A Core Language for Executable Models of Cyber Physical Systems

(Work In Progress Report)

Walid Taha
Halmstad University,
Halmstad, Sweden.
Rice University
Houston, TX, USA
Walid.Taha@hh.se

Paul Brauner
Rice University
Houston, TX, USA
polux2000@gmail.com

Robert Cartwright
Rice University
Houston, TX, USA
corky.cartwright@gmail.com

Verónica Gaspes
Halmstad University
Halmstad, Sweden
Veronica.Gaspes@hh.se

Aaron Ames
University of Texas A&M
College Station, TX, USA
aames@tamu.edu

Alexandre Chapoutot
ENSTA ParisTech
Paris, France
alexandre.chapoutot@ensta-paristech.fr

ABSTRACT

Recently we showed that an expressive class of mathematical equations can be automatically translated into simulation codes. Focusing on the expressivity of equations on continuous functions, this work considered only minimal interaction with discrete behaviors and only a static number of statically connected components. However, the interaction between continuous and hybrid components in many cyber physical domains is highly coupled, and such systems are often highly dynamic in both respects. This paper gives an overview of a proposed core language for capturing executable hybrid models of highly dynamic cyber physical systems.

Keywords

Modeling, Simulation, Cyber-Physical Systems.

1. INTRODUCTION

Systems that evolve over a dense notion of time interact in complex ways that can confound both intuition and analytical methods [7, 2]. This problem is acute for nonlinear differential systems, which include virtually all three-dimensional mechanical systems, and for which solutions rarely have closed form descriptions [5]. As a result, successful analysis and design of novel cyber physical systems invariably includes an extensive experimental component that relies either on physical prototypes or on simulation. Today, both types of experimentation can be prohibitively costly and slow. Physical experiments incur material costs and pose challenges in control and reproducibility, measurement and instrumentation, and safety. Simulation has the potential to reduce these problems and to significantly accelerate innovation. But current methods either raise questions about fidelity or are extremely labor intensive. Using mainstream tools means depending on proprietary, black-box codes that come with a fixed set of component models and that offer only limited support for building custom models. Writing simulation codes by hand requires both effort and specialized expertise in mapping the high-level analytical models to executable codes, software implementation (including debugging and testing), and dealing with issues of floating point numerical precision. These difficulties can be debilitating for designers of cyber physical systems, especially novel and creative designs.

To reduce the cost of building simulations, we recently developed and presented an automated, scalable mapping from an expressive class of mathematical equations to code, showing that this natural mathematical formalism can be viewed as an executable language for modeling mechanical systems [17]. While the examples used to illustrate the method focused on the mechanics, the formalism, called Acumen, can be used for other physical domains such as electrical, hydraulic, or heat transfer systems.

Originally, Acumen was developed as an extension of event-driven formalisms [13, 12, 8] that have a similar flavor to synchronous languages [6]. Acumen added systems of equations on functions on dense time for the purpose of describing the "physical environment" surrounding the "purely cyber controllers" that were already describable in synchronous formalisms. Although this approach seems intuitive at first, it makes an unnecessary association between physical and continuous and cyber and discrete. In reality, physical environments often exhibit both continuous and discrete behaviors. For example, the two legs of a walking robot continually and discretely change modality and role with each step forward. Dually, the use of a purely discrete model for controllers or embedded systems is overly restrictive. In early stages of design one may use purely continuous controller models

for simplicity. In later stages of design, it may be important to capture physical aspects of a digital implementation, such as dense-time behavior, delays, energy consumption, or heat emissions. Thus, the expressivity needed to model both the physical and cyber components calls for tight integration between continuous and discrete behaviors.

This paper describes the key features of a new, more uniform design for Acumen. The proposed design allows fine-grained coupling between continuous and discrete behaviors in a unified notion of a hybrid object. Such objects have time-varying state, carry hybrid laws that specify their behavior, can be dynamically created and terminated, and can include and dynamically coordinate "child" objects.

In the rest of the paper, we review closely related work (Section II), summarize the proposed design (Section III), and describe how it is simulated (Section IV).

2. RELATED WORK

Acumen [16, 17, 1] is a modeling language being developed with the goal of bridging the gap between several important efforts in modeling and simulation, hybrid systems verification, and synchronous languages. In what follows we describe how the proposed design relates to other efforts.

The purely discrete event-driven predecessors [13, 12, 8] of Acumen have their roots in Functional Reactive Programming (FRP) [11], which itself supports both continuous and discrete behaviors in a purely functional setting. In formulating the predecessors of Acumen, we narrowed the functional framework to purely discrete systems to focus on the real-time properties of embedded controllers.

Modelica's support for equation-based (or relational) modeling [4] provided the initial inspiration for Acumen's equations on functions of dense time. Going beyond Modelica and other equation-based languages, the full Acumen language supports partial derivatives that can be used to specify systems using Euler-Lagrange equations, which still can be symbolically eliminated by translation to time derivatives.

Like CHARON [3], Acumen is a hybrid systems simulation language inspired by hybrid automata [7, 2] and hybrid logics [10]. Acumen differs from CHARON in being untyped, deterministic, and built on a single, dynamic notion of object. We present a more detailed comparison after Core Acumen has been introduced (Section III).

Dynamic differential logic [10] encouraged us to explore a more "imperative" style of describing an object's state, and which is reflected in design presented in this paper. A key difference between hybrid logics and languages aimed at simulation (such as FRP, Modelica, and Acumen) is the treatment of non-determinism. Non-determinism is advantageous in formalisms used for automated reasoning, because it can be used to weaken assumptions and thus strengthen the established properties. But non-determinism is highly problematic for simulation formalisms, because it may require exploring a vast number of options for bounded domains, and is simply not possible for unbounded domains.

Allowing discrete computations to be repeated arbitrarily until there are no more additional changes is a standard way for computing a fixed point. Synchronous languages [6] such as Lustre or Esterel use a similar strategy for converging on the result of a synchronous system. In the proposed design we compute the fixed point for the state of the whole model being simulated. By the Bekic theorem, this produces the same result as computing the fixed point for all components of the system independently [14].

3. CORE ACUMEN

Acumen's semantics is defined by a series of translations from a large source language into progressively smaller subsets of the language. The purpose of the core language is to serve as the minimal subset needed to express all the features of the full source language. In this section, the proposed design for Acumen's core language is illustrated by a series of small examples.

Acumen is implemented as free software and is available along with a hands-on tutorial from the Acumen website[1]. All examples can be simulated and visualized using the online distribution, version 10.12.13. The implementation and the tutorial include more examples than we provide in this description of the core language. The tutorial also presents the grammar (BNF), explains class parameters, how to define local notions of time, how to use the graphical user interface, and describes other features of the language that natural extend the subset presented here.

3.1 Objects and Hybrid Laws

Acumen objects are introduced by defining a class for each kind of object, and then by creating instances of these classes at a particular point in model/simulation time. As an example, consider a device consisting of a battery and discrete controller that decides whether the battery should charge or drive a load. When charging, the voltage on the battery increases at a constant rate until it reaches its full capacity. Then, the battery stops charging and starts driving the load until the voltage is too low. When that happens, the battery switches back to charging. In Core Acumen, the device is modeled as follows:

```
class Contraption ()
  private v = 0; v' = 0; mode = 0 end
  switch mode
    case 0 // Charge (until high)
      if (v < 0.8)  v' [=] 1/2
      else  mode = 1   end
    case 1 // Drive (until low)
      if (v > 0.2) v' [=] -v
      else  mode = 0  end
  end
end
```

All variables (such as `v`, `v'`, and mode in the first class) are implicitly functions of time [11]. Variables introduced in the `private` section of the class are the state of any object of this class. The value assigned to each variable is the initial value it has at the instant the object is created. A prime (`'`) following a variable name denotes its derivative with respect to time. Thus, given an initial value for the variable `v`, if `v'` is defined, then the value of `v` will be automatically determined for future values as well. The `switch` statement allows different sets of rules to govern the behavior of the

system under different conditions. The `case` that applies is determined by the value of the variable `mode`. When the value of mode is `0`, then the first `if` statement is active. When the value is `1`, then the second `if` statement is active. The two `if` statements follow a parallel pattern: their true branch contains a *continuous assignment*, and the false branch contains a *discrete assignment*. The main difference between discrete and continuous assignments is that discrete assignments block the progress of logical time, in the sense that simulation cannot advance beyond a particular point in time until all discrete assignments have been performed. Continuous assignments, in contrast, happen continuously, and pose no particular constraints on how the simulator advances time.

3.2 Class `Main` and the `simulator` parameter

In any Acumen model there must be a declaration for a class called `Main`. This class always represents the entire world being modeled.

Even though the goal of Acumen is to automate building simulation codes, there are fundamental computability limitations that dictate that not all implementation details can be hidden from the user. For example, there is no single, universal method for solving a system of equations, be it linear, non-linear, time-varying, or differential. Yet bridges and airplanes need to be built, and they will be, whether or not we help engineers write their simulation codes. Thus, a pragmatic decision is made in Acumen to allow the user to include in models additional details needed to perform the simulation. To support this, each `Main` object is required to have a parameter (by convention called the `simulator`) that allows the user to express how the model should be simulated. Continuing the above example, we can write:

```
class Main (simulator)
  private mode = "Init" end
  switch mode
    case "Init"
      simulator.timeStep = 0.001;
      simulator.endTime  = 5.0;
      create Contraption ();
      mode = "Persist"
    case "Persist"
  end
end
```

The default parameters for simulation start time, end time, and step size are 0, 10, and 0.005, respectively. The rest of the definition for this class uses a variable called mode to distinguish between two different states, one for initializing (or creating) the model and the other for letting the model run.

3.3 Object Life Cycle, Migration, and Regulation

When a `create` command is encountered at a certain point in model/simulation time, an object is introduced to the model. Similarly, objects can be terminated. Initially, any new object is considered the child of the object that created it. To allow objects to regulate the behavior of their children, Acumen allows iteration over children. It is also possible to move objects dynamically from one parent object to another, thereby changing the set of "external laws" that apply to this object.

To illustrate these concepts, we will consider an artificial example that allows us to showcase all of these facilities concisely. In this example, we introduce a class for "fancy balls". The basic functionality of these balls is to bounce. In addition, they have a limited life span that is specified by the parent at the time they are created. Additionally, each fancy ball ensures that each child has a lifespan that is at least two seconds longer than that of its parent. The `Main` class will specify a world where two such objects are created at time 0, and then, at time 2, the second ball is moved from the top level world object to be a child object of the first ball. The model capturing this behavior is as follows:

```
class FancyBall (t, x,x',x'')
 private t'=0; end
 x'' [=] -9.8;
 if x<=0 x' = -0.6*x'; x  = -x end;
 t' [=] -1;
 if t<=0 terminate self end;
 for c = self.children
   c.t [=] 2
 end
end
class Main (simulator)
  private
   mode ="Init"; n = 0;  t=0; t'=1;
   a = create FancyBall (5,10,5,0);
   b = create FancyBall (3,10,7,0);
  end
   t' [=] 1;
   n [=] sum 1 for i in self.children
         if true;
  switch mode
     case "Init"
       if t>2
         move b a;
         mode = "Persist" end
     case "Persist"
  end
end
```

The code includes an additional private variable n which keeps track of the number of children in the top level world, and illustrates one iteration construct in Acumen (summation). By default, simulating this model in Acumen produces the following plot. The plot shows the objects and their variables presented in the order that they are created, and for the duration that they exist, and with vertical scales normalized by the value range:

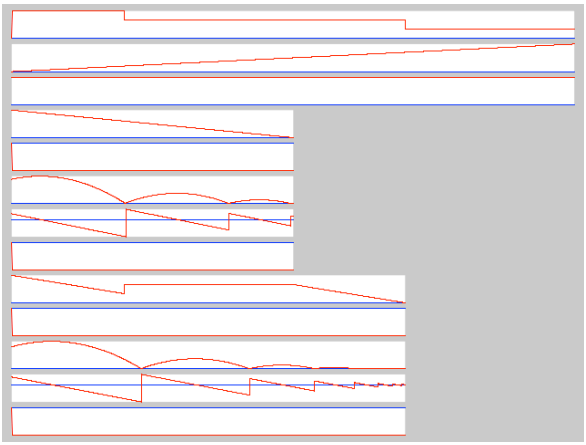

The first band plots the number of children at the top level. As expected, the number of top-level children drops at time 2, because the second ball has been moved to be a child of the first ball. The nineth band is the variable `t` of the second object stops decreasing linearly and keeps a fixed value (2) as dictated by fancy ball's rule for its children. The number of top-level children does not change when the first object dies because the default behavior is that the grand parent inherits any grandchild that survives a terminated parent.

This example also displays *Zeno behavior*, where an infinite number of discrete transitions occur in a finite amount of time (see [9] for more on formally detecting Zeno behavior and [15] for the semantics of this behavior). The finite time interval for the simulation is evidence of the existence of Zeno behavior in this example, and points to the strong need to consider simulation semantics that account for multiple discrete computations at a single time instance.

3.4 Comparison with CHARON

To illustrate the points made about the relation with CHARON in Related Work (Section II), we consider a simple model of thermostat coming from CHARON's user manual. The temperature `x` of a room is controlled to keep it in the target range of 68-82 degrees Fahrenheit. The temperature can evolve continuously over time. The heater is activated if the value of `x` is less than 70 and the evolution of `x` follows the equation `x' = -x + 100`. If the value of `x` is greater than 80 the heater is off and the temperature follows the rule `x' = -x`.

A CHARON program is made of a set of *agents* that are executed concurrently. Agents may be aggregated to form a more complex agent. Each agent is made of a set of *modes* each representing a state of hybrid automaton system in which only one mode can be activated at a time. A mode may also be made of sub-modes. A set of local or shared *variables* may be associated with agents or with modes. These variables are the main communication technique in CHARON. Each variable is declared with a type (e.g. `int` or `real`), a kind (e.g. `analog` or `discrete`) to express if it is a continuous-time or a discrete-time variable, and access restriction. A special operator `d` is used to represent derivatives (see mode `onOff`). CHARON explicitly declares guarded transitions between modes. An action may be associated in case of the transition is taken. Finally, an invariant property may be associated with each mode. If the invariant is false, then the automata must transition to another mode. If it does not, it is considered *blocked*. The implementation of the example described above is as follows:

```
agent thermostat(){
  mode top = thermostatTop()
}
mode thermostatTop(){
  private analog real x;
  mode on =
   onOff(-10000000000.0, 82.0, 100.0);
  mode off =
   onOff(68.0, 10000000000.0, 0.0);
  trans toSubMode from default
        to on when true do {x= 73}
  trans fromOnToOff from on
        to off when x > 80.0 do {}
  trans fromOffToOn from off
        to on when x < 70.0 do {}
}
mode onOff(real a, real b, real c){
  readWrite analog real x;
  inv invOnOff {x > a and x < b}
  diff dOnOff {d(x) == -x + c}
}
```

Deterministic CHARON programs can be expressed in Acumen as follows. Agents are mapped to classes. Local variables are mapped to local variables, and additional variables are introduced as needed for derivatives by looking at the rest of the agent definition. Modes are mapped to strings that can be assigned to a mode variable for that class. A switch statement is then used to capture the rules for that mode. To deal with invariants, an extra mode called Blocked is introduced, and the invariants are test at the end of the case for each switch value (mode).

```
class Thermostat(mode)
 private x = 0; x' = 0 end
 switch mode
   case "Top" x = 73; mode = "On"
   case "On"  x' [=] -x + 100;
     // Transition from on to off
     if x > 80 mode = "Off" end;
     // Invariant invOnOff
     if not (x>-10000000000.0 && x < 82.0)
      mode = "Blocked" end;
   case "Off" x' [=] -x;
     // Transition from off to on
     if x < 70 mode = "On" end;
     // Invariant invOnOff
     if not (x>68.0 && x < 10000000000.0)
        mode = "Blocked" end;
case "Blocked"
  end
end
```

Thus, Core Acumen is a smaller language that can still naturally express deterministic CHARON programs. As a hierarchical agent language, CHARON supports tree-structured models of the world. However, it is not clear that CHARON supports a notion of agent mobility similar to the one supported by core Acumen.

4. SIMULATION SEMANTICS

Acumen models are simulated by a fine interleaving of a sequences that can consist of multiple discrete computations followed by a single computation updating the values that should evolve continuously. Conceptually, we can think of any instances in time as follows:

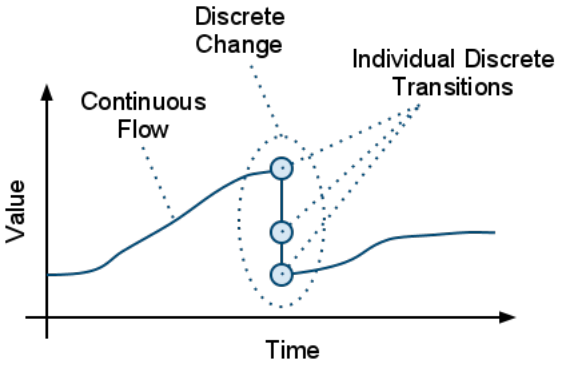

In the discrete phase of each sequence, all actions that require discrete change are performed. In the continuous phase, any range of numerical methods and tools for approximating continuous behavior can be used. Thus, simulating what is happening at any single instance in time consists of zero or more discrete steps followed by a single continuous step. The discrete steps capture sudden changes in state such as the impact of two objects, and consist of evaluating all active discrete assignments in the program until the whole system is stabilized. A system is stabilized when no more discrete steps are required. The following example illustrates how discrete assignments are handled:

```
class Main (simulator)
  private x  = 0; y  = 1; z  = 1; end
  if x<5 x = x+1 end;
end
```

The entire model is repeatedly evaluated until the condition in this statement is false. Simulation time (or logical time) is not advanced during these iterations. Acumen considers such changes to all be happening in the same instant. Using this type of global fixed point semantics allows Acumen to realize, among other things, what is sometimes called the "synchrony hypothesis", whereby the author of the model assume that certain discrete or digital events can happen "fast enough" so that we can view them in the rest of the model as happening instantaneously. In the example above, because the initial value of `x` is zero, the iteration will end when x has the value 5.

Once all discrete actions have taken place, the system moves on to performing all adjustments to the continuous state of the system. The continuous step performs all updates in parallel, meaning that all updates are based on the state that results after the sequence of discrete steps, rather than some later state that resulted from other continuous updates.

5. REFERENCES

[1] Acumen website. `http://www.acumen-language.org`, March 2011.

[2] R. Alur, C. Courcoubetis, N. Halbwachs, T. A. Henzinger, P.-H. Ho, X. Nicollin, A. Olivero, J. Sifakis, and S. Yovine. The algorithmic analysis of hybrid systems. *Theor. Comput. Sci.*, 138, February 1995.

[3] R. Alur, R. Grosu, Y. Hur, V. Kumar, and I. Lee. Modular specification of hybrid systems in charon. In *Proceedings of the Third International Workshop on Hybrid Systems: Computation and Control*, HSCC '00, 2000.

[4] D. Broman. *Meta-Languages and Semantics for Equation-Based Modeling and Simulation.* PhD thesis, Department of Computer and Information Science, Linköping University, 2010.

[5] W. M. Haddad and V. Chellaboina. *Nonlinear Dynamical Systems and Control: A Lyapunov-Based Approach.* Princeton University Press, 2008.

[6] N. Halbwachs. *Synchronous Programming of Reactive Systems.* Kluwer Academic Publishers, 1993.

[7] T. A. Henzinger. The theory of hybrid automata. In *Proceedings of the 11th Annual IEEE Symposium on Logic in Computer Science*, LICS '96, 1996.

[8] R. Kaiabachev, W. Taha, and A. Zhu. E-frp with priorities. In *Proceedings of the 7th ACM & IEEE international conference on Embedded software*, EMSOFT '07, 2007.

[9] A. Lamperski and A. D. Ames. On the existence of zeno behavior in hybrid systems with non-isolated zeno equilibria. In *47th IEEE Conference on Decision and Control*, 2008.

[10] A. Platzer and J.-D. Quesel. European train control system: A case study in formal verification. In *Proceedings of the 11th International Conference on Formal Engineering Methods: Formal Methods and Software Engineering*, ICFEM '09, 2009.

[11] Z. Wan and P. Hudak. Functional reactive programming from first principles. *SIGPLAN Not.*, 35, May 2000.

[12] Z. Wan, W. Taha, and P. Hudak. Real-time frp. *SIGPLAN Not.*, 36, October 2001.

[13] Z. Wan, W. Taha, and P. Hudak. Event-driven frp. In *Proceedings of the 4th International Symposium on Practical Aspects of Declarative Languages*, PADL '02, 2002.

[14] G. Winskel. *The Formal Semantics of Programming Languages.* The MIT Press, Cambridge, Massachusetts, 1993.

[15] H. Zheng, E. A. Lee, and A. D. Ames. Beyond zeno: Get on with it! In J. P. Hespanha and A. Tiwari, editors, *HSCC*, volume 3927 of *Lecture Notes in Computer Science*. Springer, 2006.

[16] A. Y. Zhu, J. Inoue, M. L. Peralta, W. Taha, M. K. O'Malley, and D. Powell. Implementing haptic feedback environments from high-level descriptions. In *Proceedings of the 2009 International Conference on Embedded Software and Systems*, 2009.

[17] Y. Zhu, E. Westbrook, J. Inoue, A. Chapoutot, C. Salama, M. Peralta, T. Martin, W. Taha, M. O'Malley, R. Cartwright, A. Ames, and R. Bhattacharya. Mathematical equations as executable models of mechanical systems. In *Proceedings of the 1st ACM/IEEE International Conference on Cyber-Physical Systems*, ICCPS '10.