

The Essence of Live Coding: Change the Program, Keep the State!

Manuel Bärenz
sonnen eServices GmbH
Germany

Abstract

One rarely encounters programming languages and frameworks that provide general-purpose and type-safe hot code swap. It is demonstrated here that this is entirely possible in Haskell, by faithfully following the motto of live coding: “Change the program, keep the state.”

With generic programming, one easily arrives at an automatic state migration function. The approach can be generalised to an arrowized Functional Reactive Programming framework that is parametrized by its side effects. It allows for building up complete live programs from reusable, modular components, and to separate data flow cleanly from control flow. Useful utilities for debugging and quickchecking are presented.

CCS Concepts: • **Theory of computation** → *Formal languages and automata theory*; • **Computing methodologies** → *Animation*; • **Computer systems organization** → *Data flow architectures*; • **Software and its engineering** → *State systems*; **Functional languages**; *Data types and structures*.

Keywords: Livecoding, Functional Reactive Programming

ACM Reference Format:

Manuel Bärenz. 2020. The Essence of Live Coding: Change the Program, Keep the State!. In *Proceedings of the 7th ACM SIGPLAN International Workshop on Reactive and Event-Based Languages and Systems (REBLS '20)*, November 16, 2020, Virtual, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3427763.3428312>

1 Introduction

Live coding has come to denote various concepts, from hot code swap on a production server to artistic performances with code that is written, executed and updated live. From the implementor’s perspective, the common denominator is

a framework to update the definition of a program while not interrupting its execution.

In a dynamically typed, process-oriented language, such as Erlang [4], hot code swap is conceptually simple: A new server process is started, the state of the old server is transferred to the new one, and the old server is shut down. The challenge lies in ensuring that the state is correctly migrated, as corrupt state might crash the server. This is a nontrivial task if no type checker is available to verify if the migrated state conforms to the schema required by the new server.

Domain specific live coding frameworks usually have a central server that generates the desired effect – be it audio, video or device communication – and offers an API through which users can create “cells”, such as oscillators or signal processors in the case of audio (such as [31]), and connect, re-order, update or destroy them during execution. The central server guarantees safe updates, and strongly typed user-side libraries (such as Tidal [20] for audio applications) exist, but the user is restricted to a domain specific language.

In dynamic languages like SmallTalk [16], new avenues to live coding open by interacting with the interpreter, which holds the whole program in memory. In a statically compiled language lacking a comparable depth of reflection, this is not directly possible, so one must abstract the program and its state into data structures, and implement hot code swap. One end of the design space are concurrent references to program components which can be exchanged, [2, 3, 24] which are somewhat closer to an interpreted environment. Let us restrict to the other end, though, where the possible data and control states are known at compile time, constraining the resulting framework slightly, but allowing for a very smooth development experience.

In this article, we implement a lightweight general purpose live coding framework in Haskell from scratch. It is not only type-safe, but also type-driven, in that boilerplate code for state migrations (which is hard to get right without a static type checker) is automatically derived from the type. It is not restricted to a particular domain, by virtue of being parametrised over an arbitrary monad: Any domain- or library-specific effect can be incorporated effortlessly, and handled with standard Haskell functions. The framework follows the tradition of monadic arrowized Functional Reactive Programming (FRP) as developed in [27] and [25]. To run live programs created in it, a runtime environment in the IO monad is supplied, but since the framework does not hide

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
REBLS '20, November 16, 2020, Virtual, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8188-8/20/11...\$15.00

<https://doi.org/10.1145/3427763.3428312>

IO in its abstraction (such as many other FRP frameworks do), it is an easy exercise to execute the live programs in e.g. the **STM** monad [12] or any other concurrency context such as an external main loop. The state of the live program can be inspected and debugged safely during the execution.

There is no subtly clever trick. We simply follow, dutifully and without compromise, the mantra of live coding:

Change the program. Keep the state.

This is not a new idea in itself. Hot code swap in Erlang realises this motto, and similar views are expressed about live coding in Elm¹ (a web frontend DSL inspired by Haskell). What is new about this work is the consequential application of this motto to create a general purpose, type-safe FRP framework² with *automatic state migration*.

Arriving at a simple *state migration function* by faithfully following the live coding mantra is a manageable task, carried out in Section 2. Upon the migration function, a live coding framework is constructed in Section 3. It is much more rewarding to recast this framework in the form of functional reactive programming, which allows us to reuse modular, functional components and separate data flow from control flow. Crucially, the state of our live programs is built up automatically by using FRP idioms. The result is presented in Section 4, which heavily draws inspiration from Dunai, a monadic arrowized FRP framework; and from Caspi's and Pouzet's work on synchronous stream functions [6]. After having implemented the data flow aspects of our framework, we turn to control flow in Section 5, and encode it completely algebraically within the program state, automatically getting a grip on exception handling and loops. A monadic interface to our live programs is presented. In Section 6, several useful tools such as debuggers and quickcheck utilities are shown.

This article is written in literate Haskell and supplies the library presented here. The source code is available at <https://github.com/turion/essence-of-live-coding>, while additional resources such as a presentation can be found at <https://www.manuelbaerenz.de/#computerscience>.

2 Change the program. Keep the state (as far as possible).

Our model of a live program will consist of a state and an effectful state transition function. A preliminary version is shown in Figure 1³. The program is initialised at a certain state, and from there its behaviour is defined by repeatedly applying the function `liveStep` to advance the state and produce effects. This is implemented in `stepProgram`. Since we

```
data LiveProgram m s = LiveProgram
  { liveState :: s
  , liveStep  :: s -> m s
  }

stepProgram
  :: Monad m
  => LiveProgram m s -> m (LiveProgram m s)
stepProgram liveProgram@LiveProgram { .. } = do
  liveState' <- liveStep liveState
  return liveProgram { liveState = liveState' }

stepProgramMVar
  :: MVar (LiveProgram IO s)
  -> IO ()
stepProgramMVar var = do
  currentProgram <- takeMVar var
  nextProgram <- stepProgram currentProgram
  putMVar var nextProgram

launch
  :: LiveProgram IO s
  -> IO (MVar (LiveProgram IO s))
launch liveProgram = do
  var <- newMVar liveProgram
  forkIO $ forever $ stepProgramMVar var
  return var
```

Figure 1. LiveProgramPreliminary.lhs

want to run the program in a separate thread while compiling a new version of the program in the foreground, we have to store the program in a concurrent variable, here an **MVar**. Given this variable, stepping the program it contains is a simple **IO** action, implemented in `stepProgramMVar`. To run the program, we fork a background thread and repeatedly call `stepProgramMVar` there.

In a dynamically typed language, such a setup is in principle enough to implement hot code swap. At some point, the execution will be paused, and the function `liveStep` is simply exchanged for a new one. Then the execution is resumed with the new transition function, which operates on the old state. Of course, the new step function has to take care of migrating the state to a new format, should this be necessary. The difficulties arise (apart from the practicalities of the implementation), from the inherent unsafety of this operation: Even if the old transition function behaved correctly, and the old state is in a valid format, the new transition function may crash or otherwise misbehave on the old state. It is very hard to reduce the probability of such a failure with tests since the current state constantly changes (by design). A static typechecker is missing, to guarantee the safety of this operation.

But let us return to Haskell, where we have such a type-checker. It immediately points out the unsafety of the migration: There is no guarantee that the new transition function will typecheck with the old state! In fact, in many situations,

¹Compare <https://elm-lang.org/blog/interactive-programming>.

²It shall be remarked that FRP is long past niche applications in the video and audio domains. It is possible to write web servers and frontends, simulations and games in it.

³The notation `$` may be unfamiliar. It can be read as "apply brackets until the end of the following expression". For example, `f $ g $ h a b c` is essentially the same as `f (g (h a b c))`.

```

hotCodeSwap
  :: (s -> s')
  -> LiveProgram m s'
  -> LiveProgram m s
  -> LiveProgram m s'
hotCodeSwap migrate newProgram oldProgram
  = LiveProgram
  { liveState = migrate $ liveState oldProgram
  , liveStep  = liveStep newProgram
  }

```

Figure 2. Preliminary/HotCodeSwap.lhs

the state type needs to be modified, and there is no function of type `LiveProgram m s -> LiveProgram m s'` already because there is no function of type `s -> s'`.

This kind of problem is not unknown. In the world of databases, it is commonplace that a table lives much longer than its initial schema. The services accessing the data can change, and thus also the requirements to the data format. The solution to this problem is a *schema migration*, an update to the database schema that alters the table in such a way that as little data as possible is lost, and that the table adheres to the new schema afterwards. Data loss is not entirely preventable, though. If a column has to be deleted, its data will not be recoverable. In turn, if a column is created, one has to supply a sensible default value (often `NULL` will suffice).

2.1 Migrating the State

We can straightforwardly adopt this solution by thinking of the program state as a small database table with a single row. Its schema is the type `s`. Given a *type migration* function, we can perform hot code swap, as shown in Figure 2. We need to supply the old live program, the new live program, and a suitable migration function.

This may be an acceptable solution to perform a planned, well-prepared intervention, but it does spoil the fun in a musical live coding performance if the programmer has to write a migration function after every single edit. What a live performer actually needs, is a function with this mysterious type signature:

```

hotCodeSwap
  :: LiveProgram m s'
  -> LiveProgram m s
  -> LiveProgram m s'

```

It is the same type signature as in Figure 2, but with the first argument, the manual migration function, removed. The new program, including its initial state, has just been compiled, and the old program is still stored in a concurrent variable. Can we possibly derive the new state by simply looking at the initial state of the new program and the old state? Is there a magical, universal migration function? If there were, it would have this type:

```
migrate :: s' -> s -> s'
```

A theoretician will probably invoke a free theorem [30] here, and infer that there is in fact a unique such function: `const!` But it is not what we were hoping for. `const s'` will throw away the old state and run the program with the new initial state – effectively restarting our program from a blank slate.

In this generality, we cannot hope for any other solution. But in the following, we are going to see how to tweak the live program definition by only twenty characters, and arrive at an effective migration function.

2.2 Type-Driven Migrations

In many cases, knowing the old state and the new initial state is sufficient to derive the new, migrated state safely. As an example, imagine the internal state of a simple webserver that counts the number of visitors to a page.

```
data State = State { nVisitors :: Int }
```

The server is initialised at 0, and increments the number of visitors every step. (For a full-fledged webserver, the reader is asked to patiently wait until the next section.)

```
server = LiveProgram (State 0) $ \State { .. }
  -> State $ return $ nVisitors + 1
```

We extend the state by the name of the last user agent to access the server (initially not present):

```
data State = State Int (Maybe ByteString)
initState = State 0 Nothing
```

From just comparing the two datatype definitions, it is apparent that we would want to keep the number of visitors, of type `Int`, when migrating. For the new argument of type `Maybe ByteString`, we cannot infer any sensible value from the old state, but we can take the value `Nothing` from the new initial state, and interpret it as a default value. A general state migration function should specialise to:

```
migrate (Server1.State nVisitors)
  (Server2.State _ mUserAgent)
  = Server2.State nVisitors mUserAgent
```

Our task was less obvious if we would have extended the state by the last access time, encoded as a UNIX timestamp:

```
data State = State Int Int
```

Here it is unclear to which of the `Ints` the old value should be migrated. It is obvious again if the datatype was defined as a record as well:

```
data State = State
  { nVisitors      :: Int
  , lastAccessUNIX :: Int
  }
```

We need to copy the `nVisitors` field from the old state, and initialise the `lastAccessUNIX` field from the new state. (Conversely, if we were to migrate back to the original definition, there is no way but to lose the data stored in `lastAccessUNIX`.) Clearly, the record labels enabled us to

```

userMigrate
  :: (Data a, Data b, Typeable c, Typeable d)
  => (c -> d)
  -> a -> b -> a
intToInteger :: Int -> Integer
intToInteger = toInteger

```

Figure 3. User migration

identify the correct target field. The solution lies in the datatype definition.

We can meta-program a migration function by reasoning about the structure of the type definition. This is possible with the techniques presented in the seminal, now classic article “Scrap Your Boilerplate” [19]. It supplies a typeclass **Typeable** which enables us to compare types and safely type-cast at runtime, and a typeclass **Data** which allows us to inspect constructor names and record field labels. Using the package `syb`, which supplies common utilities when working with **Data**, our migration function is implemented in under 50 lines of code, with the following signature:

```
migrate :: (Data a, Data b) => a -> b -> a
```

It handles the two previously mentioned cases: Constructors with the same names, but some mismatching arguments, and records with some coinciding field labels, but possibly a different order. In nested datatype definitions, the function recurses into all children of the data tree. Needless to say, if the types do match, then the old state is identically copied.

Sometimes it is necessary to manually migrate some part of the state. Assume, for the sake of the example, that our webserver has become wildly popular, and `nVisitors` is close to `maxInt`. We need to migrate this value to an arbitrary precision **Integer**. It is easy to extend `migrate` by a special case provided by the user, shown in Figure 3. Here, we would use `userMigrate intToInteger` to migrate the state.

To use the automatic migration function, we only need to update the live program definition to include the **Data** constraint, as shown in Figure 4. This is a small restriction. The **Data** typeclass can be automatically derived for every algebraic data type, except those that incorporate *functions*. We have to refactor our live program such that all functions are contained in `liveStep` (and can consequently not be migrated), and all data is contained in `liveState`.

Now that we have a universal migration function, it is not necessary to carry the type of the state around in the type signature. In fact it would be cumbersome in combination with **MVars** (which can’t change their type), and a real burden when later modularising the state. Consequently, the type is made existential. The only necessary information is that it is an instance of **Data**.

```

data LiveProgram m = forall s . Data s
  => LiveProgram
  { liveState :: s
  , liveStep  :: s -> m s
  }
hotCodeSwap
  :: LiveProgram m
  -> LiveProgram m
  -> LiveProgram m
hotCodeSwap
  (LiveProgram newState newStep)
  (LiveProgram oldState _)
  = LiveProgram
  { liveState = migrate newState oldState
  , liveStep  = newStep
  }

```

Figure 4. LiveProgram.lhs

3 The Runtime

3.1 Hands on Interaction

Enough declaration. Let us get semantic and run some live programs! In the preliminary version, a function `stepProgram` implemented a single execution step, and it can be reused here, up to removing the explicit state type. The runtime behaviour of a live program is defined by calling this function repeatedly. We could of course run the program in the foreground thread:

```

foreground :: Monad m => LiveProgram m -> m ()
foreground liveProgram = do
  stepProgram liveProgram
  foreground liveProgram

```

But this would leave no possibility to exchange the program with a new one. Instead, we can store the program in an **MVar** and call `stepProgramMVar` on it. Now that we can migrate any **Data**, we can follow the original plan of exchanging the live program in mid-execution:

```

update
  :: MVar (LiveProgram IO)
  -> LiveProgram IO
  -> IO ()
update var newProg = do
  oldProg <- takeMVar var
  putMVar var $ hotCodeSwap newProg oldProg

```

The old program is retrieved from the concurrent variable, migrated to the new state, and put back for further execution. And so begins our first live coding session in GHCi (line breaks added for readability):

```

> var <- newMVar $ LiveProgram 0
  $ \s -> print s >> return (s + 1)
> stepProgramMVar var
0

```



```

> stepProgramMVar var
1
> update var $ LiveProgram 0
  $ \s -> print s >> return (s - 1)
> stepProgramMVar var
2
> stepProgramMVar var
1
> stepProgramMVar var
0

```

Upon updating, the state was correctly preserved. The programs were specified in the interactive session here, but of course we will want to load the program from a file, and use GHCi's `:reload` functionality when we have edited it. But as soon as we do this, the local binding `var` is lost. The package `foreign-store` [9] offers a remedy: `var` can be stored persistently across reloads. To facilitate its usage, GHCi macros are defined for the initialisation and reload operations.

Of course, it is not intended to enter `:liverestep` repeatedly when coding. We want to launch a separate thread which executes the steps in the background. Again, we can reuse the function `launch`. (Only the type signature needs updating.) Using `ghcid` (“GHCi as a daemon” [23]), the launching and reloading operations can be automatically triggered upon starting `ghcid` and editing the code, allowing for a smooth live coding experience without any manual intervention.

In the next subsection, a full example is shown.

3.2 Live Coding a Webserver

To show that live coding can be applied to domains outside audio and video applications, let us realise the example from the previous section and create a tiny webserver using the WAI/Warp framework [33]. It is supposed to count the number of visitors, and keep this state in memory when we change the implementation.

The boiler plate code, which is suppressed here, initialises the Warp server, uses `launch` to start our live program in a separate thread and waits for user input to update it.

To save ourselves an introduction to Warp, we will communicate to it via two **MVars**, which we need to share with the live program. The textbook solution is to supply the variables through a **Reader** environment, which needs to be supplied to the live program before execution. This can be done by transporting the program along the `runReaderT` monad morphism. A function `hoistLiveProgram` does this (borrowing nomenclature from the `mmorph` [11] package).

The server logic is shown in Figure 6. It is initialised at 0 visitors. The step function receives the number of past visitors and blocks on an **MVar** until a request (which is discarded) to the server arrives. The number of visitors is incremented by 1, and baked into a response, which is in another **MVar**. Finally, the updated state (the incremented number of visitors) is returned, and passed to the next step.

```

data Env = Env
  { requestVar  :: MVar Request
  , responseVar :: MVar String
  }

```

Figure 5. DemoWai.lhs

```

data State = State
  { nVisitors :: Integer
  } deriving Data

oldServer :: LiveProgram (ReaderT Env IO)
oldServer = LiveProgram
  { liveState = State 0
  , liveStep  = \State { .. } -> do
      Env { .. } <- ask
      _ <- lift $ takeMVar requestVar
      let nVisitorsNew = nVisitors + 1
          lift $ putMVar responseVar $ unlines
              [ "Ye Olde Server greets visitor #"
              <> show nVisitorsNew <> "."
              ]
      return $ State nVisitorsNew
  }

```

Figure 6. DemoWai1.lhs

We then modify⁴ the server logic as in Figure 7. Additionally to the number of visitors, we also store the last user agent name in the state, if it was sent. For this, one more record field is added to the state type.

Let us run the old server, and switch to the new one during execution. From a console, we access the running server:

```

$ curl localhost:8080
Ye Olde Server greets visitor #1.
$ curl localhost:8080
Fancy Nu $3rv3r says HI to #2.
$ curl localhost:8080
Fancy Nu $3rv3r says HI to #3.
Last agent: curl/7.72.0

```

It correctly remembered the number of past visitors upon reload and initialised the last user agent with the value **Nothing**. When accessing the new server again, it stored the user agent as expected.

4 Live Coding as Arrowized Functional Reactive Programming

Writing out the complete state of the live program explicitly is tedious. We have to plan the whole program in advance and artificially separate its state from the step function. Such a development approach prevents us from writing programs

⁴The function `unpack` from the `bytestring` package converts between different kinds of strings. `requestHeaders` from the `wai` package extracts HTTP headers, such as the user agent name, from a request.

```

data State = State
  { nVisitors :: Integer
  , lastAgent :: Maybe String
  } deriving Data
newServer :: LiveProgram (ReaderT Env IO)
newServer = LiveProgram
  { liveState = State 0 Nothing
  , liveStep = \State { .. } -> do
    Env { .. } <- ask
    request <- lift $ takeMVar requestVar
    let nVisitorsNew = nVisitors + 1
        lastAgentStrings = case lastAgent of
          Nothing -> []
          Just str -> ["Last agent: " <> str]
        lastAgentNew = fmap unpack
            $ lookup "User-Agent"
            $ requestHeaders request
    lift $ putMVar responseVar $ unlines $
      [ "Fancy Nu $rv3r says HI to #"
      <> show nVisitorsNew <> "."
      ] ++ lastAgentStrings
    return $ State nVisitorsNew lastAgentNew
  }

```

Figure 7. DemoWai2.lhs

in a modular fashion. The purpose of this section is to show that we can develop live programs modularly by extending the approach presented so far to an arrowized FRP framework. Coincidentally, but naturally, we will end up with a coalgebraic presentation of synchronous stream functions which differs from Caspi’s and Pouzet’s work [6] only in notation and the presence of monads.

In ordinary functional programming, the smallest building blocks are functions. It stands to reason that in live coding, they should also be some flavour of functions, in fact, **Arrows** [15]. We will see that it is possible to define bigger live programs from reusable components. Crucially, the library user is disburdened from separating state and step function. The state type is built up behind the scenes, in a manner compatible with the automatic state migration.

4.1 Cells

In our definition of live programs as pairs of state and state steppers, we can generalise the step functions to an additional input and output type. Live programs are thus generalised to effectful *Mealy machines* [21]. Let us call them cells, the building blocks of everything live:

```

data Cell m a b = forall s . Data s => Cell
  { cellState :: s
  , cellStep :: s -> a -> m (b, s)
  }

```

Such a cell may progress by one step, consuming an a as input, and producing, by means of an effect in some monad m , not only the updated cell, but also an output datum b :

```

step
  :: Monad m
  => Cell m a b
  -> a -> m (b, Cell m a b)
step Cell { .. } a = do
  (b, cellState') <- cellStep cellState a
  return (b, Cell { cellState = cellState', .. })

```

As a simple example, consider the following **Cell** which adds all input and returns the delayed sum each step:

```

sumC :: (Monad m, Num a, Data a) => Cell m a a
sumC = Cell { .. }
  where
    cellState = 0
    cellStep accum a = return (accum, accum + a)

```

We recover live programs as the special case of trivial input and output:

```

liveCell
  :: Monad m
  => Cell m () ()
  -> LiveProgram m
liveCell Cell { .. } = LiveProgram
  { liveState = cellState
  , liveStep = \state -> do
    (_, state') <- cellStep state ()
    return state'
  }

```

4.2 FRP for Automata-Based Programming

Our cells are known in the literature as “Effectful Mealy Machines”, “transducers” and “resumptions” [22], [28], [13, Section 7], [1, Section 5.4]. They are known for their relevance to stream functions [6], suggesting that they offer a wide variety of applications in FRP. The essential parts of the API, which is heavily inspired by the FRP library *dunai* [27], are shown here. We will address the data flow aspects in this section, investigating control flow later in Section 5.

Composition. By being an instance of the type class **Category**, cells implement sequential composition:

```

(>>>) :: Monad m
  => Cell m a b
  -> Cell m b c
  -> Cell m a c

```

For two cells `cell1` and `cell2` with state types `state1` and `state2`, the composite `cell1 >>> cell2` holds a pair of both states:

```

data Composition state1 state2 = Composition
  { state1 :: state1
  , state2 :: state2
  } deriving Data

```

The step function executes the steps of both cells after each other. They only touch their individual state variable, the state stays encapsulated. The custom datatype is isomorphic to an ordinary Haskell tuple (`state1`, `state2`). Yet it is beneficial to introduce it, since it allows us to extend the migration function easily such that it correctly handles the common case where we live change a cell `cellMiddle` to a composition, such as `cellLeft >>> cellMiddle`, or to `cellMiddle >>> cellRight`.

The Sensor-SF-Actuator-Pattern. Composing **Cells** sequentially allows us to form live programs out of *sensors*, pure signal functions and *actuators*:

```
type Sensor a = Cell IO () a
type SF a b = forall m . Cell m a b
type Actuator b = Cell IO b ()
buildProg :: Sensor a -> SF a b -> Actuator b
  -> LiveProgram IO
buildProg sensor sf actuator = liveCell
  $ sensor >>> sf >>> actuator
```

This (optional) division of the reactive program into three such parts is inspired by Yampa [25], and was formulated in this way in [27, Section 7.1.2]. We conveniently build a whole live program from smaller components. It is never necessary to specify a big state type manually, it will be composed from basic building blocks like **Composition**.

Arrowized FRP. **Cells** are an instance of the **Arrow** type class, which allows us to lift pure functions to **Cells**:

```
arr
  :: Monad m => (a -> b)
  -> Cell m a b
```

Together with the **ArrowChoice** and **ArrowLoop** classes (discussed in the appendix), cells can be used in *arrow notation* [26] with **case**-expressions, **if then else** constructs and recursion. The next subsection gives some examples.

An essential aspect of an FRP framework is some notion of *time*. As this approach essentially uses the *dunai* API, a detailed treatment of time domains and clocks as in *rhine* [5] could be readily applied here, but this will be deferred to future work. For simplicity and explicitness, assume that we will execute all **Cells** at a certain fixed step rate, say, twenty five steps per second. Then Euler integration can be defined:

```
stepRate :: Num a => a
stepRate = 25
integrate
  :: (Data a, Fractional a, Monad m)
  => Cell m a a
integrate = arr (/ stepRate) >>> sumC
```

The time since activation of a cell is then famously [25, Section 2.4] defined as:

```
localTime
  :: (Data a, Fractional a, Monad m)
```

```
=> Cell m b a
localTime = arr (const 1) >>> integrate
```

Monads and Their Morphisms. Beyond standard arrows, a **Cell** can encode effects in a monad, so it is not surprising that Kleisli arrows can be lifted:

```
arrM
  :: Monad m => (a -> m b)
  -> Cell m a b
```

In case our **Cell** is in another monad than **IO**, one can define a function that transports a cell along a monad morphism:

```
hoistCell
  :: (forall x . m1 x -> m2 x)
  -> Cell m1 a b -> Cell m2 a b
```

For example, we may eliminate a **ReaderT** `r` context by supplying the environment through the `runReaderT` monad morphism, or lift into a monad transformer:

```
liftCell
  :: (Monad m, MonadTrans t)
  => Cell m a b
  -> Cell (t m) a b
liftCell = hoistCell lift
```

As described in [27, Section 4], we can successively handle effects (such as global state, read-only variables, logging, exceptions, and others) until we arrive at **IO**. Then we can execute the live program in the same way as before.

4.3 A Sine Generator

Making use of the **Arrows** syntax extension⁵, we can implement a harmonic oscillator that will produce a sine wave with amplitude 10 and given period length:

```
sine
  :: MonadFix m
  => Double -> Cell m () Double
sine t = proc () -> do
  rec
    let acc = - (2 * pi / t) ^ 2 * (pos - 10)
        vel <- integrate -< acc
            pos <- integrate -< vel
    returnA -< pos
```

By the laws of physics, velocity is the integral of acceleration, and position is the integral of velocity. In a harmonic oscillator, the acceleration is in the negative direction of the position, multiplied by a spring factor depending on the period length, which can be given as an argument. The integration arrow encapsulates the current position and velocity of the oscillator as internal state, and returns the position.

⁵Arrow notation – or `proc .. do` notation – is similar to monadic `do` notation, except that not only is there a dedicated binder `<-` for output values, but also an application operator `-<` for *input* values. The notation is disguised into the arrow operators, such as `arr` and the composition `>>>`.

The caller can now decide how to handle the value `e`, should it occur. This approach can be adapted to cells. A function is supplied that runs the `ExceptT e` layer:

```
runExceptC
  :: (Data e, Monad m)
  => Cell (ExceptT e m) a          b
  -> Cell          m a (Either e b)
```

To appreciate its inner workings, let us again look at the state it encapsulates:

```
data ExceptState state e
  = NotThrown state
  | Exception e
  deriving Data
```

As long as no exception occurred, `runExceptC` cell simply stores the state of cell, wrapped in the constructor `NotThrown`. The output value `b` is passed on. As soon as the exception `e` is thrown, the state switches to `Exception e`, and the exception is output forever.

As soon as the exception is thrown, we can “live bind” it to further cells as an extra input:

```
(>>>=) :: (Data e1, Monad m)
  => Cell (ExceptT e1 m) a b
  -> Cell (ExceptT e2 m) (e1, a) b
  -> Cell (ExceptT e2 m) a b

(>>>=) cell1 cell2 = proc a -> do
  eb <- liftCell $ runExceptC cell1 -< a
  case eb of
    Right b -> returnA -< b
    Left e  -> cell2  -< (e, a)
```

We run the exception effect of the first cell. Before it has thrown an exception, its output is simply forwarded. As soon as the exception is thrown, the second cell is activated and fed with the input and the thrown exception.

5.2 Control Flow Context

Inspired by [5, Section 2, “Control Flow through Exceptions”], we introduce a newtype:

```
newtype CellExcept m a b e = CellExcept
  { runCellExcept :: Cell (ExceptT e m) a b }
```

We can enter the `CellExcept` context from an exception-throwing cell, trying to execute it until the exception occurs:

```
try
  :: Cell (ExceptT e m) a b
  -> CellExcept m a b e
try = CellExcept
```

And we can leave it safely once we have proven that there are no exceptions left to throw, i.e. the exception type is empty (represented in Haskell by `Void`):

```
safely
  :: Monad m
  => CellExcept m a b Void
  -> Cell m a b
```

One way to prove the absence of further exceptions is, of course, to run an exception-free cell:

```
safe
  :: Monad m
  => Cell m a b
  -> CellExcept m a b Void
```

The Return of the Monad. Our new hope is to give `Functor`, `Applicative` and `Monad` instances to `CellExcept`. We will explore now how this allows for rich control flow.

The `Functor` instance is not too hard. When an exception is raised, we simply apply a given function to it:

```
instance Functor m
  => Functor (CellExcept m a b) where
  fmap f (CellExcept cell) = CellExcept
    $ hoistCell (withExceptT f) cell
```

The pure function of the `Applicative` class (or equivalently, return of the `Monad`), is simply throwing an exception, wrapped in the newtype:

```
pure
  :: Monad m
  => e
  -> CellExcept m a b e
pure e = CellExcept $ arr (const e) >>> throwC
```

Like the sequential application operator `<*>` from the `Applicative` class can be defined from the bind operator `>>=`, it can also be defined from the *live bind* operator `>>>=` introduced previously. As a technical tour-de-force, even a `Monad` instance for `CellExcept` can be derived with some modifications. This is shown at length in an appendix⁷.

But how can `Applicative` and `Monad` be put to use? The foreground value of `CellExcept` is the thrown exception. With `pure`, such values are created, and `Functor` allows us to perform computations with them. With `Applicative` and `Monad`, we *chain* the execution of exception throwing cells:

```
sineWait
  :: Double -> CellExcept IO () String Void
sineWait t = do
  try $ arr (const "Waiting...") >>> wait 2
  safe $ sine t >>> arr asciiArt
```

This `do`-block can be read intuitively. Initially, the first cell is executed, which returns the message “Waiting...” every second. After three seconds, it throws an exception, which is handled by activating the sine generator. Since all exceptions have been handled, we leave the `CellExcept` context and run the resulting program:

```
printSineWait :: LiveProgram IO
printSineWait = liveCell
  $ safely (sineWait 8)
  >>> printEverySecond
```

⁷Available online at <https://www.manuelbaerenz.de/essence-of-live-coding/EssenceOfLiveCodingAppendix.pdf>.

```

Waiting...
Waiting...
>>=
  >>=
    >>=
      >>=
        >>=
          >>=
            >>=
              >>=
                >>=

```

The crucial advantage of handling control flow this way is that the *control state* – that is, the information which exceptions have been thrown and which cell is currently active – is encoded completely in the overall state of the live program, and can thus be migrated automatically. Let us rerun the above example, but after the first try statement has already passed control to the sine generator we shorten the period length of the sine wave and reload:

```

Waiting...
Waiting...
>>=
  >>=
    >>=
      >>=
        >>=
          >>=
            >>=
              >>=
                >>=

```

The migrated program did not restart and wait again, but remembered to immediately continue executing the sine generator from the same phase as before. This is in contrast to simplistic approaches to live coding in which the control flow state is forgotten upon reload, and restarted each time.

In most other programming languages where control flow is builtin, this would typically require reworking the compiler or interpreter, but in Haskell, we succeed entirely within the language.

5.3 Exceptions Forever

What if we want to change between the oscillator and a waiting period indefinitely? In other words, how do we repeatedly execute this action:

```

sinesWaitAndTry
  :: MonadFix m
  => CellExcept m () String ()

```

```

sinesWaitAndTry = do
  try $ arr (const "Waiting...") >>> wait 1
  try $ sine 5 >>> arr asciiArt >>> wait 5

```

The one temptation we have to resist is to recurse in the **CellExcept** context to prove the absence of exceptions:

```

sinesForever'
  :: MonadFix m
  => CellExcept m () String Void
sinesForever' = do
  sinesWaitAndTry
  sinesForever'

```

It typechecks, but it does *not* execute correctly. As the initial state is built up, this definition inquires about the initial state of all cells in the **do**-expression, but the last one is again `sinesForever'`, and thus already initialising such a cell hangs in an infinite loop. Using the standard function `forever :: Applicative f => f a -> f ()` has the same deficiency, as it is defined in the same way.

The resolution is an explicit loop operator, and faith in the library user to remember to employ it.

```

foreverE
  :: (Monad m, Data e)
  => e
  -> Cell (ReaderT e (ExceptT e m)) a b
  -> Cell m a b

```

The loop function receives as arguments an initial exception, and a cell that is to be executed forever. Of course, the monad `m` may again contain exceptions that can be used to break from this loop. Again, it is instructive to look at the internal state of the looped cell:

```

data ForeverE e s = ForeverE
  { lastException :: e
  , initState     :: s
  , currentState  :: s
  }
deriving Data

```

`foreverE e cell` starts with the initial state of cell, and a given value `e`. Then cell is stepped, mutating `currentState`, until it encounters an exception. This new exception is stored, and the cell is restarted with the original initial state. The cell may use the additional input `e` to ask for the last thrown exception (or the initial value, if none was thrown yet). The exception is thus the only method of passing on data to the next loop iteration.⁸ In our example, we need not pass on any data, so a simpler version of the loop operator is defined:

```

foreverC
  :: (Data e, Monad m)
  => Cell (ExceptT e m) a b

```

⁸It is the user's responsibility to ensure that it does not introduce a space leak, for example through a lazy calculation that builds up bigger and bigger thunks.

wish to test them in a similar way before reloading. As a simple example, we wish to assure that `sumC` will never output negative numbers if only positive numbers are fed into it. Our test cell is thus defined as:

```
testCell :: Monad m => Cell m (Positive Int) Bool
testCell
  = arr getPositive >>> sumC >>> arr (>= 0)
```

We begin by running a cell repeatedly against a list of inputs, collecting its outputs:

```
embed
  :: Monad m
  => [a]
  -> Cell m a b
  -> m [b]
```

If the input type `a` can be generated arbitrarily, then so can a list of `as`. After running the cell with all inputs, we form the conjunction of all properties, with `QuickCheck`'s `conjoin`. Effects in `IO` can be embedded in `QuickCheck` [8] with the monad morphism `run`, and executed with `monadicIO`. Cobbling all those pieces together makes cells testable:

```
instance (Arbitrary a, Show a, Testable prop)
  => Testable (Cell IO a prop) where
  property cell = property
    $ \as -> monadicIO $ fmap conjoin
    $ embed as $ hoistCell run cell
```

Let us execute our test:

```
> quickCheck testCell
+++ OK, passed 100 tests.
```

A large class of properties can be tested this way. We can unit test all components of a new version of our live program before reloading it. To go further, one could set up *stateful property-based testing* [14] for the live coding environment.

Migration Tests. Even better, we can test *before reloading* whether the newly migrated state would be valid. Given some tests on intermediate values in the computation, we collect all test properties in a `Writer` effect:

```
logTest
  :: Monad m
  => Cell m a prop
  -> Cell (WriterT [prop] m) a ()
logTest cell
  = liftCell cell
  >>> arrM (return >>> tell)
```

Now the tests can be included in the definition of the whole live program without adding new outputs. When the program is built, we can optionally test the properties:

```
liveCheck
  :: Testable prop
  => Bool
  -> LiveProgram (WriterT [prop] IO)
  -> LiveProgram IO
```

```
liveCheck test = hoistLiveProgram performTests
  where
    performTests action = do
      (s, props) <- runWriterT action
      when test $ quickCheck $ conjoin props
      return s
```

The function `liveCheck True` will run `quickCheck` on all properties, while `liveCheck False` gives the “production” version of our program, with tests disabled. We launch two separate threads and run the test version in one of them and the production version in the other. Always reloading into the test version first, we can ensure that the migration will create valid state before migrating the live system.

7 Conclusion

General purpose live coding can be simple and free from boilerplate. It is most naturally cast in the form of a Functional Reactive Programming (FRP) framework, and conforms well with the synchronous, arrowized paradigm in the tradition of Yampa and Dunai. The state migration is type-safe, and type-driven, in that it is derived generically from the datatype definition. By parametrizing the cells over arbitrary monads, and leveraging the exception monad, we can reason about effects and separate data flow aspects from control flow. The approach is extensible as debugging and testing methods can be added easily.

Further Directions. To use the framework in any setting beyond a toy application, wrappers have to be written that explicitly integrate it in the external loops of existing frameworks, such as web frontends and backends, OpenGL, `gloss` [18], or audio libraries. As a start, the multi-clock FRP library `rhine` [5] could be adapted to this approach.

The automatic migration only guarantees that the new state will typecheck. However, if further properties beyond the reach of Haskell's type system are expected to hold for the old state, those are not guaranteed for the new state. Within Haskell, quickchecking is our only hope. An extension such as refinement types (see e.g. [29] about `LiquidHaskell`) can automatically verify certain algebraic constraints, though. It would be a great enrichment to generalise automatic migration to such a type system.

The author thanks Iván Pérez for his work on Yampa, Dunai, and numerous other projects in the FRP world; the reviewers of Haskell Symposium 2018 and REBLS 2020 for very helpful comments that enriched and streamlined this work; Paolo Capriotti for the initial idea that led to monadic exception control flow; and the `sonnen VPP` team, especially Fabian Linges, for helpful discussions about hot code swap in Erlang.

References

- [1] Samson Abramsky, Esfandiar Haghverdi, and Philip Scott. 2002. Geometry of Interaction and Linear Combinatory Algebras. *Mathematical Structures in Computer Science* 12 (10 2002), 625–665. <https://doi.org/10.1017/S0960129502003730>
- [2] Heinrich Apfelmus. 2016. HyperHaskell. <https://github.com/HeinrichApfelmus/hyper-haskell>.
- [3] Heinrich Apfelmus. 2019. Functors and music. In *Proceedings of the 7th ACM SIGPLAN International Workshop on Functional Art, Music, Modeling, and Design*. 52–55.
- [4] Joe Armstrong. 2013. *Programming Erlang: software for a concurrent world*. Pragmatic Bookshelf.
- [5] Manuel Bärenz and Ivan Perez. 2018. Rhine: FRP with type-level clocks. In *Proceedings of the 11th International Symposium on Haskell*. ACM, 145–157.
- [6] Paul Caspi and Marc Pouzet. 1998. A Co-iterative Characterization of Synchronous Stream Functions. In *Workshop on Coalgebraic Methods in Computer Science (CMCS'98) (Electronic Notes in Theoretical Computer Science)*, Vol. 11. Elsevier, Lisbon, Portugal, 1–21. [https://doi.org/10.1016/S1571-0661\(04\)00050-7](https://doi.org/10.1016/S1571-0661(04)00050-7)
- [7] Koen Claessen and John Hughes. 2000. QuickCheck: A Lightweight Tool for Random Testing of Haskell Programs. In *Proceedings of the International Conference on Functional Programming (ICFP 2000)*. ACM.
- [8] Koen Claessen and John Hughes. 2002. Testing monadic code with QuickCheck. <http://www.cse.chalmers.se/~rjmh/Papers/QuickCheckST.ps>. *ACM SIGPLAN Notices* 37, 12 (2002), 47–59.
- [9] Chris Done. 2014. foreign-store. <https://github.com/esoylemez/foreign-store>.
- [10] Andy Gill and Ross Paterson. 2004. transformers. <https://hub.darcs.net/ross/transformers/>.
- [11] Gabriel Gonzalez. 2013. mmorph. <https://github.com/Gabriel439/Haskell-MMorph-Library>.
- [12] Tim Harris, Simon Marlow, and Simon Peyton Jones. 2005. Composable memory transactions. In *PPoPP '05: Proceedings of the tenth ACM SIGPLAN symposium on Principles and practice of parallel programming (ppopp '05: proceedings of the tenth acm sigplan symposium on principles and practice of parallel programming ed.)*. ACM Press, 48–60. <https://www.microsoft.com/en-us/research/publication/composable-memory-transactions/>
- [13] Ichiro Hasuo and Bart Jacobs. 2011. Traces for coalgebraic components. *Mathematical Structures in Computer Science* 21, 2 (2011), 267–320. <https://doi.org/10.1017/S0960129510000551>
- [14] Fred Hebert. 2019. *Property-Based Testing with PropEr, Erlang, and Elixir*. Pragmatic Bookshelf.
- [15] John Hughes. 2000. Generalising monads to arrows. *Science of computer programming* 37, 1-3 (2000), 67–111.
- [16] Daniel Ingalls. 2020. The evolution of Smalltalk: from Smalltalk-72 through Squeak. *Proceedings of the ACM on Programming Languages* 4, HOPL (2020), 1–101.
- [17] Mark P Jones. 1995. Functional programming with overloading and higher-order polymorphism. In *International School on Advanced Functional Programming*. Springer, 97–136.
- [18] Ben Lippmeier. 2010. Gloss. <https://github.com/benl23x5/gloss>.
- [19] Ralf Lämmel and Simon Peyton Jones. 2003. Scrap your boilerplate: a practical approach to generic programming. ACM Press, 26–37. <https://www.microsoft.com/en-us/research/publication/scrap-your-boilerplate-a-practical-approach-to-generic-programming/>
- [20] Alex McLean. 2014. Making programming languages to dance to: live coding with tidal. In *Proceedings of the 2nd ACM SIGPLAN international workshop on Functional art, music, modeling & design*. ACM, 63–70.
- [21] George H Mealy. 1955. A method for synthesizing sequential circuits. *The Bell System Technical Journal* 34, 5 (1955), 1045–1079.
- [22] Robin Milner. 1975. Processes: A Mathematical Model of Computing Agents. In *Logic Colloquium '73*, H.E. Rose and J.C. Shepherdson (Eds.). Studies in Logic and the Foundations of Mathematics, Vol. 80. Elsevier, 157 – 173. [https://doi.org/10.1016/S0049-237X\(08\)71948-7](https://doi.org/10.1016/S0049-237X(08)71948-7)
- [23] Neil Mitchell. 2014. ghcid. <https://github.com/ndmitchell/ghcid/>.
- [24] Tom E. Murphy. 2016. A livecoding semantics for functional reactive programming. In *Proceedings of the 4th International Workshop on Functional Art, Music, Modelling, and Design*. 48–53.
- [25] 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, 51–64.
- [26] Ross Paterson. 2001. A new notation for arrows. In *ACM SIGPLAN Notices*, Vol. 36. ACM, 229–240.
- [27] Ivan Perez, Manuel Bärenz, and Henrik Nilsson. 2016. Functional Reactive Programming, Refactored. In *Proceedings of the 9th International Symposium on Haskell (Nara, Japan) (Haskell 2016)*. ACM, 33–44.
- [28] Maciej Píróg and Jeremy Gibbons. 2014. The coinductive resumption monad. *Electronic notes in theoretical computer science* 308 (2014), 273–288.
- [29] Niki Vazou, Eric L Seidel, Ranjit Jhala, Dimitrios Vytiniotis, and Simon Peyton-Jones. 2014. Refinement types for Haskell. In *ACM SIGPLAN Notices*, Vol. 49. ACM, 269–282.
- [30] Philip Wadler. 1989. Theorems for free!. In *FPCA*, Vol. 89. 347–359.
- [31] Scott Wilson, David Cottle, and Nick Collins. 2011. *The SuperCollider Book*. The MIT Press.
- [32] Daniel Winograd-Cort and Paul Hudak. 2014. Settable and non-interfering signal functions for FRP: how a first-order switch is more than enough. *ACM SIGPLAN Notices* 49, 9 (2014), 213–225.
- [33] Kazu Yamamoto, Michael Snoyman, and Andreas Voellmy. [n.d.]. Warp. <http://www.aosabook.org/en/posa/warp.html>. In *The Performance of Open Source Applications*.