

Functional Reactive Programming

Lecture 24

Tuesday, April 28, 2014

Pure programs (i.e., programs that are free from side-effects such as reading or writing memory or producing input and output) can be easier to reason about than programs with side-effects. This is because when composing two pure programs e_1 and e_2 together (e.g., in sequence, in parallel, etc.) we don't need to think about how the side effects of e_1 and e_2 might interact, and change the behavior of the programs.

However, many of the program we write *should* have side effects, because we want these programs to interact in interesting ways with the external environment. For example, we want programs to read user input, receive messages over the network, produce output on a screen, and send messages to servers. Only rarely do we want to write completely pure programs, where the only interesting computation is the final result.

Functional Reactive Programming (FRP) is a paradigm that combines pure functional programming approaches with the ability for programs to interact with their external environments. A key concept in FRP is a *signal*, which is a time-varying value. Signals can be used to represent the current value of something in the environment, such as the co-ordinates of a mouse, or the current value in a user input field. An FRP program can compute new signals from existing ones, and the output of an FRP program can be a signal, e.g., the contents of a window, that may contain images and text that change over time.

Conceptually, you can think of an FRP program as being continually recomputed so that as the value of input signals change, the value of the output signals also changes. For example, suppose that s is an input signal (for example, the current horizontal position of the mouse). The following FRP program computes a signal whose current value is two times the current value of s .

$$\text{lift } (\lambda x. x + x) s$$

The primitive `lift` takes a pure function $(\lambda x. x + x)$ and evaluates the function on each value that signal s takes on. The result of program is itself a signal, i.e., a time-varying value. As time passes and signal s takes on new values, the output signal also changes over time.

1 Elm

To further explore FRP, we will focus on the Elm programming language. Elm is a functional-reactive programming language that specifically targets graphical user interfaces (GUIs). That is, in addition to FRP language constructs, it also has a rich library focused on GUIs. Elm was designed and implemented by Evan Czaplicki '12. You can find out more about Elm at <http://elm-lang.org/>.

The following syntax describes a simple core calculus for Elm. It includes integers n , unit $()$, let expressions, and conditionals (where if e_1 e_2 e_3 evaluates e_2 if e_1 evaluates to a non-zero integer, and evaluates e_3 if e_1 evaluates to zero). We assume there is some set of primitive input signals i , representing, for example, the current position of the mouse, the current size of the window, etc. In Elm, signals can be thought of as a stream of discrete values. So, for example, if we had a signal representing the current position of the mouse, the signal would in essence be a stream of pairs of integers, with perhaps a new pair being created whenever the position of the mouse changed.

This core calculus for Elm also has primitive operations `liftn` and `foldp`, which we describe below.

Numbers	$n \in \mathbf{Int}$
Variables	$x \in \mathbf{Var}$
Input signals	$i \in \mathbf{Input}$
Expressions	$e ::= () \mid n \mid x \mid \lambda x:\eta. e \mid e_1 e_2$ $\mid \text{if } e_1 e_2 e_3 \mid \text{let } x = e_1 \text{ in } e_2 \mid i$ $\mid \text{lift}_n e e_1 \dots e_n \mid \text{foldp } e_1 e_2 e_3$

The primitive operation lift_n takes a pure function that takes n arguments, and n signals, and “lifts” the pure function to the signals, producing an output signal that varies as any of the input signals vary. (This is actually a family of primitive operations, one operation for each $n \geq 1$.) For example, in the following program, assume that i and j are input signals of integers. The expression evaluates to a signal that is the sum of the two input signals.

$$\text{lift}_2 (\lambda a. \lambda b. a + b) i j$$

The lift primitives are where pure functions interact with signals. They allow signals to be transformed (e.g., $\text{lift}_1 (\lambda x. x + x) i$) and combined (as in the example immediately above).

Elm is pure, in that the only side effects that Elm programs have are via signals to and from the external environment. However, it is useful for a program to have state. Elm provides the primitive foldp , which performs a “fold from the past” on the values of a signal. The name is by analogy with foldl and foldr which perform folds over a list from the left and right respectively. Expression $\text{foldp } e_1 e_2 e_3$ expects e_3 to be a signal, and e_2 to be the initial value of the accumulator. Expression e_1 is expected to be a function $\tau \rightarrow \tau' \rightarrow \tau'$, where the signal produces values of type τ , and the accumulator is of type τ' . The foldp expression evaluates to a signal that produces values of type τ' , i.e., it is the sequence of values of the accumulator.

For example, the following program counts the number of times that primitive signal i is equal to zero. (For example, if i is the horizontal mouse position, it counts the number of times the mouse is at the left edge of the window.) Note that the initial value of the accumulator is 0, and each time the value of i changes, if the new value is zero, then the accumulator c is incremented.

$$\text{foldp } (\lambda x. \lambda c. \text{if } x \text{ c } (c + 1)) 0 i$$

1.1 Type system

Consider the following program, where i , j , and k are signals of integer values.

$$\text{let } s = \text{lift } (\lambda x. \text{if } x \text{ j } k) i \text{ in } \dots$$

The value s is a signal whose values are themselves signals (of integers). This is not a problem in and of itself, but can cause problems in some circumstances. For example, suppose that instead of signal k above, we used a more complex signal expression, perhaps one that depended, via foldp on the entire history of some input signal width (e.g., the width of the window).

$$\text{let } s = \text{lift } (\lambda x. \text{if } x \text{ j } (\text{foldp } (\lambda w. a. a + w) 0 \text{ width})) i \text{ in } \dots$$

Suppose we run this program for 10 minutes, and then signal i becomes zero for the first time. At that time, signal s should produce the signal $\text{foldp } (\lambda w. a. a + w) 0 \text{ width}$. But what is the current value of this signal? It should be the sum of the window widths throughout the entire execution of the program up to this point. To compute this value, we would have had to save all of the history of the signal width , even though we don’t know whether signal $\text{foldp } (\lambda w. a. a + w) 0 \text{ width}$ will be used. Alternatively, we could compute the current value of the signal just using the current value of width (i.e., ignoring the history). But this would allow the possibility of having two identically defined signals that have different values, based on when they were created. We avoid these issues by ruling out signals of signals.

Elm's type system is stratified, separating “simple types” and “signal types”. Signal types include signals over simple types, and functions that produce signal types. Note that the types do not allow signals of signals, nor functions that consume signals and produce values of primitive types. Elm's type system also ensures that signals and values are used correctly.

Simple types	$\tau ::= \mathbf{unit} \mid \mathbf{int} \mid \tau \longrightarrow \tau'$
Signal types	$\sigma ::= \mathbf{signal} \tau \mid \tau \longrightarrow \sigma \mid \sigma \longrightarrow \sigma'$
Types	$\eta ::= \tau \mid \sigma$
T-UNIT	$\frac{}{\Gamma \vdash () : \mathbf{unit}}$
T-NUMBER	$\frac{}{\Gamma \vdash n : \mathbf{int}}$
T-VAR	$\frac{\Gamma(x) = \eta}{\Gamma \vdash x : \eta}$
T-INPUT	$\frac{\Gamma(i) = \mathbf{signal} \tau}{\Gamma \vdash i : \mathbf{signal} \tau}$
T-LAM	$\frac{\Gamma, x : \eta \vdash e : \eta'}{\Gamma \vdash \lambda x : \eta. e : \eta \rightarrow \eta'}$
T-APP	$\frac{\Gamma \vdash e_1 : \eta \rightarrow \eta' \quad \Gamma \vdash e_2 : \eta}{\Gamma \vdash e_1 e_2 : \eta'}$
T-COND	$\frac{\Gamma \vdash e_1 : \mathbf{int} \quad \Gamma \vdash e_2 : \eta \quad \Gamma \vdash e_3 : \eta}{\Gamma \vdash \mathbf{if} \ e_1 \ e_2 \ e_3 : \eta}$
T-LET	$\frac{\Gamma \vdash e_1 : \eta \quad \Gamma, x : \eta \vdash e_2 : \eta'}{\Gamma \vdash \mathbf{let} \ x = e_1 \ \mathbf{in} \ e_2 : \eta'}$
T-LIFT	$\frac{\Gamma \vdash e : \tau_1 \rightarrow \dots \rightarrow \tau_n \rightarrow \tau \quad \Gamma \vdash e_i : \mathbf{signal} \tau_i \ \forall i \in 1..n}{\Gamma \vdash \mathbf{lift}_n \ e \ e_1 \dots e_n : \mathbf{signal} \tau}$
T-FOLD	$\frac{\Gamma \vdash e_f : \tau \rightarrow \tau' \rightarrow \tau' \quad \Gamma \vdash e_b : \tau' \quad \Gamma \vdash e_s : \mathbf{signal} \tau}{\Gamma \vdash \mathbf{foldp} \ e_f \ e_b \ e_s : \mathbf{signal} \tau'}$

1.2 Operational semantics

The operational semantics of Elm are actually in two phases. In the first phase, execution of the program occurs according to standard call-by-value semantics, without any evaluation of signals. This first phase of execution essentially defines a “signal graph”, that is the plumbing that describes how output signals are computed by computation over the values of input signals. In the second phase of execution, as the values of the input signals vary, the computation is performed on the new input values to produce the output values, i.e., the value of the output signals. Essentially, the first phase of computation allows us to perform computation that is independent of the values of signals, and put the program in a form where it will be possible to efficiently compute new output values as the input values change.

We define an intermediate language such that a term in the intermediate language is the result of evaluating a source language program. The grammar below defines the syntax of the intermediate language.

Values	$v ::= () \mid n \mid \lambda x : \tau. e$
Signal terms	$s ::= x \mid \mathbf{let} \ x = s \ \mathbf{in} \ u \mid i \mid \mathbf{lift}_n \ v \ s_1 \dots s_n \mid \mathbf{foldp} \ v_1 \ v_2 \ s$
Final terms	$u ::= v \mid s$

We define an evaluation-context-based call-by-value semantics to reduce Elm core-language expressions to final terms.

$$\begin{aligned}
 E &::= [\cdot] \mid E \ e \mid \mathbf{if} \ E \ e_2 \ e_3 \mid \mathbf{let} \ x = E \ \mathbf{in} \ e \mid \mathbf{let} \ x = s \ \mathbf{in} \ E \mid \\
 &\quad \mathbf{lift}_n \ E \ e_1 \dots e_n \mid \mathbf{lift}_n \ v \ s_1 \dots E \dots e_n \mid \\
 &\quad \mathbf{foldp} \ E \ e_2 \ e_3 \mid \mathbf{foldp} \ v_1 \ E \ e_3 \mid \mathbf{foldp} \ v_1 \ v_2 \ E \\
 F &::= [\cdot] \ e \mid \mathbf{if} \ [\cdot] \ e_2 \ e_3 \mid \mathbf{let} \ x = [\cdot] \ \mathbf{in} \ e \mid \mathbf{lift}_n \ [\cdot] \ e_1 \dots e_n \mid \mathbf{foldp} \ [\cdot] \ e_2 \ e_3 \mid \mathbf{foldp} \ v_1 \ [\cdot] \ e_3
 \end{aligned}$$

$$\begin{array}{c}
\text{CONTEXT} \frac{e \longrightarrow e'}{E[e] \longrightarrow E[e']} \quad \text{COND-TRUE} \frac{v \neq 0}{\text{if } v \ e_2 \ e_3 \longrightarrow e_2} \quad \text{COND-FALSE} \frac{}{\text{if } 0 \ e_2 \ e_3 \longrightarrow e_3} \\
\\
\text{APP} \frac{}{(\lambda x : \eta. e_1) \ e_2 \longrightarrow \text{let } x = e_2 \text{ in } e_1} \quad \text{REDUCE} \frac{}{\text{let } x = v \text{ in } e \longrightarrow e[v/x]} \\
\\
\text{EXPAND} \frac{x \notin fv(F[\cdot])}{F[\text{let } x = s \text{ in } u] \longrightarrow \text{let } x = s \text{ in } F[u]}
\end{array}$$


The semantics has some interesting features. First, application $(\lambda x : \eta. e_1) \ e_2$ reduces to a let expression $\text{let } x = e_2 \text{ in } e_1$ (Rule APP), and then expression e_2 is evaluated (using context $\text{let } x = E \text{ in } e$). If e_2 evaluates to a non-signal value, then substitution is performed (Rule REDUCE). Otherwise, if e_2 evaluates to a signal value, then evaluation continues on expression e_1 , without performing substitution (due to context $\text{let } x = s \text{ in } E$).

Rule EXPAND expands the scope of a let expression to allow evaluation to continue. Contexts F describe where Rule EXPAND can be applied, and includes all contexts where a simple value v is required. For a context F , we write $fv(F[\cdot])$ for the set of free variables in F . We assume that expressions are equivalent up to renaming of bound variables, and an appropriate variable x can always be chosen such that $x \notin fv(F[\cdot])$.

Let's see an example of the first phase evaluation. Suppose that i is an input signal of integers.

$$\begin{aligned}
& \text{let double} = \lambda f. \lambda x. f \ (f \ x) \text{ in} \\
& \text{let } t = \text{lift} \ (\text{double} \ (\lambda y. y + 1)) \ i \text{ in} \\
& \text{lift}_2 \ (\lambda a. \lambda b. a + b) \ i \ t \\
\longrightarrow & \text{let } t = \text{lift} \ ((\lambda f. \lambda x. f \ (f \ x)) \ (\lambda y. y + 1)) \ i \text{ in} \\
& \text{lift}_2 \ (\lambda a. \lambda b. a + b) \ i \ t \\
\longrightarrow & \text{let } t = \text{lift} \ (\lambda x. (\lambda y. y + 1) \ ((\lambda y. y + 1) \ x)) \ i \text{ in} \\
& \text{lift}_2 \ (\lambda a. \lambda b. a + b) \ i \ t
\end{aligned}$$

This last expression is a final term. No more evaluation can be performed without having some values from signals. Conceptually, this final term can be thought of as a graph that indicates which signals are consumed by which computation. The graph below shows that there are 3 signals in the program. One is the input signal i . The signal t (defined as $\text{lift} \ ((\lambda f. \lambda x. f \ (f \ x)) \ (\lambda y. y + 1)) \ i$) uses the input signal i . The program evaluates to the signal $\text{lift}_2 \ (\lambda a. \lambda b. a + b) \ i \ t$, which uses both signal i and signal t .



For the second phase evaluation, the final term of the first phase is repeatedly evaluated as new values of input signals arrive. We won't dig into the semantics here, but you can't find a semantics in the paper "Asynchronous Functional Reactive Programming for GUIs."¹ In the paper, the semantics for the second phase evaluation is given by a translation to a CML program, since in that paper we are concerned with a concurrent semantics for signal evaluation.

¹E. Czaplicki and S. Chong. Asynchronous functional reactive programming for GUIs. In *Proceedings of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 411–422, New York, NY, USA, June 2013. ACM Press.

2 Continuous FRP

The Elm language assumes that signals are discrete. That is, the value produced by a signal changes discretely. For some signals, such as mouse clicks, key presses, etc., this is clearly the right choice, as the signal represents discrete events in the external environment. Indeed, any practical FRP language needs to be able to represent discrete signals. But some signals are naturally continuous. For example, the current position of the mouse could be seen as a continuous signal, as could other signals such as time, or the position of a glyph on a canvas.

Clearly, even in a language with continuous signals, in implementation such signals will be approximated by discrete values. However, continuous signals can have an elegant semantics that are amenable to “rendering” them at different resolution. That is, based on the computing power of a particular machine on which you are running an FRP, the machine may be able to sample an external signal (e.g., the mouse position, a sensor reading) at higher or lower frequency. In a FRP language with continuous signals, a higher sampling frequency means that the implementation is more closely approximating the “ideal” semantics, analogous to how a higher-resolution monitor may render a vector graphic more accurately than a lower-resolution monitor. By contrast, in a FRP language with discrete signals, the frequency of sampling may impact the meaning of the program. For example, if we are performing a `foldp` on the horizontal mouse position, summing up the values, then sampling the mouse position with higher frequency will result in larger sums. (Elm provides a primitive signal that will attempt to sample at a specified frequency, but will adjust based on the machine’s capabilities.)

Something for you to think about: what is the equivalent of `foldp` for a continuous signal?