

Mixed-Level Embedding and JIT Compilation for an Iteratively Staged DSL

George Giorgidze and Henrik Nilsson

Functional Programming Laboratory
School of Computer Science
University of Nottingham
United Kingdom
{ggg,nhn}@cs.nott.ac.uk

Abstract. This paper explores how to implement an *iteratively staged* domain-specific language (DSL) by embedding into a functional language. The domain is modelling and simulation of physical systems where models are expressed in terms of non-causal differential-algebraic equations; i.e., sets of constraints solved through numerical simulation. What distinguishes our language is that the equational constraints are *first class entities* allowing for an evolving model structure characterised by *repeated generation* of updated constraints. Hence iteratively staged. Our DSL can thus be seen as a combined functional and constraint programming language, albeit a two-level one, with the functional language chiefly serving as a meta language. However, the two levels do interact throughout the simulation. The embedding strategy we pursue is a mixture of deep and shallow, with the deep embedding enabling just-in-time (JIT) compilation of the constraints as they are generated for efficiency, while the shallow embedding is used for the remainder for maximum leverage of the host language. The paper is organised around a specific DSL, but our implementation strategy should be applicable for iteratively staged languages in general. Our DSL itself is further a novel variation of a declarative constraint programming language.

1 Introduction

Embedding is a powerful and popular way to implement domain-specific languages (DSLs) [11]. Compared with implementing a language from scratch, extending a suitable general-purpose programming language, the *host language*, with notions and vocabulary addressing a particular application or problem domain tends to save a lot of design and implementation effort. Or, to put this differently, for a given implementation effort, embedding usually results in a more capable and complete language than would an implementation from scratch.

There are two basic approaches to language embeddings: *shallow* and *deep*. In a shallow embedding, domain-specific notions are expressed directly in host-language terms, typically through a higher-order combinator library. This is a light-weight approach that makes it easy to leverage the facilities of the host

language. However, the syntactic freedom is limited, and the static semantics of the embedded language must be relatively close to that of the host language for an embedding to be successful. In contrast, a deep embedding is centred around a *representation* of embedded language terms that then are given meaning by interpretation or compilation. This is a more heavy-weight approach, but also more flexible. In particular, for optimisation or compilation, it is often necessary to inspect terms, suggesting a deep embedding. The two approaches can be combined to draw on the advantages of each. This leads to *mixed-level* embedding.

In this paper, we explore how to embed a language, Hydra [20], for modelling and simulation of physical systems into a functional language. Hydra is a *non-causal* language where systems are modelled by constraints expressed as *undirected* Differential Algebraic Equations (DAEs). These equations are solved by specialised symbolic and numerical simulation methods. A defining aspect of Hydra is that the equations are *first-class entities* at the functional level, allowing for flexible model composition and evolving model structure. Specifically, at an *event*, simulation is stopped and, *depending* on results thus far, (partly) new equations are *generated* describing a (partly) new problem to be solved. We refer to this kind of DSL as *iteratively staged* to emphasise that the *domain* is characterised by repeated program generation and execution. Iterative staging allows Hydra to handle classes of systems that current main-stream non-causal modelling languages cannot [7,19]; see Sec. 2 for an example.

Hydra can be seen as a functional and constraint or logical programming language in that it combines a functional and a relational approach to programming. However, the integration of the two approaches is less profound than in, say, functional logic languages based on residuation or narrowing [9]. Hydra is a two-level language, where the functional part to a large extent serves as a meta language. However, the two layers do interact throughout the simulation.

We have chosen Haskell with Glasgow Haskell Compiler (GHC) extensions as the host language, GHC’s quasiquoting facility [14,15] being one reason for this choice. Because performance is a primary concern in our application domain, the simulation code corresponding to the current equations has to be compiled into native code. As this code is determined *dynamically*, this necessitates *just-in-time* (JIT) compilation. For this part of the language we use a deep embedding along with the Low Level Virtual Machine (LLVM) [12]: a language-independent, optimising, compiler back-end with JIT compilation support. In contrast, we retain a shallow embedding for high-level, symbolic computation to get maximum leverage from the host language. Note that we are not concerned with (hard) *real-time* performance here: we are prepared to pay the price of brief “pauses” for symbolic processing and JIT compilation to minimise the computational cost of the actual simulation as this typically dominates the overall costs.

An alternative might have been to use a *multi-staged* host language like MetaOCaml [24] as we then in principle should get run-time code generation for free. We have so far not explored this approach as we wanted to have tight control over the generated code. Also, not predicating our approach on a multi-staged host language means that some of our ideas and implementation techniques

can be more readily deployed in other contexts; for example, to enhance the capabilities of existing implementations of non-causal languages, or to settings where code needs to be generated for co-processors of various kinds.

Compilation of Embedded DSLs (EDSLs) is today a standard tool in the DSL-implementer’s tool chest. The seminal example, at least in a functional setting, is Pan [4]. More recent examples include Paradise [1] and Flask [15]. However, in all these languages, code generation and compilation happens once, followed by execution, meaning that the host language acts as a powerful but fundamentally conventional macro language.

As Hydra is iteratively staged, the problems we are facing are in many ways different. Moreover, rather than acting merely as a macro language that is out of the picture once the generated program is ready for execution, the host language is in our case part of the dynamic semantics of the embedded language through the shallow parts of the embedding. With this paper, we thus add further tools to the DSL tool chest for embedding a class of languages that hitherto has not been studied much. Specifically, our contributions are:

- a case study of mixed-level embedding of iteratively staged DSLs;
- using JIT compilation to implement an iteratively staged EDSL efficiently.

Additionally, we consider static type checking in the context of iterative staging and quasiquoting-based embedding. While Hydra is specialised, we believe the ideas underlying the implementation are of general interest, and that Hydra itself should be of interest to programming language researchers interested in languages that combine functional and relational approaches to programming. The implementation is available on-line¹ under the open source BSD license.

The rest of the paper is organised as follows. In Sec. 2, we introduce non-causal modelling and our language Hydra in more detail. Sec. 3 explains the Haskell embedding of Hydra and Sec. 4 then describes how iteratively staged programs are executed. Related work is discussed in Sec. 5 and Sec. 6 concludes.

2 Background

This section provides an introduction to Functional Hybrid Modelling (FHM) [20] and to our specific instance Hydra [6,7]. We focus on aspects that are particularly pertinent to the present setting. The reader is referred to the earlier papers on FHM and Hydra for a more general treatment.

2.1 Functional Hybrid Modelling

Functional Hybrid Modelling (FHM) [20] is a new approach to designing *non-causal modelling languages* [3] supporting *hybrid systems*: systems that exhibit both continuous- and discrete-time behaviour. This class of languages is intended for modelling and simulation of systems that can be described by DAEs. Examples include electrical, mechanical, hydraulic, and other physical systems, as

¹ <http://cs.nott.ac.uk/~ggg/>

well as their combinations. *Non-causal*² in this context refers to treating the equations as being *undirected*: an equation can be used to solve for any of the variables occurring in it. This is in contrast to *causal* modelling languages where equations are restricted to be *directed*: only “known” variables on one side of the equal sign, and only “unknown” variables on the other. Causal modelling languages are consequently somewhat akin to functional programming languages, while non-causal ones are akin to (constraint) logic programming languages. However, note that the domain of the variables in our case is *time-varying values* or *signals*: functions of continuous time.

An advantage of causal modelling languages is that simulation is relatively straightforward thanks to equations being directed. The advantages of non-causal languages over causal ones include that models are more *reusable* (the equations can be used in many ways) and more *declarative* (the modeller can focus on *what* to model, worrying less about *how* to model it to enable simulation) [3]. These are crucial advantages in many modelling domains. As a result, a number of successful non-causal modelling languages have been developed. Modelica [17] is a prominent, state-of-the-art example.

However, one well-known weakness of current non-causal languages is that their support for modelling *structurally dynamic systems*, systems where the equations that describe the dynamic behaviour change at discrete points in time, usually is limited [21]. There are a number of reasons for this. A fundamental one is that languages like Modelica, to facilitate efficient simulation, are designed and subsequently implemented on the assumption that the model is translated into simulation code once and for all, *before* simulation starts. This severely limits the possibilities for describing structurally dynamic systems, in particular *highly* structurally dynamic ones, where the number of structural configurations, or system *modes*, is large, unbounded, or impossible to determine in advance.

The idea of FHM is to enrich a purely functional language with a few key abstractions for supporting hybrid, non-causal modelling. In particular, first-class *signal relations*, relations on signals described by undirected DAEs, provide support for non-causal modelling, and *dynamic switching* among signal relations that are *computed* at the point when they are being “switched in” provides support for describing highly structurally dynamic systems [20].

Our hypothesis is that the FHM approach will result in non-causal modelling languages that are more expressive than the current ones, yet have relatively simple, declarative semantics. Results so far have been promising. The capability to compute and use new signal relations during simulation has already allowed us to non-causally model and simulate systems that, for example, Modelica cannot handle [7,19]. We present one such example in the following. The dynamic computation of and switching among signal relations is, of course, also what makes FHM iteratively staged.

² Do not confuse this with *temporal* causality. A system is temporally causal if its output only depends on present and past input, and temporally non-causal if the output depends on future input.




Fig. 1. A pendulum subject to gravity

2.2 Hydra by Example: The Breaking Pendulum

To introduce Hydra, let us model a physical system whose structural configuration changes abruptly during simulation: a simple pendulum that can break at a specified point in time; see Fig. 1. The pendulum is modelled as a body represented by a point mass m at the end of a rigid, mass-less rod, subject to gravity $m\vec{g}$. If the rod breaks, the body will fall freely. This makes the differences between the two configurations sufficiently large that e.g. Modelica does not support non-causal modelling of this system. Instead, if simulation across the breaking point is desired, the modeller is forced to model the system in a causal, less declarative way [16, pp. 31–33].

As already mentioned, there are two levels to Hydra: the *functional level* and the *signal level*. The functional level is concerned with the definition of ordinary functions operating on time-invariant values. The signal level is concerned with the definition of relations on signals, the *signal relations*, and, *indirectly*, the definition of the *signals* themselves as solutions satisfying these relations.

Signal relations are *first-class* entities at the functional level. The type of a signal relation is parametrised on a *descriptor* of the types of the signals it relates: essentially a tuple of the types carried by the signals. For example, the type of a signal relation relating three real-valued signals is $SR (Real, Real, Real)$.

Signals, in contrast to signal relations, are *not* first-class entities at the functional level. However, crucially, *instantaneous values* of signals can be propagated back to the functional level, allowing the future system structure to depend on signal values at specific points in time, *events*, as explained in the following.

The definitions at the signal level may freely refer to entities defined at the functional level as the latter are time-invariant, known parameters as far as solving the equations are concerned. However, the opposite is not allowed: time-varying entities are confined to the signal level. The only signal-level notion that exists at the functional level is the *time-invariant* signal relation.

Hydra is currently implemented as an embedding in Haskell using *quasi-quoting* [14,15]. Haskell thus provides the functional level almost for free through shallow embedding. In contrast, the signal level is realised through deep embedding: signal relations expressed in terms of Hydra-specific syntax are, through the

```

type Coordinate = (Double, Double)
type Velocity = (Double, Double)
type Body = (Coordinate, Velocity)

g :: Double
g = 9.81

freeFall :: Body → SR Body
freeFall ((x0, y0), (vx0, vy0)) = [hydra|
  sigrel ((x, y), (vx, vy)) where
    init (x, y)      = ($x0$, $y0$/)
    init (vx, vy)   = ($vx0$, $vy0$/)
    (der x, der y)   = (vx, vy)
    (der vx, der vy) = (0, -$g$/)
  ]

pendulum :: Double → Double
           → SR Body
pendulum l phi0 = [hydra|
  sigrel ((x, y), (vx, vy)) where
    init phi      = $ phi0 $
    init der phi = 0
    init vx      = 0
    init vy      = 0
    x           = $ l $ * sin phi
    y           = - $ l $ * cos phi
    (vx, vy)    = (der x, der y)
    der (der phi)
      + ($g / l $) * sin phi = 0
  ]

```

Fig. 2. Models of the two modes of the pendulum

quasiquote machinery, turned into an internal representation, essentially an abstract syntax tree (AST), that then is compiled into simulation code. This, along with the reasons for using quasiquote, is discussed in more detail in an earlier paper [6]. However, that paper only treated *structurally static* systems.

Fig. 2 illustrates how to model the two modes of the pendulum in Hydra. The type *Body* represents the possible states of the pendulum body, its position and velocity, both 2-dimensional vectors represented by pairs of doubles. Each mode is represented by a function that maps the *parameters* of the mode to a relation on signals; that is, an instance of the defining system of DAEs for specific values of the parameters. In the unbroken mode, the parameters are the length of the rod *l* and the initial angle of deviation *phi0*. In the broken mode, the signal relation is parametrised on the initial state of the body.

The character sequences [*hydra*| and] are the open and close quasiquotes. Between them, we have signal-level definitions expressed in our custom syntax. The keyword **sigrel** starts the definition of a signal relation. It is followed by a pattern that introduces *signal variables* giving local names to the signals that are going to be constrained by the signal relation. This pattern thus specifies the *interface* of a signal relation.

Note the two kinds of variables: the functional level ones representing *time-invariant* parameters, and the signal-level ones, representing *time-varying* entities, the signals. Functional-level fragments, such as variable references, are spliced into the signal level by enclosing them between antiquotes, \$. On the other hand time-varying entities are not allowed to escape to the functional level (meaning signal-variables are not in scope between antiquotes).

After the keyword **where** follow the equations that define the relation. These equations may introduce additional signal variables as needed. Equations marked by the keyword **init** are initialisation equations used to specify initial conditions. The operator *der* indicates differentiation with respect to time of the signal-valued expression to which it is applied.

The non-causal nature of Hydra can be seen particularly clearly in the last equation of the unbroken mode that simply states a constraint on the angle of deviation and its second derivative, without making any assumption regarding which of the two time-varying entities is going to be used to solve for the other (both g and l are time-invariant functional-level variables).

To model a pendulum that breaks at some point, we need to create a composite model where the model that describes the dynamic behaviour of the unbroken pendulum is replaced, at the point when it breaks, by the model describing a free falling body. These two submodels must be suitably joined to ensure the continuity of both the position and velocity of the body of the pendulum.

To this end, the *switch*-combinator, which forms new signal relations by *temporal composition*, is used:

$$\text{switch} :: (\text{Eval } b) \Rightarrow \text{SR } (a, E b) \rightarrow (b \rightarrow \text{SR } a) \rightarrow \text{SR } a$$

The composite behaviour is governed by the first signal relation until an *event* of type b occurs ($E b$ in the type signature above). At this point, the second argument to *switch* is applied to the value carried by the event to *compute* the signal relation that is going to govern the composite behaviour from then on.

The type class *Eval* allows for translation of instantaneous values of signals into corresponding Haskell values. The translation is used in the implementation of the *switch* combinator. Event signals are *discrete-time* signals, signals that are only defined at (countably many) discrete points in time, as opposed to the continuous-time signals that (conceptually) are defined everywhere on a continuous interval of time. Each point of definition of an event signal is known as an *event occurrence*. Unlike continuous-time signals, the causality of event signals is always fixed. As we will see (Sec. 3), the realisation of this combinator involves mixed-level embedding techniques as signal relations constructed by *switch* do have an explicit representation at the abstract syntax level, but part of this representation is a host-language function.

Fig. 3 shows how *switch* is used to construct a model of a breaking pendulum. The *pendulum* model is first extended into a signal relation *pendulumBE* that also provides the event signal that defines when the pendulum is to break: see Fig. 3(a). In our example, an event is simply generated at an *a priori* specified point in time, but the condition could involve arbitrary time-varying entities. The value of the event signal is the state (position and velocity) of the pendulum at the point of the event, allowing the succeeding model to be initialised so as to ensure the continuity of the position and velocity as discussed above.

To bring the equations of *pendulum* into the definition of *pendulumBE*, *pendulum* is first applied to the length of the pendulum and the initial angle of deviation at the *functional level* (within antiquotes), thus computing a signal relation. This *relation* is then *applied*, at the signal level, using the *signal relation application operator* \diamond . This instantiates the equations of *pendulum* in the context of *pendulumBE*. Unfolding signal relation application in Hydra is straightforward: the actual arguments (signal-valued expressions) to the right of the signal relation application operator \diamond are simply substituted for the corresponding formal arguments (signal variables) in the body of the signal relation

<pre> pendulumBE :: Double → Double → Double → SR (Body, E Body) pendulumBE t l phi0 = [Hydra] sigrel (((x, y), (vx, vy)), event e) where \$ pendulum l phi0 \$ ◇ ((x, y), (vx, vy)) event e = ((x, y), (vx, vy)) when time = \$t \$ </pre>	<pre> breakingPendulum :: SR Body breakingPendulum = switch (pendulumBE 10 1 (pi / 4)) freeFall </pre>
--	--

(a) Pendulum extended with a breaking event (b) Composition using switch

Fig. 3. The breaking pendulum

to the left of \diamond , renaming as necessary to avoid name clashes. The process of unfolding signal relation applications is called *flattening*. It is in many ways similar to the transformation of hierarchical models in languages like Modelica into a flat system of equations, a process usually also referred to as flattening. See [6] for further details.

Finally, a model of the breaking pendulum can be composed by switching from *pendulumBE* to *freeFall*: see Fig. 3(b). Note that the switching event carries the state of the pendulum at the breaking point as a value of type *Body*. This value is passed to *freeFall*, resulting in a model of the pendulum body in free fall initialised to ensure the continuity of its position and velocity.

In our particular example, the pendulum is only going to break once. In other words, there is not much iteration going on, and it would in principle (with a suitable language design) be straightforward to generate code for both modes of operation prior to simulation. However, this is not the case in general. For example, given a parametrised signal relation:

$$sr1 :: Double \rightarrow SR ((Double, Double), E Double)$$

we can recursively define a signal relation *sr* that describes an overall behaviour by “stringing together” the behaviours described by *sr1*:

$$sr :: Double \rightarrow SR (Double, Double)$$

$$sr\ x = switch (sr1\ x)\ sr$$

In this case, because the number of instantiations of *sr1* in general cannot be determined statically, and because each instantiation can depend on the parameter in arbitrarily complex ways, there is no way to generate all code prior to simulation. However, the pendulum example is simple and suffice for illustrative purposes. Moreover, despite its simplicity, it is already an example with which present non-causal languages struggle, as mentioned earlier.

In practical terms, the *switch*-combinator is a somewhat primitive way of describing variable model structure. Our aim is to enrich Hydra with higher-level constructs as described in the original FHM paper [20]. The basic principles of the implementation should, however, not change much.

3 Embedding

In this section, we describe the Haskell embedding of Hydra in further detail. First, we introduce a Haskell data type that represents an embedded signal relation. This representation is untyped. We then introduce typed combinators that ensure only well-typed signal relations can be constructed.

The following data type is the central, untyped representation of signal relations. There are two ways to form a signal relation: either from a set of defining equations, or by composing signal relations temporally:

```
data SigRel =
  SigRel      Pattern [Equation]
  | SigRelSwitch SigRel (Expr → SigRel)
```

The constructor *SigRel* forms a signal relation from equations. Such a relation is represented by a pattern and the list of defining equations. The pattern serves the dual purpose of describing the *interface* of the signal relation in terms of the types of values carried by the signals it relates and their time domains (continuous time or discrete time/events), and of introducing names for these signals for use in the equations. Patterns are just nested tuples of signal variable names along with indications of which are event signals: we omit the details. The list of equations constitute a system of DAEs³ that defines the signal relation by expressing constraints on the (signal) variables introduced by the pattern and any additional local signal variables. The various kinds of equations are discussed below.

The *switch*-combinator forms a signal relation by temporal composition of two signal relations. Internally, such a temporal composition is represented by a signal relation constructed by *SigRelSwitch*. The first argument is the signal relation that is initially active. The second argument is the function that, in the case of an event occurrence from the initially active relation, is applied to the value carried by that occurrence to compute a new signal relation. Here, the *value subset* of the type *Expr* is used to represent the value. This new signal relation then becomes the active one, replacing the initial signal relation, including the switch itself.

Note the use of a mixture of shallow and deep techniques of embedding. The embedded function in a signal relation constructed by *SigRelSwitch* corresponds to the shallow part of the embedding. The rest of the data types constitute a deep embedding, providing an explicit representation of language terms for further symbolic processing and ultimately compilation, as will be discussed in more detail in the following.

As to the equations, there are four different kinds:

```
data Equation =
  EquationInit Expr Expr | EquationEq Expr Expr |
  EquationEvent String Expr Expr | EquationSigRelApp SigRel Expr
```

³ Although not necessarily a *fixed* such system as these equations may refer to signal relations that contain switches.

Initialisation equations, constructed by *EquationInit*, provide initial conditions. They are only in force when a signal relation instance first becomes active.

Equations constructed by *EquationEq* are basic equations imposing the constraint that the valuations of the two expressions have to be equal for as long as the containing signal relation instance is active (for example, equations like *der (der x) = 0*).

Equations constructed by *EquationEvent* define event signals; that is, they represent equations like **event** $e = (x, y)$ **when** $time = 3$. These equations are directed. The string is the name of the defined event signal. The first expression gives the value of the event signal at event occurrences. The second expression defines these occurrences. An event occurs whenever the signal represented by this expression *crosses* 0 (i.e., the signal's sign changes from positive to negative or from negative to positive). For the above example, the expression defining the event occurrences would thus be $time - 3$. Notation like **when** $x = y$, as in the example here, is standard practice but can be misleading in that if interpreted literally, it does not necessarily define a set of only countably many points in time. For example, $0 = 0$ is always true, implying uncountably many occurrences, which would be flawed. However, as the signal $0 - 0$ never *crosses* 0, the actual semantics is that this does not define any event occurrences at all.

The fourth kind of equation is signal relation application, *EquationSigRelApp*, that is equations like $sr \diamond (x, y + 2)$. This brings all equations of a signal relation into scope by instantiating them for the expressions to which the relation is applied (taking care to avoid accidental name clashes).

Finally, the representation of expressions is a standard first-order term representation making it easy to manipulate expressions symbolically (e.g. computing symbolic derivatives) and compiling expressions to simulation code:

```
data Expr = ExprUnit | ExprReal Double | ExprVar String | ExprTime |
          ExprTuple Expr Expr [Expr] | ExprApp Function [Expr]
data Function = FuncDer | FuncNeg | FuncAdd | FuncSub | FuncMul | ...
```

We use quasiquoting, a recent Haskell extension implemented in GHC, to provide a convenient surface syntax for signal relations. We have implemented a quasiquoter that takes a string in the concrete syntax of Hydra and generates Haskell code that builds the signal relation in the mixed-level representation described above. GHC executes the quasiquoter for each string between the quasiquotes before type checking.

While the internal representation of a signal relation is untyped, Hydra itself is *statically* typed, and we thus have to make sure that only type-correct Hydra programs are accepted. As Hydra fragments are generated dynamically, during simulation, we cannot postpone the type checking to after program generation. Nor can we do it early, at quasiquoting time, at least not completely, as no type information from the context around quasiquoted program fragments are available (for example, types of antiquoted Haskell expressions). In the current version of Hydra, only domain-specific scoping rules are checked at the stage of quasiquoting. Luckily, the type system of the present version of Hydra is fairly

simple; in particular, Hydra is simply typed, so by using the standard technique of phantom types, the part of the type checking that requires type information outside the quasiquotes is delegated to the host language type checker [23].

A phantom type is a type whose type constructor has a parameter that is not used in its definition. We define phantom type wrappers for the untyped representation as follows:

```
data SR a      = SR      SigRel
data PatternT a = PatternT Pattern
data ExprT a   = ExprT   Expr
data E a
```

Phantom types can be used to restrict a function to building only type-correct domain-specific terms. For example, a typed combinator *sigrel* can be defined in the following way:

```
sigrel :: PatternT a → [Equation] → SR a
sigrel (PatternT p) eqs = SR (SigRel p eqs)
```

As can be seen, the type of the pattern that defines the interface of the signal relation is what determines its type. Similarly, we define a typed combinator *switch*:

```
switch :: (Eval b) ⇒ SR (a, E b) → (b → SR a) → SR a
switch sr f = SigRelSwitch (forget sr) (forget ∘ f ∘ eval)
```

E is a type constructor with no constituent data constructors. It is used to type patterns that introduce event signals. The data for the event signals are constructed using event equations. Typed signal relations are translated into their untyped counterparts using the following function:

```
forget :: SR a → SigRel
forget (SR sr) = sr
```

The *Eval* class method *eval* is used to translate a value carried by an event occurrence (represented by an element of the value subset of the type *Expr*) into the corresponding Haskell value:

```
class Eval a where
  eval :: Expr → a
```

Instances of the *Eval* type class are at present provided for *Double* and arbitrarily nested tuples with fields of type *Double*.

A signal relation that is defined using the *switch* combinator is structurally dynamic. However, the type of the *switch* combinator statically guarantees that its type (i.e., its interface) remains unchanged. Thus, a structurally dynamic signal relation can be used in a signal relation application just like any other signal relation.

Well-typed equations are constructed using combinators in a similar way:

$$\begin{aligned} \text{equationEq} &:: \text{ExprT } a \rightarrow \text{ExprT } a \rightarrow \text{Equation} \\ \text{equationSigRelApp} &:: \text{SR } a \rightarrow \text{ExprT } a \rightarrow \text{Equation} \end{aligned}$$

Typed combinators for the remaining parts of the language, including *Pattern* and *Expr*, are defined using the same technique.

To summarise, under the hood the representation is still untyped. By only exposing the typed combinators for building signal relations, it is guaranteed that only well-typed terms can be constructed. Consequently, the quasiquoter of Hydra has only access to typed combinators for building signal relations. However, symbolic transformations (such as symbolic differentiation and flattening) on embedded language terms work on the untyped representation. These transformations need to be programmed with care as the Haskell type checker cannot verify that the transformations are type-preserving.

Several type system extensions of Haskell, such as generalised algebraic data types (GADTs), existential types, and type families, make alternative techniques for typing EDSLs possible. One option in our case might be to adopt a typed internal representation of the signal relations. The Haskell type checker could then hopefully be used to verify that transformations are type preserving. While this approach would require more work from the EDSL implementer, the additional correctness guarantees would be beneficial. We have not yet attempted to do this, opting for the simpler, untyped representation for now, but the work by Guillemette et al. suggests a more typed approach ought to be feasible [8].

4 Simulation

In this section we describe how an iteratively staged Hydra program is run. The process is illustrated in Fig. 4 and is conceptually divided into four stages. In the first stage, a signal relation is flattened and subsequently transformed into a mathematical representation suitable for numerical simulation. In the second stage, this representation is JIT compiled into efficient machine code. In the third stage, the compiled code is passed to a numerical solver that simulates the system until the end of simulation or an event occurrence. In the fourth stage, in the case of an event occurrence, the event is analysed, a corresponding new signal relation is computed and the process is repeated from the first stage. In the following, each stage is described in more detail.

As a first step, all signal variables are renamed to give them distinct names. This helps avoiding name clashes during unfolding of signal relation applications, flattening (see Sec. 2.2), and thus simplifies this process. Having carried out this preparatory renaming step, all signal relation applications are unfolded until the signal relation is completely flattened.

Further symbolic processing is then performed to transform the flattened signal relation into a form that is suitable for numerical simulation. In particular, derivatives of compound signal expressions are computed symbolically. In the

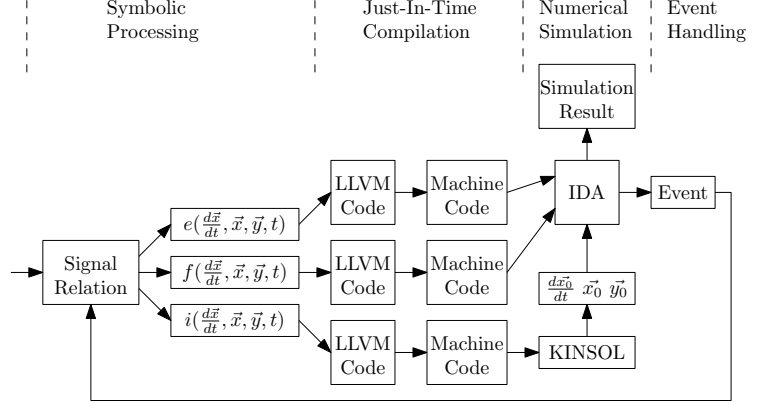


Fig. 4. Execution model of Hydra

case of higher-order derivatives, extra variables and equations are introduced to ensure that all derivatives in the flattened system are first order.

Finally, the following equations are generated at the end of the stage of symbolic processing:

$$i\left(\frac{d\vec{x}}{dt}, \vec{x}, \vec{y}, t\right) = 0 \quad (1)$$

$$f\left(\frac{d\vec{x}}{dt}, \vec{x}, \vec{y}, t\right) = 0 \quad (2)$$

$$e\left(\frac{d\vec{x}}{dt}, \vec{x}, \vec{y}, t\right) = 0 \quad (3)$$

Here, \vec{x} is a vector of differential variables, \vec{y} is a vector of algebraic variables, and t is time. Equation (1) determines the initial conditions for (2); that is, the values of $\frac{d\vec{x}}{dt}, \vec{x}$ and \vec{y} at time $t = t_0$, the starting time for the current set of equations. Equation (2) is the main DAE of the system that is integrated in time starting from the initial conditions. Equation (3) specifies the event conditions (signals crossing 0).

As the functions i , f , and e are invoked from within inner loops of the solver, they have to be compiled into machine code for efficiency: any interpretive overhead here would be considered intolerable by practitioners for most applications. However, as Hydra allows the equations to be changed in arbitrary ways *during* simulation, the equations have to be compiled whenever they change, as opposed to only prior to simulation. Our Hydra implementation employs JIT machine code generation using the compiler infrastructure provided by LLVM. The functions i , f and e are compiled into LLVM instructions that in turn are compiled by the LLVM JIT compiler into native machine code. Function pointers to the generated machine code are then passed to the numerical solver.

The numerical suite used in the current implementation of Hydra is called SUNDIALS [10]. The components we use are KINSOL, a nonlinear algebraic

equation systems solver, and IDA, a differential algebraic equation systems solver. The code for the function i is passed to KINSOL that numerically solves the system and returns initial values (at time t_0) of $\frac{d\vec{x}}{dt}, \vec{x}$ and \vec{y} . These vectors together with the code for the functions f and e are passed to IDA that proceeds to solve the DAE by numerical integration. This continues until either the simulation is complete or until one of the events defined by the function e occurs. Event detection facilities are provided by IDA.

At the moment of an event occurrence (one of the signals monitored by e crossing 0), the numerical simulator terminates and presents the following information to an event handler: name of the event variable for which an event occurrence has been detected, time t_e of the event occurrence and instantaneous values of the signal variables (i.e., values of $\frac{d\vec{x}}{dt}, \vec{x}$ and \vec{y} at time t_e).

The event handler traverses the original unflattened signal relation and finds the event value expression (a signal-level expression) that is associated with the named event variable. In the case of the breaking pendulum model, the expression is $((x, y), (vx, vy))$. This expression is evaluated by substituting the instantaneous values of the corresponding signals for the variables. The event handler applies the second argument of the *switch* combinator (i.e., the function to compute the new signal relation to switch into) to the functional-level event value. In the case of the breaking pendulum model, the function *freeFall* is applied to the instantaneous value of $((x, y), (vx, vy))$. The result of this application is a new signal relation. The part of the original unflattened signal relation is updated by replacing the old signal relation with the new one. The flat system of equations for the previous mode and the machine code that was generated for it by the LLVM JIT compiler are discarded. Then the simulation process for the updated model is repeated from the first stage and onwards.

In previous work [7], we conducted benchmarks to evaluate the performance of the proposed execution model. The initial results are encouraging. For a small system with handful of equations (such as the breaking pendulum) the total time spent on run-time symbolic processing and code generation is only a couple of hundredth of a second. To get an initial assessment of how well our approach scales, we also conducted a few large scale benchmarks (thousands of equations). These demonstrated that the overall performance of the execution model seems to scale well. In particular, time spent on run-time symbolic processing and JIT compilation increased roughly linearly in the number of equations for these tests. The results also demonstrate that the time spent on JIT compilation dominates over the time spent on run-time symbolic processing. Above all, our benchmarks indicated that the time for symbolic processing and compilation remained modest in absolute terms, and thus should be relatively insignificant compared with the time for simulation in typical applications.

In the current implementation of Hydra, a new flat system of equations is generated at each mode switch without reusing the equations of the previous mode. It may be useful to identify exactly what has changed at each mode switch, thus enabling the reuse of *unchanged* equations and associated code from the previous mode. In particular, this could reduce the burden placed on

the JIT compiler, which in our benchmarks accounted for most of the switching overheads. Using such techniques, it may even be feasible to consider our kind of approach for structurally dynamic *soft real-time* applications.

Our approach offers new functionality in that it allows non-causal modelling and simulation of structurally dynamic systems that simply cannot be handled by static approaches. Thus, when evaluating the feasibility of our approach, one should weigh the inherent overheads against the limitation and inconvenience of not being able to model and simulate such systems non-causally.

5 Related Work

The deep embedding techniques used in the Hydra implementation for domain-specific optimisations and efficient code generation draws from the extensive work on compiling staged domain-specific embedded languages. Examples include Pan by Elliott et al. [4], a DSL for image synthesis and manipulation, Paradise by Augustsson et al. [1], a DSL for pricing financial products, and Flask by Mainland et al. [15], a DSL for programming sensor networks. Pan is perhaps the setting most similar to ours as it combines leveraging the host language for program generation with compilation to speed up the resulting performance-critical computations. However, unlike Hydra, code generation and compilation only happens once in Pan, Paradise, and Flask, prior to execution, meaning that the host language is used only for meta-programming and that it is not involved during the execution of the final code.

The use of quasiquoting in the implementation of Hydra was inspired by Flask [15]. However, we had to use a different approach to type checking. A Flask program is type checked by a domain-specific type checker *after* being generated, just before the subsequent compilation into the code to be deployed on the sensor network nodes. This happens at *host language run-time*. Because Hydra is iteratively staged, we cannot use this approach: we need to move type checking back to *host language compile-time*. The Hydra implementation thus translates embedded programs into typed combinators at the stage of quasiquoting, charging the host language type checker with checking the embedded terms. This ensures that only well-typed programs are generated at run-time.

Lee et al. are developing a DSL embedded in Haskell for data-parallel array computations on a graphics processing unit (GPU) [13]. GPU programs are first-class entities. The embedded language is being designed for run-time code generation, compilation and execution, with results being fed back for use in further host language computations. Thus, this is another example of what we term iterative staging. At the time of writing, the implementation is interpreted. However, a JIT compiler for a GPU architecture is currently being developed.

The FHM design was originally inspired by Functional Reactive Programming (FRP) [5], particularly Yampa [18]. A key difference is that FRP provides *functions* on signals whereas FHM generalises this to *relations* on signals. FRP can thus be seen as a framework for *causal* simulation, while FHM supports non-causal simulation. Signal functions are first class entities in most incarnations

of FRP, and new ones can be computed and integrated into a running system dynamically. This means that these FRP versions, including Yampa, also are examples of iteratively staged languages. However, as all FRP versions supporting highly dynamic system structure so far have been interpreted, the program generation aspect is much less pronounced than what is the case for FHM.

In the area of non-causal modelling, Broman’s work on the Modelling Kernel Language (MKL) has a number of similarities to FHM [2]. MKL provides a λ -abstraction for defining functions and an abstraction similar to **sigrel** for defining non-causal models. Both functions and non-causal models are first-class entities in MKL, enabling higher-order, non-causal modelling like in FHM. However, support for structural dynamism has not yet been considered.

Non-causal languages that do support more general forms of structural dynamism than the current mainstream ones include MOSILAB [22], a Modelica extension, and Sol [25], a Modelica-like language. MOSILAB has a compiled implementation, but insists all structural configurations are predetermined to enable compilation once and for all, prior to simulation. Sol is less restrictive, but currently only has an interpreted implementation. Both MOSILAB and Sol could thus benefit from the implementation techniques described in this paper.

6 Conclusions

In this paper we presented a novel implementation approach for non-causal modelling and simulation languages supporting structural dynamism through iterative staging. Our approach is to embed the language in Haskell, using a mixed-level embedding to capitalise maximally on the host language while simultaneously enabling an efficient implementation through JIT compilation.

The resulting EDSL combines a functional and relational approach to modelling, and can thus be viewed as a kind of functional constraint logic programming language, albeit specialised to modelling and simulation and with a looser coupling between the functional and constraint logic parts than more typical such languages.

The iterative staging allows us to model systems that current, main-stream, non-causal modelling languages cannot handle without resorting to causal modelling. As far as we are aware, this is the first compiled implementation of a non-causal modelling language that supports highly structurally dynamic systems, demonstrating the practical feasibility of such a language as compilation of simulation code is usually considered essential for reasons of performance.

The use of the EDSL approach was instrumental to achieve the above. By reusing the features of the host language and its tool chain, we could focus our efforts on the problems that are specific to non-causal modelling and simulation. We also note that LLVM has worked really well for our purposes. Compilation of iteratively staged, mixed-level EDSLs does not seem to have attracted much attention thus far. We hope the implementation techniques we have developed will be useful to others who are faced with implementing such languages.

Acknowledgements. This work was supported by EPSRC grant EP/D064554/1. We would like to thank Neil Sculthorpe and the anonymous reviewers for their thorough and constructive feedback that helped improving the paper.

References

1. Augustsson, L., Mansell, H., Sittampalam, G.: Paradise: a two-stage DSL embedded in Haskell. In: Proceeding of the 13th ACM SIGPLAN International Conference on Functional Programming (ICFP 2008), pp. 225–228. ACM Press, New York (2008)
2. Broman, D., Fritzson, P.: Higher-order acausal models. In: Fritzson, P., Cellier, F., Broman, D. (eds.) Proceedings of the 2nd International Workshop on Equation-Based Object-Oriented Languages and Tools, Paphos, Cyprus. Linköping Electronic Conference Proceedings, vol. 29, pp. 59–69. Linköping University Electronic Press (2008)
3. Cellier, F.E.: Object-oriented modelling: Means for dealing with system complexity. In: Proceedings of the 15th Benelux Meeting on Systems and Control, pp. 53–64 (1996)
4. Elliott, C., Finne, S., de Moor, O.: Compiling embedded languages. In: Taha, W. (ed.) SAIG 2000. LNCS, vol. 1924, pp. 9–27. Springer, Heidelberg (2000)
5. Elliott, C., Hudak, P.: Functional reactive animation. In: Proceedings of International Conference on Functional Programming, pp. 163–173 (June 1997)
6. Giorgidze, G., Nilsson, H.: Embedding a functional hybrid modelling language in Haskell. In: Proceedings of the 20th International Symposium on the Implementation and Application of Functional Languages. University of Hertfordshire, Hatfield (September 2008)
7. Giorgidze, G., Nilsson, H.: Higher-order non-causal modelling and simulation of structurally dynamic systems. In: Proceedings of the 7th International Modelica Conference, Como, Italy. Linköping Electronic Conference Proceedings, Linköping University Electronic Press (September 2009)
8. Guillemette, L.-J., Monnier, S.: A type-preserving compiler in Haskell. In: ICFP 2008: Proceeding of the 13th ACM SIGPLAN International Conference on Functional Programming, pp. 75–86. ACM, New York (2008)
9. Hanus, M., Kuchen, H., Moreno-Navarro, J.J.: Curry: A truly functional logic language. In: Proceedings Workshop on Visions for the Future of Logic Programming, Portland, Oregon, USA, pp. 95–107 (December 1995)
10. Hindmarsh, A.C., Brown, P.N., Grant, K.E., Lee, S.L., Serban, R., Shumaker, D.E., Woodward, C.S.: Sundials: Suite of nonlinear and differential/algebraic equation solvers. *ACM Trans. Math. Softw.* 31(3), 363–396 (2005)
11. Hudak, P.: Modular domain specific languages and tools. In: Proceedings of Fifth International Conference on Software Reuse, pp. 134–142 (June 1998)
12. Lattner, C.: LLVM: An infrastructure for multi-stage optimization. Master’s thesis, Computer Science Dept., University of Illinois at Urbana-Champaign, Urbana, IL (December 2002), <http://llvm.org>
13. Lee, S., Chakravarty, M., Grover, V., Keller, G.: GPU kernels as data-parallel array computations in Haskell. In: Workshop on Exploiting Parallelism using GPUs and other Hardware-Assisted Methods (2009)
14. Mainland, G.: Why it’s nice to be quoted: quasiquoting for Haskell. In: Proceedings of the ACM SIGPLAN Workshop on Haskell Workshop, pp. 73–82. ACM, New York (2007)

15. Mainland, G., Morrisett, G., Welsh, M.: Flask: Staged functional programming for sensor networks. In: Proceedings of the Thirteenth ACM SIGPLAN International Conference on Functional Programming, Victoria, British Columbia, Canada. ACM Press, New York (September 2008)
16. The Modelica Association. Modelica – A Unified Object-Oriented Language for Physical Systems Modeling: Tutorial version 1.4 (December 2000)
17. The Modelica Association: Modelica – A Unified Object-Oriented Language for Physical Systems Modeling: Language Specification version 3.0 (September 2007), <http://www.modelica.org/documents/ModelicaSpec30.pdf>
18. Nilsson, H., Courtney, A., Peterson, J.: Functional reactive programming, continued. In: Proceedings of the 2002 ACM SIGPLAN Haskell Workshop, Pittsburgh, Pennsylvania, USA, pp. 51–64. ACM Press, New York (October 2002)
19. Nilsson, H., Giorgidze, G.: Exploiting structural dynamism in Functional Hybrid Modelling for simulation of ideal diodes. In: Proceedings of the 7th EUROSIM Congress on Modelling and Simulation, Prague, Czech Republic. Czech Technical University Publishing House (September 2010)
20. Nilsson, H., Peterson, J., Hudak, P.: Functional hybrid modeling. In: Dahl, V. (ed.) PADL 2003. LNCS, vol. 2562, pp. 376–390. Springer, Heidelberg (2002)
21. Nilsson, H., Peterson, J., Hudak, P.: Functional hybrid modeling from an object-oriented perspective. In: Fritzon, P., Cellier, F., Nytsch-Geusen, C. (eds.) Proceedings of the 1st International Workshop on Equation-Based Object-Oriented Languages and Tools (EOOLT), Berlin, Germany. Linköping Electronic Conference Proceedings, vol. 24, pp. 71–87. Linköping University Electronic Press (2007)
22. Nytsch-Geusen, C., Ernst, T., Nordwig, A., Schwarz, P., Schneider, P., Vetter, M., Wittwer, C., Nouidui, T., Holm, A., Leopold, J., Schmidt, G., Mattes, A., Doll, U.: MOSILAB: Development of a Modelica-based generic simulation tool supporting model structural dynamics. In: Proceedings of the 4th International Modelica Conference, Hamburg, Germany, pp. 527–535 (2005)
23. Rhiger, M.: A foundation for embedded languages. *ACM Trans. Program. Lang. Syst.* 25(3), 291–315 (2003)
24. Taha, W.: A gentle introduction to multi-stage programming. In: Lengauer, C., Batory, D., Blum, A., Vetta, A. (eds.) *Domain-Specific Program Generation*. LNCS, vol. 3016, pp. 30–50. Springer, Heidelberg (2004)
25. Zimmer, D.: Introducing Sol: A general methodology for equation-based modeling of variable-structure systems. In: Proceedings of the 6th International Modelica Conference, Bielefeld, Germany, pp. 47–56 (2008)