}\;\; f\;\; (v + 1))$$

The guardedness of $\mathit{higher_order}$ now depends on the definition of f, which is an unknown parameter. As in our earlier work [13], we instead use a *next-step* modality $\bullet A$ to track the times at which values are available in their types.

A value of type $\bullet A$ is a computation that will yield a value of type A when executed on the *next* clock tick. The tail function has type $\mathit{tail} : S(A) \to \bullet S(A)$, expressing that the tail of a stream only becomes available in the future. Similarly the type of cons is refined to $\mathit{cons} : A \times \bullet S(A) \to S(A)$, capturing that streams are constructed from a value today and a stream tomorrow. By giving the fixed point combinator the type $\mathit{fix} : (\bullet A \to A) \to A$, we ensure that all recursive definitions are well-founded without restricting their syntactic form.

If we care about space usage, however, this use of types to track guardedness still admits too many programs. The problem is to limit the amount of data that must be buffered to carry it from one time tick to the next. In the case of *nats* above, it is clear that the current state of the counting process can always be held in a single natural number. But consider a similar definition at a higher type:

$$\begin{array}{l} \mathit{constantly_leak} \;\;:\; \mathsf{S}(A) \to \mathsf{S}(\mathsf{S}(A)) \\ \mathit{constantly_leak}\;\; xs = \mathsf{cons}(xs,\;\; \mathit{constantly_leak}\;\; xs) \end{array}$$

The call $\mathit{constantly_leak}\;\; xs$ yields a stream of streams which is constantly xs. This is a perfectly well-founded functional definition but, as a stateful process, requires the whole accumulated history of xs to be buffered so it can be pushed forward on each time step.

Figure 1 shows the evolution of a stream as time passes. The gray nodes denote the nodes already produced, the doubly-circled nodes show the current value of the head of the stream, and the white nodes mark the elements yet to be produced. At each time step, one node moves into the past (becomes gray), and the current head advances one step farther into the stream.

As *constantly_leak xs* retains a reference to its argument, it needs to buffer all of the gray nodes to correctly enumerate the elements of the streams to be produced in future. Running *constantly_leak xs* for n time steps thus requires $O(n)$ space, which is unreasonable. To see this more clearly, consider the following function:

```
f : S(A) → S(S(A)) → ℕ → ℕ → S(A)
f xs yss n m =
  let (x, xs') = (head xs, tail xs) in
  let (ys, yss') = (head yss, tail yss) in
  if n = 0 then cons(x, f (head yss') yss' m (m+1))
           else cons(x, f xs' yss' (n−1) m)

diag : S(A) → S(A)
diag xs = f xs (constantly_leak xs) 0 1
```

Now, *diag* (*nats* 0) yields $(0, 0, 1, 0, 1, 2, 0, 1, 2, 3, \ldots)$, enumerating the prefixes of *nats* 0. As each prefix appears infinitely often, the whole history of the initial stream must be saved as the program executes. There are definitions with the same type as *constantly_leak* that *can* be implemented in constant space:

```
tails :: S(A) → S(S(A))
tails xs = cons(xs, tails (tail xs))
```

Since *tails* returns the successive tails of its stream argument, we do not need to remember gray nodes (see Figure 1), and hence can implement the function without any buffering. Substituting *tails* for *constantly_leak* turns *diag* into the (space-efficient) identity.

To account for the memory usage associated with creating and buffering stream data, we adapt a variant of the linear (affine) *resource types* of Hofmann's LFPL [9, 10]. The type $\lozenge$ represents a permission to create one new stream; the tensor product $R \otimes S$ is the permission to do both R and S; and the linear function space $R \multimap A$ builds an A, consuming the resources in R. We further refine the construction of streams of type $S(A)$ to take *three* arguments using the syntactic form $\mathsf{cons}(u, e, u'.\ e')$. The term $u : \lozenge$ is a permission to allocate a cons cell and the term $e : A$ is the head of the stream. The tail, $e' : \bullet S(A)$, is defined in scope of the variable u', which re-binds the allocation permission that will be freed up on the next time step.

We still permit sharing of stream values without restriction since dataflow programs gain efficiency precisely from the ability to share. Therefore, we also support a context of unrestricted variables and an intuitionistic[2] function space $A \to B$. (Hofmann's original language featured a strictly linear type discipline.) Function closures can also need buffering if they capture streams in their environment, so it is useful to introduce the type constructor $!A$, which classifies those A-values that need no buffering and may be freely carried forward in time.

2.1 Examples

Our language makes the intuition of bounded resource consumption explicit, while remaining close to standard functional programming style. We begin with the definition of two classic stream functions, $nats$ and *fib*, in our language:

```
nats : !ℕ → ◊ ⊸ S(!ℕ)
nats = fix loop : !ℕ → ◊ ⊸ S(!ℕ)
         λ x u. let !n = x in cons(u, !n, u'. loop !(n+1) u')

fibs : !ℕ → !ℕ → ◊ ⊸ S(!ℕ)
fibs = fix loop : !ℕ → !ℕ → ◊ ⊸ S(!ℕ).
         λ x y u.
           let !n = x in
           let !m = y in
           cons(u, !n, u'. loop !m !(n+m) u')
```

[2] We do *not* decompose the intuitionistic function space as $!A \multimap B$.

The function *nats* differs from the standard definition in two ways. First, we include an extra resource argument of type $\lozenge$, which contains a permission used as the extra argument to the cons operator. Furthermore, we require the argument of the function to be of the modal type $!\mathbb{N}$. Since the variable n is used in both the head and the tail of the cons-expression, we need to know that n can be used at different times without requiring additional space usage. This is exactly what the type $!A$ allows: an expression `let !a = a' in ...` binds the modal value a' to a variable a which can be used at any current or future time. The introduction form $!e$ for $!A$ constructs a value of type A, which may only mention resource-free variables. The definition of the function *fibs* is similar.

Similarly, we give the definition of the *constant* function, which returns a constant stream of values.

```
constant : !A → ◊ ⊸ S(!A)
constant = fix loop : !A → ◊ ⊸ S(!A).
             λ a' u. let !a = a' in cons(u, a', v. loop !a v)
```

The modal type $!A$ in the typing of *constant* prevents one defining *constantly_leak* in terms of *constant*, since stream values are never resource-free (as cons-expressions include a resource argument).

However, we can still define functions constructing higher-order streams, so long as space usage is bounded. In this, we improve upon much previous work on efficient implementations of FRP. An example is the *tails* function that we described earlier; in our language, this is programmed as follows:

```
tails : S(A) → ◊ ⊸ S(S(A))
tails =
  fix tails : S(A) → ◊ ⊸ S(S(A)).
    λ xs u.
      let xs' = tail(xs) in
      cons(u, xs, u'. tails xs' u')
```

Higher type streams also let us *define* many of the switching combinators of FRP libraries, without having to build them in as primitive operations. For example:

```
switch : S(bool) → S(S(A)) → ◊ ⊸ S(A)
switch =
  let loop =
    fix loop : S(bool) → S(S(A)) → S(A) → ◊ ⊸ S(A).
      λ bs xss current u.
        let yss = tail(xss) in
        let bs' = tail(bs) in
        if head(bs) then
          let zs = tail(head xss) in
          cons(u, head(head xss), u'. loop bs' yss zs u')
        else
          let zs = tail(current) in
          cons(u, head(current), u'. loop bs' yss zs u')
  in λ bs xss u. loop bs xss (head xss) u
```

The function $switch$ takes a stream of boolean events and a stream of streams. It then yields the elements of the head of the stream of streams until the boolean stream yields true, at which point it starts generating the elements of the current stream from the stream of streams. In this way, it is easy to go beyond simple static dataflow programs, without having to contort programs to fit a fixed set of combinators.

However, with the full resources of a higher-order programming language available, it is often convenient to define programs in terms of familiar stream functionals such as map.

```
map : !(A → B) → S(A) → ◊ ⊸ S(B)
map h =
  let !f = h in
  fix loop : S(A) → ◊ ⊸ S(B).
    λ xs u.
      let ys = tail xs in
      cons(u, f (head xs), u'. loop ys u')
```

This function illustrates a common pattern of higher-order programming in our language. We often wish to use functional arguments at many different times: in this case, we want to apply the argument to each element of the input stream. Therefore, the functional arguments in higher-order functions often need to be under the time-independence modality $!(A \to B)$.

The *unfold* function provides another nice example of higher-order programming:

```
unfold : !(X → A × •X) → X → ◊ ⊸ S(A) =
unfold h =
  let !f = h in
  fix loop : X → ◊ ⊸ S(A).
    λ x u.
      let (a, d) = f(x) in
      let •x' = d in
      cons(u, a, v. loop x' v)
```

Using *unfold*, one can directly translate deterministic state machines into stream programs, passing in the state transformer function and the initial state as the arguments. This function is also the first example with the delay modality — our state transformer takes a state, and returns a value and the state to use on the *next* timestep. The elimination form let $\bullet y = e$ in ... takes an expression of type $\bullet A$ and binds it to a variable y of type A, one tick in the future. This lets us use it in the third argument to cons, since the tail of a stream is also an expression one tick in the future.

Next, we illustrate the importance of being able to use streams intuitionistically, even as we track resources linearly.

```
zip : S(A) × S(B) → ◊ ⊸ S(A × B)
zip =
  fix zip : S(A) × S(B) → ◊ ⊸ S(A × B).
    λ (xs, ys) u.
      let xs' = tail(xs) in
      let ys' = tail(ys) in
      cons(u, (head xs, head ys), v. zip (xs', ys') v)

sum : S(ℕ) × S(ℕ) → (◊ ⊗ ◊) ⊸ S(ℕ) =
  λ (xs, ys) (u, v). map !+ (zip (xs,ys) u) v

double : S(ℕ) → (◊ ⊗ ◊) ⊸ S(ℕ)
double ns (u, v) = sum(ns, ns) (u, v)
```

We first define the function *zip*, which takes two streams and returns a stream of pairs, and the function *sum*, which takes two streams of natural numbers and pointwise sums their elements. These two functions are then used to define *double*, which takes a stream of numbers and returns a new stream of elements each of which is twice the size of the input. Note that *double* works by passing the *sum* function the *same* stream in both arguments, decisively violating linearity.

The *map* and *zip* functions (together with *unzip*, which we do not define here) witness that $S(\cdot)$ is a Cartesian functor. We can also define maps illustrating other semantic properties of streams.

```
cokleisli : !(S(A) → B) → S(A) → ◊ ⊸ S(B)
cokleisli g =
  let !f = g in
  fix loop : S(A) → ◊ ⊸ S(B).
    λ xs u.
      let ys = tail(xs) in cons(u, f xs, u'. loop ys u')

flip : S(•A) → •(◊ ⊸ S(A))
flip =
  fix flip : S(•A) → •(◊ ⊸ S(A))
    λ xs.
      let •x' = head(xs) in
      let xs' = tail(xs) in
      •(let •f = flip xs' in
          λ u. cons(u, x', u'. f u'))

unflip : •S(A) → ◊ ⊸ S(•A)
unflip =
  fix unflip : •S(A) → ◊ ⊸ S(•A).
    λ xs' u.
      let •ys = xs' in
      cons(u, •(head(ys)), u'.
              let ys' = tail(ys) in
              unflip (• ys') u')
```

The *cokleisli* function lifts a function from streams $S(A)$ to a type B to a function from streams of A to streams of B, giving a comonad structure [24] to the $S(\cdot)$ functor. The *flip* and *unflip* functions define an isomorphism between streams of delayed values and delayed streams of values. These operations make use of the ability to explicitly delay expressions e until the next tick with the $\bullet e$ introduction form for $\bullet A$.

In our earlier work, we had general delay operators $\delta_A : A \to \bullet A$, which shifted values of any type forward one tick into the future. However, a given piece of data may represent different values as time passes, and so we do *not* want delay maps of type $\delta_A : S(A) \to \bullet S(A)$, since this type does not capture the additional storage needed to move the argument forward one step in the future. However, it is possible in our system to *define* delay operators with types such as $S(!\mathbb{N}) \to \Diamond \multimap \bullet S(!\mathbb{N})$, which explicitly represent the buffering in the type:

```
buffer : !ℕ → S(!ℕ) → ◊ ⊸ S(!ℕ)
buffer b xs u =
  let !y = head(xs) in            // let head of y be usable later
  let ys = tail(xs) in            // ys at time 1
  cons(u, b, v. buffer !y ys v)   // now call buffer at time 1
```

The *buffer* function prepends a number to the front of a stream and can be used to construct a delay operator:

```
delaynats : S(!ℕ) → ◊ ⊸ •S(!ℕ)
delaynats xs u =
  let !y = head(xs) in    // let head of y be usable later
  let ys = tail(xs) in    // tail ys is at time 1
  •(buffer !y ys u)       // now call buffer at time 1 with y,ys
```

This gives a delay operator for streams, but additionally asks for a resource with which to construct the delayed stream. The definition illustrates our design philosophy that expensive crosstime operations should be programmed explicitly.

We conclude with a traditional example of stream programming, the sieve of Eratosthenes:

```
sieve : !(ℕ → bool) → S(!ℕ) → ◊ ⊸ S(!(option ℕ))
sieve p xs u =
  let !n = head xs in
  let ns = tail(xs) in
  let !pred = p in
  if pred n then
    let q = !(λ j. pred j ∧ j mod n ≠ 0) in
    cons(u, !(Some n), v. sieve q ns v)
  else
    cons(u, !None, v. sieve p ns v)

primes : ◊ ⊗ ◊ ⊸ S(ℕ)
primes (u,v) = sieve !(is_odd) (nats 2 u) v
```

The *sieve* function incrementally constructs a filter predicate that tests each element of the stream for divisibility by all the primes seen so far. Since the number of primes is infinite, the size of the sieve predicate's lambda-term grows without bound, but is nevertheless accepted by our language's type system. By design, we allow all ordinary call-by-value functional programs (which have reasonable compositional cost semantics [23]), and only use typing to track the unusual memory leaks of FRP. Hence our types only track the size of the dataflow graph (i.e., the number of live cons cells).

3. Syntax and Typing

We give the types and syntax of the programming language in Figure 2. The types include base types P, stream types $S(A)$, the next-step modality $\bullet A$, ordinary functions $A \to B$, the purity modality $!A$, and the linear function space $R \multimap A$. Resources R include the allocation permission $\Diamond$, and the tensor product $R \otimes S$. For space reasons, we do not include products $A \times B$ or sums $A + B$ here. (Since we require coproducts to define switching combinators, we emphasise that this is purely for space reasons: there are no technical complications associated with sum types.)

The typing rules are defined in Figure 3. The two type judgements of our type theory are $\Theta \vdash t :_i R$ and $\Theta; \Pi; \Gamma \vdash e :_i A$. Both judgements are time-indexed, in the sense that the type system judges a term to have a type at a particular time i. Furthermore, each hypothesis in each context is indexed by the time at which they can be used. (As usual, we take contexts to be unordered, and implicitly assume alpha-renaming to ensure that all variables are distinct.)

The judgement $\Theta \vdash t :_i R$ states that in the context of affine resource variables Θ, the term t has resource type R, at time i. The resource terms are built from affine pairs of the ground type $\Diamond$ and are permissions to allocate one or more cons cells.

The judgement $\Theta; \Pi; \Gamma \vdash e :_i A$ has three contexts: the affine resource context Θ contains again permissions to allocate cons cells; the intuitionistic context Π contains pure hypotheses (i.e., variables in Π bind non-state-dependent values); and the intuitionistic context Γ binds arbitrary value types, and permits unrestricted sharing and reuse of variables. Under these three contexts, we judge a term e to be an expression of type A at time i.

There are only two rules for the affine resource term calculus in Figure 3. The RHYP rule allows a resource to be used at any time after the context says it is available. The tensor rule R$\otimes$I lets one form an affine pair $\langle t, t' \rangle$, dividing the resources in the context between the two components. Observe that these rules allow weakening but not contraction.

The rule PHYP lets us use a pure hypothesis at any time after the time index in the variable context. In contrast, the rule EHYP only permits using a variable x at *exactly* its time index in Γ. This difference is one of the keys to accurately tracking space usage: we may substitute values which require buffering for variables in Γ, and by disallowing implicit transport of values across time, we ensure that the programmer uses explicit buffering whenever needed.

The rules $\to$I and $\to$E introduce and eliminate intuitionistic functions. The introduction rule does not permit the body to use any resources, since we can call functions multiple times. (The presence of Θ in the conclusion of $\to$I and the other value forms builds in weakening, so that we do not have to give a separate structural rule.) The elimination rule does allow expressions to use resources, since we will use a call-by-value evaluation strategy that will evaluate the two terms (using up their resource permissions) before substituting a value into a lambda-term.

The rules $\multimap$I and $\multimap$E introduce and eliminate linear functions. The introduction rule only permits the body to use the resources it receives in the argument, since we need to ensure that the function can be safely called multiple times. As a result, our typing rules do not permit currying of linear functions ($R \multimap S \multimap A \not\simeq R \otimes S \multimap A$), even though our underlying semantic model does permit it. If our type theory had the tree-structured contexts of the logic of bunched implications [20], then currying linear functions would be syntactically expressible. However, type checking for bunched calculi is still a difficult problem, and so in this work we restrict our attention to a linear fragment.

The rules SI, SE-HEAD and SE-TAIL are the introduction and elimination rules for streams. The syntactic form $\mathsf{cons}(t, e, u.\ e')$ takes three arguments: the expression t is a permission to create a cons cell, the expression e is the head of the stream, and e' is the tail of the stream. The tail subterm e' occurs underneath a binder for the resource variable u'. The intuition is that each stream takes up one unit of space at each successive time step, and u' names the permission t, after one time step has elapsed. This lets us pass the permission to use t to functions on subsequent time steps in the body of e'.

The rule SE-HEAD is straightforward: given a stream of type $S(A)$, we get a value of type A, at the same time. The rule SE-TAIL uses the form $\mathsf{let}\ x = \mathsf{tail}(e)\ \mathsf{in}\ e'$ to bind the tail of e to the variable x in e'. The tail of a stream at time i lives at time $i+1$ and we choose a binding elimination form to maintain the invariant that no term of time i contains any subterms at any earlier time.

The rules $\bullet$I and $\bullet$E introduce and eliminate delay terms. The rule $\bullet$I says that if e is a term of type A at time $i + 1$, then $\bullet e$ is a term of type $\bullet A$ at time i. Like the other value introduction forms, it prevents e from using any resources, so that e can be substituted freely. The rule $\bullet$E gives a binding elimination $\mathsf{let}\ \bullet x = e\ \mathsf{in}\ e'$ for the next-step modality. We use a binding elimination for the same reason as in the rule SE-TAIL — we do not want terms of time i to contain subterms of time $< i$.

The rule !I introduces a term of pure type $!A$. It does so by typing the body of a term $!A$ at type A with an empty resource and shared context. Since e can only refer to hypotheses in Π, which are pure, it follows that e itself must be pure. The rule !E types the elimination form $\mathsf{let}\ !x = e\ \mathsf{in}\ e'$, which binds the value e to the variable x in the pure context Π.

The rule $\otimes$E is the elimination form for the tensor. Given a term t of type $R \otimes S$, the expression $\mathsf{let}\ \langle u, v \rangle = t\ \mathsf{in}\ e$ binds the components to the variables u and v for the scope of e.

The rule LET introduces local let-bindings. The introduced binding $x = e$ must be at the same time i as the overall expression.

The rule FIX types fixed points $\mathsf{fix}\ x : A.\ e$. The body e is typed at time i with the recursive hypothesis $x :_{i+1} A$ one tick later. The one-step delay ensures the guardedness of recursive definitions. Furthermore, the typing of the expression e is derived under an empty resource context and with an empty shared context. Since e may unfold multiple times, giving e resources would violate linearity. Furthermore, the unrollings can happen at different times, which means that using any variables in the shared context might require buffering.

We now state the structural and substitution principles.

Lemma 1. *(Admissibility of Weakening and Contraction)*

1. *If* $\Theta \vdash t :_i R$ *then* $\Theta, \Theta' \vdash t :_i R$.
2. *If* $\Theta; \Pi; \Gamma \vdash e :_i A$ *then* $\Theta, \Theta'; \Pi, \Pi'; \Gamma, \Gamma' \vdash e :_i A$.
3. *If* $\Theta; \Pi; \Gamma, x :_i A, y :_i A \vdash e :_i B$ *then* $\Theta; \Pi; \Gamma, x :_i A \vdash [x/y]e :_i A$.
4. *If* $\Theta; \Pi, x :_i A, y :_j A; \Gamma \vdash e :_i B$ *and* $i \leq j$ *then* $\Theta; \Pi; \Gamma, x :_i A \vdash [x/y]e :_i A$.

Theorem 1. *(Substitution) We have that:*

1. *If* $\Theta \vdash t :_i R$ *and* $\Theta', u :_i R \vdash t' :_j R'$, *then* $\Theta, \Theta' \vdash [t/u]t' :_j R'$.
2. *If* $\Theta \vdash t :_i R$ *and* $\Theta, u :_i R; \Pi; \Gamma \vdash e :_j A$, *then* $\Theta, \Theta'; \Pi; \Gamma \vdash [t/u]e :_j A$.
3. *If* $\cdot; \Pi; \cdot \vdash e :_i A$ *and* $\Theta; \Pi, x :_j A; \Gamma \vdash e' :_k B$ *and* $i \leq j$, *then* $\Theta; \Pi; \Gamma \vdash [e/x]e' :_k B$.
4. *If* $\cdot; \Pi; \Gamma \vdash e :_i A$ *and* $\Theta; \Pi; \Gamma, x :_i A \vdash e' :_j B$, *then* $\Theta; \Pi; \Gamma \vdash [e/x]e' :_j B$.

Lemma 1 and Theorem 1 can be proved by structural inductions on the type derivations in the respective premises.

Types			
General	A	$::=$	$P \mid A \to A \mid S(A) \mid \bullet A$
		$\mid$	$!A \mid R \multimap A$
Resource	R	$::=$	$\Diamond \mid R \otimes R$

Terms			
General	e	$::=$	$x \mid \lambda x : A.\ e \mid e\ e' \mid \lambda u : R.\ e \mid e\ t$
		$\mid$	$\mathsf{cons}(t, e, u.\ e') \mid \mathsf{head}(e)$
		$\mid$	$\mathsf{let}\ x = \mathsf{tail}(e)\ \mathsf{in}\ e'$
		$\mid$	$\bullet e \mid \mathsf{let}\ \bullet x = e\ \mathsf{in}\ e'$
		$\mid$	$!e \mid \mathsf{let}\ !x = e\ \mathsf{in}\ e'$
		$\mid$	$\mathsf{let}\ \langle u, v\rangle = t\ \mathsf{in}\ e \mid \mathsf{fix}\ x : A.\ e$
		$\mid$	$\mathsf{let}\ x = e\ \mathsf{in}\ e'$
Resource	t	$::=$	$u \mid \langle t, t'\rangle$
Values	v	$::=$	$\lambda x : A.\ e \mid \lambda u : R.\ e \mid\ !v \mid \bullet\ e \mid x$

Contexts			
General	Γ	$::=$	$\cdot \mid \Gamma, x :_i A$
Pure	Π	$::=$	$\cdot \mid \Pi, x :_i A$
Resource	Θ	$::=$	$\cdot \mid \Theta, u :_i R$

Evaluation Contexts	C	$::=$	$\square \mid \mathsf{let}\ x = \mathsf{tail}(y)\ \mathsf{in}\ C$
		$\mid$	$\mathsf{let}\ x = \mathsf{cons}(u, v, u'.\ e)\ \mathsf{in}\ C$

Figure 2. Syntax

4. Operational Semantics

Theorem 1 formulates the substitution principles so that a general expression e that replaces a variables has to be typed without resource variables. That is, the substitution $[e/x]e'$ is only sound if the substituted term e uses no resource variables. On the other hand, the typing rule for cons cells demands the use of a resource, raising the question: what is the operational semantics of cons cells?

A purely substitution-based operational semantics cannot be correct, because it does not account for the sharing of cons cells. Consider the following expression that is well-typed in our system.

```
let xs = cons(u, ...) in // u is the linear allocation permission
sum (xs, xs) v
```

Here, we construct xs once, but use it twice in the call to sum. We cannot simply substitute the cons into the body of the let in our system, as that would duplicate the linear variable u.

One approach for managing permissions is to introduce a heap for cons cells, and refer to streams indirectly by reference. However, adding a heap moves us away from our functional intuitions and makes it more difficult to connect to our denotational model. Instead, we retain the idea of referring to streams by reference, but use *variables* for the indirect reference, by defining evaluation to put terms into let-normal form, such as:

```
let xs = cons(u1, e1, v1. e1') in
let ys = cons(u2, e2, v2. e2') in
let xs' = tail(xs) in
...
let zs = cons(u3, e3, v2. e3') in
let ys' = tail(ys) in
v
```

Now, the nested let-bindings act as our heap. The value v may contain many references to individual streams (such as xs), but since each stream is bound only once, we can respect the linearity constraint on the allocation permissions u_i. (Taking the tail of streams, as in the definition xs' and ys', also needs to be in let-normal form, since we cannot cut out the tail of a cons cell until

$$\boxed{\Theta \vdash t :_i R} \qquad \boxed{\Theta; \Pi; \Gamma \vdash e :_i A}$$

$$\frac{u :_i R \in \Theta \qquad i \leq j}{\Theta \vdash u :_j R}\ \textsc{RHyp}$$

$$\frac{\Theta \vdash t :_i R \qquad \Theta' \vdash t' :_i R'}{\Theta, \Theta' \vdash \langle t, t'\rangle :_i R \otimes R'}\ \mathrm{R}{\otimes}\mathrm{I}$$

$$\frac{x :_i A \in \Pi \qquad i \leq j}{\Theta; \Pi; \Gamma \vdash x :_j A}\ \textsc{PHyp} \qquad \frac{x :_i A \in \Gamma}{\Theta; \Pi; \Gamma \vdash x :_i A}\ \textsc{EHyp}$$

$$\frac{\cdot; \Pi; \Gamma, x :_i A \vdash e :_i B}{\Theta; \Pi; \Gamma \vdash \lambda x : A.\ e :_i A \to B}\ {\to}\mathrm{I}$$

$$\frac{\Theta; \Pi; \Gamma \vdash e :_i A \to B \qquad \Theta'; \Pi; \Gamma \vdash e' :_i A}{\Theta, \Theta'; \Pi; \Gamma \vdash e\ e' :_i B}\ {\to}\mathrm{E}$$

$$\frac{u :_i R; \Pi; \Gamma \vdash e :_i A}{\Theta; \Pi; \Gamma \vdash \lambda u : R.\ e :_i R \multimap A}\ {\multimap}\mathrm{I}$$

$$\frac{\Theta; \Pi; \Gamma \vdash e :_i R \multimap A \qquad \Theta' \vdash t :_i R}{\Theta, \Theta'; \Pi; \Gamma \vdash e\ t :_i A}\ {\multimap}\mathrm{E}$$

$$\frac{\Theta \vdash t :_i \Diamond \qquad \Theta'; \Pi; \Gamma \vdash e :_i A \qquad \Theta'', u :_{i+1} \Diamond; \Pi; \Gamma \vdash e' :_{i+1} S(A)}{\Theta, \Theta', \Theta''; \Pi; \Gamma \vdash \mathsf{cons}(t, e, u.\ e') :_i S(A)}\ \mathrm{SI}$$

$$\frac{\Theta; \Pi; \Gamma \vdash e :_i S(A)}{\Theta; \Pi; \Gamma \vdash \mathsf{head}(e) :_i A}\ \textsc{SE-head}$$

$$\frac{\Theta; \Pi; \Gamma \vdash e :_i S(A) \qquad \Theta'; \Pi; \Gamma, y :_{i+1} S(A) \vdash e' :_i B}{\Theta, \Theta'; \Pi; \Gamma \vdash \mathsf{let}\ y = \mathsf{tail}(e)\ \mathsf{in}\ e' :_i B}\ \textsc{SE-tail}$$

$$\frac{\cdot; \Pi; \Gamma \vdash e :_{i+1} A}{\Theta; \Pi; \Gamma \vdash \bullet e :_i \bullet A}\ \bullet\mathrm{I}$$

$$\frac{\Theta; \Pi; \Gamma \vdash e :_i \bullet A \qquad \Theta'; \Pi; \Gamma, x :_{i+1} A \vdash e' :_i B}{\Theta, \Theta'; \Pi; \Gamma \vdash \mathsf{let}\ \bullet x = e\ \mathsf{in}\ e' :_i B}\ \bullet\mathrm{E}$$

$$\frac{\cdot; \Pi; \cdot \vdash e :_i A}{\Theta; \Pi; \Gamma \vdash\ !e :_i\ !A}\ !\mathrm{I}$$

$$\frac{\Theta; \Pi; \Gamma \vdash e :_i\ !A \qquad \Theta'; \Pi, x :_i A; \Gamma \vdash e' :_i B}{\Theta, \Theta'; \Pi; \Gamma \vdash \mathsf{let}\ !x = e\ \mathsf{in}\ e' :_i B}\ !\mathrm{E}$$

$$\frac{\Theta \vdash t :_i R \otimes S \qquad \Theta', u :_i R, v :_i S; \Pi; \Gamma \vdash e :_i C}{\Theta, \Theta' \vdash \mathsf{let}\ \langle u, v\rangle = t\ \mathsf{in}\ e :_i C}\ {\otimes}\mathrm{E}$$

$$\frac{\cdot; \Pi, x :_{i+1} A; \cdot \vdash e :_i A}{\Theta; \Pi; \Gamma \vdash \mathsf{fix}\ x : A.\ e :_i A}\ \textsc{Fix}$$

$$\frac{\Theta; \Pi; \Gamma \vdash e :_i A \qquad \Theta'; \Pi; \Gamma, x :_i A \vdash e' :_i B}{\Theta, \Theta'; \Pi; \Gamma \vdash \mathsf{let}\ x = e\ \mathsf{in}\ e' :_i B}\ \textsc{Let}$$

Figure 3. Typing Rules

the next tick.) Using bindings to represent sharing will make it easier to continue using our denotational model to interpret the resulting terms. The scoping rules for let-binding also restrict us to a DAG dependency structure, an invariant that imperative reactive programming implementations based on dependency graphs must go to some lengths to implement and maintain.

Let-normalizing cons cells has a second benefit: we can advance the global clock by taking the tails of each cons cell in the context.

```
let xs = [u1/v1]e1' in
let ys = [u2/v2]e2' in
let xs' = xs in
...
let zs = [u3/v3]e3' in
let ys' = ys in
v
```

Since we know where all of the cons cells are, we can rewrite them to model the passage of time. Advancing the clock for tail expressions simply drops the tail. Intuitively, yesterday they promised a tail stream today, and after the step the binding they refer to contains that tail stream.

Our operational semantics has two phases: the within-step operational semantics, which puts an expression into let-normal form, and the step semantics, which advances the clock by one tick by rewriting the cons cells to be their tails.

4.1 Within-Step Operational Semantics

The syntax of values and evaluation contexts is given in Figure 2 and the typing and auxilliary operations are given in Figure 5. We define the reduction relation in Figure 6, in big-step style. We write $\Sigma \rhd_i \Omega_i \vdash C \dashv \Omega'_i$ for the evaluation-context typing judgement. The context Σ is a resource context Θ that consists only of $\Diamond$ hypotheses. Similar, the contexts Ω_i are restricted forms of general contexts Γ consisting only of stream variables at time i or $i+1$. Both are defined in Figure 4. The judgement $\Sigma \rhd_i \Omega_i \vdash C \dashv \Omega'_i$ reads as "the evaluation context C creates the bindings in Ω'_i, uses the resources in Σ to do so, and may refer to the bindings in Ω_i".

The context-concatenation operation $C \circ C'$ appends two evaluation contexts C and C'. It is defined in Figure 5 and satisfies the following properties.

Lemma 2. *(Context Concatenation) We have that:*

- $\circ$ *is associative with unit* $\Box$.
- *If* $\Sigma \rhd_i \Omega_i \vdash C \dashv \Omega'_i$ *and* $\Sigma' \rhd_i \Omega_i, \Omega'_i \vdash C' \dashv \Omega''_i$, *then* $\Sigma, \Sigma' \rhd_i \Omega \vdash C \circ C' \dashv \Omega'_i, \Omega''_i$.
- *If* $\Sigma \rhd_i \Omega_i \vdash C_1 \circ C_2 \dashv \Omega'_i$, *then there exist* Σ_1, Σ_2 *and* Ω^1_i, Ω^2_i *such that* $\Sigma = \Sigma_1, \Sigma_2$ *and* $\Omega'_i = \Omega^1_i, \Omega^2_i$ *and* $\Sigma_1 \rhd_i \Omega_i \vdash C_1 \dashv \Omega^1_i$ *and* $\Sigma_2 \rhd_i \Omega_i, \Omega^1_i \vdash C_2 \dashv \Omega^2_i$.

These properties all follow from routine inductions.

In Figure 6, we give the context semantics, evaluating an expression in a context into a value in a larger context. Note that the value forms for streams are variables, since we need to preserve sharing for them, and we can use variable names as pointers into the evaluation context. For most expression forms, the context semantics works as expected; it evaluates each subexpression in context, building a value in a larger context.

The rule CONSE is one of the two rules that extend the context. It evaluates a cons cell, creates a binding to a fresh variable, and returns the fresh variable as the value for that stream. The other rule that extends the context is TAILE. It adds a binding to the context naming the tail it constructs. The rule HEADE, on the other hand, uses the $C@x \Rightarrow v$ relation that is defined in Figure 5 to find the head of the cons cell bound to the variable x.

The rule FIXE looks entirely conventional. We simply unfold the fixed point and continue. We are nevertheless able to prove a normalization result for the within-step operational semantics since the fixed point substitutes for a variable at a future time.

We begin the metatheory with a type preservation proof.

Theorem 2. *(Type Preservation) If we have that*

$$\Sigma \rhd_i \cdot \vdash C \dashv \Omega_i, \quad \Sigma'; \cdot; \Omega_i \vdash e :_i A, \quad \text{and} \quad C[e] \Downarrow C''[v],$$

then there is an Ω'_i *and* C' *such that*

$$C'' = C \circ C', \;\; \Sigma' \rhd_i \Omega_i \vdash C' \dashv \Omega'_i, \;\; \text{and} \;\; \cdot; \cdot; \Omega_i, \Omega'_i \vdash v :_i A.$$

This theorem follows from a routine structural induction.

To show soundness, we will prove termination via a Kripke logical relations argument. Since we evaluate terms e in contexts C, and return a value v in some larger context C', we take our Kripke worlds to be the closed contexts. That is, worlds are those C such that $\Sigma \rhd_i \cdot \vdash C \dashv \Omega_i$. We define the ordering $C' \sqsubseteq C$ on worlds so that there should be some C_1 such that $C' \equiv C \circ C_1$. Thus, a future world is one in which more bindings are available.

In Figure 4, we define the logical relation by induction on types. There is one clause $\mathcal{V}_A(C)$ for each type A, defining a subset of the well-typed expressions of type A, closed save for the variables bound by C. The expression relation $\mathcal{E}_A(\Sigma; C)$ consists of the expressions that use resources in Σ and evaluate to a value in $\mathcal{V}_A(C)$ in the context C.

The definition of the logical relation for streams states that a variable x is in $\mathcal{V}_{S(A)}(C)$ if x binds a cons cell with v in its head, and v is in the A-relation. As expected, the relation for functions consists of lambdas such that in any future world, applying a value in the A-relation should result in a term in the expression relation at type B. In the case of the delay modality we allow any well-typed value, since the within-step evaluation relation does not evaluate any terms at time i, and the body of a delay is at time $i+1$. The $!A$ relation consists of values $!v$ such that v is in the A-relation in the empty world $\Box$, since we want values of type $!A$ to not depend on the stream values C binds. Last, the relation at $R \multimap A$ consists of those lambda-terms $\lambda u : R.\ e$ such that for any resource t of type R in context Σ, the expression $[t/u]e$ is in the expression relation for A with resources Σ.

To prove the fundamental property, we define some auxilliary predicates in Figure 4. The predicate $\mathit{Good}(\Omega_i)$ picks out those contexts in which all of the bindings in Ω_i at time i contain true streams, according to the stream relation. The $\mathcal{V}(\Gamma_{\geq i}; C)$ and $\mathcal{V}!(\Pi_{\geq i})$ sets extend the value relation to substitutions rather than single values, and $\Sigma \vdash \vartheta : \Theta$ defines linear substitutions. The notations $\Pi_{\geq i}$ and $\Gamma_{\geq i}$ mean contexts where every variable is at time i or later. In these substitutions, we only require substitutands for variables at time i to lie in the logical relation, and require only well-typededness for other variables, since within-step evaluation only affects the current tick.

Theorem 3. *(Fundamental Property of Logical Relations) Suppose* $\Theta; \Pi_{\geq i}; \Omega_i, \Gamma_{\geq i} \vdash e :_i A$. *Furthermore, suppose that* $C \in \mathit{Good}(\Omega)$ *and* $\Sigma \vdash \vartheta : \Theta$ *and* $\pi \in \mathcal{V}!(\Pi_{\geq i})$ *and* $\gamma \in \mathcal{V}(\Gamma_{\geq i}; C)$. *Then* $(\vartheta \circ \pi \circ \gamma)(e) \in \mathcal{E}_A(\Sigma; C)$.

As is usual for logical relations, this theorem follows from a structural induction on the derivation of $e : A$.

The fundamental property suffices to prove normalization, once we observe that typing derivations satisfy the following history independence property:

Lemma 3. *(History Independence)*

- *If* $\Theta; \Pi; \Gamma, x :_i A \vdash e :_j B$ *and* $i < j$, *then* $\Theta; \Pi; \Gamma \vdash e :_j B$.
- *If* $\Theta; \Pi, x :_i A; \Gamma \vdash e :_j B$ *and* $i < j$, *then* $\Theta; \Pi, x :_j A; \Gamma \vdash e :_j B$.
- *If* $\Theta, u :_i R; \Pi; \Gamma \vdash e :_j B$ *and* $i < j$, *then* $\Theta, u :_j R; \Pi; \Gamma \vdash e :_j B$.

Parameter of the logical relation: i

$$\text{Base Contexts} \quad \Sigma ::= \cdot \mid \Sigma, u :_0 \Diamond$$
$$\Omega_i ::= \cdot \mid \Omega_i, x :_i S(A) \mid \Omega_i, x :_{i+1} S(A)$$

$$\boxed{\mathcal{V}_A(\Sigma \rhd_i \cdot \vdash C \dashv \Omega_i) \subseteq \{v \mid \cdot;\cdot;\Omega_i \vdash v :_i A\}}$$

$$\boxed{\mathcal{E}_A(\Sigma;\Sigma' \rhd_i \cdot \vdash C \dashv \Omega_i) \subseteq \{e \mid \Sigma;\cdot;\Omega_i \vdash e :_i A\}}$$

$$\begin{array}{l}
\mathcal{V}_{S(A)}(C) = \{x \mid \exists v.\ C@x \Rightarrow v \wedge v \in \mathcal{V}_A(C)\} \\
\mathcal{V}_{A\to B}(C) = \\
\quad \{\lambda x : A.\ e \mid \forall C', v \in \mathcal{V}_A(C \circ C').\ [v/x]e \in \mathcal{E}_B(\cdot; C \circ C')\} \\
\mathcal{V}_{\bullet A}(\Sigma \rhd_i \cdot \vdash C \dashv \Omega_i) = \\
\quad \{\bullet e \mid \cdot;\cdot;\Omega_i \vdash \bullet e :_i \bullet A\} \\
\mathcal{V}_{!A}(C) = \\
\quad \{!v \mid v \in \mathcal{V}_A(\Box)\} \\
\mathcal{V}_{R\multimap A}(C) = \\
\quad \{\lambda u{:}R.\ e \mid \forall C', \Sigma, t.\ \Sigma \vdash t :_i R \Rightarrow [t/u]e \in \mathcal{E}_A(\Sigma; C \circ C')\}
\end{array}$$

$$\begin{array}{l}
\mathcal{E}_A(\Sigma; C) = \\
\quad \{e : A \mid \exists C', v.\ C[e] \Downarrow (C \circ C')[v] \wedge v \in \mathcal{V}_A(C \circ C')\}
\end{array}$$

$$\mathit{Good}(\Omega_i) = \{C \mid \forall x :_i S(A) \in \Omega_i.\ x \in \mathcal{V}_{S(A)}(C)\}$$

$$\begin{array}{l}
\mathcal{V}(\cdot; C) = \{\cdot\} \\
\mathcal{V}(\Gamma, x :_i A; C) = \\
\quad \{(\gamma, [v/x]) \mid \gamma \in \mathcal{V}(\Gamma; C) \wedge v \in \mathcal{V}_A(C)\} \\
\mathcal{V}(\Gamma, x :_{j>i} A; \Sigma \rhd_i \cdot \vdash C \dashv \Omega_i) = \\
\quad \{(\gamma, [e/x]) \mid \gamma \in \mathcal{V}(\Gamma; C) \wedge \cdot;\cdot;\Omega_i \vdash e :_j A\}
\end{array}$$

$$\begin{array}{l}
\mathcal{V}!(\cdot) = \{\cdot\} \\
\mathcal{V}!(\Pi, x :_i A) = \{(\gamma, [v/x]) \mid \gamma \in \mathcal{V}!(\Gamma) \wedge v \in \mathcal{V}_A(\Box)\} \\
\mathcal{V}!(\Pi, x :_{j>i} A) = \{(\gamma, [e/x]) \mid \gamma \in \mathcal{V}!(\Gamma) \wedge \cdot;\cdot;\cdot \vdash e :_j A\}
\end{array}$$

$$\boxed{\Sigma \vdash \vartheta : \Theta}$$

$$\frac{}{\cdot \vdash \cdot : \cdot} \qquad \frac{\Sigma \vdash t :_i R \qquad \Sigma' \vdash \vartheta : \Theta}{\Sigma, \Sigma' \vdash (\vartheta, [t/u]) : (\Theta, u :_i R)}$$

Figure 4. Logical Relation for Termination

These properties are all proved by structural induction.

The syntax ensures expressions at time $j > i$ do not depend on a variable of time i. As a result, we only need to consider contexts in which Π and Γ contain variables no younger than the current time. Normalization immediately follows:

Corollary 1. *(Normalization) Suppose* $\Sigma;\cdot;\cdot \vdash e :_i A$. *Then* $\Box[e] \Downarrow C[v]$.

Finally note that we are considering normalization of open terms, since we have no constants of type $\Diamond$. The non-existence of such constants is, of course, what ensures that the language respects space bounds.

Theorem 4. *(Space Bounded Evaluation) Suppose* $\Sigma;\cdot;\cdot \vdash e :_i A$ *and* $\Box[e] \Downarrow C[v]$. *Then the* size *of* C *— the number of cons cells it binds — is bounded by the size of* Σ.

Given type preservation, this theorem is straightforward. Each cons cell in the context needs a distinct resource variable, so the number of cons cells in C is clearly bounded by the size of Σ.

4.2 Next-Step Operational Semantics

Recall that when we advance time, we want to replace the tails of stream variables with just the variable. We define the necessary operations in Figure 7. To tick the clock, we define the relation $C[x] \leadsto e$, which takes a stream in context $C[x]$ and constructs a new expression e by cutting the tails of the streams in C, and sending each tail expressions $\mathsf{tail}(y)$ in C to y. This models the effect of advancing time by one tick, so that all the streams in the expression context become their tails, and all references to tails become references to the updated stream variable.

To show the type-correctness of this operation, we introduce the $\mathsf{Step}(\Omega_i)$ operation, which tells us how the typing changes: it sends all streams at time i to time $i+1$ (and leaves streams at time $i+1$ alone). Now, we can prove the following type soundness theorem.

Theorem 5. *(Ticking the Clock) If* $\Sigma;\cdot;\Omega_i \vdash C[x] :_i S(A)$ *and* $C[x] \leadsto e$, *then* $\Sigma;\cdot;\mathsf{Step}(\Omega_i) \vdash e :_{i+1} S(A)$.

This theorem follows from a routine structural induction, and establishes that our semantics is space-bounded. Since Σ is the same before and after the tick from $C[x] \leadsto e$, it follows that evaluating e will never construct more streams than Σ permits. Together with the within-step type preservation property, our clock-ticking theorem gives a (purely syntactic) proof of the space-boundedness of our language.

However, while our definition of advancing the clock is intuitively plausible, and even makes proving memory safety easy, it is still unclear in what sense it is *correct*. To illustrate the issues involved, observe that our language contains lambda-expressions (which may contain free variable references), and that the tick operator essentially imperatively updates stream values. It is important to prove that ticking the clock cannot change the meaning of a lambda-term in a way which violates the purely functional character of functional reactive programming.

To show this, we will first give a denotational semantics of stream programs. Then we will show that if $C[x] \leadsto e$, then the meaning of e does in fact equal the tail of the denotational meaning of $C[x]$, showing that ticking the clock really does advance the clock in the way we expect. This is morally a standard adequacy proof, applied to an unusual operational semantics.

5. Semantic Intuitions

5.1 Causality and Ultrametric Spaces

The intuition underpinning reactive programming is the *stream transformer*, a function which takes a stream of inputs and generates a stream of outputs. But not all functions on streams are implementable reactive programs — in order to be implementable at all, reactive programs must respect the *causality* condition. That is, the first n outputs of a stream function may depend on at most its first n inputs. Writing $\lfloor xs \rfloor_n$ for the n-element prefix of the stream xs, we formalize causality as follows:

Definition 1. *(Causality) A stream function* $f : A^\omega \to B^\omega$ *is* causal, *when for all* n *and all streams* as *and* as', *we have that if* $\lfloor as \rfloor_n = \lfloor as' \rfloor_n$ *then* $\lfloor f\ as \rfloor_n = \lfloor f\ as' \rfloor_n$.

Furthermore, reactive programs often define streams by feedback. If a stream transformer can produce the first value of its output without looking at its input, then we can constructing a fixed point via feedback, taking the n-th output and supplying it as the input at time $n+1$. So as long as we can generate *more* than n outputs from the first n inputs, we can find a fixed point. Formalizing this gives us a definition of guardedness for defining fixed points:

Definition 2. *(Guardedness) A function* $f : A^\omega \to B^\omega$ *is* guarded, *when there exists* $k > 0$ *such that for all* n *and all streams* as *and* as', *if* $\lfloor as \rfloor_n = \lfloor as' \rfloor_n$ *then* $\lfloor f\ as \rfloor_{n+k} = \lfloor f\ as' \rfloor_{n+k}$.

$$\boxed{\Sigma \rhd_i \Omega_i \vdash C \dashv \Omega_i'}$$

$$\frac{}{\Sigma \rhd_i \Omega_i \vdash \Box \dashv \Omega_i} \text{ExtNil}$$

$$\frac{\cdot;\cdot;\Omega_i \vdash v :_i A \qquad \Sigma', u' :_{i+1} \Diamond; \cdot; \Omega_i \vdash e :_{i+1} S(A) \qquad \Sigma \rhd_i \Omega_i, x :_i S(A) \vdash C \dashv \Omega_i'}{\Sigma, u :_i \Diamond, \Sigma' \rhd_i \Omega_i \vdash \mathsf{let}\ x = \mathsf{cons}(u, v, u'.\ e)\ \mathsf{in}\ C \dashv \Omega_i', x :_i S(A)} \text{ExtCons}$$

$$\frac{y :_i S(A) \in \Omega_i \qquad \Sigma \rhd_i \Omega_i, x :_{i+1} S(A) \vdash C \dashv \Omega_i'}{\Sigma \rhd_i \Omega_i \vdash \mathsf{let}\ x = \mathsf{tail}(y)\ \mathsf{in}\ C \dashv \Omega_i', x :_{i+1} S(A)} \text{ExtTail}$$

$$\boxed{C \circ C'}$$

$$\begin{array}{l} \Box \circ C' \triangleq C' \\ (\mathsf{let}\ x = \mathsf{tail}(y)\ \mathsf{in}\ C) \circ C' \triangleq \mathsf{let}\ x = \mathsf{tail}(y)\ \mathsf{in}\ (C \circ C') \\ (\mathsf{let}\ x = \mathsf{cons}(u, v, u'.\ e)\ \mathsf{in}\ C) \circ C' \\ \quad \triangleq \mathsf{let}\ x = \mathsf{cons}(u, v, u'.\ e)\ \mathsf{in}\ (C \circ C') \end{array}$$

$$\boxed{C@x \Rightarrow v}$$

$$\frac{C \equiv C' \circ \mathsf{let}\ y = \mathsf{tail}(z)\ \mathsf{in}\ \Box \qquad C'@x \Rightarrow v}{C@x \Rightarrow v} \text{LookupTail}$$

$$\frac{C \equiv C' \circ \mathsf{let}\ y = \mathsf{cons}(u, v', u'.\ e)\ \mathsf{in}\ \Box \qquad C'@x \Rightarrow v}{C@x \Rightarrow v} \text{LookupNext}$$

$$\frac{C \equiv C' \circ \mathsf{let}\ x = \mathsf{cons}(u, v, u'.\ e)\ \mathsf{in}\ \Box}{C@x \Rightarrow v} \text{LookupCurr}$$

Figure 5. Context Typing and Operations

$$\boxed{C[e] \Downarrow C'[v]}$$

$$\frac{}{C[v] \Downarrow C[v]} \text{ValE}$$

$$\frac{C[e] \Downarrow C'[\lambda x : A.\ e''] \qquad C'[e'] \Downarrow C''[v] \qquad C''[[v/x]e''] \Downarrow C'''[v']}{C[e\ e'] \Downarrow C'''[v']} \text{AppE}$$

$$\frac{C[e] \Downarrow C'[\lambda u : R.\ e'] \qquad C'[[t/u]e'] \Downarrow C''[v]}{C[e\ t] \Downarrow C''[v]} \text{LAppE}$$

$$\frac{C[e] \Downarrow C'[v] \qquad x \text{ fresh} \qquad C'' = C' \circ \mathsf{let}\ x = \mathsf{cons}(t, v, u.\ e')\ \mathsf{in}\ \Box}{C[\mathsf{cons}(t, e, u.\ e')] \Downarrow C''[x]} \text{ConsE}$$

$$\frac{C[e] \Downarrow C'[x] \qquad C'@x \Rightarrow \mathsf{cons}(t, v, u.\ e)}{C[\mathsf{head}(e)] \Downarrow C'[v]} \text{HeadE}$$

$$\frac{C[e] \Downarrow C'[y] \qquad (C' \circ \mathsf{let}\ x = \mathsf{tail}(y)\ \mathsf{in}\ \Box)[e'] \Downarrow C''[v]}{C[\mathsf{let}\ x = \mathsf{tail}(e)\ \mathsf{in}\ e'] \Downarrow C''[v]} \text{TailE}$$

$$\frac{C[e] \Downarrow C'[\bullet e_1] \qquad C'[[e_1/x]e'] \Downarrow C''[v'']}{C[\mathsf{let}\ \bullet x = e\ \mathsf{in}\ e'] \Downarrow C''[v'']} \bullet\text{E}$$

$$\frac{C[[\mathsf{fix}\ x : A.\ e/x]e] \Downarrow C'[v]}{C[\mathsf{fix}\ x : A.\ e] \Downarrow C'[v]} \text{FixE}$$

$$\frac{C[e] \Downarrow C'[v] \qquad C'[[v/x]e'] \Downarrow C''[v'']}{C[\mathsf{let}\ x = e\ \mathsf{in}\ e'] \Downarrow C''[v'']} \text{LetE}$$

$$\frac{C[[t_1/u, t_2/v]e'] \Downarrow C'[v']}{C[\mathsf{let}\ \langle u, v \rangle = \langle t_1, t_2 \rangle\ \mathsf{in}\ e'] \Downarrow C'[v']} \otimes\text{E} \qquad \frac{C[e] \Downarrow C'[v]}{C[!e] \Downarrow C'[!v]} !\text{I}$$

$$\frac{C[e] \Downarrow C'[!v] \qquad C'[[v/x]e'] \Downarrow C''[v'']}{C[\mathsf{let}\ !x = e\ \mathsf{in}\ e'] \Downarrow C''[v'']} !\text{E}$$

Figure 6. Within-Step Operational Semantics

However, these definitions apply only to stream functions, and real programs need more types than just the stream type. So we need generalisations of causality which work at other types such as streams of streams and higher-order functions. To generalize these definitions, we follow our earlier work [13] by moving to a category of metric spaces. A *complete 1-bounded bisected ultrametric space* A (which we will simply call "ultrametric space") is a pair $(|A|, d)$, where $|A|$ is a set and $d \in |A| \times |A| \to [0, 1]$ is a distance function satisfying the following properties:

1. $d(x, y) = 0$ iff $x = y$
2. $d(x, y) = d(y, x)$
3. $d(x, z) \leq \max(d(x, y), d(y, z))$
4. $d(x, y) = 0$ or 2^{-n} for some n
5. All Cauchy sequences have limits

We take the morphisms between ultrametric spaces to be the *non-expansive maps* $f : A \to B$. These are the set-theoretic functions $f \in |A| \to |B|$ such that:

For all $a, a' \in |A|$, we have $d_B(f\ a, f\ a') \leq d_A(a, a')$

That is, a morphism between A and B is a function f such that it takes any two points in A to two points in B that are at least as close — it is a non-distance-increasing function.

The category of ultrametric spaces is useful for two reasons. First, the causal stream functions are *exactly* the nonexpansive maps between spaces of streams with the Cantor metric (i.e., the distance between two streams is 2^{-n}, where n is the first position at which they disagree). Since the category of 1-bounded complete ultrametric spaces is Cartesian closed, we have our higher-type generalisation of causality — one which would be very difficult to find from purely operational considerations. Second, nonempty metric spaces satisfy Banach's theorem, which lets us define fixed points at arbitrary types:

Proposition 1. *(Banach's Contraction Map Theorem) If A is a nonempty complete metric space, and $f : A \to A$ is a strictly contractive function, then f has a unique fixed point.*

5.2 Modeling Space Bounds with Length Spaces

As noted earlier, simple casuality still allows undesirable functions. Requiring a stream function to depend only on its history does not prevent it from depending on its *whole* history.

$$\boxed{C[z] \rightsquigarrow e}$$

$$\frac{}{\Box[z] \rightsquigarrow x}$$

$$\frac{C[z] \rightsquigarrow e}{\mathsf{let}\ x = \mathsf{tail}(y)\ \mathsf{in}\ C[z] \rightsquigarrow \mathsf{let}\ x = y\ \mathsf{in}\ e}$$

$$\frac{C[z] \rightsquigarrow e}{\mathsf{let}\ x = \mathsf{cons}(u, v, u'.\ e')\ \mathsf{in}\ C[z] \rightsquigarrow \mathsf{let}\ x = [u/u']e'\ \mathsf{in}\ e}$$

$$\boxed{\mathsf{Step}(\Omega_i)}$$

$$\begin{array}{lcl} \mathsf{Step}(\cdot) & = & \cdot \\ \mathsf{Step}(\Omega_i, x :_i S(A)) & = & \mathsf{Step}(\Omega_i), x :_{i+1} S(A) \\ \mathsf{Step}(\Omega_i, x :_{i+1} S(A)) & = & \mathsf{Step}(\Omega_i), x :_{i+1} S(A) \end{array}$$

$$\boxed{\mathsf{Step}_{\Omega_i}(\omega) \in [\![\Omega_i]\!]_i \to [\![\mathsf{Step}(\Omega_i)]\!]_i}$$

$$\begin{array}{lcl} \mathsf{Step}_{\cdot}(\langle\rangle) & = & \langle\rangle \\ \mathsf{Step}_{\Omega_i, x :_i S(A)}(\omega, vs) & = & (\mathsf{Step}_{\Omega_i}(\omega), \mathit{tail}(vs)) \\ \mathsf{Step}_{\Omega_i, x :_{i+1} S(A)}(\omega, vs) & = & (\mathsf{Step}_{\Omega_i}(\omega), vs) \end{array}$$

Figure 7. The Next Step Operator

To deal with this issue, we adapt the *length spaces* of Hofmann [10], which give a model of space-bounded computation. The idea behind this model is to start with a partially ordered resource monoid R representing space resources ($\mathbb{N}$ in the original work). One then constructs the category of length spaces as follows.

A *length space* A is a pair $(|A|, \sigma_A : |A| \to R)$, consisting of a set of elements $|A|$ and a *size function* σ which assigns a size $\sigma_A(a)$ to each element $a \in |A|$. A morphism of length spaces $f : A \to B$ is a *non-size-increasing function*. That is, it is a set-theoretic function $f \in |A| \to |B|$ with the property that:

$$\forall a \in |A|.\ \sigma_B(f\ a) \le \sigma_A(a)$$

The programming language intuition is that a morphism $A \to B$ is a term of type B with a free variable in A, and so a term cannot use more memory than it receives from its environment.

To model the permission to allocate, we can define a length space of type $\Diamond \triangleq (1, \lambda\langle\rangle.\ 1)$. The space $\Diamond$ is uninteresting computationally (its set only has the unit in it), but it brings a permission to allocate with it. So we can model computations which do allocation by giving them permission elements, thereby controlling the allocation performed.

6. The Denotational Semantics

6.1 The Resource Model

In the synchronous dataflow model, there is a global, ambient notion of time. Furthermore, higher-order reactive programs can create a dataflow graph dynamically, by waiting for an event before choosing to build cons cells to do some computation. So we need a resource structure capable of modelling space usage over time.

Therefore we take resources to be the monoidal lattice $R = (\mathrm{Time} \to \mathrm{Space}, \bot, \max, \top, \min, 0, \oplus, \le)$, where $\mathrm{Time} = \mathbb{N}$, and $\mathrm{Space} = \mathbb{N} \uplus \{\infty\}$ (the vertical natural numbers with a topmost element). Intuitively, time is discrete, and measured in ticks. Space counts the number of cons cells used in the program, and may be infinite (obviously, we cannot implement such programs). We define the lattice operations as follows:

1. $\bot = \lambda k.\ 0$
2. $\top = \lambda k.\ \infty$
3. $0 = \lambda k.\ 0$
4. $\max(c, d) = \lambda k.\ \max(c_k, d_k)$
5. $\min(c, d) = \lambda k.\ \min(c_k, d_k)$
6. $c \oplus d = \lambda k.\ c_k + d_k$
7. $c \le d$ iff $\forall k \in \mathrm{Time}$, we have $c_k \le d_k$

Essentially, we lift the lattice structure of the vertical natural numbers pointwise across time (with $(0, +)$ as the monoidal structure), so that a resource $c \in R$ describes the number of cons cells that are used at each time step.

We then turn R into an ultrametric space by equipping it with the Cantor metric:

$$d_R(c, d) = 2^{-n} \text{ where } n = \min\{k \in \mathrm{Time} \mid c_k \ne d_k\}$$

6.2 The Category of Complete Ultrametric Length Spaces

A *complete 1-bounded bisected ultrametric length space* A (which we will gloss as "metric length space") is a tuple $(|A|, d, \sigma)$, where $(|A|, d)$ is a complete 1-bounded bisected ultrametric space, and $\sigma_A : |A| \to R$ is a size function giving each element of $|A|$ a size drawn from R.

Furthermore, the size function $\sigma : |A| \to R$ must be a nonexpansive map between $(|A|, d)$ and (R, d_R). Nonexpansiveness ensures that we cannot tell if the memory usage requirements of two elements of $|A|$ differs until we know that the elements themselves differ. In addition to being intuitively reasonable, this requirement ensures that limits of Cauchy sequences will be well-behaved with respect to size, which we need to ensure the completeness of the size-0 subspace of A that we use to interpret $!A$.

The morphisms of this category are the *nonexpansive size-preserving maps* $f : A \to B$, which are the set-theoretic functions $f \in |A| \to |B|$ such that:

- For all $a, a' \in |A|$, we have $d_B(f\ a, f\ a') \le d_A(a, a')$
- For all $a \in |A|$, we have $\sigma_B(f\ a) \le \sigma_A(a)$

That is, the morphisms we consider are the functions which are both causal and space-bounded.

6.3 Categorical Structure

Metric length spaces and nonexpansive size-preserving maps form a category that we use to interpret our programming language. First, it forms an intuitionistic bicartesian BI category, which is a doubly-closed category with both cartesian and monoidal closed structure, as well as supporting coproduct structure. Second, this category also models the resource types $\Diamond$ of Hofmann [10], as well as a resource-freedom modality $!A$, which is comonadic in the usual fashion of linear logic. Third, it supports a version of the delay modality of our earlier work [13], which lets us interpret guarded recursion via Banach's fixed point theorem.

We give the definitions of all of these objects below. In Figure 9, we define the distance and size functions, and in Figure 8, we give the natural transformations associated with the objects.

- $1 = (\{*\}, d_1, \sigma_1)$
- $A + B = (|A| + |B|, d_{A+B}, \sigma_{A+B})$
- $A \times B = (|A| \times |B|, d_{A\times B}, \sigma_{A\times B})$
- $A \Rightarrow B = (|A| \to |B|, d_{A\to B}, \sigma_{A\to B})$
- $A + B = (|A| + |B|, d_{A+B}, \sigma_{A+B}$

- $A \star B = (|A| \times |B|, d_{A \star B}, \sigma_{A \star B})$
- $A -\!\!\star B = (|A| \to |B|, d_{A -\!\star B}, \sigma_{A -\!\star B})$
- $\bullet A = (|A|, d_{\bullet A}, \sigma_{\bullet A})$
- $S(A) = (|A|^{\omega}, d_{S(A)}, \sigma_{S(A)})$
- $!A = (\{a \in A \mid \sigma_A(a) = 0\}, \sigma_{!A})$
- $\Diamond = (\{*\}, d_1, \sigma_{\Diamond})$
- $\diamond = (\{*\}, d_1, \sigma_{\diamond})$

The construction of Cartesian and monoidal products closely follows that of Hofmann [10]. The Cartesian product is a "sharing product", in which the associated resources are available to both components (this explains the use of $\max$), and the monoidal product is a "disjoint product", in which the resources are divided between the two components (explaining the use of $\oplus$ in the size function). The best intuition for the closed structure comes from implementing first-class functions as closures: the monoidal exponential $A -\!\!\star B$ takes an argument which does not share with the captured environment, and the Cartesian exponential $A \Rightarrow B$ which does.

A difference between our work and earlier work on length spaces is our heavy use of the category's Cartesian closed structure. Indeed, dal Lago and Hofmann [15] use a realizability model to remove the Cartesian closed structure from their semantics — they wished to prevent duplicated variables in lambda-terms from enabling large increases in the size of a lambda-term under reduction, since this makes establishing strict resource bounds more difficult. As we only want to track the allocation of cells, but wish to allow free sharing otherwise, the CCC structure takes on a central role in our model.

Our next-step modality's metric $d_{\bullet A}$ is the same as in Krishnaswami and Benton [13], but the size function $\sigma_{\bullet A}$ (which shifts all sizes 1 timestep into the future relative to σ_A), means that $\bullet(A \to B) \not\simeq \bullet A \to \bullet B$. This breaks with Krishnaswami and Benton [13] and Nakano [18], significantly changing the elimination rules.

As mentioned earlier, this limitation is intentional: we do not want delay operators at types for which delay would be expensive. Our semantics rules out such maps with the size function for streams plus the requirement that morphisms are non-size-increasing. The size function for streams gives a size for the stream as 1 to account for the size of the stream itself, plus the maximum space usage of all the values the stream takes on. Intuitively, a stream can seen as taking the space for an infinitary Cartesian product $A \times \bullet A \times \bullet^2 A \times \ldots$, plus a constant for the stream cell itself. This is the only place where we increment the size of a value relative to its components, which justifies the idea that sizes measure the number of cons cells. Since delaying a stream shifts its time usage by one step, we have no *a priori* reason to expect that a delay map will exist at all types.

However, for types such as $\mathbb{N}, !A$, and $\Diamond$, there do exist maps $A \to \bullet A$, which is why time subsumption is justified for the linear and pure contexts. Furthermore, all types whose values are all of size zero have maps $A \to !A$. As a result, we can introduce constants corresponding to such maps for these types, allowing types such as numbers and booleans be promoted to the pure fragment. (In fact, our implementation applies these coercions implicitly, providing a slightly friendlier syntax than presented here.)

Our semantic model contains a space $\Diamond$ to interpret the resource type $\Diamond$, which gives 1 unit of space at every time tick. Our model uses the additional metric length space $\diamond$, which gives 1 unit of space at time 0, and no units of space at any other time. This lets us give a nice type to $cons : \diamond \star (A \times \bullet S(A)) \to S(A)$. Note that the type $A \times \bullet S(A)$ lacks the space to form a stream — we need 1 unit of space at time 0, which neither the A nor the $\bullet S(A)$ provide.

1	$: A \to I$	$=$	$\lambda a.\ *$
π_1	$: A \times B \to A$	$=$	$\lambda(a,b).\ a$
π_2	$: A \times B \to B$	$=$	$\lambda(a,b).\ b$
$\langle f, g \rangle$	$: A \to B \times C$	$=$	$\lambda a.\ (f\ a, g\ a)$
	where $f : A \to B, g : A \to C$		
$\lambda \langle f \rangle$	$: A \to B \Rightarrow C$	$=$	$\lambda a.\ \lambda b.\ f(a,b)$
	where $f : A \times B \to C$		
$eval$	$: (A \Rightarrow B) \times A \to B$	$=$	$\lambda(f,a).\ f\ a$
$f \star g$	$: A \star B \to C \star D$	$=$	$\lambda(a,b).\ (f\ a, g\ b)$
	where $f : A \to C, g : B \to D$		
α	$: (A \star B) \star C \to A \star (B \star C)$	$=$	$\lambda((a,b),c).(a,(b,c))$
α^{-1}	$: A \star (B \star C) \to (A \star B) \star C$	$=$	$\lambda(a,(b,c)).((a,b),c)$
γ	$: A \star B \to B \star A$	$=$	$\lambda(a,b).\ (b,a)$
ρ	$: A \star I \to A$	$=$	$\lambda(a,*).\ a$
ρ^{-1}	$: A \star I \to A$	$=$	$\lambda a.\ (a,*)$
$\hat{\lambda} \langle f \rangle$	$: A \to B -\!\!\star C$	$=$	$\lambda a.\ \lambda b.\ f(a,b)$
	where $f : A \star B \to C$		
$eval_{-\!\star}$	$: (A -\!\!\star B) \star A \to B$	$=$	$\lambda(f,a).\ f\ a$
ϵ	$: !A \to A$	$=$	$\lambda a.\ a$
$f^{\dagger}$	$: !A \to !B$	$=$	$\lambda a.\ f\ a$
	where $f : !A \to B$		
$!f$	$: !A \to !B$	$=$	$\lambda a.\ f\ a$
	where $f : A \to B$		
δ	$: !A \to \bullet !A$	$=$	$\lambda a.\ a$
$\bullet f$	$: \bullet A \to \bullet B$	$=$	$\lambda a.\ f\ a$
	where $f : A \to B$		
η	$: ! \bullet A \to \bullet !A$	$=$	$\lambda a.\ a$
η^{-1}	$: \bullet !A \to \bullet A$	$=$	$\lambda a.\ a$
$head$	$: S(A) \to A$	$=$	$\lambda(x \cdot xs).\ x$
$tail$	$: S(A) \to \bullet S(A)$	$=$	$\lambda(x \cdot xs).\ xs$
$cons$	$: \diamond \star (A \times \bullet S(A) \to S(A)$	$=$	$\lambda(*, (x, xs)).\ x \cdot xs$
$split$	$: \Diamond \to \diamond \star \bullet \Diamond$	$=$	$\lambda *.\ (*, *)$
$split^{-1}$	$: \diamond \star \bullet \Diamond \to \Diamond$	$=$	$\lambda(*,*).\ *$
fix	$: !(\bullet !A \Rightarrow !A) \to !A$	$=$	$\lambda f.\ \mu(f)$
ι	$: \bullet A \times \bullet B \to \bullet(A \times B)$	$=$	$\lambda(a,b).\ (a,b)$
ι^{-1}	$: \bullet(A \times B) \to \bullet A \times \bullet B$	$=$	$\lambda(a,b).\ (a,b)$
$\iota_{\star}$	$: \bullet A \star \bullet B \to \bullet(A \star B)$	$=$	$\lambda(a,b).\ (a,b)$
$\iota_{\star}^{-1}$	$: \bullet(A \star B) \to \bullet A \star \bullet B$	$=$	$\lambda(a,b).\ (a,b)$
ξ	$: A \star B \to A \times B$	$=$	$\lambda(a,b).\ (a,b)$
σ	$: A \times !B \to A \star !B$	$=$	$\lambda(a,b).\ (a,b)$
ψ	$: !(A \times B) \to !A \star !B$	$=$	$\lambda(a,b).\ (a,b)$
ψ^{-1}	$: !A \star !B \to !(A \times B)$	$=$	$\lambda(a,b).\ (a,b)$

Figure 8. Categorical Combinators

6.4 Denotational Interpretation

We give the interpretation of types and contexts in Figure 10. The interpretation of types offers no surprises, but the interpretation of contexts is relative to the current time. The interpretations of Θ and Δ keeps hypotheses at times earlier than the current time, but Γ simply drops all earlier hypotheses. This corresponds to the difference between the type rules RHYP and PHYP on the one hand, and the rule EHYP on the other. In all three cases, future hypotheses are interpreted with the delay modality.

In Figure 11, we give a time-indexed interpretation function for expressions, $[\![\Theta; \Pi; \Gamma \vdash e :_i A]\!]_i$ which has the type $[\![\Theta]\!]_i \star ![\![\Pi]\!]_i \star [\![\Gamma]\!]_i \to [\![A]\!]$. The interpretation of $\bullet$I makes use of the functoriality of $\bullet A$ to interpret the body of the delay in the future, and then bring it back to the past, with the necessary action on contexts defined in Figure 10. The other rules are as expected, with the resource context managed in a single-threaded way and the other contexts duplicated freely. We can then show that this semantics is sound with respect to substitution.

$$
\begin{array}{lcl}
d_{A+B} & = & \lambda(v,v').\ \begin{cases} d_A(x,x') & \text{if } v = \mathsf{inl}\ x \wedge v' = \mathsf{inl}\ x' \\ d_B(y,y') & \text{if } v = \mathsf{inr}\ y \wedge v' = \mathsf{inr}\ y' \\ 1 & \text{otherwise} \end{cases} \\
d_1 & = & \lambda(\langle\rangle, \langle\rangle).\ 0 \\
d_{A\times B} & = & \lambda((a,b),(a',b')).\ \max(d_A(a,a'), d_B(b,b')) \\
d_{A\star B} & = & \lambda((a,b),(a',b')).\ \max(d_A(a,a'), d_B(b,b')) \\
d_{A\Rightarrow B} & = & \lambda(f,g).\ \max\{d_B(f\ a, g\ a) \mid a \in |A|\} \\
d_{A\multimap\!\star B} & = & \lambda(f,g).\ \max\{d_B(f\ a, g\ a) \mid a \in |A|\} \\
d_{\bullet A} & = & \lambda(a,a').\ \frac{1}{2} d_A(a,a') \\
d_{S(A)} & = & \lambda(xs,ys).\ \max\left\{2^{-n} \cdot d_A(xs_n, ys_n) \mid n \in \mathbb{N}\right\} \\
d_{!A} & = & \lambda(a,a').\ d_A(a,a') \\
d_{\Diamond} & = & \lambda(\langle\rangle, \langle\rangle).\ 0 \\
d_{\diamond} & = & \lambda(\langle\rangle, \langle\rangle).\ 0
\end{array}
$$

$$
\begin{array}{lcl}
\sigma_{A+B} & = & \lambda v.\ \begin{cases} \sigma_A(a) & \text{if } v = \mathsf{inl}\ a \\ \sigma_B(b) & \text{if } v = \mathsf{inr}\ b \end{cases} \\
\sigma_1 & = & \lambda\langle\rangle.\ 0 \\
\sigma_{A\times B} & = & \lambda(a,b).\ \max(\sigma_A(a), \sigma_B(b)) \\
\sigma_{A\star B} & = & \lambda(a,b).\ \sigma_A(a) \oplus \sigma_B(b) \\
\sigma_{A\Rightarrow B} & = & \lambda f.\ \min\{k \in R \mid \forall a \in |A|.\ \sigma_B(f\ a) \leq \max(k, \sigma_A(a))\} \\
\sigma_{A\multimap\!\star B} & = & \lambda f.\ \min\{k \in R \mid \forall a \in |A|.\ \sigma_B(f\ a) \leq k \oplus \sigma_A(a)\} \\
\sigma_{\bullet A} & = & \bullet\sigma_A \\
\sigma_{S(A)} & = & \lambda xs.\ \lambda k.\ 1 \oplus \max\left\{(\bullet^i \sigma_A)(xs_i) \mid i \in \mathrm{Time}\right\} \\
\sigma_{\Diamond} & = & \lambda\langle\rangle.\ \lambda k.\ 1 \\
\sigma_{\diamond} & = & \lambda\langle\rangle.\ \lambda k.\ \text{if } k = 0 \text{ then } 1 \text{ else } 0 \\
\sigma_{!A} & = & \lambda a.\ 0 \\
\\
\bullet\sigma & = & \lambda a.\ \lambda k.\ \text{if } k = 0 \text{ then } 0 \text{ else } \sigma(a)(k-1)
\end{array}
$$

Figure 9. The Distance and Size Functions

Theorem 6. *(Semantic Substitution) Suppose* $\theta \in [\![\Theta]\!]_i$, $\theta' \in [\![\Theta']\!]_i$, $\pi \in [\![\Pi]\!]_i$ *and* $\gamma \in [\![\Gamma]\!]_i$. *Then*

1. *If* $\Theta \vdash t :_i R$ *and* $\Theta', u :_i R \vdash t' :_j R'$, *then* $[\![\Theta, \Theta' \vdash [t/u]t' :_j R']\!]\ (\theta, \theta')$ *is equal to* $[\![\Theta', u :_i R \vdash t' :_j R']\!]\ (\theta', [\![\Theta \vdash t :_i R]\!]\ \theta)$.
2. *If* $\Theta \vdash t :_i R$ *and* $\Theta', u :_i R; \Pi; \Gamma \vdash e :_j A$, *then* $[\![\Theta, \Theta'; \Pi; \Gamma \vdash [t/u]e :_j A]\!]\ ((\theta, \theta'), \pi, \gamma)$ *equals* $[\![\Theta', u :_i R; \Pi; \Gamma \vdash e :_j A]\!]\ ((\theta', [\![\Theta \vdash t :_i R]\!]\ \theta), \pi, \gamma)$
3. *If* $\cdot; \Pi; \cdot \vdash e :_i A$ *and* $\Theta; \Pi, x :_j A; \Gamma \vdash e' :_k B$ *and* $i \leq j$, *then* $[\![\Theta; \Pi; \Gamma \vdash [e/x]e' :_k B]\!]\ (\theta, \pi, \gamma)$ *equals* $[\![\Theta; \Pi, x :_j A; \Gamma \vdash e' :_k B]\!]\ (\theta, (\pi, [\![\cdot; \Pi; \cdot \vdash e :_i A]\!]\ (\langle\rangle, \pi, \langle\rangle)), \gamma)$
4. *If* $\cdot; \Pi; \Gamma \vdash e :_i A$ *and* $\Theta; \Pi; \Gamma, x :_i A \vdash e' :_j B$, *then* $\Theta; \Pi; \Gamma \vdash [e/x]e' :_j B$ *equals* $[\![\Theta; \Pi; \Gamma, x :_i A \vdash e' :_j B]\!]\ (\theta, \pi, (\gamma, [\![\cdot; \Pi; \Gamma \vdash e :_i A]\!]\ (\langle\rangle, \pi, \gamma)))$

These theorems follow from structural induction on the typing derivation of the term being substituted into.

Figure 12, gives the interpretation of contexts, using the expression semantics to define the meaning of each stream bound by the context. This interpretation is sound with respect to the "hole-filling" of terms in contexts:

Theorem 7. *(Context Soundness) If* $\Sigma \rhd_i \Omega_i \vdash C \dashv \Omega'_i$ *and* $\Sigma'; \cdot; \Omega_i, \Omega'_i \vdash e :_i A$ *and* $\sigma \in [\![\Sigma]\!]_i$ *and* $\sigma' \in [\![\Sigma']\!]_i$ *and* $\omega \in [\![\Omega_i]\!]_i$, *then* $[\![\Sigma, \Sigma'; \cdot; \Omega_i \vdash C[e] :_i A]\!]\ ((\sigma, \sigma'), \omega)$ *is equal to* $[\![\Sigma'; \cdot; \Omega_i, \Omega'_i \vdash e :_i A]\!]\ (\sigma', \langle\rangle, (\omega, [\![\Sigma \rhd_i \Omega_i \vdash C \dashv \Omega'_i]\!]\ (\sigma, \omega))$.

Now we can show that the within-step operational semantics is sound with respect to the denotational model:

Theorem 8. *(Soundness of Within-Step Semantics) Let* $\Sigma; \cdot; \cdot \vdash C[e] :_i A$ *and* $C[e] \Downarrow C'[v]$. *Then* $[\![\Sigma; \cdot; \cdot \vdash C[e] :_i A]\!]$ *equals* $[\![\Sigma; \cdot; \cdot \vdash C'[v] :_i A]\!]$.

Finally, we can show that advancing the clock has the expected semantics:

Theorem 9. *(Soundness of Advancing the Clock) Let* $\Sigma; \cdot; \Omega_i \vdash C[x] :_i S(A)$, *and* $\sigma \in [\![\Sigma]\!]_i$ *and* $\omega \in [\![\Omega_i]\!]_i$, *and suppose* $C[x] \rightsquigarrow e$. *Then we have that*

$$\mathsf{tail}([\![\Sigma; \cdot; \Omega_i \vdash C[x] :_i S(A)]\!]\ (\sigma, \langle\rangle, \omega))$$

is equal to $\bullet[\![\Sigma; \cdot; \mathsf{Step}(\Omega_i) \vdash e :_{i+1} S(A)]\!]\,(\sigma', \langle\rangle, \omega'))$ *where* $\sigma' = \mathrm{Next}^i_\Sigma\ \sigma$ *and* $\omega' = \mathsf{Step}_{\Omega_i}(\omega)$.

This theorem connects the operation of stepping the heap with the denotational interpretation — each time we advance the clock, each stream in a closed context C will become its tail. So advancing the clock successively enumerates the elements of the stream.

7. Discussion

In this paper, we have introduced an expressive type system for writing stream programs, and given an operational semantics respecting the space-efficiency claims of the type system. Our semantic model is one of the primary contributions of this paper, since it lets us reason about space usage without surrendering the sets-and-functions view of FRP. Also, our model contains many operations which are not currently expressible in our language: for example, in the future we might want richer types in the affine context and function space, so that operations like in-place map can be typed $!(A \to B) \to S(A) \multimap\!\star S(B)$.

Our language contains an explicit delay modality (as in logical presentations of step-indexing [2, 7, 18]) and an update-based operational semantics. The reader may find it surprising that, although our operational semantics does make use of a form of mutable higher-order store, the logical relation we give is not step-indexed. The reason this is possible is essentially that FRP is a *synchronous* model of computation, in which all the updates happen in a separate phase from ordinary functional evaluation. This explains why we were able to present a two-phase semantics, and since no heap modifications take place during ticks, there is no need for step-indexing.

Wan et al. [25] introduced real-time FRP; a restricted subset of FRP sharing many of the same design choices of synchronous dataflow languages. It is essentially first-order (streams can carry general values of the host language, but these values can not themselves refer to streams), and makes use of a novel continuation typing to ensure that all recursive signals are tail-recursive. As a result, the language requires only constant-stack and constant-time FRP reductions. Event-driven FRP [26] is similar, but relaxes FRP's timing constraints by dropping the global clock.

None of our examples using higher-order functions or nested streams can be programmed in real-time FRP, which is carefully engineered to avoid exposing a first-class behaviour type, and so cannot express higher-order operations on streams. All productive real-time FRP programs can be written in our language, since we can define all of real-time FRP's stream operators (switching, delays, multimodal recursion) as ordinary user-level programs.

Liu et al.'s causal commutative arrows [17] are another attempt to control the memory and space usage of reactive programs. This work takes advantage of the fact that pure arrowized FRP (i.e., without switching) builds fixed dataflow graphs, allowing programs to be optimized into single-loop code via a transformation reminiscent of the Bohm-Jacopini theorem. Our language does not seem to support such an elegant normal form, because of the presence

$$
\begin{array}{lcl}
[\![A \to B]\!] & = & [\![A]\!] \Rightarrow [\![B]\!] \\
[\![S(A)]\!] & = & S([\![A]\!]) \\
[\![\bullet A]\!] & = & \bullet [\![A]\!] \\
[\![!A]\!] & = & ![\![A]\!] \\
[\![R \multimap A]\!] & = & [\![R]\!] \twoheadrightarrow\!\!\star [\![A]\!] \\
[\![\Diamond]\!] & = & \Diamond \\
[\![R \otimes S]\!] & = & [\![R]\!] \star [\![S]\!]
\end{array}
$$

$\boxed{[\![\Theta]\!]_i}$ $\boxed{[\![\Pi]\!]_i}$ $\boxed{[\![\Gamma]\!]_i}$

$$
\begin{array}{lcll}
[\![\Theta]\!]_i & \in & \mathrm{Obj} & \\
[\![\cdot]\!]_i & = & I & \\
[\![\Theta, x :_n R]\!]_i & = & [\![\Theta]\!]_i \star \bullet^{n-i}[\![R]\!] & \text{when } n \geq i \\
[\![\Theta, x :_n R]\!]_i & = & [\![\Theta]\!]_i \star [\![R]\!] & \text{when } n < i
\end{array}
$$

$$
\begin{array}{lcll}
[\![\Pi]\!]_i & \in & \mathrm{Obj} & \\
[\![\cdot]\!]_i & = & I & \\
[\![\Pi, x :_n A]\!]_i & = & [\![\Pi]\!]_i \times \bullet^{n-i}[\![A]\!] & \text{when } n \geq i \\
[\![\Pi, x :_n A]\!]_i & = & [\![\Pi]\!]_i \times [\![A]\!] & \text{when } n < i
\end{array}
$$

$$
\begin{array}{lcll}
[\![\Gamma]\!]_i & \in & \mathrm{Obj} & \\
[\![\cdot]\!]_i & = & I & \\
[\![\Gamma, x :_n A]\!]_i & = & [\![\Gamma]\!]_i \times \bullet^{n-i}[\![A]\!] & \text{when } n \geq i \\
[\![\Gamma, x :_n A]\!]_i & = & [\![\Gamma]\!]_i & \text{when } n < i
\end{array}
$$

$\boxed{\mathrm{Next}^i_\Theta \in [\![\Theta]\!]_i \to \bullet[\![\Theta]\!]_{i+1}}$ $\boxed{\mathrm{Next}^i_\Pi \in [\![\Pi]\!]_i \to \bullet[\![\Pi]\!]_{i+1}}$

$\boxed{\mathrm{Next}^i_\Gamma \in [\![\Gamma]\!]_i \to \bullet[\![\Gamma]\!]_{i+1}}$

$$
\begin{array}{lcll}
\mathrm{Next}^i_\cdot \ \langle\rangle & = & \langle\rangle & \\
\mathrm{Next}^i_{\Theta,x:_nR}\ (\theta, r) & = & (\mathrm{Next}^i_\Theta\ \theta, \delta(r)) & \text{if } n \leq i \\
\mathrm{Next}^i_{\Theta,x:_nR}\ (\theta, r) & = & (\mathrm{Next}^i_\Theta\ \theta, r) & \text{if } n > i
\end{array}
$$

$$
\begin{array}{lcll}
\mathrm{Next}^i_\cdot \ \langle\rangle & = & \langle\rangle & \\
\mathrm{Next}^i_{\Pi,x:_nA}(\pi, v) & = & (\mathrm{Next}^i_\Pi\ \pi, \delta(v)) & \text{if } n \leq i \\
\mathrm{Next}^i_{\Pi,x:_nA}(\pi, v) & = & (\mathrm{Next}^i_\Pi\ \pi, v) & \text{if } n > i
\end{array}
$$

$$
\begin{array}{lcll}
\mathrm{Next}^i_\cdot \ \langle\rangle & = & \langle\rangle & \\
\mathrm{Next}^i_{\Gamma,x:_nA}(\gamma, v) & = & \mathrm{Next}^i_\Gamma & \text{if } n \leq i \\
\mathrm{Next}^i_{\Gamma,x:_nA}(\gamma, v) & = & (\mathrm{Next}^i_\Gamma\ \gamma, v) & \text{if } n > i
\end{array}
$$

Figure 10. Interpretation of Types and Contexts

of higher-order streams and functions, but it would nonetheless be interesting to investigate related optimizations in our setting.

Sculthorpe and Nilsson's work [22] on safe functional programming uses Agda's types to ensure productivity, by having dependent types track whether signal processors have delays before permitting feedback. Our guardedness modality is simpler but less flexible, since it cannot depend on the *values* a signal produces. However, the advantage of our modality is that it works smoothly at higher-order.

Cooper and Krishnamurthi [6] described the FrTime system, which embeds FRP into the PLT Scheme (now Racket) implementation. One commonality between FrTime and our work is that switching does not come from special primitives, but from ordinary conditionals and case statements. Unlike our denotational model, Cooper's operational semantics [5] exemplifies the "imperative FRP" tradition, in which reactivity is modelled explicitly as

$\boxed{[\![\Theta \vdash t :_i R]\!]_i \in [\![\Theta]\!]_i \to [\![R]\!]}$

$\boxed{[\![\Theta; \Pi; \Gamma \vdash e :_i A]\!]_i \in [\![\Theta]\!]_i \star ![\![\Pi]\!]_i \star [\![\Gamma]\!]_i \to [\![A]\!]}$

$[\![\Theta \vdash u :_j R]\!]_i\ \theta = \theta(u)$
$[\![\Theta, \Theta' \vdash (t, t') :_j R \otimes S]\!]_i\ (\theta, \theta') = ([\![t]\!]\ \theta, [\![t']\!]\ \theta')$

$[\![\Theta; \Pi; \Gamma \vdash x :_j A]\!]_i\ (\theta, \pi, \gamma) = \pi(x)$ (if $x :_i A \in \Pi$)

$[\![\Theta; \Pi; \Gamma \vdash x :_i A]\!]_i\ (\theta, \pi, \gamma) = \gamma(x)$ (if $x :_i A \in \Gamma$)

$[\![\Theta; \Pi; \Gamma \vdash \lambda x : A.\ e :_i A \to B]\!]_i\ (\theta, \pi, \gamma) =$
$\quad \lambda v.\ [\![\cdot; \Pi; \Gamma, x :_i A \vdash e :_i B]\!]_i\ (\langle\rangle, \pi, (\gamma, v))$

$[\![\Theta, \Theta'; \Pi; \Gamma \vdash e\ e' :_i B]\!]_i\ ((\theta, \theta'), \pi, \gamma) =$
$\quad \text{let } f = [\![e : A \to B]\!]_i\ (\theta, \pi, \gamma) \text{ in}$
$\quad \text{let } v = [\![e' : A]\!]_i\ (\theta', \pi, \gamma) \text{ in}$
$\quad \mathit{eval}(f, v)$

$[\![\Theta; \Pi; \Gamma \vdash !e :_i !A]\!]_i\ (\theta, \pi, \gamma) = [\![\cdot; \Pi; \cdot \vdash e :_i A]\!]^\dagger_i\ (\langle\rangle, \pi, \langle\rangle)$

$[\![\Theta, \Theta'; \Pi; \Gamma \vdash \mathsf{let}\ !x = e\ \mathsf{in}\ e' :_i B]\!]_i\ (\theta, \pi, \gamma) =$
$\quad [\![e']\!]_i\ (\theta', \psi^{-1}(\pi, [\![e]\!]_i\ (\theta', \pi, \gamma)), \gamma)$

$[\![\Theta, \Theta', \Theta''; \Pi; \Gamma \vdash \mathsf{cons}(t, e, u'.e') :_i S(A)]\!]_i\ ((\theta, \theta', \theta''), \pi, \gamma) =$
$\quad \text{let } (d, r) = \mathit{split}([\![t]\!]_i\ \theta) \text{ in}$
$\quad \text{let } h = [\![e]\!]_i\ (\theta', \pi, \gamma) \text{ in}$
$\quad \text{let } t = \bullet[\![e']\!]_{i+1}\ ((\mathrm{Next}^i_{\Theta''}\ \theta'', r), \mathrm{Next}^i_\Pi\ \pi, \mathrm{Next}^i_\Gamma\ \gamma) \text{ in}$
$\quad \mathit{cons}(d, (h, t))$

$[\![\Theta; \Pi; \Gamma \vdash \mathsf{head}(e) :_i A]\!]_i\ (\theta, \pi, \gamma) =$
$\quad \mathit{head}([\![\Theta; \Pi; \Gamma \vdash e :_i S(A)]\!]_i\ (\theta, \pi, \gamma))$

$[\![\Theta, \Theta'; \Pi; \Gamma \vdash \mathsf{let}\ x = \mathsf{tail}(e)\ \mathsf{in}\ e' :_i B]\!]_i\ ((\theta, \theta'), \pi, \gamma) =$
$\quad \text{let } vs = [\![\Theta]\!]_i \Pi \Gamma e S(A) i \text{ in } (\theta, \pi, \gamma)$
$\quad [\![\Theta'; \Pi; \Gamma, x :_{i+1} S(A) \vdash e' :_i B]\!]_i\ (\theta', \pi, (\gamma, \mathit{tail}(vs)))$

$[\![\Theta; \Pi; \Gamma \vdash \bullet e :_i \bullet A]\!]_i\ (\theta, \pi, \gamma) =$
$\quad \bullet[\![\cdot; \Pi; \Gamma \vdash e :_{i+1} A]\!]_{i+1}\ (\langle\rangle, \mathrm{Next}^i_\Pi\ \pi, \mathrm{Next}^i_\Gamma\ \gamma)$

$[\![\Theta, \Theta'; \Pi; \Gamma \vdash \mathsf{let}\ \bullet x = e\ \mathsf{in}\ e' :_i B]\!]_i\ ((\theta, \theta'), \pi, \gamma) =$
$\quad \text{let } v = [\![\Theta; \Pi; \Gamma \vdash e :_i \bullet A]\!]_i\ (\theta, \pi, \gamma) \text{ in}$
$\quad [\![\Theta'; \Pi; \Gamma, x :_{i+1} A \vdash e' :_i B]\!]_i\ (\theta', \pi, (\gamma, v))$

$[\![\Theta; \Pi; \Gamma \vdash \lambda u : R.\ e :_i R \multimap A]\!]_i\ (\theta, \pi, \gamma) =$
$\quad \lambda r.\ [\![u :_i R; \Pi; \Gamma \vdash e :_i A]\!]\ (r, \pi, \gamma)$

$[\![\Theta, \Theta'; \Pi; \Gamma \vdash e\ t :_i A]\!]_i\ ((\theta, \theta'), \pi, \gamma) =$
$\quad \mathit{eval}_{\twoheadrightarrow\!\star}([\![e]\!]\ (\theta, \pi, \gamma), [\![t]\!]\ \theta')$

$[\![\Theta, \Theta'; \Pi; \Gamma \vdash \mathsf{let}\ (u, v) = t\ \mathsf{in}\ e :_i A]\!]_i\ ((\theta, \theta'), \pi, \gamma) =$
$\quad \text{let } (r, s) = [\![\Theta \vdash t :_i R \otimes S]\!]\ \theta \text{ in}$
$\quad [\![\Theta', u :_i R, v :_i S; \Pi; \Gamma \vdash e :_i A]\!]\ ((\theta', r, s), \pi, \gamma)$

$[\![\Theta; \Pi; \Gamma \vdash \mathsf{fix}\ x : A.\ e :_i A]\!]_i\ (\theta, \pi, \gamma) =$
$$\text{let } f = \lambda v. \left(\begin{array}{l} \text{let } \pi' = \psi^{-1}(\pi, \eta^{-1}\ v) \text{ in} \\ [\![\cdot; \Pi, x :_{i+1} A; \cdot \vdash e :_i A]\!]_i\ (\langle\rangle, \pi', \langle\rangle) \end{array} \right) \text{ in}$$
$\quad \mathit{fix}(f)$

$[\![\Theta, \Theta'; \Pi; \Gamma \vdash \mathsf{let}\ x = e\ \mathsf{in}\ e' :_i B]\!]_i\ ((\theta, \theta'), \pi, \gamma) =$
$\quad \text{let } v = [\![\Theta; \Pi; \Gamma \vdash e :_i A]\!]_i\ (\theta, \pi, \gamma) \text{ in}$
$\quad [\![\Theta'; \Pi; \Gamma, x :_i A \vdash e' :_i B]\!]_i\ (\theta', (\pi, v), \gamma)$

Figure 11. Denotational Semantics of Terms

$$\boxed{[\![\Sigma \rhd_i \Omega_i \vdash C \dashv \Omega_i']\!] \in [\![\Sigma]\!]_i \star [\![\Omega_i]\!] \to [\![\Omega_i']\!]}$$

$$
\begin{array}{l}
[\![\Sigma \rhd_i \Omega_i \vdash \Box \dashv \cdot]\!]\ (\theta, \gamma) = \langle\rangle \\
[\![\Sigma \rhd_i \Omega_i \vdash \mathsf{let}\ y = \mathsf{tail}(x)\ \mathsf{in}\ C \dashv y :_{i+1} S(A), \Omega_i']\!]\ (\theta, \gamma) = \\
\quad \text{let } v = \mathit{tail}(\gamma(x)) \text{ in} \\
\quad \text{let } \gamma' = [\![\Sigma \rhd_i \Omega_i, y :_{i+1} S(A) \vdash C \dashv \Omega_i']\!]\ (\theta, (\gamma, v)) \text{ in} \\
\quad (\gamma', v) \\
[\![\Sigma, \Sigma', \Sigma'' \rhd_i \Omega_i \vdash \mathsf{let}\ x = \mathsf{cons}(u, v, u'.\ e)\ \mathsf{in}\ C \dashv \Omega_i', x :_{i+1} S(A)]\!] \\
\ ((\theta, \theta', \theta''), \gamma) = \\
\quad \text{let } (r, d) = [\![\Sigma' \vdash u :_i \Diamond]\!]\ \theta' \text{ in} \\
\quad \text{let } h = [\![\cdot; \cdot; \Omega_i \vdash v :_i A]\!]\ (\langle\rangle, \langle\rangle, \gamma) \text{ in} \\
\quad \text{let } \theta_1'' = (\mathrm{Next}^i_{\Sigma''}\ \theta'', d) \text{ in} \\
\quad \text{let } \gamma_1 = \mathrm{Next}^{\Omega_i}_{\gamma} \text{ in} \\
\quad \text{let } t = \bullet[\![\Sigma'', u' :_{i+1} \Diamond; \cdot; \Omega_i \vdash e :_{i+1} S(A)]\!]\ (\theta_1'', \langle\rangle, \gamma_1') \text{ in} \\
\quad \text{let } vs = \mathit{cons}(r, (h, t)) \text{ in} \\
\quad \text{let } \gamma' = [\![\Sigma \rhd_i \Omega_i, x :_i S(A) \vdash C \dashv \Omega_i']\!]\ (\theta, (\gamma, vs)) \text{ in} \\
\quad (\gamma', vs)
\end{array}
$$

Figure 12. Interpretation of Contexts

a kind of heap update in the operational semantics. We think the operational semantics in this paper, which is quasi-imperative in flavour, is close enough to both a denotational model and Cooper's semantics that it makes sense to study how to reunify the pure and the imperative flavours of FRP.

Krishnaswami and Benton [14] have also presented a language for writing reactive GUI programs. This language also makes use of linear types, but in this case they are used to track the allocation of new GUI widgets. It is not yet clear to us how one might combine space-boundedness with this kind of dynamism: we may need to add a dynamic allocation monad to our model to integrate the two lines of work.

Supporting other dynamic data structures (not necessarily streams) suggests looking at the work of Acar et al. [1], who have studied adding user-controlled incrementalization to self-adjusting computation, which shares many implementation issues with FRP. Conversely, it will be worth investigating whether our denotational semantics can be adapted to provide a useful categorical cost semantics for self-adjusting computation[16].

References

[1] U. A. Acar, G. E. Blelloch, R. Ley-Wild, K. Tangwongsan, and D. Türkoglu. Traceable data types for self-adjusting computation. In *PLDI*, pages 483–496, 2010.

[2] A. W. Appel, P.-A. Melliès, C. D. Richards, and J. Vouillon. A very modal model of a modern, major, general type system. In M. Hofmann and M. Felleisen, editors, *POPL*, pages 109–122. ACM, 2007. ISBN 1-59593-575-4.

[3] G. Berry and L. Cosserat. The ESTEREL synchronous programming language and its mathematical semantics. In *Seminar on Concurrency*, pages 389–448. Springer, 1985.

[4] P. Caspi, D. Pilaud, N. Halbwachs, and J. Plaice. LUSTRE: A declarative language for real-time programming. In *POPL*, 1987.

[5] G. Cooper. *Integrating dataflow evaluation into a practical higher-order call-by-value language*. PhD thesis, Brown University, 2008.

[6] G. Cooper and S. Krishnamurthi. Embedding dynamic dataflow in a call-by-value language. *Programming Languages and Systems*, pages 294–308, 2006.

[7] D. Dreyer, A. Ahmed, and L. Birkedal. Logical step-indexed logical relations. In *LICS*, pages 71–80. IEEE, 2009.

[8] C. Elliott and P. Hudak. Functional reactive animation. In *ICFP*, 1997.

[9] M. Hofmann. A type system for bounded space and functional in-place update. *Nordic Journal of Computing*, 7(4), 2000.

[10] M. Hofmann. Linear types and non-size-increasing polynomial time computation. *Information and Computation*, 183(1), 2003.

[11] P. Hudak, A. Courtney, H. Nilsson, and J. Peterson. Arrows, robots and functional reactive programming. In *Advanced Functional Programming*, volume 2638 of *LNCS*. Springer, 2003.

[12] J. Hughes. Generalizing monads to arrows. *Sci. Comput. Program.*, 37(1-3), 2000.

[13] N. Krishnaswami and N. Benton. Ultrametric semantics of reactive programs. In *LICS*. IEEE, 2011.

[14] N. Krishnaswami and N. Benton. A semantic model of graphical user interfaces. In *ICFP*, 2011.

[15] U. D. Lago and M. Hofmann. Realizability models and implicit complexity. *Theor. Comput. Sci.*, 412(20):2029–2047, 2011.

[16] R. Ley-Wild, U. A. Acar, and M. Fluet. A cost semantics for self-adjusting computation. In Z. Shao and B. C. Pierce, editors, *POPL*, pages 186–199. ACM, 2009. ISBN 978-1-60558-379-2.

[17] H. Liu, E. Cheng, and P. Hudak. Causal commutative arrows and their optimization. In *ICFP*, 2009.

[18] H. Nakano. A modality for recursion. In *LICS*, pages 255–266, 2000.

[19] H. Nilsson, A. Courtney, and J. Peterson. Functional reactive programming, continued. In *ACM Haskell Workshop*, pages 51–64. ACM, 2002.

[20] P. W. O'Hearn and D. J. Pym. The logic of bunched implications. *Bulletin of Symbolic Logic*, 5(2), 1999.

[21] M. Pouzet. *Lucid Synchrone, version 3. Tutorial and reference manual*. Université Paris-Sud, LRI, 2006.

[22] N. Sculthorpe and H. Nilsson. Safe functional reactive programming through dependent types. In *ICFP*, 2009.

[23] D. Spoonhower, G. E. Blelloch, R. Harper, and P. B. Gibbons. Space profiling for parallel functional programs. In J. Hook and P. Thiemann, editors, *ICFP*, pages 253–264. ACM, 2008. ISBN 978-1-59593-919-7.

[24] T. Uustalu and V. Vene. The essence of dataflow programming. In K. Yi, editor, *APLAS*, volume 3780 of *Lecture Notes in Computer Science*, pages 2–18. Springer, 2005. ISBN 3-540-29735-9.

[25] Z. Wan, W. Taha, and P. Hudak. Real-time FRP. In *ICFP*, pages 146–156, 2001.

[26] Z. Wan, W. Taha, and P. Hudak. Event-driven FRP. In *PADL*, pages 155–172, 2002.