Functional Implementations of Continuous Modeled Animation

Conal Elliott

Microsoft Research
Graphics Group
http://www.research.microsoft.com/~conal

(To appear in the proceedings of PLILP/ALP '98)

Abstract. Animation is a temporally *continuous* phenomenon, but is typically programmed in terms of a discrete sequence of changes. The use of discreteness serves to accommodate the machine that is *presenting* an animation, rather than the person *modeling* an animation with the help of a computer. Using a continuous model of time for animation allows for natural specification, avoiding some artificial details, but is difficult to implement with generality, robustness and efficiency. This paper presents and motivates *continuous modeled animation*, and sketches out a naive functional implementation for it. An examination of some of the practical problems with this implementation leads to several alternate representations, all of which have difficulties in themselves, some quite subtle. We hope that the insights and techniques discussed in this paper lead to still better representations, so that animation may be specified in natural terms without significant loss of performance.

1 Introduction

A functional approach to animation offers the possibility to make animations much easier and more natural to express, by removing the need to direct the details of *presentation* and allowing the programmer instead to say what an animation is [1, 8]. Following the evolution of 3D graphics, we have termed this approach "modeling", as opposed to "presentation" [4].

Given that we want to model animations, what notion of time should we adopt? The first fundamental choice is discrete vs continuous, that is do we think of time as moving forward in a (discrete) sequence of "clock ticks", or a (continuous) flow?

A discrete model of time fits more easily into popular *computational* concepts, because modern computers operate in temporally discrete steps. More specifically, animations are presented as a sequence of frames, delivered to the user at a finite rate.

On the other hand, a continuous model of time seems a more natural fit with our human perception of time. Moreover, we have a rich heritage of mathematical, scientific, and engineering tools for understanding and describing the

basic animation concepts of motion, growth, etc, and these tools are based on the continuous notion of time. These tools include polynomials, trigonometry, and calculus, together with their applications to physical motion, governed by Newton's law of motion and its applications to gravitational attraction, springs, etc.

Just as continuous image models such as Postscript [11] naturally give rise to spatial resolution-independence and hence scalability, the continuous time model for animation yields *temporal* resolution-independence and scalability. For instance, one can stretch or squeeze a given animation to fit a desired duration.

Fran (*Functional reactive animation*) is a theory of animation, and implementation as a Haskell library, that is based on a continuous time model. As such, it offers the possibility of allowing people to express animations in ways that match their intuitive understandings and that leverage the wealth of mathematical tools at their disposal.

The cost of offering a continuous time model is a more challenging implementation problem, since continuous animation descriptions must be translated into discrete lower-level software and hardware structures.[1] In fact, we have implemented Fran many times, and have not yet achieved a satisfactory combination of generality, robustness and efficiency. Several subtle difficulties became apparent to us only through implementation and experimentation.

The purpose of this paper is to present the implementation techniques we have explored, together with some insight into practical difficulties, in the hope of facilitating work leading to still better representations.

An expanded version of this paper presents a few more implementation alternatives that we have investigated, but have implemented only partially or not at all [7].

Previous papers on Fran have presented its vocabulary and semantics, its role as an "embedded language", and examples of its use [8, 5, 4, 14]. This paper gives only minimal treatment of these issues.

2 A User Perspective on Fran

Fran is a declarative animation library (or "embedded language") that embodies the continuous time model of animation, and is implemented in Haskell [9].

2.1 Behaviors

For any type `ty`, the Fran type `Behavior ty` represents `ty`-valued animations, i.e., functions from continuous time to `ty`. As an example, consider the following

[1] The translation from continuous to discrete inevitably introduces error. Usually these errors are negligible, but sometimes, as in some systems of differential equations or condition-based events, errors can become significant. These problems are inherent in applying computers to simulate continuous phenomena, regardless of the programming paradigm.

definition of a real-valued animation, which has value 1 at time 0 and grows quadratically:

```
type RealB = Behavior Double   -- helpful synonym

growQuad :: RealB
growQuad = 1 + time ** 2
```

There is some notational magic going on here. We are using behavior-specific overloadings of familiar numerical operators for addition and exponentiation, and even numeric literals, using Haskell's type class mechanism. The types used in this example are as follows:

```
1, 2      :: RealB
(+), (**) :: RealB -> RealB -> RealB
time      :: RealB
```

Literals like 1 and 2 are implicitly converted into behaviors through application of the polymorphic `constantB` function:

```
constantB :: a -> Behavior a
```

For any value `x`, the behavior `constantB x` has value `x` at every time.

The behavior versions of functions like `+` and `**`, as well as `sin`, `log`, etc., are all defined in terms of the non-behavior versions, by using lifting operations.

```
lift1 :: (a -> b) ->
         Behavior a -> Behavior b
lift2 :: (a -> b -> c) ->
         Behavior a -> Behavior b -> Behavior c
lift3 :: (a -> b -> c -> d) ->
         Behavior a -> Behavior b -> Behavior c -> Behavior d
-- etc
```

Lifting is hidden inside of type class instances, with declarations like the following:

```
instance  Num a => Num (Behavior a)  where
  (+)        =  lift2 (+)
  (*)        =  lift2 (*)
  negate     =  lift1 negate
  abs        =  lift1 abs
  fromInt n  =  constantB (fromInt n)
```

The first line says that for any "number type" `a`, the type of `a`-valued behaviors is also a number type. The second line says that the behavior version of `+` is the binary-lifted version of the unlifted `+` operation. (The definitions appear self-referential, but are not. Overload resolution distinguishes the two different versions of "`+`".) The function `fromInt` is used to resolve literal integer overloading. The reason for this definition, instead of the more obvious one involving

`lift1`, is that `fromInt` must still work on numbers, not number-valued behaviors. This requirement is partly desirable, but is also partly due to the restricted type of `fromInt`. In such cases, Fran provides additional definitions with names formed by adding "`B`" or "`*`" to the unlifted version, such as the following.

```
fromIntB :: Num a => Behavior Int -> Behavior a
fromIntB = lift1 fromInt

(==*) :: Eq a => Behavior a -> Behavior a -> Behavior Bool
(==*) =  lift2 (==)
```

Because `Behavior` is a type constructor, we can apply it to any type. For instance, when applied to a type `Point2` of 2D static points, the result is the type of motionsn in 2D space.

It is often quite natural to express behaviors in terms of time-varying velocities. For this reason, Fran supports integration over a variety of types, as long as they implement vector space operations.

```
integral :: VectorSpace v => Behavior v -> User -> Behavior v
```

Examples of vector space types are reals, 2D vectors, and 3D vectors. The "user" argument to `integral` has two roles. One is to choose the integral's start time, and the other is to assist choice of step-size in the numerical solution. If a user `u` has start time t_0, then `integral b u` is a behavior that at time t has approximately the value $\int_{t_0}^{t} b(t')dt'$.

For modularity, it is useful to construct a behavior and then, separately, transform it in time. Fran, therefore, has a time transformation operation. Since a time transform remaps time, it is specified as a time-valued behavior, which semantically is a function from time to time.

```
timeTransform :: Behavior a -> Behavior Time -> Behavior a
```

The lifting operators can all be expressed in terms of a single operator:

```
($*) :: Behavior (a -> b) -> Behavior a -> Behavior b
```

The name "`$*`" comes from the fact that it is the lifted version of function application, whose Haskell infix operator is "`$`".

The lifting operators are defined simply in terms of `constantB` and "`$*`", as follows.

```
lift0            = constantB
lift1 f b1       = lift0 f $* b1
lift2 f b1 b2    = lift1 f b1 $* b2
lift3 f b1 b2 b3 = lift2 f b1 b2 $* b3
-- etc
```

Note that the basic combinators "`$*`", `constantB`, `time`, and `timeTransform` correspond semantically to type-specialized versions of the classic S, K, I, and B combinators.

2.2 Events

Besides behaviors, the other principle notion in Fran is the *event*. For any type `ty`, the Fran type `Event ty` is the type of `ty`-valued events, each of which is semantically a time-sorted sequence of time/`ty` pairs, also refered to as "occurrences" of the event.

This paper is chiefly concerned with the implementation of continuous *behaviors*, rather than events (which are intrinsically discrete). We will informally describe event operators as they arise.

2.3 Reactive Behaviors

While one can define an infinite set of behaviors using just the behavior combinators given above, such behaviors are not very dynamic. Fran greatly enriches these behaviors by supporting reactivity, via the following primitive, which uses behavior-valued events.

```
untilB :: Behavior a -> Event (Behavior a) -> Behavior a
```

Given behavior `b` and an event `e`, the behavior `b 'untilB' e` acts like `b` until `e` occurs, and then switches to acting like the behavior that accompanies the event occurrence. Although `untilB` only needs the *first* occurrence of an event, the other occurrences are used by some of the event combinators. This meaning of events turned out to be more convenient than the one in [8], since it enables higher level combinators for constructing reactive behaviors, as illustrated in [4].

3 Implementing Continuous Behaviors

We now turn to the main thrust of our paper, which is an exploration of how to implement continuous behaviors.

Representation A: time-to-value functions. The semantics of behaviors suggest a very simple representation:

```
newtype Behavior a = Behavior (Time -> a)

-- Sample a behavior "at" a given time
at :: Behavior a -> Time -> a
Behavior f `at` t = f t
```

This definition introduces both a new type constructor and a value constructor, both called `Behavior`. (A simple `type` definition in Haskell would not have allowed overloading.) Using this representation, it is easy to define our simple behavior combinators:

```
constantB x = Behavior (const x)

fb $* xb = Behavior (\ t -> (fb `at` t) (xb `at` t))

time = Behavior (\ t -> t)

timeTransform b tt = Behavior (at b . at tt)
```

To sample a behavior constructed by `b `untilB` e` at a time t, first check whether the event `e` occurs before t.[2] If so, sample the new behavior, `b'`, that is part of the event occurrence, and if not, sample `b`.

```
b `untilB` e = Behavior sample
 where
   sample t = case (e `occ` t) of
                Nothing  ->  b  `at` t
                Just b'  ->  b' `at` t
```

Here we presume a function `occ`, with the following signature, for checking for an event occurrence before a given time.

```
occ :: Event a -> Time -> Maybe a
```

3.1 The Problem of Non-Incremental Sampling

While Representation A given above is appealing in terms of simplicity, it has a serious performance problem. It allows nothing to be remembered from one sampling to another. In animation, however, behaviors are typically sampled at a sequence of display times separated by small differences. For instance, given an integral behavior, it is vital for efficiency that intermediate results are carried from one sampling to the next, so that only a small amount of extra work is required. Similarly, for a *reactive* behavior, i.e., one constructed with `untilB`, it is important to make incremental progress in event detection.

Consider the implementation of `untilB` above, and suppose we want to sample a reactive behavior with times $t_1, t_2, \ldots$. For every t_i, sampling must consider whether the event has occurred, and choose to sample the old behavior or the new one. Moreover, it is frequently the case that the new behavior it itself reactive, and so on, in an infinite chain. In such a case, the time it takes to compute each sample will increase without bound, which is clearly unacceptable. Moreover, consider what is involved in computing "`e `occ` t`" in `untilB`. There are two possibilities: either event occurrences are being rediscovered, or they are cached somehow. If they are rediscovered, then the cost of sampling increases even more so with each sampling. If, however, they are cached, then the cache itself becomes a large space leak. Together these problems cause what we call a "space-time leak".

[2] By choosing *before*, rather than before *or at*, the sampling time t, animations may be self- or mutually-reactive.

Representation B: residual behaviors. A crucial observation is that typically, the sampling times t_1, t_2, ..., are monotonically increasing. If we assume that this typical case holds, then we can remove the space-time leak.[3] The idea is to have sampling yield not only a value, but also a "residual behavior".

```
newtype Behavior a = Behavior (Time -> (a, Behavior a))

at :: Behavior a -> Time -> (a, Behavior a)
Behavior f `at` t = f t
```

The combinator implementations are not quite as simple as before, but still reasonable. Constant behaviors always return the same pair:

```
constantB x = b
 where
   b = Behavior (const (x, b))
```

The `time` behavior is also simple.

```
time = Behavior (\ t -> (t, time))
```

The "`$*`" combinator samples its argument behaviors, applies one resulting value to the other, and applies "`$*`" to the residual behaviors.

```
fb $* xb = Behavior sample
 where sample t = (f x, fb' $* xb')
         where (f, fb') = fb `at` t
               (x, xb') = xb `at` t
```

Time transformation can be implemented much like "`$*`".

```
timeTransform b tt = Behavior sample
 where sample t = (x, timeTransform b' tt')
         where (t', tt') = tt `at` t
               (x , b' ) = b  `at` t'
```

Time transformation can violate our monotonicity assumption for time streams, if the time transform `tt` is not itself monotonic. It would be possible to check for monotonic sampling dynamically, at least for reactive behaviors, though Fran does not do so. Checking monotonicity "statically", i.e., when a `timeTransform` or `untilB` behavior is constructed does not seem to be feasible.

The `occ` function, used for checking event occurrences, is changed in a way much like behaviors, so that it now yields a residual event along with a possible occurrence.

```
occ :: Event a -> Time -> (Maybe a, Event a)
```

[3] Unfortunately, the cost of this assumption is a significant restriction in the time transforms that may be applied to reactive behaviors.

Next, consider reactivity. Sampling a reactive behavior with a time that is after the event's first occurrence, the newly constructed behavior (a) no longer checks for the event occurrence, thus eliminating the time leak, and (b) no longer holds onto the old behavior, thus eliminating the space leak.

```
b `untilB` e  = Behavior sample
 where
   sample t =
     case (e `occ` t) of
       -- No occurrence before t; keep looking
       (Nothing, eNext) -> let (x, bNext) = b `at` t in
                           (x, bNext `untilB` eNext)

       -- Found it; sample new behavior
       (Just bNew, _)  -> bNew `at` t
```

Representation C: stream functions. An alternative solution to the problem of non-incremental sampling is to map time streams to value streams. We will see in Section 3.2 that this representation has advantages over Representation B.

```
newtype Behavior a = Behavior ([Time] -> [a])

at :: Behavior a -> [Time] -> [a]
Behavior f `at` ts = f ts
```

Constant behaviors always return the same list containing an infinite repetition of a value:

```
constantB x = Behavior (const (repeat x))
```

The "`$*`" combinator samples the function and argument behaviors, and applies resulting functions to corresponding arguments, using the binary mapping functions `zipWith`.

```
fb $* xb =
  Behavior (\ ts -> zipWith ($) (fb `at` ts) (xb `at` ts))
```

Time is the identity as it was in Representation A, but of a different type:

```
time = Behavior (\ ts -> ts)
```

Reactivity is implemented by scanning through a list of possible event occurrences, while enumerating behavior samples. For convenience, assume a stream sampler function for events:

```
occs :: Event a -> [Time] -> [Maybe a]
```

The implementation of `untilB`:

```
b `untilB` e =
  Behavior (\ ts -> loop ts (b `at` ts) (e `occs` ts))
 where
   -- Non-occurrence.  Emit first b sample and continue looking
   loop (_:ts') (x:xs') (Nothing:mbOccs') =
      x : loop ts' xs' mbOccs'
   -- First event occurrence.  Discard the rest of the b values
   -- and possible event occurrences, and continue with the
   -- new behavior
   loop ts      _       (Just bNew : _  ) = bNew `at` ts
```

A weakness of Representations B and C is that they cause a great deal of construction with each sampling. For this reason, it is important to have a garbage collector that deals very efficiently with the rapid production of short-lived structures, as in generational garbage collection. For instance, the Glasgow Haskell Compiler [12] has such a collector.

3.2 The Problem of Redundant Sampling

Another serious problem with all the representations preceding is that they lead to redundant sampling. As a simple example, consider the following behavior that linearly interpolates between two numerical behaviors `b1` and `b2`, according to the interpolation parameter `a`, so that the resulting behavior is equal to `b1` when `a` is zero and `b2` when `a` is one.

```
interp :: RealB -> RealB -> RealB -> RealB
interp b1 b2 a  =  (1 - a) * b1 + a * b2
```

The problem is that sampling the behavior generated by `interp` at some time t ends up sampling `a` at t twice, due to the repeated use of `a` in the body of `interp`. If `a` is a complex behavior the redundant sampling is costly. Worse yet, composition of functions like `interp` multiplies the redundancy, as in the following example.

```
doubleInterp :: RealB -> RealB -> RealB -> RealB -> RealB -> RealB
doubleInterp b1 b2 b3 b4 a = interp b1 b2 a'
  where a' = interp b3 b4 a
```

When the result of `doubleInterp` is sampled with a time t, the behavior `a'` will be sampled twice at t, which means `a` will be sampled four times, and the interpolation work for `a'` done twice. If we were to cube the result of `doubleInterp`, via the definition

```
cube :: Num a => a -> a
cube x = x * x * x
```

then sampling work would be multiplied by three, causing `a` to be evaluated twelve times for each sample time t.

One approach to solving the problem of redundant sampling is applying lazy memoization [10, 3], which may be supported with a function of the form type.

```
memo :: Eq a => (a -> b) -> (a -> b)
```

Semantically, `memo` is the identity. Operationally, the closure returned contains a mutable "memo table".

Any of the three representations discussed so far may be memoized. For instance, in Representation B one could simply memoize the constructed sampling functions. By way of example, here is an appropriately modified "`$*`". The only change is that the created `sample` function is memoized.

```
fb $* xb = Behavior (memo sample)
 where sample t = (f x, fb' $* xb')
         where (f, fb') = fb `at` t
               (x, xb') = xb `at` t
```

A drawback to this implementation is that the memoization overhead must be paid for each sample time of each component behavior, and so slows down sampling, rather than speeding it up, except in circumstances of extreme redundant sampling.

A more efficient alternative would to start with Representation C, so that memoization works at the level of *lists* of times rather than individual times. Rather than base memoization on the usual elementwise notion of list equality, which would be particularly problematic because our time lists are infinite, it suffices to use pointer equality on the list representations, as recommended in [10].

This is the representation used in the current version of Fran (1.11), except that memo tables are managed explicitly, rather than through a higher-order `memo` function. The reason for this exception is that memo tables need to be "aged", as will explained below.

Our algebra of behaviors is related to, and in some ways inspired by, Backus' language FP [2]. In FP, programs are always expressed at the function level, with application to "object-level" values kept implicit. This property leads to redundant applications of a function to the same argument, similar to the problem discussed in this section. The Illinois Functional Programming Interpreter [13] addressed this problem by using an "expression cache." For some recursive algorithms, expression caching reduced the asymptotic running time. Normally this caching had a more modest effect, speeding up some computations and slowing down others.

3.3 The Problem of Space-Leaking Memo Tables

A subtle but important consideration for any use of memoization is garbage collection. The memoized functions contain tables of domain and range values, and

these tables typically grow without bound. When all references to a memoized function are lost, the contents of its memo table are reclaimed, as described in [3]. For example, consider the following behavior, which uses the `cube` and `interp` functions defined above to stretch a given picture. (The Fran type `ImageB` represents "image-valued behaviors", i.e., two-dimensional animations.)

```
anim1 :: ImageB
anim1 = stretch s1 pic
 where
   s1 = cube (interp 0.5 1.5 (cube time))
```

If `anim1` is displayed, it will be sampled with a time list that depends on the speed of the machine and the time-varying load from other processes. As long as there is no reference to `anim1` besides the one given to the display function, the representations of `anim1` and `s1`, including memo tables will be reclaimed. The list cells in the time and value lists will also be eligible for reclamation in an efficient manner as animation display progresses.

Unfortunately, memory use is not always so well-behaved. One problematic case is that in which the definition of `anim1` is a top-level definition, also known as a "CAF" (constant applicative form). In that case, in the current Haskell implementations we know of, the behaviors will never be reclaimed.[4]

Another problematic situation for reclamation of memo tables arises when a behavior is retained in order to be restarted, as in the following example, which makes a looping behavior out of `s1`, restarting every two seconds.[5]

```
anim2 :: ImageB
anim2 = stretch s2 pic
 where
   s1 = cube (interp 0.5 1.5 (cube time))
   s2 = s1 `untilB` timeIs 2 -=> later 2 s2
```

Here we have used the `later` operator to shift `s2` by two seconds. (Alternatively, the semantics of `untilB` could do this kind of shifting automatically. Although there are some technical difficulties, we intend to alter Fran in this way. The implementation issue discussed below remains, however.) In this case, `s1` cannot released, because it will be used every two seconds to generate a new list of scale values. For each memo table, an entry is added every two seconds, and the evaluated size of each entry grows through lazy evaluation by a time, a scale, and two cons cells every sample, e.g., 10 per second.

For these reasons, it seems necessary for memo tables to have special support from the garbage collector, so that when there are no live pointers to an object other than memo table keys, then the object gets reclaimed and the memo table entries get deleted. This idea is described in [10] and has been in use in some Smalltalk, Lisp and Scheme implementations for quite a while, but as far as we

[4] An upcoming release of the Glasgow Haskell Compiler fixes this problem.

[5] The event "`e -=> v`" occurs whenever `e` occurs and has value `v` at each occurrence. Syntactically, it binds more tightly than `` `untilB` ``.

know it has not yet been implemented in a Haskell run-time system. Some Fran programs leak space for exactly this reason.

A more serious, and more subtle, problem with memoization comes from Fran's `afterE` combinator. This combinator is essentially a dual to `snapshot`. The two signatures are as follows:

```
snapshot :: Event a -> Behavior b -> Event (a, b)
afterE   :: GBehavior bv => Event a -> bv -> Event (a, bv)
```

The idea of `e ‘snapshot‘ b` is to pair up the event data from each occurrence of `e` with snapshots of `b` values at the corresponding occurrence time. Dually, `e ‘afterE‘ b` gives access to the *residual* of `b` at each occurrence of `e`. Its purpose is to be able to coordinate with a concurrently running behavior after an event occurrence, but fortunately its use can often be hidden inside of other event and behavior combinators. Because `afterE` need not actually sample, it can be used with "generalized behaviors", a notions that subsumes the actual behavior types.

Fran implements `afterE` in terms of the following more primitive "aging" function:

```
afterTimes :: GBehavior bv => bv -> [Time] -> [bv]
```

The idea is to create a list of times corresponding not only to occurrences of an event, but to closely spaced *non-occurrences* as well. This list is fed into `afterTimes`, which then produces a stream of updated versions of its generalized behavior argument. (Note: it may well be possible and desirable to hide `afterE` beneath a set of more abstract combinators, but the implementation issues remain.)

Now consider the interaction between `afterTimes` and memoization. Suppose we have a behavior `bv` with a memo entry for the time stream $t_1, t_2, \ldots$ and corresponding value stream $x_1, x_2, \ldots$, and we have another time stream $t'_1, t'_2, \ldots$ as an argument to `afterTimes`. How should we initialize the aged behaviors' memo tables? If we use `bv`'s memo table, we will have a space leak, and, moreover, these entries are not likely to get cache hits. If, on the other hand, we start with empty memo tables, we will end up repeating a lot of work. A crucial observation is that these aged behaviors are often sampled with *aged sample time streams*, i.e., suffixes of the time streams that have already been used to sample `bv` itself. Rather than reusing or discarding memo tables in their entirety, the current Fran implementation "ages" the tables, replacing each time- and value-stream pair with corresponding stream suffixes. For instance, if $t_2 < t'_1 \leq t_3$, then the first aged memo pair would be $t_3, t_4, \ldots$ and $x_3, x_4, \ldots$.

Note that if we were to memoize Representation A or B instead of C, then it would become trickier to use a garbage collector to trim the memo tables. If the `Time` type is something like `Float` or `Double` (as in Fran's implementation), then we could easily keep a time/sample pair alive in a memo table by accident. If indeed these accidental retentions happen, a solution would be to introduce a data type to wrap around the time values.

3.4 The Problem of Synchrony

Memoization solves the problem of redundant sampling, but only when the different uses of a behavior are sampled with exactly the same time streams. There are, however, at least three situations in which we would want a behavior to be sampled with different time streams.

The first situation is the application of a time transformation. Consider the following example.

```
b :: RealB
b = b1 + timeTransform b1 (time / pi)
 where
   b1 = cube (interp 0.5 1.5 (cube time))
```

There is a second situation in which we might like to sample a behavior with two different rates, namely when different degrees of precision are needed. As an example, suppose we have an image `im1` moving in a complex motion path `mot`, with an overlaid, much smaller copy of itself:

```
im :: ImageB
im = stretch 0.1 im2 `over` im2
 where
   im2 = move mot im1
```

If we were sampling in order to stay within a error bound measured in 2D distance, rather than *temporal* rate, then the first use of `im2` would require less accuracy from `mot` and `im1` than the second use, because the spatial inaccuracies are reduced by a factor of ten.

A third situation calling for variable sampling rate is detection of events defined by boolean behaviors. As discussed in [8], interval analysis (IA) can be used in a robust nand efficient algorithm to detect occurrences of such events. The sampling patterns are adaptive to the nature of the condition's constituent behaviors.

In all of these cases, lack of synchrony disallows sharing of work between different sampling patterns. We do not have a solution to this problem. Note, however, that the values used for display in between samplings are necessarily approximate. (Fran uses a linearly interpolating engine [6].) As explored in [7], this observation suggests sharing of work among non-synchronous sampling patterns.

3.5 Structural Optimizations

Fran's algebra of behaviors and events satisfies several algebraic properties that are exploited for optimization. Roughly speaking, these identities fall into the categories of "static" and "behavioral" properties.

By a "static" property, we mean one that is lifted to the level of behaviors directly from a property on static values. Examples include the following identity

and distributive properties on images, where `over` is the (associative but not commutative) image overlay operation and `*%` applies a 2D transform to an image.

```
emptyImage `over` im  ==  im

im `over` emptyImage  ==  im

xf *% (im1 `over` im2)  ==  xf *% im1 `over` xf *% im2
```

Another useful property is following one for `condB`, which is the lifted form of if-then-else:

```
condB (constantB True ) b c  ==  b

condB (constantB False) b c  ==  c
```

By a "behavioral" property, we mean one that applies to behaviors over all types. For example, when "`$*`" is applied to constant behaviors, the result is a constant behavior, i.e.,

```
constantB f $* constantB x == constantB (f x)
```

As a consequence, a lifted n-ary function applied to n constant behaviors yields a constant behavior.

Another useful property is distributivity of lifted functions over reactivity. First consider a simple case.

```
lift1 f (b `untilB` e)  ==  lift1 f b `untilB` e ==> lift1 f
```

The event `e ==> h` occurs when `e` occurs, and its occurrence values result from applying the function `h` to the corresponding value from `e`. Syntactically, it binds more tightly than `` `untilB` ``.

Here is an obvious candidate for the general case:

```
fb $* (xb `untilB` e)  ==
  fb $* xb  `untilB` e ==> \ xb' ->
  fb $* xb'
```

and similarly for the case that `fb` is reactive. If both argument behaviors are reactive, then both rules may be applied sequentially (in either order). Similarly, if in the first rule, `xb` itself is reactive, applying the rule will give rise to another rule application as well.

There is, however, an operational problem with a rule like the one above, namely that it holds onto the behavior `fb` while waiting for the event `e` to occur, thus causing a space leak. Then when `e` finally occurs, `fb` will get sampled starting at the occurrence time. If `fb` has meanwhile undergone many transitions, there may be a lot of catching up to do, which amounts to a "time leak". To fix both of these problems, use `afterE` to give quick access to the residual of `fb`.

```
fb $* (xb `untilB` e)  ==
  fb  $* xb  `untilB` e `afterE` fb ==> \ (xb',fb') ->
  fb' $* xb'
```

The identities above, while widely applicable, do not improve performance by themselves. Their merit is that they often enable other optimizations. For instance, consider the following animation:

```
b :: RealB
b = (time `untilB` timeIs 5 -=> 0) + b2
```

Applying the `$*/untilB` identity yields the following:

```
time + b2 `untilB` (timeIs 5 -=> 0) `afterE` b2 ==>
   \ (b1', b2') -> b1' + b2'
```

When the transition occurs at time 5, the new behavior will be `0 + b2'`, which can be simplified to `b2'`.

4 Conclusions

Modern software and hardware technology are temporally discrete in nature and so encourage discrete software models. In the context of animation, a continuous approach is more natural because it more closely reflects the real-world behaviors being modeled. In this paper, we have explored several functional implementations of continuous animation and some problems that arise. Some of these problems are rather subtle and became apparent only through costly trial and error. See [7] for more alternatives that we have partially explored. We hope that their discussion here will motivate further work in pursuit of the goals of efficiently-executing, naturally-specified interactive animation.

5 Acknowledgements

Discussions with Sigbjorn Finne and Simon Peyton Jones helped to clarify operational properties of the representations discussed in this paper. Extensive comments from anonymous reviewers greatly improved the presentation.

References

1. Kavi Arya. A functional approach to animation. *Computer Graphics Forum*, 5(4):297–311, December 1986.
2. John Backus. Can programming be liberated from the von Neumann style? *Comm. ACM*, 8:613–641, 1978.
3. Byron Cook and John Launchbury. Disposable memo functions (extended abstract). In *Proceedings of the 1997 ACM SIGPLAN International Conference on Functional Programming*, page 310, Amsterdam, The Netherlands, June 1997.

4. Conal Elliott. Modeling interactive 3D and multimedia animation with an embedded language. In *The Conference on Domain-Specific Languages*, pages 285–296, Santa Barbara, California, October 1997. USENIX. WWW version at `http://-www.research.microsoft.com/~conal/papers/dsl97/dsl97.html`.
5. Conal Elliott. Composing reactive animations. *Dr. Dobb's Journal*, pages 18–33, July 1998. Expanded version with animated GIFs: `http://www.research.-microsoft.com/~conal/fran/{tutorial.htm,tutorialArticle.zip}`.
6. Conal Elliott. From functional animation to sprite-based display (expanded version). Technical Report MSR-TR-98-28, Microsoft Research, July 1998. `http://-www.research.microsoft.com/~conal/papers/spritify/long.ps`.
7. Conal Elliott. Functional implementations of continuous modeled animation (expanded version). Technical Report MSR-TR-98-25, Microsoft Research, June 1998. `http://www.research.microsoft.com/~conal/papers/plilpalp98/long.ps`.
8. Conal Elliott and Paul Hudak. Functional reactive animation. In *Proceedings of the 1997 ACM SIGPLAN International Conference on Functional Programming*, pages 263–273, Amsterdam, The Netherlands, 9–11 June 1997.
9. John Peterson et. al. Haskell 1.3: A non-strict, purely functional language. Technical Report YALEU/DCS/RR-1106, Department of Computer Science, Yale University, May 1996. Current WWW version at `http://haskell.org/-report/index.html`.
10. John Hughes. Lazy memo functions. In J. P. Jouannaud, editor, *Functional Programming Languages and Computer Architecture*, volume 201 of *LNCS*, pages 129–146. Springer Verlag, September 1985.
11. Adobe Systems Incorporated. *POSTSCRIPT Language: Tutorial and CookBook*. Addison-Wesley Publishing Company, Inc, 1991.
12. Simon L. Peyton Jones, Cordelia V. Hall, Kevin Hammond, Will Partain, and Philip Wadler. The Glasgow Haskell compiler: a technical overview. In *Proc. UK Joint Framework for Information Technology (JFIT) Technical Conference*, July 93.
13. Arch D. Robinson. The Illinois functional programming interpreter. In *Proceedings SIGPLAN '87 Symposium on Interpreters and Interpretive Techniques*, pages 64–73. ACM, ACM, June 1987.
14. Simon Thompson. A functional reactive animation of a lift using Fran. Technical Report 5-98, University of Kent, Computing Laboratory, May 1998. `http://www.cs.ukc.ac.uk/pubs/1998/583/index.html`.