

Modal FRP For All

Functional Reactive Programming Without Space Leaks in Haskell

PATRICK BAHR, IT University of Copenhagen, Denmark

Functional reactive programming (FRP) provides a high-level interface for implementing reactive systems in a declarative manner. However, this high-level interface has to be carefully reigned in to ensure that programs can in fact be executed in practice. Specifically, one must ensure that FRP programs are productive, causal and can be implemented without introducing space leaks. In recent years, modal types have been demonstrated to be an effective tool to ensure these operational properties.

In this paper, we present RATTUS, a modal FRP language that extends and simplifies previous modal FRP calculi while still maintaining the operational guarantees for productivity, causality, and space leaks. The simplified type system makes RATTUS a practical programming language that can be integrated with existing functional programming languages. To demonstrate this, we have implemented a shallow embedding of RATTUS in Haskell that allows the programmer to write RATTUS code in familiar Haskell syntax and seamlessly integrate it with regular Haskell code. Thus RATTUS combines the benefits enjoyed by FRP libraries such as Yampa, namely access to a rich library ecosystem (e.g. for graphics programming), with the strong operational guarantees offered by a bespoke type system. The soundness proof based on a logical relations argument has been formalised using the Coq proof assistant.

Additional Key Words and Phrases: Functional reactive programming, Modal types, Haskell, Type systems

1 INTRODUCTION

Reactive systems perform an ongoing interaction with their environment, receiving inputs from the outside, changing their internal state and producing some output. Examples of such systems include GUIs, web applications, video games, and robots. Programming such systems with traditional general-purpose imperative languages can be very challenging: The components of the reactive system are put together via a complex and often confusing web of callbacks and shared mutable state. As a consequence, individual components cannot be easily understood in isolation, which makes building and maintaining reactive systems in this manner difficult and error-prone.

Functional reactive programming (FRP), introduced by Elliott and Hudak [1997], tries to remedy this problem by introducing time-varying values (called *behaviours* or *signals*) and *events* as a means of communication between components in a reactive system instead of shared mutable state and callbacks. Crucially, signals and events are first-class values in FRP and can be freely combined and manipulated. These high-level abstractions not only provide a rich and expressive programming model. They also make it possible for us to reason about FRP programs by simple equational methods.

Elliott and Hudak's original conception of FRP is an elegant idea that allows for direct manipulation of time-dependent data but also immediately leads to the question of what the interface for signals and events should be. A naive approach would be to model signals as streams defined by the following Haskell data type¹

```
data Str a = a ::: (Str a)
```

¹Here `:::` is a data constructor written as a binary infix operator.

Author's address: Patrick Bahr, IT University of Copenhagen, Denmark, paba@itu.dk.

2018. 2475-1421/2018/1-ART1 \$15.00
<https://doi.org/>

A stream of type $Str\ a$ thus consists of a head of type a and a tail of type $Str\ a$. The type $Str\ a$ encodes a discrete signal of type a , where each element of a stream represents the value of that signal at a particular time.

Combined with the power of higher-order functional programming we can easily manipulate and compose such signals. For example, we may apply a function to the values of a signal:

$$\begin{aligned} map &:: (a \rightarrow b) \rightarrow Str\ a \rightarrow Str\ b \\ map\ f\ (x \::: xs) &= f\ x \::: map\ f\ xs \end{aligned}$$

However, this representation is too permissive and allows the programmer to write *non-causal* programs, i.e. programs where the present output depends on future input such as the following:

$$\begin{aligned} clairvoyance &:: Str\ Int \rightarrow Str\ Int \\ clairvoyance\ (x \::: xs) &= map\ (+1)\ xs \end{aligned}$$

This function takes the input n of the *next* time step and returns $n + 1$ in the *current* time step. In practical terms, this reactive program cannot be effectively executed since we cannot compute the current value of the signal that it defines.

Much of the research in FRP has been dedicated to addressing this problem by adequately restricting the interface that the programmer can use to manipulate signals. This can be achieved by exposing only a carefully selected set of combinators to the programmer or by using a more sophisticated type system. The former approach has been very successful in practice, not least because it can be readily implemented as a library in existing languages. This library approach also immediately integrates the FRP language with a rich ecosystem of existing libraries and inherits the host language's compiler and tools. The most prominent example of this approach is Arrowised FRP [Nilsson et al. 2002], as implemented in the Yampa library for Haskell [Hudak et al. 2004], which takes signal functions as primitive rather than signals themselves. However, this library approach forfeits some of the simplicity and elegance of the original FRP model as it disallows direct manipulation of signals.

More recently, an alternative to this approach has been developed [Bahr et al. 2019, 2021; Jeffrey 2014; Jeltsch 2013; Krishnaswami 2013; Krishnaswami and Benton 2011; Krishnaswami et al. 2012] that uses a *modal* type operator \bigcirc that captures the notion of time. Following this idea, an element of type $\bigcirc a$ represents data of type a arriving in the next time step. Signals are then modelled by the type of streams defined instead as follows:

$$\mathbf{data}\ Str\ a = a \::: (\bigcirc(Str\ a))$$

That is, a stream of type $Str\ a$ is an element of type a now and a stream of type $Str\ a$ later, thus separating consecutive elements of the stream by one time step. Combining this modal type with guarded recursion [Nakano 2000] in the form of a fixed point operator of type $(\bigcirc a \rightarrow a) \rightarrow a$ gives a powerful type system for reactive programming that guarantees not only causality, but also *productivity*, i.e. the property that each element of a stream can be computed in finite time.

Causality and productivity of an FRP program means that it can be effectively implemented and executed. However, for practical purposes it is also important whether it can be implemented with given finite resources. If a reactive program requires an increasing amount of memory or computation time, it will eventually run out of resources to make progress or take too long to react to input. It will grind to a halt. Since FRP programs operate on a high level of abstraction, it can be very difficult to reason about their space and time cost. A reactive program that exhibits a gradually slower response time, i.e. computations take longer and longer as time progresses, is said to have a *time leak*. Similarly, we say that a reactive program has a *space leak*, if its memory use is gradually increasing as time progresses, e.g. if it holds on to memory while continually allocating more.

Within both lines of work – the library approach and the modal types approach – there has been an effort to devise FRP languages that avoid *implicit* space leaks, i.e. space leaks that are caused by the implementation of the FRP language rather than explicit memory allocations intended by the programmer. For example, [Ploeg and Claessen \[2015\]](#) devised an FRP library for Haskell that avoids implicit space leaks by carefully restricting the API to manipulate events and signals. Based on the modal operator \bigcirc described above, [Krishnaswami \[2013\]](#) has devised a *modal* FRP calculus that permit an aggressive garbage collection strategy that rules out implicit space leaks.

The absence of space leaks is an operational property that is notoriously difficult to reason about in higher-level languages. For a simple example, consider the following innocuously looking function *const* that takes an element of type *a* and repeats it indefinitely as a stream:

```
const :: a → Str a
const x = x ::: const x
```

In particular, this function can be instantiated at type $const :: Str\ Int \rightarrow Str\ (Str\ Int)$, which has an inherent space leak with its memory usage growing linearly with time: At each time step *n* it has to store all previously observed input values from time step 0 to *n*. On the other hand, instantiated with the type $const :: Int \rightarrow Str\ Int$, the function can be efficiently implemented. To distinguish between these two scenarios, [Krishnaswami \[2013\]](#) introduced the notion of *stable types*, i.e. types such as *Int* that are time invariant and whose values can thus be transported into the future without causing space leaks.

Contributions. In this paper, we present RATTUS, a practical modal FRP language that takes its ideas from the modal FRP calculi of [Krishnaswami \[2013\]](#) and [Bahr et al. \[2019\]](#) but with a simpler and less restrictive type system that makes it attractive to use in practice. Like the Simply RaTT calculus of [Bahr et al.](#), we use a Fitch-style type system [[Clouston 2018](#)], which extends typing contexts with *tokens* to avoid the syntactic overhead of the dual-context-style type system of [Krishnaswami \[2013\]](#). In addition, we further simplify the typing system by (1) only requiring one kind of *token* instead of two, (2) allowing tokens to be introduced without any restrictions, and (3) simplifying the guarded recursion scheme. The resulting calculus is simpler and more expressive, yet still retains the operational guarantees of the earlier calculi, namely productivity, causality and admissibility of an aggressive garbage collection strategy that prevents implicit space leaks. We have proved these properties by a logical relations argument formalised using the Coq theorem prover (see supplementary material).

To demonstrate its use as a practical programming language, we have implemented RATTUS as an embedded language in Haskell. This implementation consists of a library that implements the primitives of the language along with a plugin for the GHC Haskell compiler. The latter is necessary to check the more restrictive variable scope rules of RATTUS and to ensure the eager evaluation strategy that is central to the operational properties. Both components are bundled in a single Haskell library that allows the programmer to seamlessly write RATTUS code alongside Haskell code. We further demonstrate the usefulness of the language with a number of case studies, including an FRP library based on streams and events as well as an arrowized FRP library in the style of Yampa. We then use both FRP libraries to implement a primitive game. The two libraries implemented in RATTUS also demonstrate different approaches to FRP libraries: discrete time (streams) vs. continuous time (Yampa); and first-class signals (streams) vs. signal functions (Yampa). The RATTUS Haskell library and all examples are included in the supplementary material.

Overview of Paper. [Section 2](#) gives an overview of the RATTUS language introducing the main concepts and their intuitions through examples. [Section 3](#) presents a case study of a simple FRP library based on streams and events, as well as an arrowized FRP library. [Section 4](#) presents the

underlying core calculus of RATTUS including its type system, its operational semantics, and our main metatheoretical results: productivity, causality and absence of implicit space leaks. [Section 5](#) gives an overview of the proof of our metatheoretical results. [Section 6](#) describes how RATTUS has been implemented as an embedded language in Haskell. [Section 7](#) reviews related work and [Section 8](#) discusses future work.

2 RATTUS NORVEGICUS DOMESTICA

To illustrate RATTUS we will use example programs written in the embedding of the language in Haskell. The type of streams is at the centre of these example programs:

```
data Str a = a ::: (○(Str a))
```

The simplest stream one can define just repeats the same value indefinitely. Such a stream is constructed by the *const* function below, which takes an integer and produces a constant stream that repeats that integer at every step:

```
const :: Int → Str Int
```

```
const x = x ::: delay (const x)
```

Because the tail of a stream of integers must be of type $\bigcirc(\text{Str Int})$, we have to use *delay*, which is the introduction form for the type modality \bigcirc . Intuitively speaking, *delay* moves a computation one time step into the future. We could think of *delay* having type $a \rightarrow \bigcirc a$, but this type is too permissive as it can cause space leaks. It would allow us to move arbitrary computations – and the data they depend on – into the future. Instead, the typing rule for *delay* is formulated as follows:

$$\frac{\Gamma, \checkmark \vdash t :: A}{\Gamma \vdash \text{delay } t :: \bigcirc A}$$

This is a characteristic example of a Fitch-style typing rule: It introduces the *token* \checkmark (pronounced “tick”) in the typing context Γ . A typing context consists of type assignments of the form $x :: A$, but it can also contain several occurrences of \checkmark . We can think of \checkmark as denoting the passage of one time step, i.e. all variables to the left of \checkmark are one time step older than those to the right. In the above typing rule, the term t does not have access to these “old” variables in Γ . There is, however, an exception: If a variable in the typing context is of a type that is time-independent, we still allow t to access them – even if the variable is one time step old. We call these time-independent types *stable* types, and in particular all base types such as *Int* and *Bool* are stable. We will discuss stable types in more detail in [Section 2.1](#).

Formally, the variable introduction rule of RATTUS reads as follows:

$$\frac{\Gamma' \text{ tick-free or } A \text{ stable}}{\Gamma, x :: A, \Gamma' \vdash x :: A}$$

That is, if x is not of a stable type and appears to the left of a \checkmark , then it is no longer in scope.

Turning back to our definition of the *const* function, we can see that the recursive call *const x* must be of type *Str Int* in the context Γ, \checkmark , where Γ contains $x :: \text{Int}$. So x remains in scope because it is of type *Int*, which is a stable type. This would not be the case if we were to generalise *const* to arbitrary types:

```
leakyConst :: a → Str a
```

```
leakyConst x = x ::: delay (leakyConst x) -- the rightmost occurrence of x is out of scope
```

In this example, x is of type a and therefore goes out of scope under *delay*: Since a is not necessarily stable, $x :: a$ is blocked by the \checkmark introduced by *delay*.

The definition of *const* also illustrates the *guarded* recursion principle used in RATTUS. For a recursive definition to be well-typed, all recursive calls have to occur in the presence of a \checkmark – in other words, recursive calls have to be guarded by delay. This restriction ensures that all recursive functions are productive, which means that each element of a stream can be computed in finite time. If we did not have this restriction, we could write the following obviously unproductive function:

```
loop :: Str Int
loop = loop -- unguarded recursive call to loop is not allowed
```

The recursive call to *loop* does not occur under a delay, and is thus rejected by the type checker.

The function *inc* below takes a stream of integers as input and increments each integer by 1:

```
inc :: Str Int → Str Int
inc (x ::: xs) = (x + 1) ::: delay (inc (adv xs))
```

Here we have to use *adv*, the elimination form for \bigcirc , to convert the tail of the input stream from type $\bigcirc(\text{Str Int})$ into type *Str Int*. Again we could think of *adv* having type $\bigcirc a \rightarrow a$, but this general type would allow us to write non-causal functions such as the following:

```
tomorrow :: Str Int → Str Int
tomorrow (x ::: xs) = adv xs -- adv is not allowed here
```

This function skips one time step so that the output at time n depends on the input at time $n + 1$.

To ensure causality, *adv* is restricted to contexts with a \checkmark :

$$\frac{\Gamma \vdash t :: \bigcirc A \quad \Gamma' \text{ tick-free}}{\Gamma, \checkmark, \Gamma' \vdash \text{adv } t :: A}$$

Not only does *adv* require a \checkmark , it also causes all bound variables to the right of \checkmark to go out of scope. Intuitively speaking delay looks ahead one time step and *adv* then allows us to go back to the present. Variable bindings made in the future are therefore not accessible once we returned to the present.

Note that *adv* causes the variables to the right of \checkmark to go out of scope *forever*, whereas it brings variables back into scope that were previously blocked by the \checkmark . That is, variables that go out of scope due to delay can be brought back into scope by *adv*

2.1 Stable types

We haven't yet made precise what stable types are. To a first approximation, types are stable if they do not contain \bigcirc or function types. The intuition here is that \bigcirc expresses a temporal aspect and thus types containing \bigcirc are not time-invariant. Moreover, functions can implicitly have temporal values in their closure and are therefore also excluded from stable types.

However, that means we cannot not implement the *map* function that takes a function $f :: a \rightarrow b$ and applies it to each element of a stream of type *Str a*, because it would require us to apply the function *f* at any time in the future. We cannot do this because $a \rightarrow b$ is not a stable type (even if *a* and *b* were stable) and therefore *f* cannot be transported into the future. However, RATTUS has the type modality \square , pronounced “box”, that turns any type *A* into a stable type $\square A$. Using the \square modality we can implement *map* as follows:

```
map :: □(a → b) → Str a → Str b
map f (x ::: xs) = unbox f x ::: delay (map f (adv xs))
```

Instead of a function of type $a \rightarrow b$, *map* takes a *boxed* function *f* of type $\square(a \rightarrow b)$ as argument. That means, *f* is still in scope under the delay because it is of a stable type. To use *f*, it has to

be unboxed using `unbox`, which is the elimination form for the \Box modality and has simply type $\Box a \rightarrow a$, this time without any side conditions.

On the other hand, the corresponding introduction form for \Box has to make sure that boxed values do not refer to non-stable variables:

$$\frac{\Gamma^\Box \vdash t :: A}{\Gamma \vdash \text{box } t :: \Box A}$$

Here, Γ^\Box denotes the typing context that is obtained from Γ by removing all variables of non-stable types and all \checkmark tokens:

$$\cdot^\Box = \cdot \quad (\Gamma, x :: A)^\Box = \begin{cases} \Gamma^\Box, x :: A & \text{if } A \text{ stable} \\ \Gamma^\Box & \text{otherwise} \end{cases} \quad (\Gamma, \checkmark)^\Box = \Gamma^\Box$$

Thus, for a well-typed term `box t`, we know that t only accesses variables of stable type.

For example, we can implement the `inc` function using `map` as follows:

```
inc :: Str Int → Str Int
inc = map (box (+1))
```

Using the \Box modality we can also generalise the constant stream function to arbitrary boxed types:

```
constBox :: Box a → Str a
constBox a = unbox a ::: delay (constBox a)
```

Alternatively, we can make use of the `Stable` type class, to constrain type variables to stable types:

```
const :: Stable a ⇒ a → Str a
const x = x ::: delay (const x)
```

So far, we have only looked at recursive definitions at the top level. Recursive definitions can also be nested, but we have to be careful how such nested recursion interacts with the typing environment. Below is an alternative definition of `map` that takes the boxed function f as an argument and then calls the `run` function that recurses over the stream:

```
map :: Box (a → b) → Str a → Str b
map f = run
  where run :: Str a → Str b
        run (x ::: xs) = unbox f x ::: delay (run (adv xs))
```

Here `run` is type checked in a typing environment Γ that contains $f :: \Box (a \rightarrow b)$. Since `run` is defined by guarded recursion, we require that its definition must type check in the typing context Γ^\Box . Because f is of a stable type, it remains in Γ^\Box and is thus in scope in the definition of `run`. That is, guarded recursive definitions interact with the typing environment in the same way as `box`, which ensures that such recursive definitions are stable and can thus safely be executed at any time in the future. As a consequence, the type checker will prevent us from writing the following leaky version of `map`.

```
leakyMap :: (a → b) → Str a → Str b
leakyMap f = run
  where run :: Str a → Str b
        run (x ::: xs) = f x ::: delay (run (adv xs)) -- f is no longer in scope here
```

The type of f is not stable, and thus it is not in scope in the definition of `run`.

Note that top-level defined identifiers such as `map` and `const` are in scope in any context after they are defined regardless of whether there is a \checkmark or whether they are of a stable type. One can

think of top-level definitions being implicitly boxed when they are defined and implicitly unboxed when they are used later on.

2.2 Operational Semantics

As we have seen in the examples above, the purpose of the type modalities \bigcirc and \square is to ensure that RATTUS programs are causal and productive. Furthermore, the typing rules also ensure that RATTUS has no implicit space leaks. In simple terms, the latter means that temporal values, i.e. values of type $\bigcirc A$, are safe to be garbage collected after two time steps. In particular, input from a stream can be safely garbage collected one time step after it has arrived. This memory property is made precise later in [Section 4](#).

This is the same memory property [Krishnaswami \[2013\]](#) and [Bahr et al. \[2019, 2021\]](#) established for their modal FRP calculi. To achieve this, these calculi restrict guarded recursive definitions to only “look ahead” at most one time step. In the Fitch-style calculi of [Bahr et al. \[2019, 2021\]](#) this can be seen in the restriction to allow at most one \checkmark in the typing context. For the same reason these two calculi also disallow function definitions in the context of a \checkmark . As a consequence, terms like $\text{delay}(\text{delay } 0)$ and $\text{delay}(\lambda x.x)$ do not type check in the calculi of [Bahr et al. \[2019, 2021\]](#).

RATTUS lifts these restrictions on ticks and instead refines the operational semantics of the language. At first glance one might think that allowing multiple ticks can be accommodated by extending the time one has to keep temporal values in memory accordingly, i.e. we may safely garbage collect temporal values after $n + 1$ time steps, if we allow at most n ticks. However, this turns out to be neither enough nor necessary. On the one hand, even allowing just two ticks would require us to keep temporal values in memory indefinitely, i.e. it permits implicit space leaks. On the other hand, if we change the evaluation strategy, we can still garbage collect all temporal values after two time steps, no matter how many ticks were involved in type checking the program.

Like the abovementioned earlier calculi of [Krishnaswami](#) and [Bahr et al.](#), RATTUS uses an eager evaluation strategy except for `delay` and `box`. That is, arguments are evaluated to values before they are passed on to functions other than `delay` and `box`. Where RATTUS differs is the evaluation of `adv`. Simply put, RATTUS follows the *temporal* semantics of `adv`: Recall that `delay t` delays the computation t by one time step and that `adv` reverses such a delay. The operational semantics of RATTUS reflects this intuition by first evaluating every term t that occurs as `delay (... adv t ...)` before evaluating `delay`. To enforce this evaluation strategy, RATTUS programs are transformed before execution employing two rewrite rules (see [Section 4.2.1](#)).

The extension in expressive power afforded by RATTUS’ more aggressive eager evaluation strategy has immediate practical benefits. Most importantly, there are no restrictions on where one can define functions. Secondly, we can write recursive functions that look several steps into the future:

```
stutter :: Int → Str Int
stutter n = n :: delay (n :: delay (stutter (n + 1)))
```

Applying `stutter` to 0 would construct a stream of numbers 0, 0, 1, 1, 2, 2, ... In order to implement `stutter` in the more restrictive language of [Krishnaswami \[2013\]](#) and [Bahr et al. \[2019, 2021\]](#) we would need to decompose it into two mutually recursive functions [[Bahr et al. 2021](#)].

The operational semantics is made more precise in [Section 4](#). In the Haskell embedding of the language, this evaluation strategy is enforced by using strict data structures and strict evaluation. The latter is achieved by a compiler plug-in that transforms all RATTUS functions so that arguments are always evaluated to weak head normal form (cf. [Section 6](#)).

3 REACTIVE PROGRAMMING IN RATTUS

3.1 Programming with streams and events

In this section we showcase how RATTUS can be used for reactive programming. To this end we use a small library of combinators for programming with streams and events defined in [Figure 1](#).

The *map* function should be familiar by now. The *zip* function combines two streams similar to Haskell's *zip* function on lists. Note however that instead of the normal pair type we use a strict pair type:

```
data a ⊗ b = !a ⊗ !b
```

It is like the normal pair type (a, b) , but when constructing a strict pair $s \otimes t$, the two components s and t are evaluated to weak head normal form.

The *scan* function is similar to Haskell's *scanl* function on lists: given a stream of values v_0, v_1, v_2, \dots , the expression *scan l (box f) v* computes the stream

$$f \ v \ v_0, \quad f \ (f \ v \ v_0) \ v_1, \quad f \ (f \ (f \ v \ v_0) \ v_1) \ v_2, \quad \dots$$

If one would want a variant of *scan* that is closer to Haskell's *scanl*, i.e. the result starts with the value v instead of $f \ v \ v_0$, one can simply replace the first occurrence of acc' in the definition of *scan* with acc . Note that the type b has to be stable in the definition of *scan* so that $acc' :: b$ is still in scope under delay.

A central component of functional reactive programming is that it must provide a way to react to events. In particular, it must support the ability to *switch* behaviour as reaction to the occurrence of an event. There are different ways to represent events. The simplest is to define events of type a as streams of type *Maybe a*. However, we will use the strict variant of the *Maybe* type:

```
data Maybe' a = Just' ! a | Nothing'
```

We can then devise a *switch* combinator that reacts to events. Given an initial stream xs and an event e that may produce a stream, *switch xs e* initially behaves as xs but changes to the new stream provided by the occurrence of an event. Note that the behaviour changes *every time* an event occurs, not only the first time.

In the definition of *switch* we use the applicative operator \otimes defined as follows

```
(⊗) :: ○(a → b) → ○a → ○b
```

```
f ⊗ x = delay ((adv f) (adv x))
```

Instead of using \otimes , we could have also written `delay (switch (adv xs) (adv fas))` instead.

Finally, *switchTrans* is a variant of *switch* that switches to a new stream function rather than just a stream. It is implemented using the variant *switchTrans'* where the initial stream function is rather just a stream.

3.2 A simple reactive program

To put our bare-bones FRP library to use, let's implement a simple single player variant of the classic game Pong: The player has to move a paddle at the bottom of the screen to bounce a ball and prevent it from falling.² The core behaviour is described by the following stream function:

```
pong :: Str Input → Str (Pos ⊗ Float)
```

```
pong inp = zip ball pad where
```

```
  pad :: Str Float
```

```
  pad = padPos inp
```

²So it is rather like Breakout, but without the bricks.

```

map :: □(a → b) → Str a → Str b
map f (x ::: xs) = unbox f x ::: delay (map f (adv xs))
zip :: Str a → Str b → Str (a ⊗ b)
zip (a ::: as) (b ::: bs) = (a ⊗ b) ::: delay (zip (adv as) (adv bs))
scan :: Stable b ⇒ □(b → a → b) → b → Str a → Str b
scan f acc (a ::: as) = acc' ::: delay (scan f acc' (adv as))
  where acc' = unbox f acc a
type Event a = Str (Maybe' a)
switch :: Str a → Event (Str a) → Str a
switch (x ::: xs) (Nothing'      ::: fas) = x ::: (delay switch ⊗ xs ⊗ fas)
switch _      (Just' (a ::: as) ::: fas) = a ::: (delay switch ⊗ as ⊗ fas)
switchTrans :: (Str a → Str b) → Event (Str a → Str b) → (Str a → Str b)
switchTrans f es as = switchTrans' (f as) es as
switchTrans' :: Str b → Event (Str a → Str b) → Str a → Str b
switchTrans' (b ::: bs) (Nothing' ::: fs) as = b ::: (delay switchTrans' ⊗ bs ⊗ fs ⊗ tail as)
switchTrans' _      (Just' f      ::: fs) as = b' ::: (delay switchTrans' ⊗ bs' ⊗ fs ⊗ tail as)
  where (b' ::: bs') = f as

```

Fig. 1. Small library for streams and events.

```

ball :: Str Pos
ball = ballPos (zip pad inp)

```

It receives a stream of inputs (button presses and how much time has passed since the last input) and produces a stream of pairs consisting of the 2D position of the ball and the x coordinate of the paddle. Its implementation uses two helper functions to compute these two components. The position of the paddle only depends on the input whereas the position of the ball also depends on the position of the paddle (since it may bounce off it):

```

padPos :: Str (Input) → Str Float
padPos = map (box fst') ∘ scan (box padStep) (0 ⊗ 0)
padStep :: (Float ⊗ Float) → Input → (Float ⊗ Float)
padStep (pos ⊗ vel) inp = ...
ballPos :: Str (Float ⊗ Input) → Str Pos
ballPos = map (box fst') ∘ scan (box ballStep) ((0 ⊗ 0) ⊗ (20 ⊗ 50))
ballStep :: (Pos ⊗ Vel) → (Float ⊗ Input) → (Pos ⊗ Vel)
ballStep (pos ⊗ vel) (pad ⊗ inp) = ...

```

Both auxiliary functions follow the same structure. They use a *scan* to keep track of some internal state, e.g. the position and velocity of the ball, while consuming the input stream. The internal state is then projected away using *map*. Here *fst'* is the first projection for the strict pair type. We can see that the ball starts at the centre of the screen (at coordinates (0, 0)) and moves towards the upper right corner.

Let's change the implementation of *pong* so that it allows the player to reset the game, e.g. after ball has fallen off the screen:

```
pong' :: Str Input → Str (Pos ⊗ Float)
pong' inp = zip ball pad where
  pad = padPos inp
  ball = switchTrans ballPos           -- starting ball behaviour
        (map (box ballTrig) inp)      -- trigger restart on pressing reset button
        (zip pad inp)                 -- input to the switch

ballTrig :: Input → Maybe' (Str (Float ⊗ Input) → Str Pos)
ballTrig inp = if reset inp then Just' ballPos else Nothing'
```

To achieve this behaviour we use the *switchTrans* combinator, which we initialise with the original behaviour of the ball. The event that will trigger the switch is constructed by mapping *ballTrig* over the input stream, which will create an event of type *Event (Str (Float ⊗ Input) → Str Pos)*, which will be triggered every time the player hits the reset button.

3.3 Arrowized FRP

The benefit of a modal FRP language is that we can directly interact with signals and events in a way that guarantees causality. A popular alternative to ensure causality is arrowized FRP [Nilsson et al. 2002], which takes *signal functions* as primitive and uses Haskell's arrow notation [Paterson 2001] to construct them. By implementing an arrowized FRP library in RATTUS, we can not only guarantee causality but also productivity and the absence of implicit space leaks. Furthermore, this demonstrates that RATTUS can also be used to implement a continuous-time FRP library, in contrast to the FRP library from Section 3.1, which is discrete-time.

The basis for such an arrowized RATTUS library is fairly easily obtained, simply by instantiating Haskell's *Arrow* type class for a suitable signal function type *SF*. We can almost copy the definition of *SF* from Nilsson et al. [2002], but we have to insert the \bigcirc modality to make it a guarded recursive type and use a strict pair type:

```
data SF a b = SF (Float → a → (⊙(SF a b) ⊗ b))
```

We can then use Haskell's convenient arrow notation for constructing signal functions. For example, assuming we have signal functions *ballPos* :: *SF (Float ⊗ Input) Pos* and *padPos* :: *SF Input Float* describing the positions of the ball and the paddle from our game in Section 3.2, we can combine these as follows:

```
pong :: SF Input (Pos ⊗ Float)
pong = proc inp → do pad ← padPos <- inp
                    ball ← ballPos <-(pad ⊗ inp)
                    returnA <-(ball ⊗ pad)
```

The *Arrow* type class only provides a basic interface for constructing *static* signal functions. To permit dynamic behaviour we need to provide additional combinators, e.g. for switching signals and for recursive definitions. The *rSwitch* combinator corresponds to the *switchTrans* combinator from Figure 1:

```
rSwitch :: SF a b → SF (a ⊗ Maybe' (SF a b)) b
```

This combinator allows us to implement our game so that it resets to its start position if we hit the reset button:

```
pong' :: SF Input (Pos ⊗ Float)
pong' = proc inp → do pad ← padPos <- inp
      let event = if reset inp then Just' ballPos else Nothing'
          ball ← rSwitch ballPos <- ((pad ⊗ inp) ⊗ event)
          returnA <- (ball ⊗ pad)
```

Arrows can be endowed with a very general recursion principle by instantiating the *ArrowLoop* type class. However, this would require us to implement a function of type $SF (b, d) (c, d) \rightarrow SF b c$. Such a recursion principle is too general for guarded recursion. Instead, we implement a more restricted recursion principle that corresponds to guarded recursion:

```
loopPre :: c → SF (a ⊗ c) (b ⊗ ⊙c) → SF a b
```

Intuitively speaking, this combinator constructs a signal function from a to b with the help of an internal state of type c . The first argument initialises the state, and the second argument is a signal function that turns input of type a into output of type b while also updating the internal state. Apart from the addition of the \odot modality and strict pair types, this definition has the same type as Yampa's *loopPre*.

Using the *loopPre* combinator we can implement the signal function of the ball:

```
ballPos :: SF (Float ⊗ Input) Pos
ballPos = loopPre (20 ⊗ 50) run where
  run :: SF ((Float ⊗ Input) ⊗ Vel) (Pos ⊗ ⊙Vel)
  run = proc ((pad ⊗ inp) ⊗ v) → do p ← integral (0 ⊗ 0) <- v
      returnA <- (p ⊗ delay (calculateNewVelocity pad p v))
```

Here we also use the *integral* combinator that computes the integral of a signal using a simple approximation that sums up rectangles under the curve:

```
integral :: (Stable a, VectorSpace a s) ⇒ a → SF a a
```

The signal function for the paddle can be implemented in a similar fashion. The complete code of the case studies presented in this section can be found in the supplementary material.

4 CORE CALCULUS

In this section we present the core calculus of RATTUS, which we call $\lambda_{\mathcal{R}}$. The purpose of $\lambda_{\mathcal{R}}$ is to formally present the language's Fitch-style typing rules, its operational semantics, and to formally prove the central operational properties, i.e. productivity, causality, and absence of implicit space leaks. To this end, the calculus is stripped down to its essence: simply typed lambda calculus extended with guarded recursive types $\text{Fix } \alpha.A$ and the two type modalities \square and \odot . Since general inductive types and polymorphic types are orthogonal to the issue of operational properties in reactive programming, we have omitted these for the sake of clarity.

4.1 Type System

Figure 2 defines the syntax of $\lambda_{\mathcal{R}}$. Besides guarded recursive types and the two type modalities, we include standard sum and product types along with unit and integer types. The type of streams of type A is represented as $\text{Fix } \alpha.A \times \alpha$. Note the absence of \odot in this type. When unfolding guarded recursive types such as $\text{Fix } \alpha.A \times \alpha$, the \odot modality is inserted implicitly: $\text{Fix } \alpha.A \times \alpha \cong A \times \odot(\text{Fix } \alpha.A \times \alpha)$. This ensures that guarded recursive types are by construction always guarded by the \odot modality.

Types	$A, B ::= \alpha \mid 1 \mid \text{Int} \mid A \times B \mid A + B \mid A \rightarrow B \mid \square A \mid \bigcirc A \mid \text{Fix } \alpha.A$
Stable Types	$S, S' ::= 1 \mid \text{Int} \mid \square A \mid S \times S' \mid S + S'$
Values	$v, w ::= \langle \rangle \mid \bar{n} \mid \lambda x.t \mid \langle v, w \rangle \mid \text{in}_i v \mid \text{box } t \mid \text{into } v \mid \text{fix } x.t \mid l$
Terms	$s, t ::= \langle \rangle \mid \bar{n} \mid \lambda x.t \mid \langle s, t \rangle \mid \text{in}_i t \mid \text{box } t \mid \text{into } t \mid \text{fix } x.t \mid l \mid x \mid t_1 t_2 \mid t_1 + t_2$ $\mid \text{adv } t \mid \text{delay } t \mid \text{case } t \text{ of } \text{in}_1 x.t_1; \text{in}_2 x.t_2 \mid \text{let } x = s \text{ in } t \mid \text{unbox } t \mid \text{out } t$

Fig. 2. Syntax of (stable) types, terms, and values. In typing rules, only closed types (no free α) are considered.

Typing contexts, defined in Figure 3, consist of variable typings $x : A$ and \surd tokens. If a typing context contains no \surd , we call it *tick-free*. The complete set of typing rules for λ_{\surd} is given in Figure 4. The typing rules for RATTUS presented in Section 2 appear in the same form also here, except for replacing Haskell's $::$ operator with the more standard notation. The remaining typing rules are entirely standard, except for the typing rule for the guarded fixed point combinator fix .

The typing rule for fix follows Nakano's fixed point combinator and ensures that the calculus is productive. In addition, following Krishnaswami [2013], the rule enforces the body t of the fixed point to be stable by strengthening the typing context to Γ^\square . Moreover, we follow Bahr et al. [2021] and assume x to be of type $\square(\bigcirc A)$ instead of $\bigcirc A$. As a consequence, recursive calls may occur at any time in the future, i.e. not necessarily in the very next time step. In conjunction with the more general typing rule for delay , this allows us to write recursive function definitions that, like *stutter* in section 2.2, look several steps into the future.

To see how the recursion syntax of RATTUS translates into the fixed point combinator of λ_{\surd} , let us reconsider the *const* function:

```
const :: Int → Str Int
const x = x :: delay (const x)
```

Such a recursive definition is simply translated into a fixed point $\text{fix } r.t$ where the recursive occurrence of *const* is replaced by $\text{adv } (\text{unbox } r)$.

$$\text{const} = \text{fix } r.\lambda x.x \text{ :: delay}(\text{adv } (\text{unbox } r) x)$$

where the stream cons operator $s \text{ :: } t$ is shorthand for $\text{into } \langle s, t \rangle$. The variable r is of type $\square(\bigcirc(\text{Int} \rightarrow \text{Str Int}))$ and applying unbox followed by adv turns it into type $\text{Int} \rightarrow \text{Str Int}$. Moreover, the restriction that recursive calls must occur in a context with \surd makes sure that this transformation from recursion notation to fixed point combinator results in a well-typed term.

The typing rule for $\text{fix } x.t$ also explains the treatment of recursive definitions that are nested inside a top-level definition. The typing context Γ is turned into Γ^\square when type checking the body t of the fixed point.

For example, reconsider the following ill-typed definition of *leakyMap*:

```
leakyMap :: (a → b) → Str a → Str b
leakyMap f = run
  where run :: Str a → Str b
        run (x :: xs) = f x :: delay (leakyMap (adv xs))
```

Translated into λ_{\surd} , it looks like this:

$$\text{leakyMap} = \lambda f.\text{fix } r.\lambda s.\text{let } x = \text{head } s \text{ in let } xs = \text{tail } s \text{ in } f x \text{ :: delay}((\text{adv } (\text{unbox } r)) (\text{adv } xs))$$

Here the pattern matching syntax is translated into projection functions head and tail that decompose a stream into its head and tail, respectively. More importantly, the variable f bound by the

$$\frac{}{\emptyset \vdash_{\mathcal{W}}} \quad \frac{\Gamma \vdash_{\mathcal{W}}}{\Gamma, x : A \vdash_{\mathcal{W}}} \quad \frac{\Gamma \vdash_{\mathcal{W}}}{\Gamma, \checkmark \vdash_{\mathcal{W}}}$$

Fig. 3. Well-formed contexts

$$\frac{\Gamma, x : A, \Gamma' \vdash_{\mathcal{W}} \quad \Gamma' \text{ tick-free or } A \text{ stable}}{\Gamma, x : A, \Gamma' \vdash_{\mathcal{W}} x : A} \quad \frac{\Gamma \vdash_{\mathcal{W}}}{\Gamma \vdash_{\mathcal{W}} \langle \rangle : 1} \quad \frac{n \in \mathbb{Z}}{\Gamma \vdash_{\mathcal{W}} \bar{n} : \text{Int}}$$

$$\frac{\Gamma \vdash_{\mathcal{W}} s : \text{Int} \quad \Gamma \vdash_{\mathcal{W}} t : \text{Int}}{\Gamma \vdash_{\mathcal{W}} s + t : \text{Int}} \quad \frac{\Gamma, x : A \vdash_{\mathcal{W}} t : B}{\Gamma \vdash_{\mathcal{W}} \lambda x. t : A \rightarrow B} \quad \frac{\Gamma \vdash_{\mathcal{W}} s : A \quad \Gamma, x : A \vdash_{\mathcal{W}} t : B}{\Gamma \vdash_{\mathcal{W}} \text{let } x = s \text{ in } t : B}$$

$$\frac{\Gamma \vdash_{\mathcal{W}} t : A \rightarrow B \quad \Gamma \vdash_{\mathcal{W}} t' : A}{\Gamma \vdash_{\mathcal{W}} t t' : B} \quad \frac{\Gamma \vdash_{\mathcal{W}} t : A \quad \Gamma \vdash_{\mathcal{W}} t' : B}{\Gamma \vdash_{\mathcal{W}} \langle t, t' \rangle : A \times B}$$

$$\frac{\Gamma \vdash_{\mathcal{W}} t : A_1 \times A_2 \quad i \in \{1, 2\}}{\Gamma \vdash_{\mathcal{W}} \pi_i t : A_i} \quad \frac{\Gamma \vdash_{\mathcal{W}} t : A_i \quad i \in \{1, 2\}}{\Gamma \vdash_{\mathcal{W}} \text{in}_i t : A_1 + A_2}$$

$$\frac{\Gamma, x : A_i \vdash_{\mathcal{W}} t_i : B \quad \Gamma \vdash_{\mathcal{W}} t : A_1 + A_2 \quad i \in \{1, 2\}}{\Gamma \vdash_{\mathcal{W}} \text{case } t \text{ of } \text{in}_1 x. t_1; \text{in}_2 x. t_2 : B} \quad \frac{\Gamma, \checkmark \vdash_{\mathcal{W}} t : A}{\Gamma \vdash_{\mathcal{W}} \text{delay } t : \bigcirc A}$$

$$\frac{\Gamma \vdash_{\mathcal{W}} t : \bigcirc A \quad \Gamma' \text{ tick-free}}{\Gamma, \checkmark, \Gamma' \vdash_{\mathcal{W}} \text{adv } t : A} \quad \frac{\Gamma \vdash_{\mathcal{W}} t : \square A}{\Gamma \vdash_{\mathcal{W}} \text{unbox } t : A} \quad \frac{\Gamma^\square \vdash_{\mathcal{W}} t : A}{\Gamma \vdash_{\mathcal{W}} \text{box } t : \square A}$$

$$\frac{\Gamma \vdash_{\mathcal{W}} t : A[\bigcirc(\text{Fix } \alpha. A)/\alpha]}{\Gamma \vdash_{\mathcal{W}} \text{into } t : \text{Fix } \alpha. A} \quad \frac{\Gamma \vdash_{\mathcal{W}} t : \text{Fix } \alpha. A}{\Gamma \vdash_{\mathcal{W}} \text{out } t : A[\bigcirc(\text{Fix } \alpha. A)/\alpha]} \quad \frac{\Gamma^\square, x : \square(\bigcirc A) \vdash_{\mathcal{W}} t : A}{\Gamma \vdash_{\mathcal{W}} \text{fix } x. t : A}$$

Fig. 4. Typing rules.

outer lambda abstraction is of a function type and thus not stable. Therefore, it is not in scope in the body of the fixed point.

4.2 Operational Semantics

The operational semantics is given in two steps: To execute a program of $\lambda_{\mathcal{W}}$ it is first translated into a more restrictive variant of $\lambda_{\mathcal{W}}$, which we call λ_{\checkmark} . The resulting λ_{\checkmark} program is then executed using an abstract machine that ensures the absence of implicit space leaks by construction.

4.2.1 Translation to λ_{\checkmark} . The λ_{\checkmark} calculus restricts typing contexts to contain at most one \checkmark and restricts the typing rules for delay and lambda abstractions as follows:

$$\frac{|\Gamma|, x : A \vdash_{\checkmark} t : B}{\Gamma \vdash_{\checkmark} \lambda x. t : A \rightarrow B} \quad \frac{|\Gamma|, \checkmark \vdash_{\checkmark} t : A}{\Gamma \vdash_{\checkmark} \text{delay } t : \bigcirc A} \quad |\Gamma| = \Gamma \quad \text{if } \Gamma \text{ is tick-free}$$

$$|\Gamma, \checkmark, \Gamma'| = \Gamma^\square, \Gamma'$$

The construction $|\Gamma|$ turns Γ into a tick-free context, which ensures that we have at most one \checkmark – even for nested occurrences of delay – and that lambda abstractions are not in the scope of a \checkmark . All other typing rules are the same as for $\lambda_{\mathcal{W}}$.

Any closed $\lambda_{\mathcal{L}}$ term can be translated into a corresponding $\lambda_{\mathcal{V}}$ term by exhaustively applying the following rewrite rules:

$$\begin{aligned} \text{delay}(C[\text{adv } t]) &\longrightarrow \text{let } x = t \text{ in } \text{delay}(C[\text{adv } x]) \quad \text{if } t \text{ is not a variable} \\ \lambda x.C[\text{adv } t] &\longrightarrow \text{let } y = \text{adv } t \text{ in } \lambda x.(C[y]) \end{aligned}$$

where C is a term with a single hole that does not occur in the scope of delay , adv , box , fix , or λ abstraction.

In a well-typed $\lambda_{\mathcal{L}}$ term, each subterm $\text{adv } t$ must occur in the scope of a corresponding delay . The above rewrite rules make sure that the subterm t is evaluated before the delay . This corresponds to the intuition that delay moves ahead in time and adv moves back in time – thus the two cancel out one another.

One can show that the rewrite rules are strongly normalising and type-preserving (in $\lambda_{\mathcal{L}}$). More importantly, exhaustively applying the rewrite rules to a closed term of $\lambda_{\mathcal{L}}$ transforms it into a closed term of $\lambda_{\mathcal{V}}$ (the full proof is given in the appendix):

THEOREM 4.1. *For each $\vdash_{\mathcal{L}} t : A$, we can effectively construct a term t' with $t \longrightarrow^* t'$ and $\vdash_{\mathcal{V}} t' : A$.*

4.2.2 Abstract machine for $\lambda_{\mathcal{V}}$. To prove the absence of implicit space leaks, we devise an abstract machine that after each time step deletes all data from the previous time step. This characteristic makes the operational semantics *by construction* free of implicit space leaks. This approach, pioneered by Krishnaswami [2013], allows us to reduce the proof of no implicit space leaks to a proof of type soundness.

At the centre of this approach is the idea to execute programs in a machine that has access to a store consisting of up to two separate heaps: A ‘now’ heap from which we can retrieve delayed computations, and a ‘later’ heap where we must store computations that should be performed in the next time step. Once the machine advances to the next time step, it will delete the ‘now’ heap and the ‘later’ heap will become the new ‘now’ heap.

The machine consists of two components: the *evaluation semantics*, presented in Figure 5, which describes the operational behaviour of $\lambda_{\mathcal{V}}$ within a single time step; and the *step semantics*, presented in Figure 6, which describes the behaviour of a program over time, e.g. how it consumes and constructs streams.

The evaluation semantics is given as a big-step operational semantics, where we write $\langle t; \sigma \rangle \Downarrow \langle v; \sigma' \rangle$ to indicate that starting with the store σ , the term t evaluates to the value v and the new store σ' . A store σ can be of one of two forms: either it consists of a single heap η_L , i.e. $\sigma = \eta_L$, or it consists of two heaps η_N and η_L , written $\sigma = \eta_N \checkmark \eta_L$. The ‘later’ heap η_L contains delayed computations that may be retrieved and executed in the next time step, whereas the ‘now’ heap η_N contains delayed computations from the previous time step that can be retrieved and executed now. We can only write to η_L and only read from η_N . However, when one time step passes, the ‘now’ heap η_N is deleted and the ‘later’ heap η_L becomes the new ‘now’ heap. This shifting of time is part of the step semantics in Figure 6, which we turn to shortly.

Heaps are simply finite mappings from *heap locations* to terms. Given a store σ of the form η_L or $\eta_N \checkmark \eta_L$, we write $\text{alloc}(\sigma)$ for a heap location l that is not in the domain of η_L . Given such a fresh heap location l and a term t , we write $\sigma, l \mapsto t$ to denote the store η'_L or $\eta_N \checkmark \eta'_L$, respectively, where $\eta'_L = \eta_L, l \mapsto t$, i.e. η'_L is obtained from η_L by extending it with a new mapping $l \mapsto t$.

Applying delay to a term t stores t on the later heap and returns its location on the heap. Conversely, if we apply adv to such a delayed computation, we retrieve the term from the now heap and evaluate it.

In principle, the abstract machine could also be used for $\lambda_{\mathcal{L}}$ terms directly, without transforming them to $\lambda_{\mathcal{V}}$ first. However, applied to a suitable $\lambda_{\mathcal{L}}$ term, the machine will try to dereference a

$$\begin{array}{c}
 \frac{}{\langle v; \sigma \rangle \Downarrow \langle v; \sigma \rangle} \qquad \frac{\langle t; \sigma \rangle \Downarrow \langle \bar{m}; \sigma' \rangle \quad \langle t'; \sigma' \rangle \Downarrow \langle \bar{n}; \sigma'' \rangle}{\langle t + t'; \sigma \rangle \Downarrow \langle \bar{m} + \bar{n}; \sigma'' \rangle} \\
 \\
 \frac{\langle t; \sigma \rangle \Downarrow \langle u; \sigma' \rangle \quad \langle t'; \sigma' \rangle \Downarrow \langle u'; \sigma'' \rangle}{\langle \langle t, t' \rangle; \sigma \rangle \Downarrow \langle \langle u, u' \rangle; \sigma'' \rangle} \qquad \frac{\langle t; \sigma \rangle \Downarrow \langle \langle v_1, v_2 \rangle; \sigma' \rangle \quad i \in \{1, 2\}}{\langle \pi_i(t); \sigma \rangle \Downarrow \langle v_i; \sigma' \rangle} \\
 \\
 \frac{\langle t; \sigma \rangle \Downarrow \langle v; \sigma' \rangle \quad i \in \{1, 2\}}{\langle \text{in}_i(t); \sigma \rangle \Downarrow \langle \text{in}_i(v); \sigma' \rangle} \qquad \frac{\langle t; \sigma \rangle \Downarrow \langle \text{in}_i(u); \sigma' \rangle \quad \langle t_i[v/x]; \sigma' \rangle \Downarrow \langle u_i; \sigma'' \rangle \quad i \in \{1, 2\}}{\langle \text{case } t \text{ of in}_1 x.t_1; \text{in}_2 x.t_2; \sigma \rangle \Downarrow \langle u_i; \sigma'' \rangle} \\
 \\
 \frac{\langle t; \sigma \rangle \Downarrow \langle \lambda x.s; \sigma' \rangle \quad \langle t'; \sigma' \rangle \Downarrow \langle v; \sigma'' \rangle \quad \langle s[v/x]; \sigma'' \rangle \Downarrow \langle v'; \sigma''' \rangle}{\langle t t'; \sigma \rangle \Downarrow \langle v'; \sigma''' \rangle} \\
 \\
 \frac{l = \text{alloc } (\sigma)}{\langle \text{delay } t; \sigma \rangle \Downarrow \langle l; \sigma, l \mapsto t \rangle} \qquad \frac{\langle t; \eta_N \rangle \Downarrow \langle l; \eta'_N \rangle \quad \langle \eta'_N(l); \eta'_N \checkmark \eta_L \rangle \Downarrow \langle v; \sigma' \rangle}{\langle \text{adv } t; \eta_N \checkmark \eta_L \rangle \Downarrow \langle v; \sigma' \rangle} \\
 \\
 \frac{\langle t; \sigma \rangle \Downarrow \langle \text{box } t'; \sigma' \rangle \quad \langle t'; \sigma' \rangle \Downarrow \langle v; \sigma'' \rangle}{\langle \text{unbox } t; \sigma \rangle \Downarrow \langle v; \sigma'' \rangle} \qquad \frac{\langle t; \sigma \rangle \Downarrow \langle v; \sigma' \rangle}{\langle \text{into } t; \sigma \rangle \Downarrow \langle \text{into } v; \sigma' \rangle} \qquad \frac{\langle t; \sigma \rangle \Downarrow \langle \text{into } v; \sigma' \rangle}{\langle \text{out } t; \sigma \rangle \Downarrow \langle v; \sigma' \rangle} \\
 \\
 \frac{\langle t[\text{box } (\text{delay } (\text{fix } x.t))/x]; \sigma \rangle \Downarrow \langle v; \sigma' \rangle}{\langle \text{fix } x.t; \sigma \rangle \Downarrow \langle v; \sigma' \rangle}
 \end{array}$$

Fig. 5. Evaluation semantics.

$$\frac{\langle t; \eta \checkmark \rangle \Downarrow \langle v \text{ :: } l; \eta_N \checkmark \eta_L \rangle}{\langle t; \eta \rangle \xRightarrow{v} \langle \text{adv } l; \eta_L \rangle} \qquad \frac{\langle t; \eta, l^* \mapsto v \text{ :: } l^* \checkmark l^* \mapsto \langle \rangle \rangle \Downarrow \langle v' \text{ :: } l; \eta_N \checkmark \eta_L, l^* \mapsto \langle \rangle \rangle}{\langle t; \eta \rangle \xRightarrow{v/v'} \langle \text{adv } l; \eta_L \rangle}$$

Fig. 6. Step semantics for streams.

heap location that has previously been garbage collected. In other words, programs in λ_{\checkmark} may have implicit space leaks.

Such implicit space leaks can occur in λ_{\checkmark} for two reasons: (1) a lambda abstraction that appears in the scope of a \checkmark , and (2) a term that requires more than one \checkmark to type check. Bahr et al. [2019] give an example of (1). The following term of type $\text{O}(\text{Str Int}) \rightarrow \text{O}(\text{Str Int})$ is an example of (2):

$$\text{fix } r.\lambda x.\text{delay}(\text{head}(\text{adv } x) \text{ :: } \text{adv}(\text{unbox } r)(\text{delay}(\text{adv}(\text{tail}(\text{adv } x))))))$$

Here the abstract machine would get stuck trying to dereference a heap location that was previously garbage collected. The problem is that the second occurrence of x is nested below two occurrences of delay , which means when $\text{adv } x$ is evaluated, the heap location bound to x is two time steps old and has been garbage collected already. Importantly, this problem would occur even if we increased the number of heaps from two to any other fixed upper bound!

On the other hand, if we apply the rewrite rules from Section 4.2.1, we obtain the following λ_{\checkmark} term that is safe to execute on the abstract machine and thus does not suffer from space leaks:

$$\text{fix } r.\lambda x.\text{let } r' = \text{unbox } r \text{ in } \text{delay}(\text{head}(\text{adv } x) \text{ :: } \text{adv } r'(\text{let } y = \text{tail}(\text{adv } x) \text{ in } \text{delay}(\text{adv } y)))$$

In the next section, we present our main theoretical result, namely that all closed λ_{\checkmark} terms of the appropriate type are safe to execute on the abstract machine.

4.3 Main results

The step semantics describes the behaviour of reactive programs. Here we consider two kinds of reactive programs: terms of type $Str\ A$ and terms of type $Str\ A \rightarrow Str\ B$. The former just produces an infinite stream of values of type A whereas the latter is reactive process that produces a value of type B for each input value of type A .

4.3.1 Productivity of the step semantics. The step semantics \xRightarrow{v} from Figure 6 describes the unfolding of streams of type $Str\ A$. Given a closed term $\vdash_{\checkmark} t : Str\ A$, it produces an infinite reduction sequence

$$\langle t; \emptyset \rangle \xRightarrow{v_0} \langle t_1; \eta_1 \rangle \xRightarrow{v_1} \langle t_2; \eta_2 \rangle \xRightarrow{v_2} \dots$$

where \emptyset denotes the empty heap and each v_i has type A . In each step we have a term t_i and the corresponding heap η_i of delayed computations. According to the definition of the step semantics, we evaluate $\langle t_i; \eta_i \checkmark \rangle \Downarrow \langle v_i \ :: \ l; \eta'_i \checkmark \eta_{i+1} \rangle$, where η'_i is η_i but possibly extended with some additional delayed computations and η_{i+1} is the new heap with delayed computations for the next time step. Crucially, the old heap η'_i is thrown away. That is, by construction, old data is not implicitly retained but garbage collected immediately after we completed the current time step.

As an example consider the following definition of the stream of consecutive numbers starting from some given number:

```
from :: Int → Str Int
from n = n :: delay (from (n + 1))
```

This definition translates to the λ_{\checkmark} term $\text{fix } r. \lambda n. n \ :: \ \text{delay}(\text{adv}(\text{unbox } r) (n + \bar{1}))$, which, in turn, rewrites into the following λ_{\checkmark} term:

$$\text{from} = \text{fix } r. \lambda n. n \ :: \ \text{let } r' = \text{unbox } r \text{ in } \text{delay}(\text{adv } r' (n + \bar{1}))$$

Let's see how the term $\text{from } \bar{0}$ of type $Str\ Int$ is executed on the machine:

$$\begin{aligned} \langle \text{from } \bar{0}; \emptyset \rangle &\xRightarrow{\bar{0}} \langle \text{adv } l'_1; l_1 \mapsto \text{from}, l'_1 \mapsto \text{adv } l_1 (\bar{0} + \bar{1}) \rangle \\ &\xRightarrow{\bar{1}} \langle \text{adv } l'_2; l_2 \mapsto \text{from}, l'_2 \mapsto \text{adv } l_2 (\bar{1} + \bar{1}) \rangle \\ &\xRightarrow{\bar{2}} \langle \text{adv } l'_3; l_3 \mapsto \text{from}, l'_3 \mapsto \text{adv } l_3 (\bar{2} + \bar{1}) \rangle \\ &\vdots \end{aligned}$$

In each step, the heap contains at location l_i the fixed point from and at location l'_i the delayed computation produced by the occurrence of delay in the body of the fixed point. The old versions of the delayed computations are garbage collected after each step and only the most recent version survives.

Our main result is that the execution of λ_{\checkmark} terms by the machine described in Figure 5 and 6 is safe. To describe the type of the produced values precisely, we need to restrict ourselves to streams over types whose evaluation is not suspended, which excludes function and modal types. This idea is expressed in the notion of *value types*, defined by the following grammar:

$$\text{Value Types } V, W ::= 1 \mid \text{Int} \mid U \times W \mid U + W$$

We can then prove the following theorem, which both expresses the fact that the aggressive garbage collection strategy of RATTUS is safe, and that stream programs are productive:

THEOREM 4.2 (PRODUCTIVITY). *Given a term $\vdash_{\checkmark} t : \text{Str } A$ with A a value type, there is an infinite reduction sequence*

$$\langle t; \emptyset \rangle \xRightarrow{v_0} \langle t_1; \eta_1 \rangle \xRightarrow{v_1} \langle t_2; \eta_2 \rangle \xRightarrow{v_2} \dots \quad \text{such that } \vdash_{\checkmark} v_i : A \text{ for all } i \geq 0.$$

The restriction to value types is only necessary for showing that each output value v_i has the correct type.

4.3.2 Causality of the step semantics. The step semantics $\xRightarrow{v/v'}$ from Figure 6 describes how a term of type $\text{Str } A \rightarrow \text{Str } B$ transforms a stream of inputs into a stream of outputs in a step-by-step fashion. Given a closed term $\vdash_{\checkmark} t : \text{Str } A \rightarrow \text{Str } B$, and an infinite stream of input values $\vdash_{\checkmark} v_i : A$, it produces an infinite reduction sequence

$$\langle t; \emptyset \rangle \xRightarrow{v_0/v'_0} \langle t_1; \eta_1 \rangle \xRightarrow{v_1/v'_1} \langle t_2; \eta_2 \rangle \xRightarrow{v_2/v'_2} \dots$$

where each output value v'_i has type B .

The definition of $\xRightarrow{v/v'}$ assumes that we have some fixed heap location l^* , which acts both as interface to the currently available input value and as a stand-in for future inputs that are not yet available. In each step, we evaluate the current term t_i in the current heap η_i

$$\langle t_i; \eta_i, l^* \mapsto v_i \text{ :: } l^* \checkmark l^* \mapsto \langle \rangle \rangle \Downarrow \langle v'_i \text{ :: } l; \eta'_i \checkmark \eta_{i+1}, l^* \mapsto \langle \rangle \rangle$$

which produces the output v'_i and the new heap η_{i+1} . Again the old heap η'_i is simply dropped. In the ‘later’ heap, the operational semantics maps l^* to the placeholder value $\langle \rangle$, which is safe since the machine never reads from the later heap. Then in the next reduction step, we replace that placeholder value with $v_{i+1} \text{ :: } l^*$ which contains the newly received input value v_{i+1} .

For an example, consider the following function that takes a stream of integers and produces the stream of prefix sums:

$sum :: \text{Str Int} \rightarrow \text{Str Int}$

$sum = run\ 0$ **where**

$run :: \text{Int} \rightarrow \text{Str Int} \rightarrow \text{Str Int}$

$run\ acc\ (x \text{ :: } xs) = \mathbf{let}\ acc' = acc + x$

$\mathbf{in}\ acc' \text{ :: } \mathbf{delay}\ (run\ acc'\ (\mathbf{adv}\ xs))$

This function definition translates to the following term sum in the λ_{\checkmark} calculus:

$$run = \mathbf{fix}\ r.\lambda acc.\lambda s.\mathbf{let}\ x = \mathbf{head}\ s\ \mathbf{in}\ \mathbf{let}\ xs = \mathbf{tail}\ s\ \mathbf{in}\ \mathbf{let}\ acc' = acc + x\ \mathbf{in}\ \mathbf{let}\ r' = \mathbf{unbox}\ r\ \mathbf{in}\ acc' \text{ :: } \mathbf{delay}\ (\mathbf{adv}\ r'\ acc'\ (\mathbf{adv}\ xs))$$

$$sum = run\ \bar{0}$$

Let’s look at the first three steps of executing the sum function with 2, 11, and 5 as its first three input values:

$$\begin{aligned} & \langle sum; \emptyset \rangle \\ & \xRightarrow{\bar{2}/\bar{2}} \langle \mathbf{adv}\ l'_1; l_1 \mapsto run, l'_1 \mapsto \mathbf{adv}\ l_1\ (\bar{0} + \bar{2})\ (\mathbf{adv}\ l^*) \rangle \\ & \xRightarrow{\bar{11}/\bar{13}} \langle \mathbf{adv}\ l'_2; l_2 \mapsto run, l'_2 \mapsto \mathbf{adv}\ l_2\ (\bar{2} + \bar{11})\ (\mathbf{adv}\ l^*) \rangle \\ & \xRightarrow{\bar{5}/\bar{18}} \langle \mathbf{adv}\ l'_3; l_3 \mapsto run, l'_3 \mapsto \mathbf{adv}\ l_3\ (\bar{13} + \bar{5})\ (\mathbf{adv}\ l^*) \rangle \\ & \vdots \end{aligned}$$

in each step of the computation the location l_i stores the fixed point run and l'_i stores the computation that calls that fixed point with the new accumulator value (0 + 2, 2 + 11, and 13 + 5, respectively) and the tail of the current input stream.

We can prove the following theorem, which again expresses the fact that the garbage collection strategy of RATTUS is safe, and that stream processing functions are both productive and causal:

THEOREM 4.3 (CAUSALITY). *Given a term $\vdash_{\checkmark} t : \text{Str } A \rightarrow \text{Str } B$ with B a value type, and an infinite sequence of values $\vdash_{\checkmark} v_i : A$, there is an infinite reduction sequence*

$$\langle t; \emptyset \rangle \xRightarrow{v_0/v'_0} \langle t_1; \eta_1 \rangle \xRightarrow{v_1/v'_1} \langle t_2; \eta_2 \rangle \xRightarrow{v_2/v'_2} \dots \quad \text{such that } \vdash_{\checkmark} v'_i : B \text{ for all } i \geq 0.$$

Since the operational semantics is deterministic, in each step $\langle t_i; \eta_i \rangle \xRightarrow{v_i/v'_i} \langle t_{i+1}; \eta_{i+1} \rangle$ the resulting output v'_{i+1} and new state of the computation $\langle t_{i+1}; \eta_{i+1} \rangle$ are uniquely determined by the previous state $\langle t_i; \eta_i \rangle$ and the input v_i . Thus, v'_i and $\langle t_{i+1}; \eta_{i+1} \rangle$ are independent of future inputs v_j with $j > i$.

4.4 Limitations

Now that we have formally precise statements about the operational properties of RATTUS' core calculus, we should make sure that we understand what they mean in practice and what their limitations are. In simple terms, the productivity and causality properties established by Theorem 4.2 and Theorem 4.3 state that reactive programs in RATTUS can be executed effectively – they always make progress and never depend on data that is not yet available. In the Haskell embedding of the language this has to be of course qualified as we can use Haskell functions that loop or crash.

In addition, by virtue of the operational semantics, the two theorems also imply that programs can be executed without implicitly retaining memory – thus avoiding *implicit space leaks*. This follows from the fact that in each step the step semantics (in Figure 6) discards the ‘now’ heap and only retains the ‘later’ heap for the next step.

However, similarly to the calculi of Krishnaswami [2013] and Bahr et al. [2019, 2021], RATTUS still allows *explicit space leaks* (we may still construct data structures to hold on to an increasing amount of memory) as well as *time leaks* (computations may take an increasing amount of time). Below we give some examples of these behaviours.

Given a strict list type

```
data List a = Nil | !a :!(List a)
```

we can construct a function that buffers the entire history of an input stream

```
buffer :: Stable a => Str a -> Str (List a)
buffer = scan (box (\xs x -> x :!(xs)) Nil
```

Given that we have a function $sum :: List Int \rightarrow Int$ that computes the sum of a list of numbers, we can write the following alternative implementation of the $sums$ function using $buffer$:

```
leakySums1 :: Str Int -> Str Int
leakySums1 = map (box sum) o buffer
```

At each time step this function adds the current input integer to the buffer of type $List Int$ and then computes the sum of the current value of that buffer. This function exhibits both a space leak (buffering a steadily growing list of numbers) and a time leak (the time to compute each element of the resulting stream increases at each step). However, these leaks are explicit.

An example of a time leak is found in the following alternative implementation of the $sums$ function:

$leakySums2 :: Str\ Int \rightarrow Str\ Int$

$leakySums2\ (x :: xs) = x :: delay\ (map\ (box\ (+x))\ (leakySums2\ (adv\ xs)))$

In each step we add the current input value x to each future output. The closure $(+x)$, which is Haskell shorthand notation for $\lambda y \rightarrow y + x$, stores each input value x .

None of the above space and time leaks are prevented by RATTUS. The space leaks in *buffer* and *leakySums1* are explicit since the desire to buffer the input is explicitly stated in the program. The other example is more subtle as the leaky behaviour is rooted in a time leak as the program constructs an increasing computation in each step. This shows that the programmer still has to be careful about time leaks. Note that these leaky functions can also be implemented in the calculi of Krishnaswami [2013] and Bahr et al. [2019, 2021], although some reformulation is necessary for the latter calculus. For more details we refer to the discussion on related work in Section 7.

5 META THEORY

Our goal is to show that RATTUS's core calculus λ_{\surd} enjoys the three central operational properties: productivity, causality and absence of implicit space leaks. These properties are stated in Theorem 4.2 and Theorem 4.3, and we show in this section how these are proved. Note that the absence of space leaks follows from these theorems because the operational semantics already ensures this memory property by means of garbage collecting the 'now' heap after each step. Since the proof is fully formalised in the accompanying Coq proofs, we only give a high-level overview here.

We prove the abovementioned theorems by establishing a semantic soundness property. For productivity, our soundness property must imply that the evaluation semantics $\langle t; \sigma \rangle \Downarrow \langle v; \sigma' \rangle$ converges for each well-typed term t , and for causality, the soundness property must imply that this is also the case if t contains references to heap locations in σ .

To obtain such a soundness result, we construct a *Kripke logical relation* that incorporates these properties. Generally speaking a Kripke logical relation constructs for each type A a relation $\llbracket A \rrbracket_w$ indexed over some world w with some closure conditions when the index w changes. In our case, $\llbracket A \rrbracket_w$ is a set of terms. Moreover, the index w consists of three components: a number ν to act as a step index [Appel and McAllester 2001], a store σ to establish the safety of garbage collection, and an infinite sequence $\bar{\eta}$ of future heaps in order to capture the causality property.

A crucial ingredient of a Kripke logical relation is the ordering on the indices. The ordering on the number ν is the standard ordering on numbers. For heaps we use the standard ordering on partial maps: $\eta \sqsubseteq \eta'$ iff $\eta(l) = \eta'(l)$ for all $l \in \text{dom}(\eta)$. Infinite sequences of heaps are ordered pointwise according to \sqsubseteq . Moreover, we extend the ordering to stores in two different ways:

$$\frac{\eta_N \sqsubseteq \eta'_N \quad \eta_L \sqsubseteq \eta'_L}{\eta_N \surd \eta_L \sqsubseteq \eta'_N \surd \eta'_L} \qquad \frac{\sigma \sqsubseteq \sigma'}{\sigma \sqsubseteq_{\surd} \sigma'} \qquad \frac{\eta \sqsubseteq \eta'}{\eta \sqsubseteq_{\surd} \eta'' \surd \eta'}$$

That is, \sqsubseteq is the pointwise extension of the order on heaps to stores, and \sqsubseteq_{\surd} is more general and permits introducing an arbitrary 'now' heap if none is present.

Given these orderings we define two logical relations, the value relation $\mathcal{V}_{\nu} \llbracket A \rrbracket_{\sigma}^{\bar{\eta}}$ and the term relation $\mathcal{T}_{\nu} \llbracket A \rrbracket_{\sigma}^{\bar{\eta}}$. Both are defined in Figure 7 by well-founded recursion according to the lexicographic ordering on the triple $(\nu, |A|, e)$, where $|A|$ is the size of A defined below, and $e = 1$ for the term relation and $e = 0$ for the value relation.

$$\begin{aligned} |\alpha| &= |\circ A| = |\text{Int}| = |1| = 1 \\ |A \times B| &= |A + B| = |A \rightarrow B| = 1 + |A| + |B| \\ |\square A| &= |\text{Fix } \alpha.A| = 1 + |A| \end{aligned}$$

In the definition of the logical relation, we use the notation $\eta; \bar{\eta}$ to denote an infinite sequence of heaps that starts with the heap η and then continues as the sequence $\bar{\eta}$. Moreover, we use the notation $\sigma(l)$ to denote $\eta_L(l)$ if σ is of the form η_L or $\eta_N \checkmark \eta_L$.

The crucial part of the logical relation that ensures both causality and the absence of space leaks is the case for $\bigcirc A$. The value relation of $\bigcirc A$ at store index σ is defined as all heap locations that map to computations in the term relation of A but at the store index $\text{gc}(\sigma) \checkmark \eta$. Here $\text{gc}(\sigma)$ denotes the garbage collection of the store σ as defined in Figure 7. It simply drops the ‘now’ heap if present. To see how this definition captures causality we have to look at the index $\eta; \bar{\eta}$ of future heaps. It changes to the index $\bar{\eta}$, i.e. all future heaps are one time step closer, and the very first future heap η becomes the new ‘later’ heap in the store index $\text{gc}(\sigma) \checkmark \eta$, whereas the old ‘later’ heap in σ becomes the new ‘now’ heap.

The central theorem that establishes type soundness is the so-called *fundamental property* of the logical relation. It states that well-typed terms are in the term relation. For the induction proof of this property we also need to consider open terms and to this end, we also need a corresponding context relation $C_v \llbracket \Gamma \rrbracket_{\sigma}^{\bar{\eta}}$, which is given in Figure 7.

THEOREM 5.1 (FUNDAMENTAL PROPERTY). *Given $\Gamma \vdash_{\checkmark} t : A$, and $\gamma \in C_v \llbracket \Gamma \rrbracket_{\sigma}^{\bar{\eta}}$, then $t\gamma \in \mathcal{T}_v \llbracket A \rrbracket_{\sigma}^{\bar{\eta}}$*

The proof of the fundamental property is a lengthy but entirely standard induction on the typing relation $\Gamma \vdash_{\checkmark} t : A$. Both Theorem 4.2 and Theorem 4.3 are then proved using the above theorem.

6 EMBEDDING RATTUS IN HASKELL

Our goal with RATTUS is to combine the operational guarantees provided modal FRP with the practical benefits of FRP libraries. Because of the Fitch-style typing rules we cannot implement RATTUS as a straightforward library of combinators. Instead we rely on a combination of a very simple library that implements the primitives of the language together with a compiler plugin that performs additional checks. We start with a description of the implementation followed by an illustration how the implementation is used in practice.

6.1 Implementation of RATTUS

At its core, our implementation consists of a very simple library that implements the primitives of RATTUS (delay, adv, box, and unbox) so that they can be readily used in Haskell code. The library is given in its entirety in Figure 8. Both \bigcirc and \square are simple wrapper types, each with their own wrap and unwrap function. The constructors *Delay* and *Box* are not exported by the library, i.e. \bigcirc and \square are treated as abstract types.

If we were to use these primitives as provided by the library we would end up with the problems illustrated in Section 2. Such an implementation of RATTUS would enjoy none of the operational properties we have proved. To make sure that programs use these primitives according to the typing rules of RATTUS, our implementation has a second component: a plugin for the GHC Haskell compiler that enforces the typing rules of RATTUS.

The design of this plugin follows the simple observation that any RATTUS program is also a Haskell program but with more restrictive rules for variable scope and when RATTUS’s primitives may be used. So type checking a RATTUS program boils down to first typechecking it as a Haskell program and then checking that it follows the stricter variable scope rules. That means, we must keep track of when variables fall out of scope due to the use of delay, adv and box, but also due to guarded recursion. Additionally, we must make sure that both guarded recursive calls and adv only appear in a context where \checkmark is present.

$$\begin{aligned}
 \mathcal{V}_v \llbracket \text{Int} \rrbracket_{\sigma}^{\bar{\eta}} &= \{\bar{n} \mid n \in \mathbb{Z}\}, \\
 \mathcal{V}_v \llbracket 1 \rrbracket_{\sigma}^{\bar{\eta}} &= \{\langle \rangle\}, \\
 \mathcal{V}_v \llbracket A \times B \rrbracket_{\sigma}^{\bar{\eta}} &= \{\langle v_1, v_2 \rangle \mid v_1 \in \mathcal{V}_v \llbracket A \rrbracket_{\sigma}^{\bar{\eta}} \wedge v_2 \in \mathcal{V}_v \llbracket B \rrbracket_{\sigma}^{\bar{\eta}}\}, \\
 \mathcal{V}_v \llbracket A + B \rrbracket_{\sigma}^{\bar{\eta}} &= \{\text{in}_1 v \mid v \in \mathcal{V}_v \llbracket A \rrbracket_{\sigma}^{\bar{\eta}}\} \cup \{\text{in}_2 v \mid v \in \mathcal{V}_v \llbracket B \rrbracket_{\sigma}^{\bar{\eta}}\}, \\
 \mathcal{V}_v \llbracket A \rightarrow B \rrbracket_{\sigma}^{\bar{\eta}} &= \left\{ \lambda x. t \mid \forall v' \leq v, \sigma' \sqsupseteq \text{gc}(\sigma), \bar{\eta}' \sqsupseteq \bar{\eta}. \forall u \in \mathcal{V}_v \llbracket A \rrbracket_{\sigma'}^{\bar{\eta}'}. t[u/x] \in \mathcal{T}_v \llbracket B \rrbracket_{\sigma'}^{\bar{\eta}'} \right\}, \\
 \mathcal{V}_v \llbracket \Box A \rrbracket_{\sigma}^{\bar{\eta}} &= \{\text{box } t \mid \forall \bar{\eta}' . t \in \mathcal{T}_v \llbracket A \rrbracket_{\emptyset}^{\bar{\eta}'}\}, \\
 \mathcal{V}_0 \llbracket \bigcirc A \rrbracket_{\sigma}^{\bar{\eta}} &= \{l \mid l \in \text{Loc}\} \\
 \mathcal{V}_{v+1} \llbracket \bigcirc A \rrbracket_{\sigma}^{\eta; \bar{\eta}} &= \{l \mid \sigma(l) \in \mathcal{T}_v \llbracket A \rrbracket_{\text{gc}(\sigma) \checkmark \eta}^{\bar{\eta}}\}, \\
 \mathcal{V}_v \llbracket \text{Fix } \alpha. A \rrbracket_{\sigma}^{\bar{\eta}} &= \left\{ \text{into}(v) \mid v \in \mathcal{V}_v \llbracket A[\bigcirc(\text{Fix } \alpha. A)/\alpha] \rrbracket_{\sigma}^{\bar{\eta}} \right\} \\
 \mathcal{T}_v \llbracket A \rrbracket_{\sigma}^{\bar{\eta}} &= \left\{ t \mid \forall \sigma' \sqsupseteq \checkmark \sigma. \exists \sigma'', v. \langle t; \sigma' \rangle \Downarrow \langle v; \sigma'' \rangle \wedge v \in \mathcal{V}_v \llbracket A \rrbracket_{\sigma''}^{\bar{\eta}} \right\}
 \end{aligned}$$

$$\begin{aligned}
 C_v \llbracket \cdot \rrbracket_{\sigma}^{\bar{\eta}} &= \{\star\} & \text{GARBAGE COLLECTION:} \\
 C_v \llbracket \Gamma, x : A \rrbracket_{\sigma}^{\bar{\eta}} &= \left\{ \gamma[x \mapsto v] \mid \gamma \in C_v \llbracket \Gamma \rrbracket_{\sigma}^{\bar{\eta}}, v \in \mathcal{V}_v \llbracket A \rrbracket_{\sigma}^{\bar{\eta}} \right\} & \text{gc}(\eta_L) = \eta_L \\
 C_v \llbracket \Gamma, \checkmark \rrbracket_{\eta_N \checkmark \eta_L}^{\bar{\eta}} &= C_{v+1} \llbracket \Gamma \rrbracket_{\eta_N}^{\eta_L; \bar{\eta}} & \text{gc}(\eta_N \checkmark \eta_L) = \eta_L
 \end{aligned}$$

Fig. 7. Logical relation.

data $\bigcirc a = \text{Delay } a$	data $\Box a = \text{Box } a$	class <i>StableInternal</i> a where
<code>delay</code> $:: a \rightarrow \bigcirc a$	<code>box</code> $:: a \rightarrow \Box a$	class <i>StableInternal</i> $a \Rightarrow \text{Stable } a$ where
<code>delay</code> $x = \text{Delay } x$	<code>box</code> $x = \text{Box } x$	
<code>adv</code> $:: \bigcirc a \rightarrow a$	<code>unbox</code> $:: \Box a \rightarrow a$	
<code>adv</code> (<i>Delay</i> x) = x	<code>unbox</code> (<i>Box</i> d) = d	

Fig. 8. Implementation of RATTUS primitives.

To enforce these additional simple scope rules we make use of GHC's plugin API which allows us to customise part of GHC's compilation pipeline. The different phases of GHC are illustrated in Figure 9 with the additional passes performed by the RATTUS plugin highlighted in bold.

After type checking the Haskell abstract syntax tree (AST), GHC passes the resulting typed AST on to the scope checking component of the RATTUS plugin, which checks the abovementioned stricter scoping rules. GHC then desugars the typed AST into the intermediate language *Core*. GHC then performs a number of transformation passes on this intermediate representation, the first two of these are provided by the RATTUS plugin: First, we exhaustively apply the two rewrite rules from Section 4.2.1 to transform the program into a single-tick form according to the typing rules of λ_{\checkmark} . Then we transform the resulting code so that RATTUS programs adhere to the call-by-value semantics. To this end, the plugin's *strictness* pass transforms all function applications so that

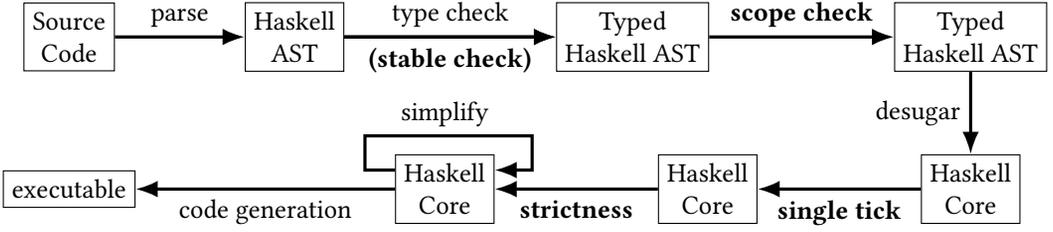


Fig. 9. Compiler phases of GHC (simplified) extended with RATTUS plugin (highlighted in bold).

arguments are evaluated to weak head normal form before they are passed to the function. In addition, this *strictness* pass also checks that RATTUS code only uses strict data types and issues a warning if lazy data types are used, e.g. Haskell’s standard list and pair types.

Not pictured in the diagram is a second scope checking pass that is performed after the *strictness* pass. After the *single tick* pass and thus also after the *strictness* pass, we expect that the code is typable according to the more restrictive typing rules of λ_{\perp} . This second scope checking pass checks this invariant for the purpose of catching implementation bugs in the RATTUS plugin. The Core intermediate language is *much* simpler than the full Haskell language, so this second scope checking pass is much easier to implement and much less likely to contain bugs. In principle, we could have saved ourselves the trouble of implementing the much more complicated scope checking at the level of the typed Haskell AST. However, by checking at this earlier stage of the compilation pipeline, we can provide the user much more helpful type error messages.

One important component of checking variable scope is checking whether types are stable. This is a simple syntactic check: a type τ is stable if all occurrences of \circ or function types in τ are nested under a \square . However, RATTUS also supports polymorphic types with type constraints such as in the *const* combinator:

```

const :: Stable a  $\Rightarrow$  a  $\rightarrow$  Str a
const x = x ::: delay (const x)

```

The *Stable* type class is defined as a primitive in the RATTUS library (see Figure 8). The library does not export the underlying *StableInternal* type class so that the user cannot declare any instances for *Stable*. Our library does not declare instances of the *Stable* class either. Instead, such instances are derived by the RATTUS plugin that uses GHC’s typechecker plugin API, which allows us to provide limited customisation to GHC’s type checking phase (see Figure 9). Using this API one can give GHC a custom procedure for resolving type constraints. Whenever GHC’s type checker finds a constraint of the form *Stable* τ , it will send it to the RATTUS plugin, which will resolve it by performing the abovementioned syntactic check on τ .

6.2 Using RATTUS

To write RATTUS code inside Haskell one must use GHC with the flag `-fplugin=Rattus.Plugin`, which enables the RATTUS plugin described above. Figure 10 shows a complete program that illustrates the interaction between Haskell and RATTUS. The language is imported via the *Rattus* module, with the *Rattus.Stream* module providing a stream library (of which we have seen an excerpt in Figure 1). The program contains only one RATTUS function, *summing*, which is indicated by an annotation. This function uses the *scan* combinator to define a stream transducer that sums up its input stream. Finally, we use the *runTransducer* function that is provided by the *Rattus.ToHaskell*

```

{-# OPTIONS -fplugin=Rattus.Plugin #-}
import Rattus
import Rattus.Stream
import Rattus.ToHaskell
{-# ANN sums Rattus #-}
sums :: Str Int → Str Int
sums = scan (box (+)) 0

main = loop (runTransducer sums)
  where loop (Trans t) = do
    input ← readLn
    let (result, next) = t input
        print result
        loop next

```

Fig. 10. Complete RATTUS program.

module. It turns a stream function of type $Str\ a \rightarrow Str\ b$ into a Haskell value of type $Trans\ a\ b$ defined as follows:

```
data Trans a b = Trans (a → (b, Trans a b))
```

This allows us to run the stream function step by step as illustrated in the main function: It reads an integer from the console passes it on to the stream function, prints out the response, and then repeats the process.

Alternatively, if a module contains only RATTUS definitions we can use the annotation

```
{-# ANN module Rattus #-}
```

to declare that all definitions in a module are to be interpreted as RATTUS code.

7 RELATED WORK

The central ideas of functional reactive programming were originally developed for the language Fran [Elliott and Hudak 1997] for reactive animation. These ideas have since been developed into general purpose libraries for reactive programming, most prominently the Yampa library [Nilsson et al. 2002] for Haskell, which has been used in a variety of applications including games, robotics, vision, GUIs, and sound synthesis.

More recently Ploeg and Claessen [2015] have developed the *FRPNow!* library for Haskell, which – like Fran – uses behaviours and events as FRP primitives (as opposed to signal functions), but carefully restricts the API to guarantee causality and the absence of implicit space leaks. To argue for the latter, the authors construct a denotational model and show using a logical relation that their combinators are not “inherently leaky”. The latter does not imply the absence of space leaks, but rather that in principle it can be implemented without space leaks.

Modal FRP calculi. The idea of using modal type operators for reactive programming goes back to Jeffrey [2012], Krishnaswami and Benton [2011], and Jeltsch [2013]. One of the inspirations for Jeffrey [2012] was to use linear temporal logic [Pnueli 1977] as a programming language through the Curry-Howard isomorphism. The work of Jeffrey and Jeltsch has mostly been based on denotational semantics, and Bahr et al. [2019, 2021]; Cave et al. [2014]; Krishnaswami [2013]; Krishnaswami and Benton [2011]; Krishnaswami et al. [2012] are the only works to our knowledge giving operational guarantees. In addition the calculi of Bahr et al. [2021] and Cave et al. [2014] can be used to encode liveness properties in modal FRP.

The guarded recursive types and fixed point combinator originate with Nakano [2000], but have since been used for constructing logics for reasoning about advanced programming languages [Birkedal et al. 2011] using an abstract form of step-indexing [Appel and McAllester 2001]. The Fitch-style approach to modal types [Fitch 1952] has been used for guarded recursion in Clocked

Type Theory [Bahr et al. 2017], where contexts can contain multiple, named ticks. Ticks can be used for reasoning about guarded recursive programs. The denotational semantics of Clocked Type Theory [Mannaa and Møgelberg 2018] reveals the difference from the more standard dual context approaches to modal logics, such as Dual Intuitionistic Linear Logic [Barber 1996]: In the latter, the modal operator is implicitly applied to the type of all variables in one context, in the Fitch-style, placing a tick in a context corresponds to applying a *left adjoint* to the modal operator to the context. Guatto [2018] introduced the notion of time warp and the warping modality, generalising the delay modality in guarded recursion, to allow for a more direct style of programming for programs with complex input-output dependencies.

Space leaks. The work by Krishnaswami [2013] and Bahr et al. [2019, 2021] is the closest to the present work. Both present a modal FRP language with a garbage collection result similar to ours. Krishnaswami [2013] pioneered this approach to prove the absence of implicit space leaks, but also implemented a compiler for his language, which translates FRP programs into JavaScript.

Like the present work, the Simply RaTT calculus of Bahr et al. [2019] uses a Fitch-style type system, which provides lighter syntax to interact with the \Box and \bigcirc modality compared to Krishnaswami's use of qualifiers in his calculus. The latter is closely related to dual context systems and requires the use of pattern matching as elimination forms of the modalities. Bahr et al. [2021] extend Simply RaTT with inductive types to express liveness properties. But otherwise their calculus is similar to Simply RaTT and the discussion below also applies to this extension.

As discussed in Section 2.2, Simply RaTT restricts where ticks may occur. In addition, Simply RaTT has a more complicated typing rule for guarded fixed points (cf. Figure 11). It uses a token $\#$ (in addition to \checkmark) to serve the role that stabilisation of a context Γ to Γ^\square serves in RATTUS. Moreover, fixed points produce terms of type $\Box A$ rather than just A . Taken together, this makes the syntax and typing for guarded recursive function definitions more complicated and less intuitive. For example, the *map* function would be defined as follows in Simply RaTT:

$$\begin{aligned} \text{map} &: \Box(a \rightarrow b) \rightarrow \Box(\text{Str } a \rightarrow \text{Str } b) \\ \text{map } f \# (a :: as) &= \text{unbox } f \ a :: (\text{map } f \ \otimes \ as) \end{aligned}$$

Here, the $\#$ is used to indicate that the argument f is to the left of the $\#$ token and only because of the presence of this token can we use the unbox combinator on f (cf. Figure 11). Additionally, the typing of recursive definitions is somewhat confusing: *map* has return type $\Box(\text{Str } a \rightarrow \text{Str } b)$ but when used in a recursive call as seen above *map* f is of type $\bigcirc(\text{Str } a \rightarrow \text{Str } b)$ instead. Moreover, we cannot call *map* recursively on its own. All recursive calls must be of the form *map* f , the exact pattern that appears to the left of the $\#$.

We argue that our typing system and syntax is simpler than both the work of Krishnaswami [2013] and Bahr et al. [2019, 2021], combining the simpler syntax of fixed points with the more streamlined syntax afforded by Fitch-style typing.

In addition, RATTUS permits recursive functions that look more than one time step into the future (e.g. *stutter* from Section 2.2), which is not possible in Krishnaswami [2013] and Bahr et al. [2019, 2021]. However, we suspect that Krishnaswami's calculus can be adapted to allow this by changing the typing rule for *fix* in the same as we did here.

We should note that that Simply RaTT will reject some programs with time leaks, e.g. *leakySums2* from Section 4.4, because of the condition in the rule for unbox requiring that Γ' be token-free. However, we can easily write a program that is equivalent to *leakySums2*, that is well-typed in Simply RaTT using tupling (essentially defining the function mutually recursively with *map*).

$$\frac{\Gamma, \sharp, x : \bigcirc A \vdash t : A}{\Gamma \vdash \text{fix } x.t : \square A} \quad \frac{\Gamma \vdash t : \square A \quad \Gamma' \text{ token-free}}{\Gamma, \sharp, \Gamma' \vdash \text{unbox } t : A} \quad \frac{\Gamma, x : A \vdash t : B \quad \Gamma' \text{ tick-free}}{\Gamma \vdash \lambda x.t : A \rightarrow B}$$

 Fig. 11. Selected typing rules from [Bahr et al. \[2019\]](#).

Finally, typing in Simply RaTT is closed under certain eta expansions involving \square , which are no longer well-typed in RaTTUS because of the typing rule for unbox. For example, we have

$$f : A \rightarrow \square B \vdash \lambda x.\text{box}(\text{unbox}(f x)) : A \rightarrow \square B$$

in Simply RaTT but not λ_{\surd} . However, because it lifts Simply RaTT's restrictions on tokens, RaTTUS is closed under several types of eta expansions that are not well-typed in Simply RaTT. For example,

$$\begin{aligned} x : \square(\square A) &\vdash_{\surd} \text{box}(\text{box}(\text{unbox}(\text{unbox } x))) : \square(\square A) \\ x : \bigcirc(\bigcirc A) &\vdash_{\surd} \text{delay}(\text{delay}(\text{adv}(\text{adv } x))) : \bigcirc(\bigcirc A) \\ f : \bigcirc(A \rightarrow B) &\vdash_{\surd} \text{delay}(\lambda x.\text{adv } f x) : \bigcirc(A \rightarrow B) \end{aligned}$$

8 DISCUSSION AND FUTURE WORK

We have shown that modal FRP can be seamlessly integrated into the Haskell programming language. Two main ingredients are central to achieving this integration: (1) the use of Fitch-style typing to simplify the syntax for interacting with the two modalities and (2) lifting some of the restrictions found in previous work on Fitch-style typing systems.

This paper opens up many avenues for future work both on the implementation side and the underlying theory. We chose Haskell as our host language as it has a compiler extension API that makes it easy for us to implement RaTTUS and convenient for programmers to start using RaTTUS with little friction. However, we think that implementing RaTTUS in call-by-value languages like OCaml or F# should be easily achieved by a simple post-processing step that checks the Fitch-style variable scope. This can be done by an external tool (not unlike a linter) that does not need to be integrated into the compiler. Moreover, while the use of the type class *Stable* is convenient, it is not necessary as we can always use the \square modality instead (cf. *const* vs. *constBox*). When a program transformation approach is not desirable, one can also use λ_{\surd} rather than λ_{\surd} as the underlying calculus. Only few programs need the flexibility of λ_{\surd} and the few that do can be easily transformed by the programmer with only little syntactic clutter.

FRP is not the only possible application of Fitch-style type systems. However, most of the interest in Fitch-style system has been in logics and dependent type theory [[Bahr et al. 2017](#); [Birkedal et al. 2018](#); [Borghuis 1994](#); [Clouston 2018](#)] as opposed to programming languages. RaTTUS is to our knowledge the first implementation of a Fitch-style programming language. We would expect that programming languages for information control flow [[Kavvos 2019](#)] and recent work on modalities for pure computations [Chaudhury and Krishnaswami \[2020\]](#) admit a Fitch-style presentation and could be implemented similarly to RaTTUS.

Part of the success of FRP libraries such as Yampa and FRPNow! is due to the fact that they provide a rich and highly optimised API that integrates well with its host language. In this paper, we have shown that RaTTUS can be seamlessly embedded in Haskell, but more work is required to design a good library and to perform the low-level optimisations that are often necessary to obtain good real-world performance. For example, our definition of signal functions in [Section 3.3](#) resembles the semantics of Yampa's signal functions, but in Yampa signal functions are defined as a GADT that can handle some special cases much more efficiently.

REFERENCES

- Andrew W. Appel and David McAllester. 2001. An Indexed Model of Recursive Types for Foundational Proof-carrying Code. *ACM Trans. Program. Lang. Syst.* 23, 5 (Sept. 2001), 657–683. <https://doi.org/10.1145/504709.504712> 00283.
- Patrick Bahr, Hans Bugge Grathwohl, and Rasmus Ejlers Møgelberg. 2017. The clocks are ticking: No more delays!. In *32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20-23, 2017*. IEEE Computer Society, Washington, DC, USA, 1–12. <https://doi.org/10.1109/LICS.2017.8005097>
- Patrick Bahr, Christian Uldal Graulund, and Rasmus Ejlers Møgelberg. 2019. Simply RaTT: a fitch-style modal calculus for reactive programming without space leaks. *Proceedings of the ACM on Programming Languages* 3, ICFP (2019), 1–27.
- Patrick Bahr, Christian Uldal Graulund, and Rasmus Ejlers Møgelberg. 2021. Diamonds are not Forever: Liveness in Reactive Programming with Guarded Recursion. (Jan. 2021). POPL 2021, to appear.
- Andrew Barber. 1996. *Dual intuitionistic linear logic*. Technical Report. University of Edinburgh, Edinburgh, UK.
- Lars Birkedal, Ranald Clouston, Bassel Mannaa, Rasmus Ejlers Møgelberg, Andrew M. Pitts, and Bas Spitters. 2018. Modal Dependent Type Theory and Dependent Right Adjoints. *arXiv:1804.05236 [cs]* (April 2018). <http://arxiv.org/abs/1804.05236> 00000 arXiv: 1804.05236.
- Lars Birkedal, Rasmus Ejlers Møgelberg, Jan Schwinghammer, and Kristian Støvring. 2011. First steps in synthetic guarded domain theory: Step-indexing in the topos of trees. In *In Proc. of LICS*. IEEE Computer Society, Washington, DC, USA, 55–64. [https://doi.org/10.2168/LMCS-8\(4:1\)2012](https://doi.org/10.2168/LMCS-8(4:1)2012)
- Valentijn Anton Johan Borghuis. 1994. *Coming to terms with modal logic: on the interpretation of modalities in typed lambda-calculus*. PhD Thesis. Technische Universiteit Eindhoven. <http://repository.tue.nl/427575> 00034.
- Andrew Cave, Francisco Ferreira, Prakash Panangaden, and Brigitte Pientka. 2014. Fair Reactive Programming. In *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '14)*. ACM, San Diego, California, USA, 361–372. <https://doi.org/10.1145/2535838.2535881>
- Vikraman Chaudhury and Neel Krishnaswami. 2020. Recovering Purity with Comonads and Capabilities. (2020). ICFP 2020, to appear.
- Ranald Clouston. 2018. Fitch-style modal lambda calculi. In *Foundations of Software Science and Computation Structures*, Christel Baier and Ugo Dal Lago (Eds.), Vol. 10803. Springer, Springer International Publishing, Cham, 258–275.
- Conal Elliott and Paul Hudak. 1997. Functional Reactive Animation. In *Proceedings of the Second ACM SIGPLAN International Conference on Functional Programming (Amsterdam, The Netherlands) (ICFP '97)*. ACM, New York, NY, USA, 263–273. <https://doi.org/10.1145/258948.258973>
- Frederic Benton Fitch. 1952. *Symbolic logic, an introduction*. Ronald Press Co., New York, NY, USA.
- Adrien Guatto. 2018. A generalized modality for recursion. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science*. ACM, 482–491.
- Paul Hudak, Antony Courtney, Henrik Nilsson, and John Peterson. 2004. Arrows, Robots, and Functional Reactive Programming. In *Advanced Functional Programming (Lecture Notes in Computer Science, Vol. 2638)*. Springer Berlin / Heidelberg. https://doi.org/10.1007/978-3-540-44833-4_6
- Alan Jeffrey. 2012. LTL types FRP: linear-time temporal logic propositions as types, proofs as functional reactive programs. In *Proceedings of the sixth workshop on Programming Languages meets Program Verification, PLPV 2012, Philadelphia, PA, USA, January 24, 2012*, Koen Claessen and Nikhil Swamy (Eds.). ACM, Philadelphia, PA, USA, 49–60. <https://doi.org/10.1145/2103776.2103783>
- Alan Jeffrey. 2014. Functional Reactive Types. In *Proceedings of the Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS) (Vienna, Austria) (CSL-LICS '14)*. ACM, New York, NY, USA, Article 54, 9 pages. <https://doi.org/10.1145/2603088.2603106>
- Wolfgang Jeltsch. 2013. Temporal Logic with "Until", Functional Reactive Programming with Processes, and Concrete Process Categories. In *Proceedings of the 7th Workshop on Programming Languages Meets Program Verification (Rome, Italy) (PLPV '13)*. ACM, New York, NY, USA, 69–78. <https://doi.org/10.1145/2428116.2428128>
- G. A. Kavvos. 2019. Modalities, Cohesion, and Information Flow. *Proc. ACM Program. Lang.* 3, POPL (Jan. 2019), 20:1–20:29. <https://doi.org/10.1145/3290333> 00000.
- Neelakantan R. Krishnaswami. 2013. Higher-order Functional Reactive Programming Without Spacetime Leaks. In *Proceedings of the 18th ACM SIGPLAN International Conference on Functional Programming (ICFP '13)*. ACM, Boston, Massachusetts, USA, 221–232. <https://doi.org/10.1145/2500365.2500588>
- Neelakantan R. Krishnaswami and Nick Benton. 2011. Ultrametric Semantics of Reactive Programs. In *2011 IEEE 26th Annual Symposium on Logic in Computer Science*. IEEE Computer Society, Washington, DC, USA, 257–266. <https://doi.org/10.1109/LICS.2011.38>
- Neelakantan R. Krishnaswami, Nick Benton, and Jan Hoffmann. 2012. Higher-order functional reactive programming in bounded space. In *Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2012, Philadelphia, Pennsylvania, USA, January 22-28, 2012*, John Field and Michael Hicks (Eds.). ACM, Philadelphia, PA, USA, 45–58. <https://doi.org/10.1145/2103656.2103665>

- Bassel Manna and Rasmus Ejlers Møgelberg. 2018. The Clocks They Are Adjunctions: Denotational Semantics for Clocked Type Theory. In *3rd International Conference on Formal Structures for Computation and Deduction, FSCD 2018, July 9-12, 2018, Oxford, UK*. New York, NY, USA, 23:1–23:17. <https://doi.org/10.4230/LIPIcs.FSCD.2018.23>
- Hiroshi Nakano. 2000. A modality for recursion. In *Proceedings Fifteenth Annual IEEE Symposium on Logic in Computer Science (Cat. No. 99CB36332)*. IEEE Computer Society, Washington, DC, USA, 255–266. <https://doi.org/10.1109/LICS.2000.855774>
- Henrik Nilsson, Antony Courtney, and John Peterson. 2002. Functional Reactive Programming, Continued. In *Proceedings of the 2002 ACM SIGPLAN Workshop on Haskell (Pittsburgh, Pennsylvania) (Haskell '02)*. ACM, New York, NY, USA, 51–64. <https://doi.org/10.1145/581690.581695>
- Ross Paterson. 2001. A new notation for arrows. *ACM SIGPLAN Notices* 36, 10 (Oct. 2001), 229–240. <https://doi.org/10.1145/507669.507664> 00234.
- Atze van der Ploeg and Koen Claessen. 2015. Practical principled FRP: forget the past, change the future, FRPNow!. In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming (ICFP 2015)*. Association for Computing Machinery, Vancouver, BC, Canada, 302–314. <https://doi.org/10.1145/2784731.2784752> 00019.
- Amir Pnueli. 1977. The Temporal Logic of Programs. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science*. IEEE Computer Society, Washington, DC, USA, 46–57. <https://doi.org/10.1109/SFCS.1977.32>

A MULTI-TICK CALCULUS

In this appendix, we give the proof of [Theorem 4.1](#), i.e. we show that the program transformation described in [Section 4.2.1](#) indeed transforms any closed λ_{\checkmark} term into a closed λ_{\surd} term.

Figure 12 gives the context formation rules of λ_{\surd} ; the only difference is the rule for adding ticks, which has a side condition so that there may be no more than one tick. Figure 13 lists the full set of typing rules of λ_{\surd} . Compared to λ_{\checkmark} (cf. Figure 4), the only rules that have changed are the rule for lambda abstraction, and the rule for delay. Both rules transform the context Γ to $|\Gamma|$, which removes the \surd in Γ if it has one:

$$\begin{aligned} |\Gamma| &= \Gamma \quad \text{if } \Gamma \text{ is tick-free} \\ |\Gamma, \surd, \Gamma'| &= \Gamma^{\square}, \Gamma' \end{aligned}$$

We define the rewrite relation \longrightarrow as the least relation that is closed under congruence and the following rules:

$$\begin{aligned} \text{delay}(C[\text{adv } t]) &\longrightarrow \text{let } x = t \text{ in } \text{delay}(C[\text{adv } x]) \quad \text{if } t \text{ is not a variable} \\ \lambda x. C[\text{adv } t] &\longrightarrow \text{let } y = \text{adv } t \text{ in } \lambda x. (C[y]) \end{aligned}$$

where C is a term with a single occurrence of a hole $[\]$ that is not in the scope of `delay`, `adv`, `box`, `fix`, or a lambda abstraction. Formally, C is generated by the following grammar.

$$\begin{aligned} C ::= & [\] \mid C t \mid t C \mid \text{let } x = C \text{ in } t \mid \text{let } x = t \text{ in } C \mid \langle C, t \rangle \mid \langle t, C \rangle \mid \text{in}_1 C \mid \text{in}_2 C \mid \pi_1 C \mid \pi_2 C \\ & \mid C + t \mid t + C \mid \text{case } C \text{ of } \text{in}_1 x.s; \text{in}_2 x.t \mid \text{case } s \text{ of } \text{in}_1 x.C; \text{in}_2 x.t \mid \text{case } s \text{ of } \text{in}_1 x.t; \text{in}_2 x.C \\ & \mid \text{into } C \mid \text{out } C \mid \text{unbox } C \end{aligned}$$

We write $C[t]$ to substitute the unique hole $[\]$ in C with the term t .

In the following, we show that for each $\vdash_{\checkmark} t : A$, if we exhaustively apply the above rewrite rules to t , we obtain a term $\vdash_{\surd} t' : A$. We prove this by proving each of the following properties in turn:

- (1) Subject reduction: If $\Gamma \vdash_{\checkmark} s : A$ and $s \longrightarrow t$, then $\Gamma \vdash_{\surd} t : A$.
- (2) Exhaustiveness: If t is a normal form for \longrightarrow , then $\vdash_{\checkmark} t : A$ implies $\vdash_{\surd} t : A$.
- (3) Strong normalisation: There is no infinite \longrightarrow -reduction sequence.

A.1 Subject reduction

We first show subject reduction (cf. [Proposition A.4](#) below). To this end, we need a number of lemmas:

LEMMA A.1 (WEAKENING). *Let $\Gamma_1, \Gamma_2 \vdash_{\checkmark} t : A$ and Γ tick-free. Then $\Gamma_1, \Gamma, \Gamma_2 \vdash_{\checkmark} t : A$.*

PROOF. By straightforward induction on $\Gamma_1, \Gamma_2 \vdash_{\checkmark} t : A$. □

LEMMA A.2. *Given $\Gamma, \surd, \Gamma' \vdash_{\checkmark} C[\text{adv } t] : A$ with Γ' tick-free, then there is some type B such that $\Gamma \vdash_{\checkmark} t : \bigcirc B$ and $\Gamma, x : \bigcirc B, \surd, \Gamma' \vdash_{\checkmark} C[\text{adv } x] : A$.*

PROOF. We proceed by induction on the structure of C .

- $[\]$: $\Gamma, \surd, \Gamma' \vdash_{\checkmark} \text{adv } t : A$ and Γ' tick-free implies that $\Gamma \vdash_{\checkmark} t : \bigcirc A$. Moreover, given a fresh variable x , we have that $\Gamma, x : \bigcirc A \vdash_{\checkmark} x : \bigcirc A$, and thus $\Gamma, x : \bigcirc A, \surd, \Gamma' \vdash_{\checkmark} \text{adv } x : A$ and Γ' .
- Cs : $\Gamma, \surd, \Gamma' \vdash_{\checkmark} C[\text{adv } t] s : A$ implies that there is some A' with $\Gamma, \surd, \Gamma' \vdash_{\checkmark} s : A'$ and $\Gamma, \surd, \Gamma' \vdash_{\checkmark} C[\text{adv } t] : A' \rightarrow A$. By induction hypothesis, the latter implies that there is some B with $\Gamma \vdash_{\checkmark} t : \bigcirc B$ and $\Gamma, x : \bigcirc B, \surd, \Gamma' \vdash_{\checkmark} C[\text{adv } x] : A' \rightarrow A$. Hence, $\Gamma, x : \bigcirc B, \surd, \Gamma' \vdash_{\checkmark} s : A'$, by [Lemma A.1](#), and thus $\Gamma, x : \bigcirc B, \surd, \Gamma' \vdash_{\checkmark} C[\text{adv } x] s : A$.

$$\frac{}{\emptyset \vdash_{\checkmark}} \quad \frac{\Gamma \vdash_{\checkmark}}{\Gamma, x : A \vdash_{\checkmark}} \quad \frac{\Gamma \vdash_{\checkmark} \quad \Gamma \text{ tick-free}}{\Gamma, \checkmark \vdash_{\checkmark}}$$

 Fig. 12. Well-formed contexts of λ_{\checkmark} .

$$\frac{\Gamma, x : A, \Gamma' \vdash_{\checkmark} \quad \Gamma' \text{ tick-free or } A \text{ stable}}{\Gamma, x : A, \Gamma' \vdash_{\checkmark} x : A} \quad \frac{\Gamma \vdash_{\checkmark}}{\Gamma \vdash_{\checkmark} \langle \rangle : 1} \quad \frac{n \in \mathbb{Z}}{\Gamma \vdash_{\checkmark} \bar{n} : \text{Int}}$$

$$\frac{\Gamma \vdash_{\checkmark} s : \text{Int} \quad \Gamma \vdash_{\checkmark} t : \text{Int}}{\Gamma \vdash_{\checkmark} s + t : \text{Int}} \quad \frac{|\Gamma|, x : A \vdash_{\checkmark} t : B}{\Gamma \vdash_{\checkmark} \lambda x. t : A \rightarrow B} \quad \frac{\Gamma \vdash_{\checkmark} s : A \quad \Gamma, x : A \vdash_{\checkmark} t : B}{\Gamma \vdash_{\checkmark} \text{let } x = s \text{ in } t : B}$$

$$\frac{\Gamma \vdash_{\checkmark} t : A \rightarrow B \quad \Gamma \vdash_{\checkmark} t' : A}{\Gamma \vdash_{\checkmark} t t' : B} \quad \frac{\Gamma \vdash_{\checkmark} t : A \quad \Gamma \vdash_{\checkmark} t' : B}{\Gamma \vdash_{\checkmark} \langle t, t' \rangle : A \times B} \quad \frac{\Gamma \vdash_{\checkmark} t : A_1 \times A_2 \quad i \in \{1, 2\}}{\Gamma \vdash_{\checkmark} \pi_i t : A_i}$$

$$\frac{\Gamma \vdash_{\checkmark} t : A_i \quad i \in \{1, 2\}}{\Gamma \vdash_{\checkmark} \text{in}_i t : A_1 + A_2} \quad \frac{\Gamma, x : A_i \vdash_{\checkmark} t_i : B \quad \Gamma \vdash_{\checkmark} t : A_1 + A_2 \quad i \in \{1, 2\}}{\Gamma \vdash_{\checkmark} \text{case } t \text{ of in}_1 x. t_1; \text{in}_2 x. t_2 : B}$$

$$\frac{|\Gamma|, \checkmark \vdash_{\checkmark} t : A}{\Gamma \vdash_{\checkmark} \text{delay } t : \text{O}A} \quad \frac{\Gamma \vdash_{\checkmark} t : \text{O}A \quad \Gamma' \text{ tick-free}}{\Gamma, \checkmark, \Gamma' \vdash_{\checkmark} \text{adv } t : A} \quad \frac{\Gamma \vdash_{\checkmark} t : \square A}{\Gamma \vdash_{\checkmark} \text{unbox } t : A} \quad \frac{\Gamma^{\square} \vdash_{\checkmark} t : A}{\Gamma \vdash_{\checkmark} \text{box } t : \square A}$$

$$\frac{\Gamma \vdash_{\checkmark} t : A[\text{O}(\text{Fix } \alpha. A)/\alpha]}{\Gamma \vdash_{\checkmark} \text{into } t : \text{Fix } \alpha. A} \quad \frac{\Gamma \vdash_{\checkmark} t : \text{Fix } \alpha. A}{\Gamma \vdash_{\checkmark} \text{out } t : A[\text{O}(\text{Fix } \alpha. A)/\alpha]} \quad \frac{\Gamma^{\square}, x : \square(\text{O}A) \vdash_{\checkmark} t : A}{\Gamma \vdash_{\checkmark} \text{fix } x. t : A}$$

 Fig. 13. Typing rules of λ_{\checkmark} .

- $sC: \Gamma, \checkmark, \Gamma' \not\vdash_{\checkmark} C[s \text{ adv } t] : A$ implies that there is some A' with $\Gamma, \checkmark, \Gamma' \not\vdash_{\checkmark} s : A' \rightarrow A$ and $\Gamma, \checkmark, \Gamma' \not\vdash_{\checkmark} C[\text{adv } t] : A'$. By induction hypothesis, the latter implies that there is some B with $\Gamma \not\vdash_{\checkmark} t : \text{O}B$ and $\Gamma, x : \text{O}B, \checkmark, \Gamma' \not\vdash_{\checkmark} C[\text{adv } x] : A'$. Hence, $\Gamma, x : \text{O}B, \checkmark, \Gamma' \not\vdash_{\checkmark} s : A' \rightarrow A$, by Lemma A.1, and thus $\Gamma, x : \text{O}B, \checkmark, \Gamma' \not\vdash_{\checkmark} s C[\text{adv } x] : A$.
- $\text{let } y = s \text{ in } C: \Gamma, \checkmark, \Gamma' \not\vdash_{\checkmark} \text{let } y = s \text{ in } C[\text{adv } t] : A$ implies that there is some A' with $\Gamma, \checkmark, \Gamma' \not\vdash_{\checkmark} s : A'$ and $\Gamma, \checkmark, \Gamma', y : A' \not\vdash_{\checkmark} C[\text{adv } t] : A$. By induction hypothesis, the latter implies that there is some B with $\Gamma \not\vdash_{\checkmark} t : \text{O}B$ and $\Gamma, x : \text{O}B, \checkmark, \Gamma', y : A' \not\vdash_{\checkmark} C[\text{adv } x] : A$. Hence, $\Gamma, x : \text{O}B, \checkmark, \Gamma' \not\vdash_{\checkmark} s : A'$, by Lemma A.1, and thus $\Gamma, x : \text{O}B, \checkmark, \Gamma' \not\vdash_{\checkmark} \text{let } y = s \text{ in } C[\text{adv } x] : A$.
- $\text{let } y = C \text{ in } s: \Gamma, \checkmark, \Gamma' \not\vdash_{\checkmark} \text{let } y = C[\text{adv } t] \text{ in } s : A$ implies that there is some A' with $\Gamma, \checkmark, \Gamma', y : A' \not\vdash_{\checkmark} s : A$ and $\Gamma, \checkmark, \Gamma' \not\vdash_{\checkmark} C[\text{adv } t] : A'$. By induction hypothesis, the latter implies that there is some B with $\Gamma \not\vdash_{\checkmark} t : \text{O}B$ and $\Gamma, x : \text{O}B, \checkmark, \Gamma' \not\vdash_{\checkmark} C[\text{adv } x] : A'$. Hence, $\Gamma, x : \text{O}B, \checkmark, \Gamma', y : A' \not\vdash_{\checkmark} s : A$, by Lemma A.1, and thus $\Gamma, x : \text{O}B, \checkmark, \Gamma' \not\vdash_{\checkmark} \text{let } y = C[\text{adv } x] \text{ in } s : A$.

The remaining cases follow by induction hypothesis and Lemma A.1 in a manner similar to the cases above. \square

LEMMA A.3. *Let $\Gamma, \Gamma' \not\vdash_{\checkmark} C[\text{adv } t] : A$ and Γ' tick-free. Then there is some type B such that $\Gamma \not\vdash_{\checkmark} \text{adv } t : B$ and $\Gamma, x : B, \Gamma' \not\vdash_{\checkmark} C[x] : A$.*

PROOF. We proceed by induction on the structure of C :

- $\llbracket \cdot \rrbracket$: $\Gamma, \Gamma' \vdash_{\not\sim} \text{adv } t : A$ and Γ' tick-free implies that there must be Γ_1 and Γ_2 such that Γ_2 is tick-free, $\Gamma = \Gamma_1, \checkmark, \Gamma_2$, and $\Gamma_1 \vdash_{\not\sim} t : \bigcirc A$. Hence, $\Gamma_1, \checkmark, \Gamma_2 \vdash_{\not\sim} \text{adv } t : A$. Moreover, $\Gamma, x : A, \Gamma' \vdash_{\not\sim} x : A$ follows immediately by the variable introduction rule.
- $\underline{C}s$: $\Gamma, \Gamma' \vdash_{\not\sim} C[\text{adv } t] s : A$ implies that there is some A' with $\Gamma, \Gamma' \vdash_{\not\sim} s : A'$ and $\Gamma, \Gamma' \vdash_{\not\sim} C[\text{adv } t] : A' \rightarrow A$. By induction hypothesis, the latter implies that there is some B with $\Gamma \vdash_{\not\sim} \text{adv } t : B$ and $\Gamma, x : B, \Gamma' \vdash_{\not\sim} C[x] : A' \rightarrow A$. Hence, $\Gamma, x : B, \Gamma' \vdash_{\not\sim} s : A'$, by Lemma A.1, and thus $\Gamma, x : B, \Gamma' \vdash_{\not\sim} C[x] s : A$.
- $\text{let } y = s \text{ in } C$: $\Gamma, \Gamma' \vdash_{\not\sim} \text{let } y = s \text{ in } C[\text{adv } t] : A$ implies that there is some A' with $\Gamma, \Gamma' \vdash_{\not\sim} s : A'$ and $\Gamma, \Gamma', y : A' \vdash_{\not\sim} C[\text{adv } t] : A$. By induction hypothesis, the latter implies that there is some B with $\Gamma \vdash_{\not\sim} t : B$ and $\Gamma, x : B, \Gamma', y : A' \vdash_{\not\sim} C[x] : A$. Hence, $\Gamma, x : B, \Gamma' \vdash_{\not\sim} s : A'$, by Lemma A.1, and thus $\Gamma, x : B, \Gamma' \vdash_{\not\sim} \text{let } y = s \text{ in } C[\text{adv } x] : A$.

The remaining cases follow by induction hypothesis and Lemma A.1 in a manner similar to the cases above. \square

PROPOSITION A.4 (SUBJECT REDUCTION). *If $\Gamma \vdash_{\not\sim} s : A$ and $s \rightarrow t$, then $\Gamma \vdash_{\not\sim} t : A$.*

PROOF. We proceed by induction on $s \rightarrow t$.

- Let $s \rightarrow t$ be due to congruence closure. Then $\Gamma \vdash_{\not\sim} t : A$ follows by the induction hypothesis. For example, if $s = s_1 s_2, t = t_1 s_2$ and $s_1 \rightarrow t_1$, then we know that $\Gamma \vdash_{\not\sim} s_1 : B \rightarrow A$ and $\Gamma \vdash_{\not\sim} s_2 : B$ for some type B . By induction hypothesis, we then have that $\Gamma \vdash_{\not\sim} t_1 : B \rightarrow A$ and thus $\Gamma \vdash_{\not\sim} t : A$.
- Let $\text{delay}(C[\text{adv } t]) \rightarrow \text{let } x = t \text{ in } \text{delay}(C[\text{adv } x])$ and $\Gamma \vdash_{\not\sim} \text{delay}(C[\text{adv } t]) : A$. That is, $A = \bigcirc A'$ and $\Gamma, \checkmark \vdash_{\not\sim} C[\text{adv } t] : A'$. Then by Lemma A.2, we obtain some type B such that $\Gamma \vdash_{\not\sim} t : \bigcirc B$ and $\Gamma, x : \bigcirc B, \checkmark \vdash_{\not\sim} C[\text{adv } x] : A'$. Hence, $\Gamma \vdash_{\not\sim} \text{let } x = t \text{ in } \text{delay}(C[\text{adv } x]) : A$.
- Let $\lambda x. C[\text{adv } t] \rightarrow \text{let } y = \text{adv } t \text{ in } \lambda x. (C[y])$ and $\Gamma \vdash_{\not\sim} \lambda x. C[\text{adv } t] : A$. Hence, $A = A_1 \rightarrow A_2$ and $\Gamma, x : A_1 \vdash_{\not\sim} C[\text{adv } t] : A_2$. Then, by Lemma A.3, there some type B such that $\Gamma \vdash_{\not\sim} \text{adv } t : B$ and $\Gamma, y : B, x : A_1 \vdash_{\not\sim} C[y] : A_2$. Hence, $\Gamma \vdash_{\not\sim} \text{let } y = \text{adv } t \text{ in } \lambda x. (C[y]) : A$.

\square

A.2 Exhaustiveness

Secondly, we show that any closed $\lambda_{\not\sim}$ term that cannot be rewritten any further is also a closed λ_{\checkmark} term (cf. Proposition A.9 below).

Definition A.5.

- We say that a term t is *weakly adv-free* iff whenever $t = C[\text{adv } s]$ for some C and s , then s is a variable.
- We say that a term t is *strictly adv-free* iff there are no C and s such that $t = C[\text{adv } s]$.

Clearly, any strictly adv-free term is also weakly adv-free.

In the following we use the notation $t \twoheadrightarrow$ to denote the fact that there is no term t' with $t \rightarrow t'$; in other words, t is a normal form.

LEMMA A.6.

- If $\text{delay } t \twoheadrightarrow$, then t is weakly adv-free.
- If $\lambda x. t \twoheadrightarrow$, then t is strictly adv-free.

PROOF. Immediate, by the definition of weakly/strictly adv-free and \twoheadrightarrow . \square

LEMMA A.7. *Let Γ' be not tick-free, $\Gamma, \checkmark, \Gamma' \vdash_{\not\sim} t : A, t \twoheadrightarrow$, and t weakly adv-free. Then $\Gamma^{\square}, \Gamma' \vdash_{\not\sim} t : A$*

PROOF. We proceed by induction on $\Gamma, \surd, \Gamma' \vdash_{\ll} t : A$:

- $\Gamma, \surd, \Gamma' \vdash_{\ll} \text{adv } t : A$: Then there are Γ_1, Γ_2 such that Γ_2 tick-free, $\Gamma' = \Gamma_1, \surd, \Gamma_2$, and $\Gamma, \surd, \Gamma_1 \vdash_{\ll} t : \bigcirc A$. Since $\text{adv } t$ is by assumption weakly adv-free, we know that t is some variable x . Since $\bigcirc A$ is not stable we thus know that $x : \bigcirc A \in \Gamma_1$. Hence, $\Gamma^\square, \Gamma_1 \vdash_{\ll} t : \bigcirc A$, and therefore $\Gamma^\square, \Gamma' \vdash_{\ll} \text{adv } t : A$.
- $\Gamma, \surd, \Gamma' \vdash_{\ll} \text{delay } t : \bigcirc A$: Hence, $\Gamma, \surd, \Gamma', \surd \vdash_{\ll} t : A$. Moreover, since $\text{delay } t \dashrightarrow$, we have by Lemma A.6 that t is weakly adv-free. We may thus apply the induction hypothesis to obtain that $\Gamma^\square, \Gamma', \surd \vdash_{\ll} t : A$. Hence, $\Gamma^\square, \Gamma' \vdash_{\ll} \text{delay } t : \bigcirc A$.
- $\Gamma, \surd, \Gamma' \vdash_{\ll} \text{box } t : \square A$: Hence, $(\Gamma, \surd, \Gamma')^\square \vdash_{\ll} t : A$, which is the same as $(\Gamma^\square, \Gamma')^\square \vdash_{\ll} t : A$. Hence, $\Gamma^\square, \Gamma' \vdash_{\ll} \text{box } t : \square A$.
- $\Gamma, \surd, \Gamma' \vdash_{\ll} \lambda x.t : A \rightarrow B$: That is, $\Gamma, \surd, \Gamma', x : A \vdash_{\ll} t : B$. Since, by assumption $\lambda x.t \dashrightarrow$, we know by Lemma A.6 that t is strictly adv-free, and thus also weakly adv-free. Hence, we may apply the induction hypothesis to obtain that $\Gamma^\square, \Gamma', x : A \vdash_{\ll} t : B$, which in turn implies $\Gamma^\square, \Gamma' \vdash_{\ll} \lambda x.t : A \rightarrow B$.
- $\Gamma, \surd, \Gamma' \vdash_{\ll} \text{fix } x.t : A$: Hence, $(\Gamma, \surd, \Gamma')^\square, x : \square \bigcirc A \vdash_{\ll} t : A$, which is the same as $(\Gamma^\square, \Gamma')^\square, x : \square \bigcirc A \vdash_{\ll} t : A$. Hence, $\Gamma^\square, \Gamma' \vdash_{\ll} \text{fix } x.t : A$.
- $\Gamma, \surd, \Gamma' \vdash_{\ll} x : A$: Then, either $\Gamma' \vdash_{\ll} x : A$ or $\Gamma^\square \vdash_{\ll} x : A$. In either case, $\Gamma^\square, \Gamma' \vdash_{\ll} x : A$ follows.
- The remaining cases follow by the induction hypothesis in a straightforward manner. For example, if $\Gamma, \surd, \Gamma' \vdash_{\ll} s t : A$, then there is some type B with $\Gamma, \surd, \Gamma' \vdash_{\ll} s : B \rightarrow A$ and $\Gamma, \surd, \Gamma' \vdash_{\ll} t : B$. Since $s t$ is weakly adv-free, so are s and t , and we may apply the induction hypothesis to obtain that $\Gamma^\square, \Gamma' \vdash_{\ll} s : B \rightarrow A$ and $\Gamma^\square, \Gamma' \vdash_{\ll} t : B$. Hence, $\Gamma^\square, \Gamma' \vdash_{\ll} s t : A$.

□

LEMMA A.8. *Let $\Gamma, \surd, \Gamma' \vdash_{\ll} t : A, t \dashrightarrow$ and t strictly adv-free. Then $\Gamma^\square, \Gamma' \vdash_{\ll} t : A$. (Note that this Lemma is about λ_{\ll} .)*

PROOF. We proceed by induction on $\Gamma, \surd, \Gamma' \vdash_{\ll} t : A$.

- $\Gamma, \surd, \Gamma' \vdash_{\ll} \text{adv } t : A$: Impossible since $\text{adv } t$ is not strictly adv-free.
- $\Gamma, \surd, \Gamma' \vdash_{\ll} \text{delay } t : \bigcirc A$: Hence, $\Gamma^\square, \Gamma', \surd \vdash_{\ll} t : A$, which in turn implies that $\Gamma^\square, \Gamma' \vdash_{\ll} \text{delay } t : \bigcirc A$.
- $\Gamma, \surd, \Gamma' \vdash_{\ll} \text{box } t : \square A$: Hence, $(\Gamma, \surd, \Gamma')^\square \vdash_{\ll} t : A$, which is the same as $(\Gamma^\square, \Gamma')^\square \vdash_{\ll} t : A$. Hence, $\Gamma^\square, \Gamma' \vdash_{\ll} \text{box } t : \square A$.
- $\Gamma, \surd, \Gamma' \vdash_{\ll} \lambda x.t : A \rightarrow B$: That is, $\Gamma, \surd, \Gamma', x : A \vdash_{\ll} t : B$. Since, by assumption $\lambda x.t \dashrightarrow$, we know by Lemma A.6 that t is strictly adv-free. Hence, we may apply the induction hypothesis to obtain that $\Gamma^\square, \Gamma', x : A \vdash_{\ll} t : B$, which in turn implies $\Gamma^\square, \Gamma' \vdash_{\ll} \lambda x.t : A \rightarrow B$.
- $\Gamma, \surd, \Gamma' \vdash_{\ll} \text{fix } x.t : A$: Hence, $(\Gamma, \surd, \Gamma')^\square, x : \square \bigcirc A \vdash_{\ll} t : A$, which is the same as $(\Gamma^\square, \Gamma')^\square, x : \square \bigcirc A \vdash_{\ll} t : A$. Hence, $\Gamma^\square, \Gamma' \vdash_{\ll} \text{fix } x.t : A$.
- $\Gamma, \surd, \Gamma' \vdash_{\ll} x : A$: Then, either $\Gamma' \vdash_{\ll} x : A$ or $\Gamma^\square \vdash_{\ll} x : A$. In either case, $\Gamma^\square, \Gamma' \vdash_{\ll} x : A$ follows.
- The remaining cases follow by the induction hypothesis in a straightforward manner. For example, if $\Gamma, \surd, \Gamma' \vdash_{\ll} s t : A$, then there is some type B with $\Gamma, \surd, \Gamma' \vdash_{\ll} s : B \rightarrow A$ and $\Gamma, \surd, \Gamma' \vdash_{\ll} t : B$. Since $s t$ is strictly adv-free, so are s and t , and we may apply the induction hypothesis to obtain that $\Gamma^\square, \Gamma' \vdash_{\ll} s : B \rightarrow A$ and $\Gamma^\square, \Gamma' \vdash_{\ll} t : B$. Hence, $\Gamma^\square, \Gamma' \vdash_{\ll} s t : A$. □

In order for the induction argument to go through we generalise the exhaustiveness property from closed λ_{\ll} terms to open λ_{\ll} terms in a context with at most one tick.

PROPOSITION A.9 (EXHAUSTIVENESS). *Let $\Gamma \vdash_{\ll}, \Gamma \vdash_{\ll} t : A$ and $t \dashrightarrow$. Then $\Gamma \vdash_{\ll} t : A$.*

PROOF. We proceed by induction on the structure of t and by case distinction on the last typing rule in the derivation of $\Gamma \vdash_{\not\sim} t : A$.

- $$\frac{\Gamma, \surd \vdash_{\not\sim} t : A}{\Gamma \vdash_{\not\sim} \text{delay } t : \bigcirc A}$$
- $\Gamma \vdash_{\not\sim} \text{delay } t : \bigcirc A$:
We consider two cases:
 - Γ is tick-free: Hence, $\Gamma, \surd \vdash_{\surd}$ and thus $\Gamma, \surd \vdash_{\surd} t : A$ by induction hypothesis. Since $|\Gamma| = \Gamma$, we thus have that $\Gamma \vdash_{\surd} \text{delay } t : \bigcirc A$.
 - $\Gamma = \Gamma_1, \surd, \Gamma_2$ and Γ_2 tick-free: By definition, $t \dashrightarrow$ and, by Lemma A.6, t is weakly adv-free. Hence, by Lemma A.7, we have that $\Gamma_1^\square, \Gamma_2, \surd \vdash_{\not\sim} t : A$. We can thus apply the induction hypothesis to obtain that $\Gamma_1^\square, \Gamma_2, \surd \vdash_{\surd} t : A$. Since $|\Gamma| = \Gamma_1^\square, \Gamma_2$, we thus have that $\Gamma \vdash_{\surd} \text{delay } t : \bigcirc A$.
 - $\Gamma, x : A \vdash_{\not\sim} t : B$

$$\frac{\Gamma, x : A \vdash_{\not\sim} t : B}{\Gamma \vdash_{\not\sim} \lambda x. t : A \rightarrow B}$$
 By induction hypothesis, we have that $\Gamma, x : A \vdash_{\surd} t : B$. There are two cases to consider:
 - Γ is tick-free: Then $|\Gamma| = \Gamma$ and we thus obtain that $\Gamma \vdash_{\surd} \lambda x. t : A \rightarrow B$.
 - $\Gamma = \Gamma_1, \surd, \Gamma_2$ with Γ_1, Γ_2 tick-free: Since $t \dashrightarrow$ by definition and t strictly adv-free by Lemma A.6, we may apply Lemma A.8 to obtain that $\Gamma_1^\square, \Gamma_2, x : A \vdash_{\surd} t : B$. Since $|\Gamma| = \Gamma_1^\square, \Gamma_2$, we thus have that $\Gamma \vdash_{\surd} \lambda x. t : A \rightarrow B$.
 - All remaining cases follow immediately from the induction hypothesis, because all other rules are either the same for both calculi or the λ_{\surd} typing rule has an additional side condition that Γ have at most one tick, which holds by assumption. \square

A.3 Strong Normalisation

PROPOSITION A.10 (STRONG NORMALISATION). *The rewrite relation \longrightarrow is strongly normalising.*

PROOF. To show that \longrightarrow is strongly normalising, we define for each term t a natural number $d(t)$ such that, whenever $t \longrightarrow t'$, then $d(t) > d(t')$. A *redex* is a term of the form $\text{delay } C[\text{adv } t]$ with t not a variable, or a term of the form $\lambda x. C[\text{adv } s]$. For each redex occurrence in a term t , we can calculate its *depth* as the length of the unique path that goes from the root of the abstract syntax tree of t to the occurrence of the redex. Define $d(t)$ as the sum of the depth of all redex occurrences in t . Since each rewrite step $t \longrightarrow t'$ removes a redex or replaces a redex with a new redex at a strictly smaller depth, we have that $d(t) > d(t')$. \square

THEOREM A.11. *For each $\vdash_{\not\sim} t : A$, we can effectively construct a term t' with $t \longrightarrow^* t'$ and $\vdash_{\surd} t' : A$.*

PROOF. We construct t' from t by repeatedly applying the rewrite rules of \longrightarrow . By Proposition A.10 this procedure will terminate and we obtain a term t' with $t \longrightarrow^* t'$ and $t' \dashrightarrow$. According to Proposition A.4, this means that $\vdash_{\not\sim} t' : A$, which in turn implies by Proposition A.9 that $\vdash_{\surd} t' : A$. \square