

Optimisation of Dynamic, Hybrid Signal Function Networks

Neil Sculthorpe and Henrik Nilsson

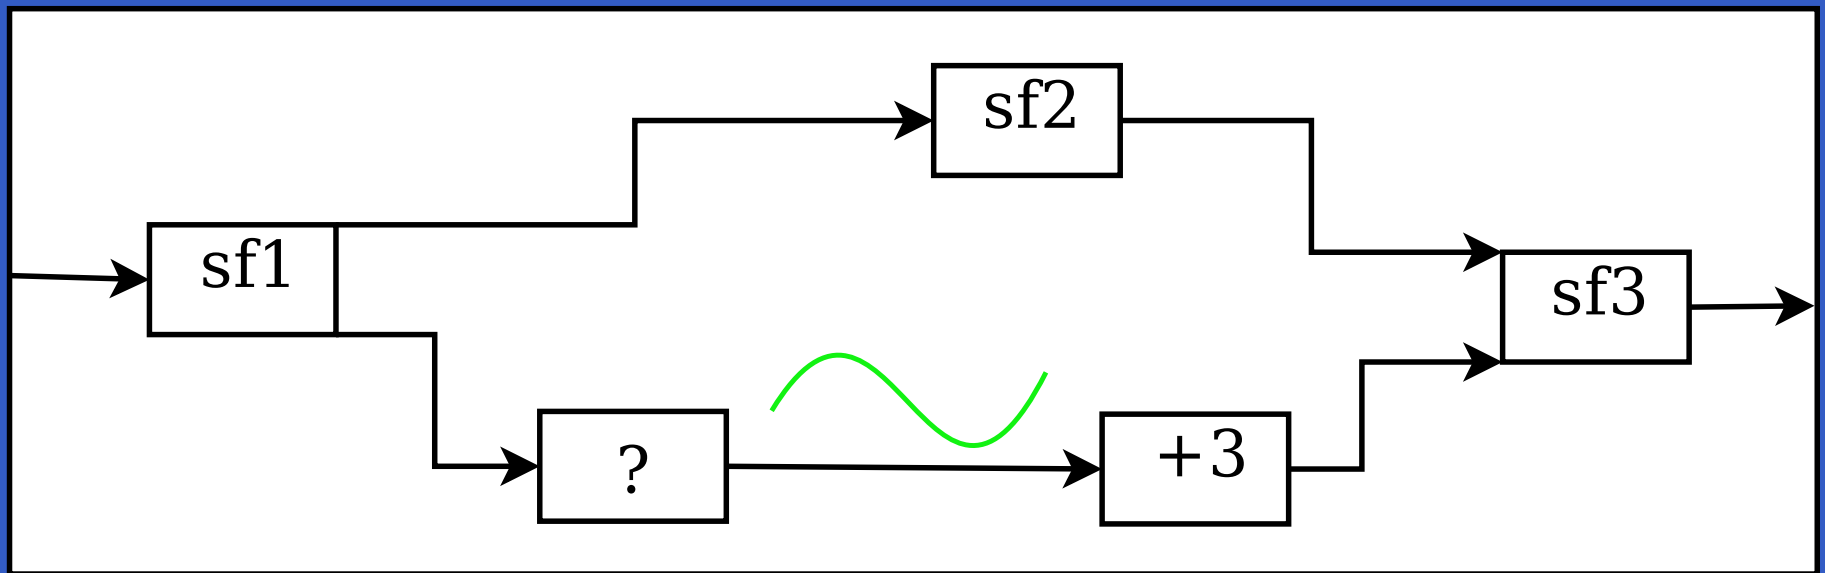
School of Computer Science
University of Nottingham, United Kingdom

Trends in Functional Programming, 27th May 2008

Outline

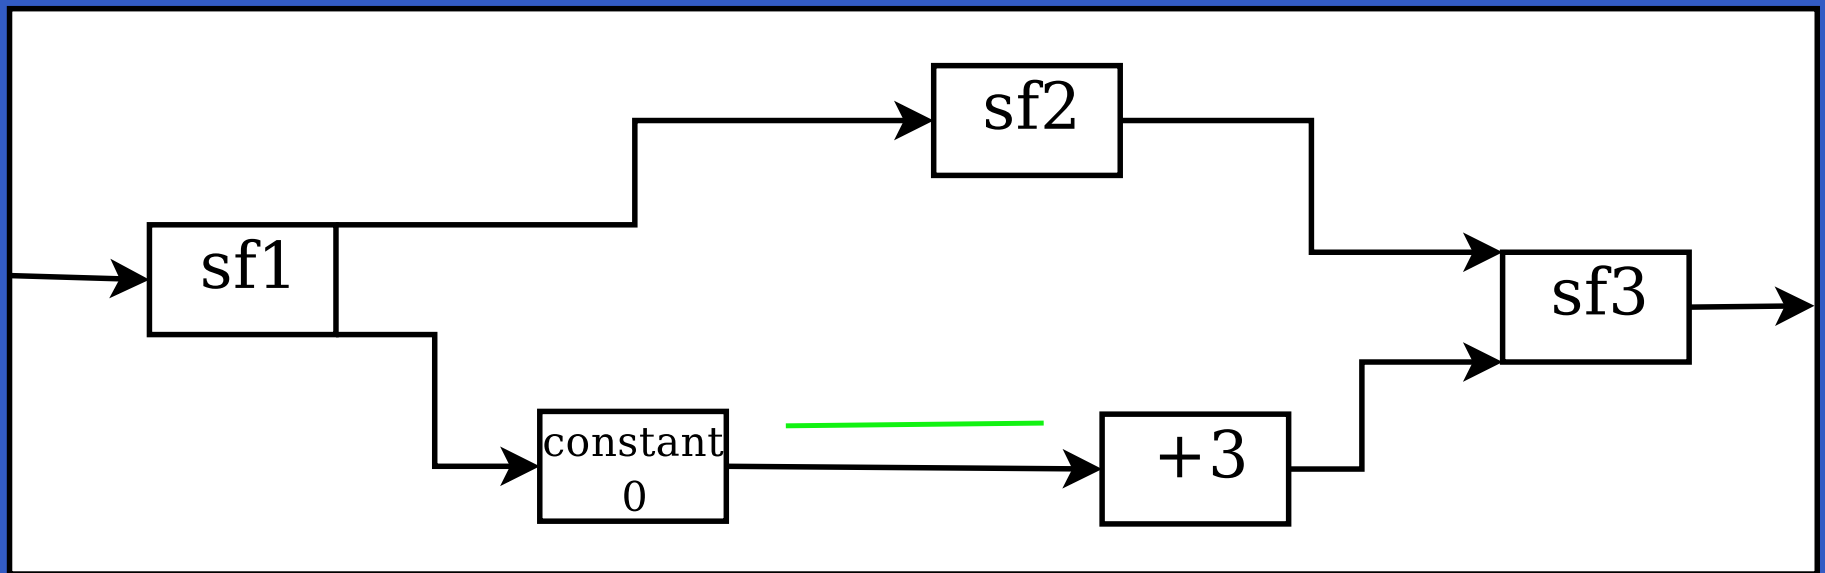
- A Brief Example
- Functional Reactive Programming and Yampa
- Our New Conceptual Framework
- A Notion of Change in that Framework
- Optimisation Examples

Example Signal Function Network



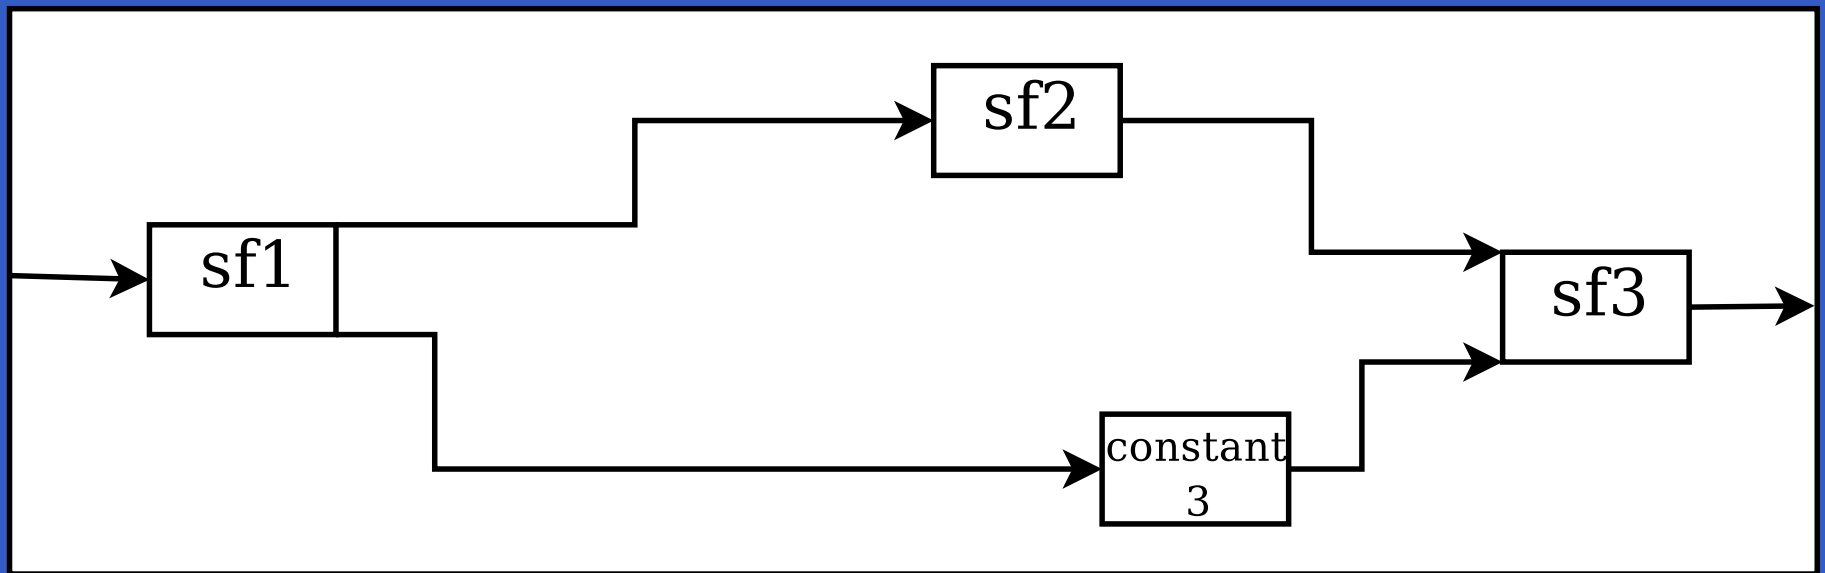
A synchronous data-flow network with hybrid and dynamic aspects.

Example Signal Function Network



A synchronous data-flow network with hybrid and dynamic aspects.

Example Signal Function Network



A synchronous data-flow network with hybrid and dynamic aspects.

Reactive Programming

- *Reactive program*: one that continually interacts with its environment, interleaving input and output in a *timely* manner.

Reactive Programming

- *Reactive program*: one that continually interacts with its environment, interleaving input and output in a *timely* manner.
- Examples include robot controllers, video games, and aeroplane control systems.

Reactive Programming

- *Reactive program*: one that continually interacts with its environment, interleaving input and output in a *timely* manner.
- Examples include robot controllers, video games, and aeroplane control systems.
- Contrast this with a *transformational program*: one that takes all input at the start of execution, and produces all output at the end (e.g. a compiler).

Reactive Programming

- *Reactive program*: one that continually interacts with its environment, interleaving input and output in a *timely* manner.
- Examples include robot controllers, video games, and aeroplane control systems.
- Contrast this with a *transformational program*: one that takes all input at the start of execution, and produces all output at the end (e.g. a compiler).
- Functional Reactive Programming (FRP) is a functional approach to reactive programming.

Yampa: A recent FRP Implementation

- *Signals* are time-varying values.

Signal $a \approx \text{Time} \rightarrow a$

Yampa: A recent FRP Implementation

- *Signals* are time-varying values.

$$\textit{Signal } a \approx \textit{Time} \rightarrow a$$

- *Signal Functions* are functions mapping signals to signals.

$$\textit{SF } a b \approx \textit{Signal } a \rightarrow \textit{Signal } b$$

Yampa: A recent FRP Implementation

- *Signals* are time-varying values.

$$\textit{Signal } a \approx \textit{Time} \rightarrow a$$

- *Signal Functions* are functions mapping signals to signals.

$$\textit{SF } a \ b \approx \textit{Signal } a \rightarrow \textit{Signal } b$$

- Signal functions can be *stateful* or *stateless*.

Yampa: Hybrid and Dynamic

- *Hybrid*: Conceptually both discrete-time (events) and continuous-time signals.

Yampa: Hybrid and Dynamic

- *Hybrid*: Conceptually both discrete-time (events) and continuous-time signals.
- In practice, events (conceptually sparse occurrences) are embedded in continuous-time signals.

$$\text{data } Event\ a = NoEvent$$
$$| \quad Event\ a$$

Yampa: Hybrid and Dynamic

- *Hybrid*: Conceptually both discrete-time (events) and continuous-time signals.
- In practice, events (conceptually sparse occurrences) are embedded in continuous-time signals.

$$\text{data } Event\ a = NoEvent$$
$$| \quad Event\ a$$

- *Dynamic*: The structure of the network can change at run-time.


Yampa: Hybrid and Dynamic

- *Hybrid*: Conceptually both discrete-time (events) and continuous-time signals.
- In practice, events (conceptually sparse occurrences) are embedded in continuous-time signals.


$$\text{data } Event\ a = NoEvent$$
$$| Event\ a$$

- *Dynamic*: The structure of the network can change at run-time.
- Signal functions (not signals!) are first class.


Some Yampa Primitives




$pure :: (a \rightarrow b) \rightarrow SF\ a\ b$



$(\gg\gg) :: SF\ a\ b \rightarrow SF\ b\ c \rightarrow SF\ a\ c$



$(**) :: SF\ a\ c \rightarrow SF\ b\ d \rightarrow$
 $SF\ (a, b)\ (c, d)$



$switch :: SF\ a\ (b, Event\ e) \rightarrow$
 $(e \rightarrow SF\ a\ b) \rightarrow SF\ a\ b$

A New Conceptual Framework

- Our new framework better supports optimisation than Yampa.

A New Conceptual Framework

- Our new framework better supports optimisation than Yampa.
- We stress that this is a *conceptual* framework, not an actual implementation.

A New Conceptual Framework

- Our new framework better supports optimisation than Yampa.
- We stress that this is a *conceptual* framework, not an actual implementation.
- The key differences from Yampa are:

A New Conceptual Framework

- Our new framework better supports optimisation than Yampa.
- We stress that this is a *conceptual* framework, not an actual implementation.
- The key differences from Yampa are:
 - a type system that makes a precise distinction between *discrete-time* and *continuous-time* signals,

A New Conceptual Framework

- Our new framework better supports optimisation than Yampa.
- We stress that this is a *conceptual* framework, not an actual implementation.
- The key differences from Yampa are:
 - a type system that makes a precise distinction between *discrete-time* and *continuous-time* signals,
 - using (heterogeneous) vectors of signals instead of nested tuples.

Continuous and Discrete Time

We give a new conceptual definition of signals to make a clear distinction between continuous and discrete time.

$\text{type } CSignal\ a \approx Time \rightarrow a$

$\text{type } ESignal\ a \approx Time \rightarrow Maybe\ a$

$\text{data } Signal\ a = C\ (CSignal\ a)$
 $\quad \quad \quad | E\ (ESignal\ a)$

Signal Vectors

Instead of tuples, we introduce *signal vectors* (a type level construct) to combine signals.

$$\begin{aligned} \text{SigVec} = & \langle \rangle \\ & | \langle C \ t \rangle \\ & | \langle E \ t \rangle \\ & | \text{SigVec} \text{ :+:+} \text{ SigVec} \end{aligned}$$



Signal Vectors

Instead of tuples, we introduce *signal vectors* (a type level construct) to combine signals.

$$\begin{array}{l} \text{SigVec} = \langle \rangle \\ | \langle C \ t \rangle \\ | \langle E \ t \rangle \\ | \text{SigVec} \text{ :+ : } \text{SigVec} \end{array}$$

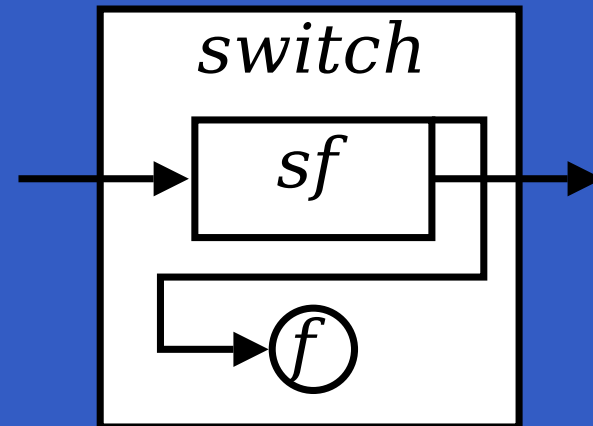

$$\text{type } (\text{SigVec } as, \text{SigVec } bs) \Rightarrow SF \ as \ bs \approx as \rightarrow bs$$

Primitives in this framework



$pure :: (a \rightarrow b) \rightarrow SF \langle td \ a \rangle \langle td \ b \rangle$

$(\gg\gg) :: SF \ as \ bs \rightarrow SF \ bs \ cs \rightarrow SF \ as \ cs$



$(**) :: SF \ as \ cs \rightarrow SF \ bs \ ds \rightarrow$
 $SF \ (as \ ::\ bs) \ (cs \ ::\ ds)$

$switch :: SF \ as \ (\langle E \ e \rangle \ ::\ bs) \rightarrow$
 $(e \rightarrow SF \ as \ bs) \rightarrow SF \ as \ bs$

Why is this useful for Optimisation?

- While we model time as continuous, at the implementation level a signal function network is executed over a discrete sequence of sample times.

Why is this useful for Optimisation?

- While we model time as continuous, at the implementation level a signal function network is executed over a discrete sequence of sample times.
- Many signal functions will produce the same output at many (if not all) sample times.

Why is this useful for Optimisation?

- While we model time as continuous, at the implementation level a signal function network is executed over a discrete sequence of sample times.
- Many signal functions will produce the same output at many (if not all) sample times.
- We would like to avoid re-computation of unchanged signals.

Change

We need a precise notion of change, which we define as follows:

Change

We need a precise notion of change, which we define as follows:

- A continuous-time signal is changing *iff* its value at the current time sample differs from its value at the preceding time sample.

Change

We need a precise notion of change, which we define as follows:

- A continuous-time signal is changing *iff* its value at the current time sample differs from its value at the preceding time sample.
- An event signal is changing *iff* there is an event occurrence at the current time sample.

Signal Function Change Classifications

- *Unchanging* (U) signal functions produce unchanging output.

Signal Function Change Classifications

- *Unchanging* (U) signal functions produce unchanging output.
- *Input-Dependent* (I) signal functions where unchanging input \Rightarrow unchanging output.

Signal Function Change Classifications

- *Unchanging* (U) signal functions produce unchanging output.
- *Input-Dependent* (I) signal functions where unchanging input \Rightarrow unchanging output.
- *Varying* (V) signal functions where the output may always change, regardless of input.

Examples

constant :: $c \rightarrow SF_U \text{ as } \langle C \ c \rangle$

never :: $SF_U \text{ as } \langle E \ e \rangle$

pure :: $(a \rightarrow b) \rightarrow SF_I \langle td \ a \rangle \langle td \ b \rangle$

edge :: $SF_I \langle C \ Bool \rangle \langle E \ () \rangle$

iPre :: $a \rightarrow SF_V \langle C \ a \rangle \langle C \ a \rangle$

Combining Change Classifications

$$(\ggg) \quad :: SF_x \text{ as } bs \rightarrow SF_y \text{ bs } cs \rightarrow SF_{(x \ggg y)} \text{ as } cs$$

$$(\ast) \quad :: SF_x \text{ as } cs \rightarrow SF_y \text{ bs } ds \rightarrow SF_{(x \sqcup y)} (as \ast bs) (cs \ast ds)$$

$$\text{switch} \quad :: SF_x \text{ as } (\langle E \ e \rangle \ast bs) \rightarrow (e \rightarrow SF_y \text{ as } bs) \rightarrow SF_{(x \text{ 'sw' } y)} \text{ as } bs$$

data *ChangeClass* = *U* | *I* | *V* **deriving** (*Eq*, *Ord*)

$$(\ggg) \quad :: \text{ChangeClass} \rightarrow \text{ChangeClass} \rightarrow \text{ChangeClass}$$

$$x \ggg U = U$$

$$x \ggg V = V$$

$$x \ggg I = x$$

$$\text{sw} \quad :: \text{ChangeClass} \rightarrow \text{ChangeClass} \rightarrow \text{ChangeClass}$$

$$U \text{ 'sw' } y = U$$

$$x \text{ 'sw' } y = x \sqcup y$$

Some Useful Optimisation Properties

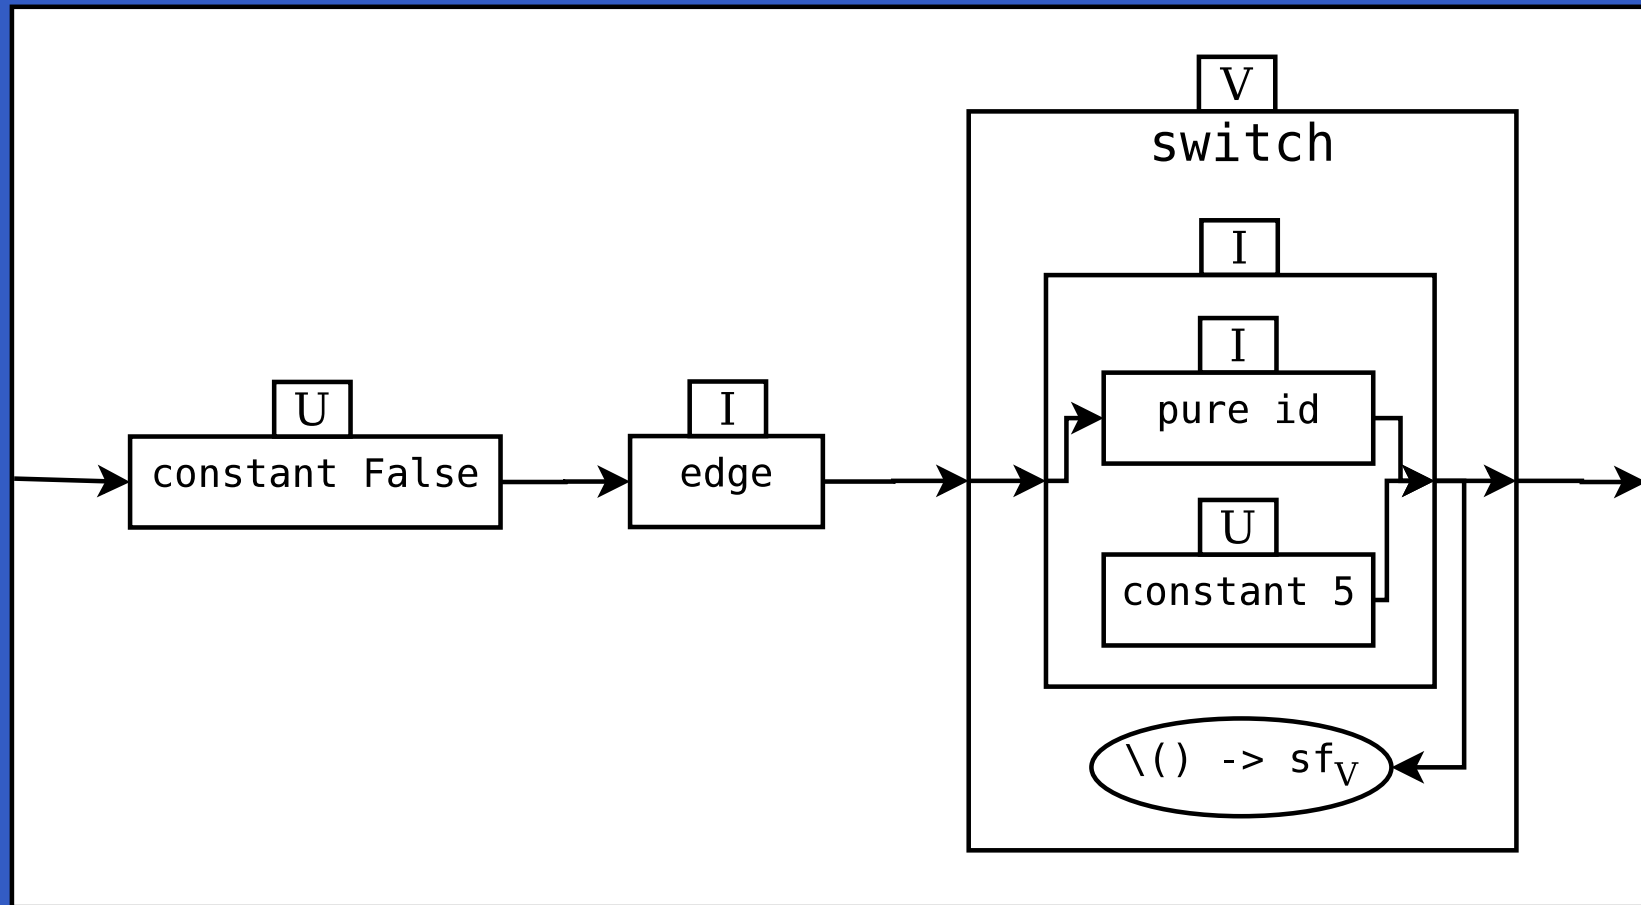
- Any composite *unchanging* signal function can be compressed into a single signal function.

Some Useful Optimisation Properties

- Any composite *unchanging* signal function can be compressed into a single signal function.
- Unchanging signal functions distribute into switches over sequential composition:

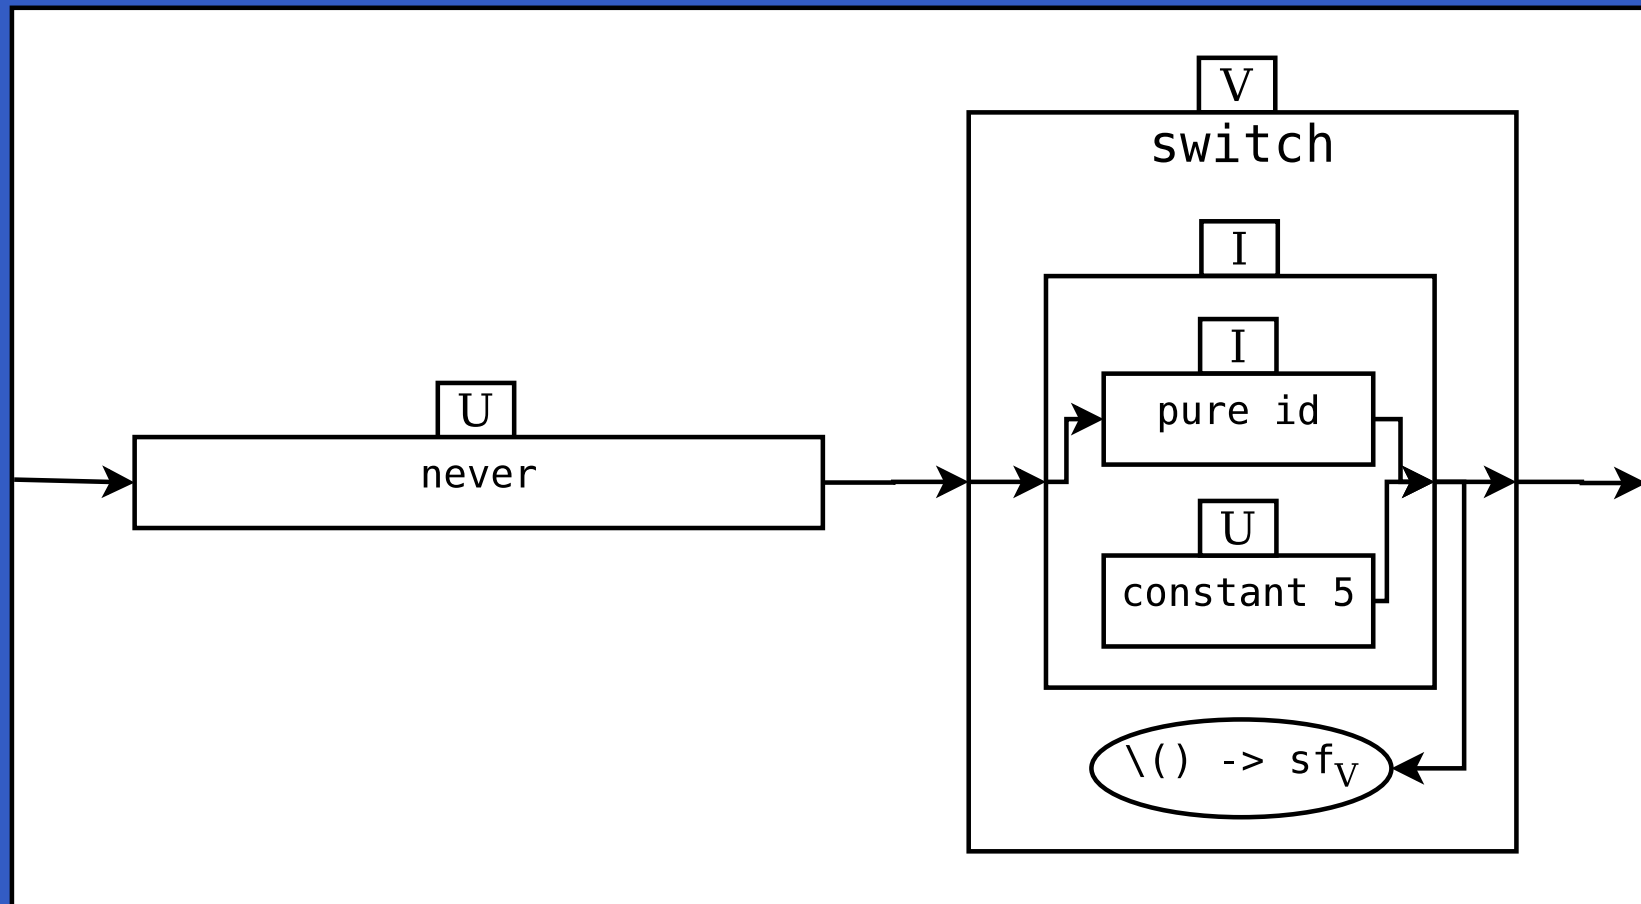
$$\begin{aligned} & sf1_U \ggg switch\ sf2\ f \\ \equiv & \\ & switch\ (sf1_U \ggg sf2)\ (\lambda a \rightarrow sf1_U \ggg f\ a) \end{aligned}$$

Example Optimisation



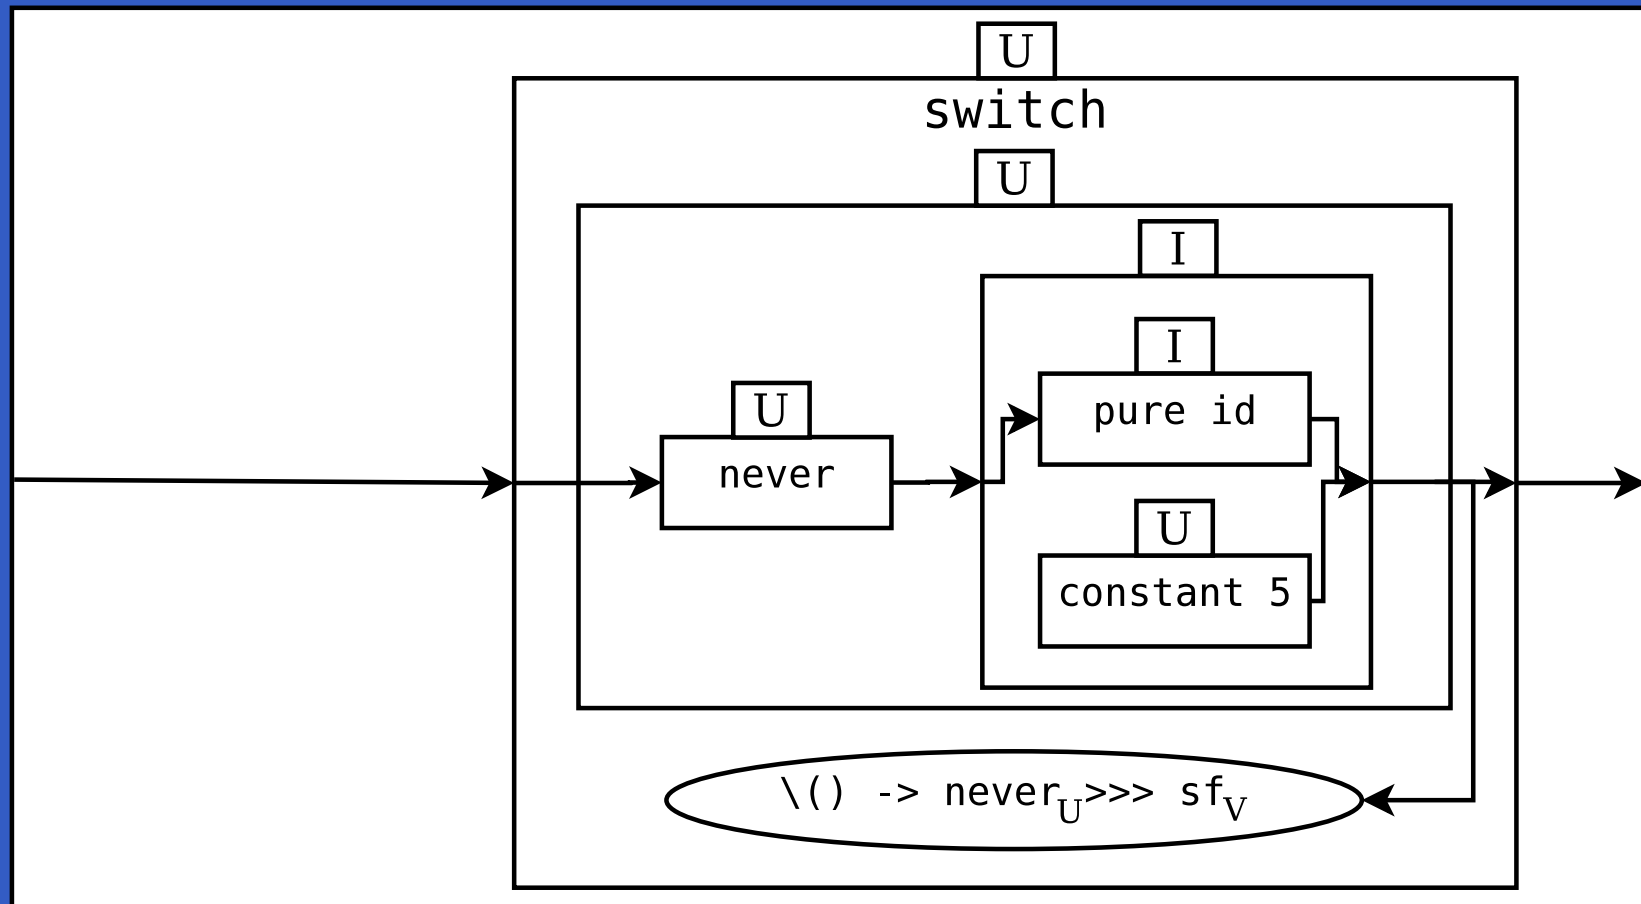
constant False \ggg *edge* \ggg *switch* (*pure id* \ast *constant 5*) ($\lambda() \rightarrow sf$)

Example Optimisation



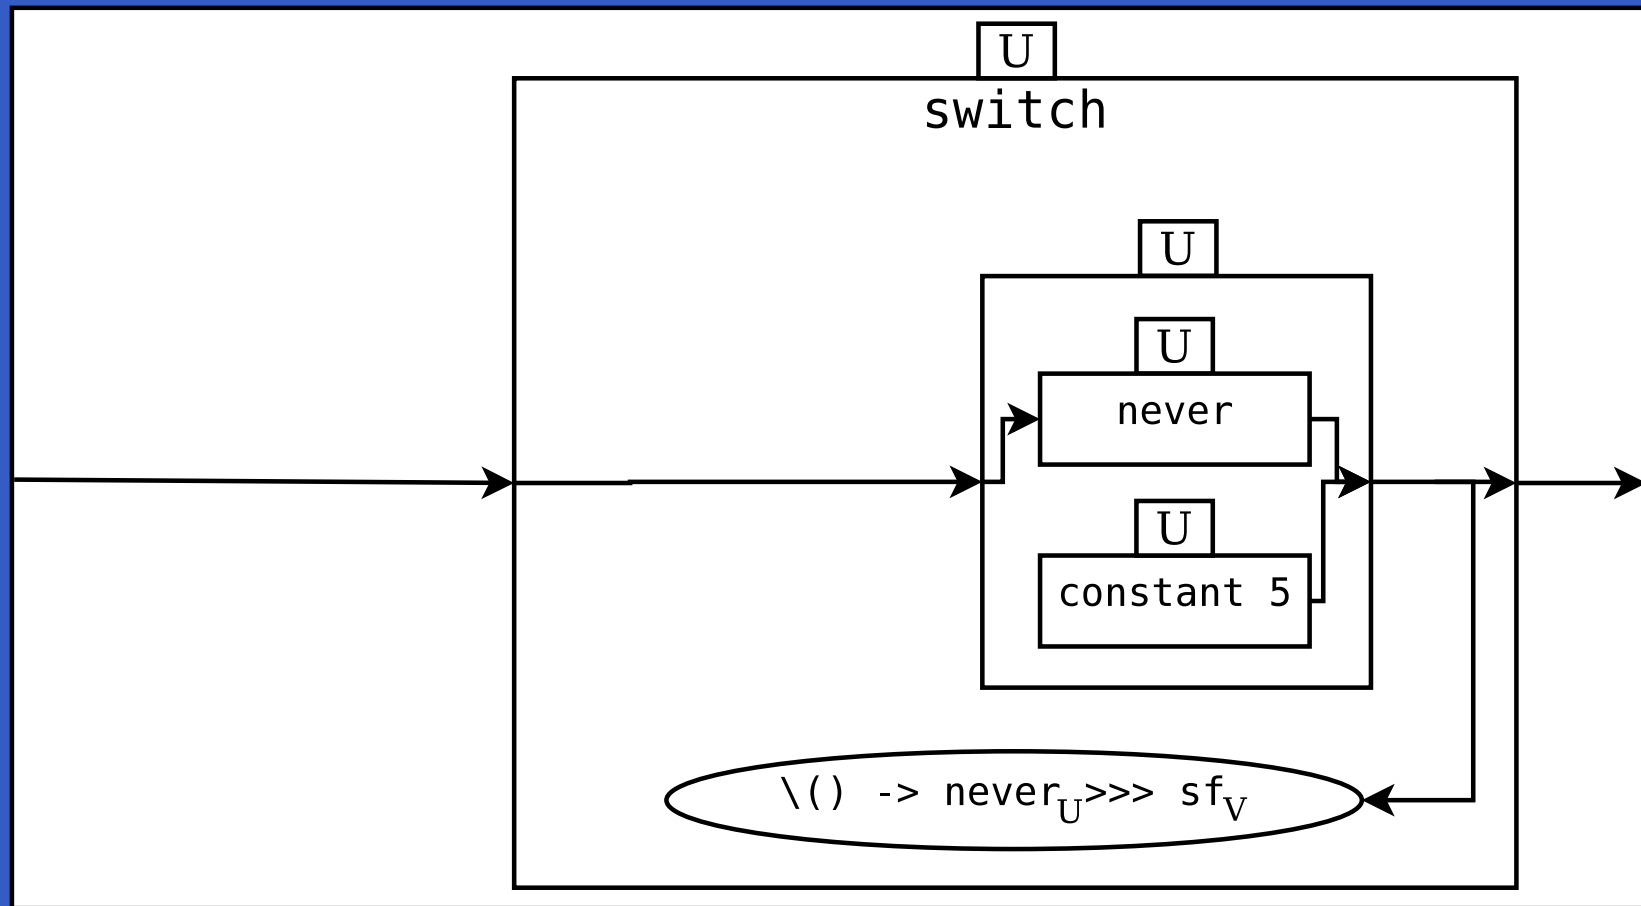
$never \ggg switch (pure\ id \ast\ constant\ 5) (\lambda() \rightarrow sf_v)$

Example Optimisation



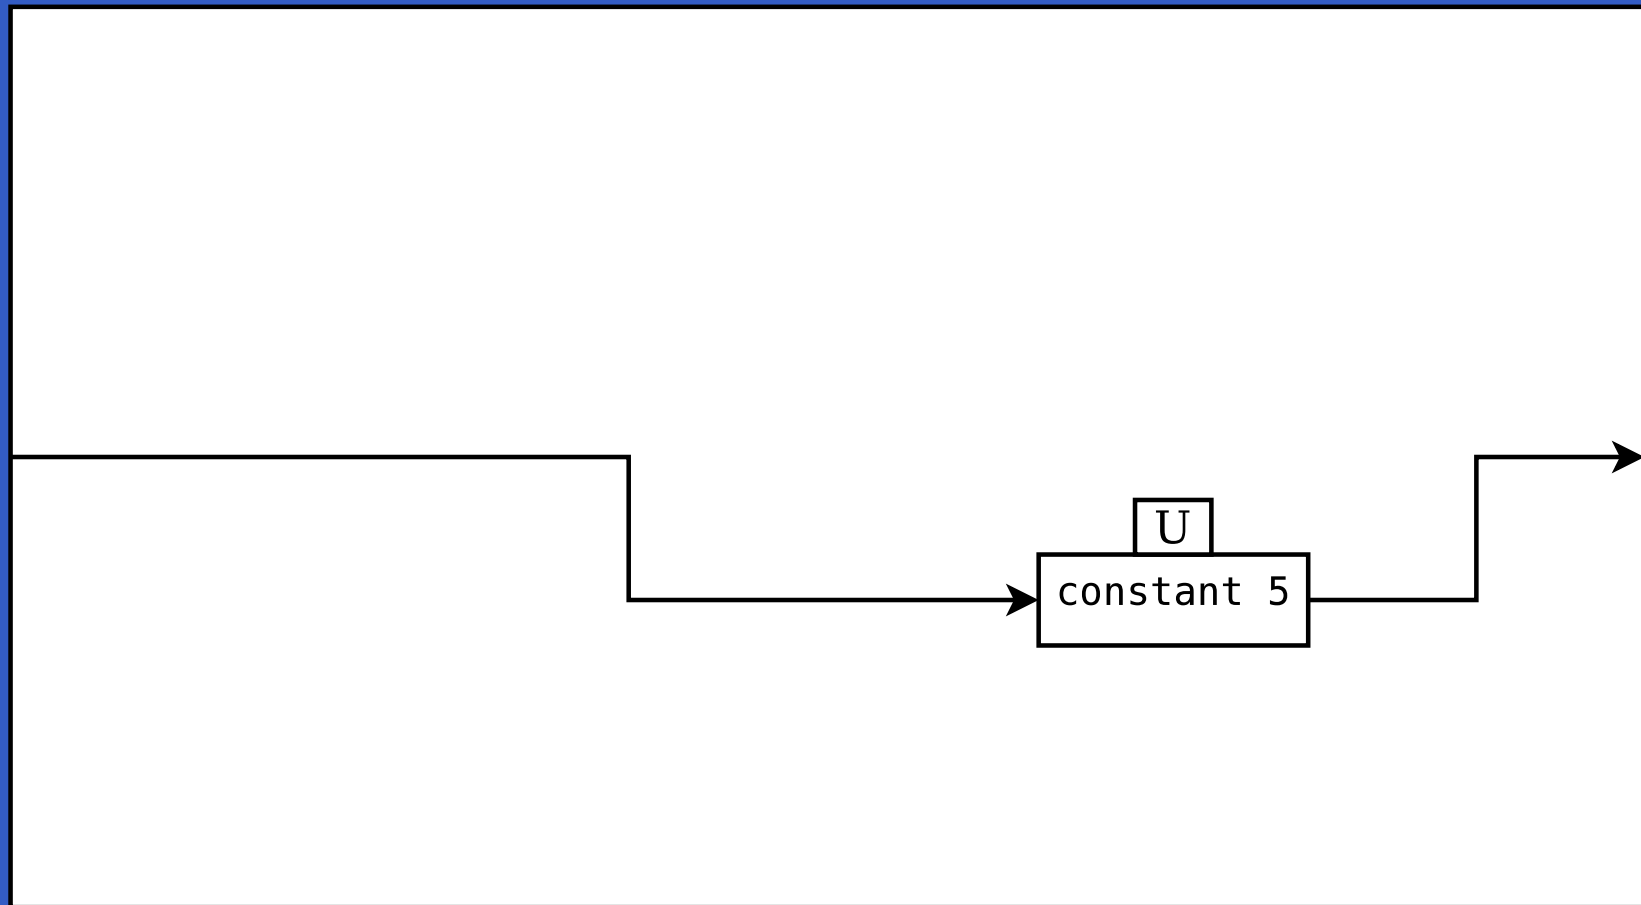
*switch (never \gg pure id $**$ constant 5) ($\lambda() \rightarrow$ never \gg sf)*

Example Optimisation



*switch (never ** constant 5) ($\lambda() \rightarrow sf$)*

Example Optimisation



constant 5

Summary

- We have discussed:
 - Our conceptual framework
 - A notion of change within that framework
 - Some optimisations that exploit that notion of change

Summary

- We have discussed:
 - Our conceptual framework
 - A notion of change within that framework
 - Some optimisations that exploit that notion of change
- Our framework also supports run-time *change propagation* optimisations.