

# Using Analytic CLP to Model and Analyze Hybrid Systems

Timothy J. Hickey    David K. Wittenberg

*Technical Report CS-03-240*  
*Computer Science Department*  
*Brandeis University*  
*{tim|dkw}@cs.brandeis.edu*

October 24, 2003

### **Abstract**

We use CLP(F), an Analytic Constraint Logic Programming language, to model hybrid systems. Analytic CLP languages combine intervals, constraints, and ODEs in a clean and natural way. CLP(F) provides an implementation of an ACLP language based on Interval Arithmetic. The semantics of CLP(F) rigorously handle non-linear ODEs and round-off error. The ODEs describing a hybrid system need only a minor change of syntax to become a CLP(F) program. This simple transformation from a physical description of a hybrid system to a program which can be used to provide a proof of safety properties of the system bridges the gap between practical tools and formal models, and allows one to easily prove statements about real-world systems. The combination of Interval Arithmetic with Analytic Constraint Logic Programming makes it easy to pose and answer many sorts of queries about a system. For example, “At what point does the system change from one state to another?”, or “What control settings result in a cycle with period  $t$ ?”

# 1 Introduction

Interval Arithmetic [Moo66] is an obvious choice for modeling hybrid systems [HHMWT00], as the interface between the analog and the digital part involves imperfect hardware whose description must include error bars. Interval arithmetic fits naturally with Constraint Logic Programming [JM94]. as an interval describing a real value  $x$  corresponds to two constraints on  $x$ , one for the upper bound, and one for the lower bound. [BO97] were the first to combine intervals with CLP.

We use CLP(F) [Hic00], an analytic CLP (Constraint Logic Programming) language which combines CLP with interval arithmetic. CLP(F) uses interval arithmetic techniques to rigorously solve constraints involving both real and function variables, constrained via arithmetic and functional equations. CLP(F) is related to QSIM [Kui93] in that each attempts to find an over approximation of the possible states of a system of ODEs (QSIM uses the term “qualitative behaviours”).

CLP(F) is an especially good fit for modeling hybrid systems described by ODEs (Ordinary Differential Equations) because it handles round-off errors, approximation errors (we use a Taylor series to calculate the ODEs), and measurement errors in a consistent natural way. These are the primary sources of computational difficulties in modeling ODEs. Because Interval Arithmetic provides constraints on the range of values that each variable can take on, it is well suited to proving that certain values are not reached (*ie.*, safety properties). Deville *et al.* [DJVH02] explore another approach to rigorously solving ODEs using constraints and intervals, which might also be fruitfully applied to the analysis of hybrid systems.

A hybrid system is a system composed of a digital part (typically a small computer) and an analog part (typically a physical system with sensors and actuators). All computer controlled or monitored processes in the real world are hybrid systems. Recent work on hybrid systems includes defining models [LSVW99] [LSV01] [ACH<sup>+</sup>95] [MMP91][GJS96] and calculating the behaviour of the analog parts [HHMWT00]. The Intelligent Highway Group at Berkeley has developed the SHIFT programming language [DGS] for describing evolving hybrid systems. Mosterman [Mos99] provides a survey of a dozen simulation packages to determine how much support they provide for simulating hybrid systems.

## 2 Advantages Over Previous Models of Hybrid Systems

There has been much work in using Constraint Logic Programming [JM94][JL87] to analyze various aspects of hybrid systems [CF00], [Urb96], [Pod00], [GJS96]. One problem with these conventional CLP approaches to modeling hybrid systems is that they must deal with the ODEs describing the continuous part of the system using some sort of approximation, such as discretization into difference equations or restriction to ODEs that have a closed form solution. This introduces a “modeling error” so that the systems are not computationally sound. One must then reason about the modeling error outside of the CLP program. Many systems ignore these errors, and leave it up to the user to understand the numerical instabilities. This sort of numerical analysis [Act96] is very tricky. To require users to understand numerical analysis under pain of getting a wrong answer is to invite error. For example, the SHIFT language [DGS] is very expressive, but it solves non-linear ODEs by using a fourth order Runge-Kutta algorithm without bounding the error term, and hence is not rigorous.

Hickey [Hic00] describes CLP(F), an Analytic Constraint Logic Programming (ACLP) language over the domain of differentiable functions. In this paper, we show how CLP(F) allows one to overcome this “modeling error” by allowing the ODE to be expressed explicitly as a constraint on function variables. The resulting ACLP program has the property that the results computed using the CLP(F) system are guaranteed to contain all solutions of the ODEs modeled by the constraints.

One of the major benefits of this approach is that the problem of analyzing the hybrid system is transformed into the problem of analyzing the corresponding CLP(F) program. In principle, one should be able to apply well understood program analysis techniques to CLP(F) and directly infer provable properties of the corresponding hybrid system. In this paper we describe only the simpler types of analysis that one can do by directly solving CLP(F) constraints related to the hybrid system. The primary disadvantage of this approach is that it is very resource intensive and hence can not currently model systems over a long

modeling period.

To demonstrate the ACLP approach to hybrid system modeling, we consider the hybrid system of a thermostat introduced in Henzinger *et al.* [HHWT98] [HHMWT00]. This is a system consisting of a stirred pot of water with a temperature sensor and a heater in it. When the measured temperature goes above a threshold, the logic circuit shuts off the heater (after a small delay). Similarly, when the measured temperature goes below a threshold, the logic circuit turns on the heater (after a small delay). The safety property in question is to establish upper and lower bounds for the temperature of the water. The state diagram is given in Figure 2.

Henzinger *et al.* take a major step towards reliability of their results by using interval arithmetic in solving the differential equations which describe the system. We improve on this by modeling the system declaratively as an ACLP program (written CLP(F) [Hic01] [Hic00]) in which the differential equations appear directly as constraints in the program and the system is modeled using intervals for all measurements (to model the inevitable error-bars of instruments) as well as to provide over-approximations to deal with rounding error.

### 3 CLP(F)

CLP(F) allows one to constrain functions by functional equations involving the same arithmetic operators and mathematical functions as discussed above. In addition, one can constrain a function to take certain values at certain points and to have a range that lies within an interval. For example, the following constraint specifies that  $F$  is a function on  $[0, 1]$  such that  $F' = F$  and  $F(0) = 1$  and  $F(A) = 2$  and  $F(1) = E$  and  $F([0, 1]) \subset [-1000, 1000]$ :

```
| ?- type([F],function(0,1)),
      {[ ddt(F,1)=F, eval(F,0)=1,
         eval(F,A)=2, eval(F,1)=E,
         F in [-1000,1000] ]}.
A = 0.6931471... E = 2.7182818...
(760 ms) no
| ?-
```

The `type` predicate is used to declare that  $F$  is an infinitely differentiable function on the interval  $[0, 1]$ . The CLP(F) system solves analytic constraints by soundly approximating analytic functions by power series and introducing arithmetic constraints among the Taylor coefficients of the functions at the endpoints, at points in the interval, and over the entire range. The function  $F$  is then constrained to be equal to its first derivative, and to take the value 1 at 0 and to take values in  $[-1000, 1000]$  for all  $x \in [0, 1]$ . The variables  $A$  and  $E$  are not declared to be functions and hence are real constants by default. They are constrained so that  $F(A) = 2$  and  $F(1) = E$ . The constraint solver finds  $A$  and  $E$  to 7 decimal digits of precision and also finds an interval for  $F$  (not shown here) that specifies intervals for its first 10 derivatives at 0 and 1, and intervals for the range of its first 10 derivatives over  $[0, 1]$ .

In this paper we will use CLP(F) to define higher order constraints which specify that two points lie on a trajectory defined by an ODE. For example, in the simplest model of a thermostat we use the following CLP(F) procedure, where  $T0, T1$  are times,  $A0, A1$  are temperatures at those times,  $A$  is the temperature function (so  $A(T0) = A0$ ) and  $\text{Alpha}, \text{Beta}$  are the heat loss and the heater element components of the ODE for  $A$ . The parameter  $I$  is a bound on the width of the interval on which  $A$  is defined. This bound is mandatory because in CLP(F) all functions must be defined on finite intervals.

```
ode((T0,A0), [I, [Alpha,Beta]],
      A, (T1,A1)) :-
  type([A],function(0,I)),
  {[ ddt(A,1) = Alpha * A + Beta,
     eval(A,0)=A0,   eval(A,T)=A1,
```

```

A in [-1.0E100,1.0E100],
T=T1-T0, T in [0,I] ]}.

```

In this case the ODE is  $f' = af + b, f(0) = a_0, f(t) = a_1$  which can be solved exactly. Thus, for this example, we can actually define the ode procedure (almost) equivalently in CLP(Intervals) without using the sophisticated analytic constraint solving machinery, as follows:

```

odeSolveForm((T0,A0), [I, [Alpha,Beta]],
              _,(T1,A1)) :-
{A1 = C1*exp(Alpha*(T1-T0)) - C2,
 C1=(A0 + C2),
 C2= Beta/Alpha}.

```

The difference between these two procedures is that the first one returns the solution function  $A$  of the ODE, which can be used to compute  $f$  at any point in the region, or to return the range, or to further constrain the trajectory, while the second merely returns the range for  $T$ .

The CLP(F) solver can also handle very complex non-linear differential equations as it is based on a “brute force” reduction of the analytic constraints into arithmetic constraints which are solved with a simple interval arithmetic constraint solver. In the more complex examples in this paper, we model a system consisting of a fluid with temperature  $A(t)$  which is heated by a heating element whose temperature  $B(t)$  has a non-linear component  $\sin(B(t))$  in its defining ODE. The component  $\sin(B(t))$  is not natural, but shows that we can handle non-linear components which are common in real systems. We omit a more realistic non-linear example for lack of space. This system is modeled by the following procedure:

```

ode2((T0,A0), [I, [Alpha,Beta,Gamma,
                  Delta]]
      ,A,(T1,A1)) :-
type([A,B],function(0,I)),
{[ ddt(A,1) = Alpha* A + Beta +Gamma*B,
   ddt(B,1) = Delta*(B + 0.1*sin(B)),
   eval(A,0)=A0,   eval(A,T)=A1,
   eval(B,0)=1,
   A in [-1.0E100,1.0E100],
   B in [-1.0E100,1.0E100],
   T=T1-T0, T in [0,I]
]}.

```

We are unaware of a closed form for the solutions of this ODE and yet the CLP(F) system is easily able to use this definition to compute  $(T1, A1)$  from  $(T0, A0)$ , or, as we will see below, to use this procedure to find values of the parameters **Alpha**, **Beta**, **Gamma**, and **Delta** which make the system behave in some desired fashion.

### 3.1 Programs

The program in Figure 1 is one way of implementing a general hybrid system simulator in CLP(F). The first parameter of the `evolve` procedure is the initial state of the hybrid system, which consists of a discrete state  $S$  and a continuous state  $X$ . The second parameter is a list of values used to specify the particular hybrid system. The third parameter is the final (or ending) state of the hybrid system. The last parameter is a list of the discrete state changes needed to go from the initial to the final state. The continuous part of these states are the values of the continuous variables at the moment of state change.

Observe that the procedure is simple. It looks for a state change from  $(S0, X0)$  to state  $S1$ . Such a change may constrain the values of  $X0$  to lie within narrow intervals. Then it looks for a new trajectory represented by the continuous variable  $X1$ . Typically, the continuous state will be represented by a pair of dependent variables  $(T, A)$  where  $A$  is the value of some ideal system sensor at time  $T$ . The `in_trajectory`

```

evolve(H,C,H, []).

evolve((S0,X0),C,(S,X),
      [(S1,X1,Range)|R]) :-
statechange((S0,X0),C,S1),
in_trajectory((S1,X0),Range,C,X1),
evolve((S1,X1),C,(S,X),R).

```

Figure 1: A general simulator for hybrid systems

procedure looks up the ODE,  $C$ , that should hold in this state and applies that ODE to the initial values  $X0$  to get the new values  $X1$ . The range  $R$  of the continuous variable during this trajectory is also returned. This is used for determining safety properties. In general other range data could also be returned (e.g. the range of the derivative of the temperature or of some other system parameter).

## 4 Henzinger’s Model and Analysis

In this section, we present the model of a thermostat with a delay in switching used by Henzinger *et al.* [HHMWT00]. Henzinger’s model consists of a finite state controller with an analog input measuring the temperature in the tank. The controller has a 1-bit output to control a heater in the tank. The tank always loses heat at a rate directly proportional to the temperature, and, while the heater is on, is heated at 4 degrees/second. Mathematically, after the heater reaches equilibrium in the on position  $A' = -A + 4$  and at equilibrium in the off position  $A' = -A$ . The controller switches the heater off within one second of the temperature going above a pre-set value, and turns the heater on within one second of the temperature dropping below another threshold. Note that this model assumes that the thermometer is perfect, the heater produces a constant and perfectly known heat output, the element heats and cools instantly and the physics of the tank are perfectly modeled by the ODE. Given those assumptions, they then use interval techniques to eliminate round off errors in proving safety properties.




Figure 2: State model allowing thermostat to shut off before element is warm

## 5 CLP(F) Model of the Thermostat

In this section we present two models of a thermostat. The first simple model demonstrates the key ideas. The second illustrates how one can easily extend the simple model to a model that more faithfully represents the real hybrid system by more closely approximating the physics of the system.

```

!!! must shorten long lines here

evolve(H,C,H).

evolve((S0,X0),C,(S,X)) :-
  statechange((S0,X0),C,S1),
  in_trajectory((S1,X0),C,X1),
  evolve((S1,X1),C,(S,X)).

in_trajectory((S0,(T0,A0)),[_I, Min, Max, ODEs],(T1,A1)) :-
  member(S0=ODE,ODEs), {T=T1-T0,T<I},
  ode((T0,A0),[T,ODE],A,(T1,A1)),
  ( (S0=on,      {[A in [-1000,Max] ]});
    (S0=off,    {[A in [Min,1000]]})).

statechange((S0,(T0,A0)),[_I, Min, Max, ODEs],S1) :-
  ( (S0=on,      {A0= Max},      S1=off);
    (S0=off,    {A0=Min },      S1=on) ).

ode((T0,A0),[I,[Alpha,Beta]],A,(T1,A1)) :-
  type([A],function(0,I)),
  {[ ddt(A,I) = Alpha * A + Beta,
    eval(A,0)=A0,  eval(A,T)=A1,
    A in [-1.0E100,1.0E100],
    T=T1-T0,  T in [0,I]
  ]}.
% call as ?- test(S,(T,A)),{A=2}. to find full/half cycle times.
test(S,X) :-
  % Step, Min, Max,  ODES for on/off states
  C=[ 2.0,  1.8, 2.3, [on=[-1,4],off=[-1,0]]],
  in_trajectory((on,(0,2)),C,X0),
  evolve((on,X0),C,(S,X)).

```

Figure 3: Simplest CLP(F) model of a thermostat

```

| ?- test(S,(T,A)),{A=2}.
A = 2, S = on, T = 0 ?
A = 2, S = off, T = 0.3022808718... ? ;
A = 2, S = on, T = 0.5029515673... ?

```

Figure 4: Query to Simple Model - When is Temp = 2?

## 5.1 The Simple Model

Our first CLP(F) model of a thermostat is shown in Figure 3. To clarify the key concepts, this first model assumes there are only two states: on and off. When the system state is on, the ODE governing the temperature  $A$  is  $A' = -A + 4$ . When the system state is off, the ODE is  $A' = -A$ . The system switches from on to off when the temperature rises above 2.3 and it switches from off to on when it drops below 1.8. The `in_trajectory` procedure models the trajectory by looking up the proper ODE for the current state and then calling the ODE procedure to constrain the new state variables  $(T1,A1)$ . It also adds the constraint that the temperature range is contained in  $[-1000, 2.3]$  (resp.  $[1.8,1000]$ ). This is not needed for our simple example because the temperature rises monotonically and then falls monotonically and then rises again. With more complex models, the temperature might not behave so nicely so this constraint states that no point in the trajectory has passed the threshold for switching. The `statechange` procedure simply indicates the condition that signals a state change and provides the new state. The `ode` procedure models the specified ODE as we have described above. Finally the `test` procedure shows how this program can be used to model the behavior of the system. It initializes the list describing the system to be analyzed and then invokes the `evolve` procedure. A simple query (“at what times is the temperature 2?”) to this system and the resulting answer is shown in Figure 4.

One subtle point about this model is that the CLP(F) solver will only work effectively if a finite step size is explicitly given (this is the `I` parameter appearing in the `in_trajectory` and `ode` procedures. If the step size is too large, then the CLP(F) solver will return very wide, unhelpful intervals for all variables. One approach to handling this is to introduce pseudo states `(on,n)`, `(off,n)`, where  $n$  is an integer representing the number of full steps that have been taken on the current trajectory in the current state. The continuous part can be modeled as `(t,a,z)` where  $t$  is the total elapsed time,  $a$  is the temperature at time  $t$ , and  $z$  is the time relative to the current step. Such an extension of the current technique is straightforward and we do not show it here due to space limitations.

```

evolve(H,C,H,[]).
evolve((SO,XO),C,(S,X),[(S1,X1,Range)|R]) :-
  statechange((SO,XO),C,S1),
  in_trajectory((S1,XO),Range,C,X1),
  evolve((S1,X1),C,(S,X),R).

in_trajectory((SO,(TO,AO,ZO)),R,[Step,Min,Max,Delay,Stime,
ODEs],[T1,A1,Z1]) :-
  member(SO=ODE,ODEs),{Z1=T,T=T1-TO,T=<Step},
  ode((TO,AO),[T,ODE],A,(T1,A1)),
  ((SO=on, {[A in [-1000,Max] ]});
  (SO=sw0, {[T < Delay]});
  (SO=cooling, {[T < Stime, A in [Min,1000] ]});
  (SO=off, {[A in [Min,1000] ]});
  (SO=sw1, {[T < Delay]});
  (SO=heating, {[T < Stime, A in [-1000,Max] ]})),
  {[V in [0,Z1], eval(A,V)=R]}.
statechange((SO,(TO,AO,T)),[Step,Min,Max,Delay,Stime,ODEs],S1) :-
  ((SO=on, {AO=Max}, S1 = sw0);
  (SO=sw0, {T=Delay}, S1=cooling);
  (SO=cooling, {T=Stime}, S1=off);
  (SO=cooling, {AO=Min}, S1=sw1);
  (SO=off, {AO=Min }, S1=sw1);
  (SO=sw1, {T=Delay}, S1=heating);
  (SO=heating, {T=Stime}, S1 = on);
  (SO=heating, {AO=Max}, S1 = sw0)).
ode((TO,AO),[T,[Alpha,Beta,Gamma,Delta]],A,(T1,A1)) :-
  type([A,B],function(0,I)),
  {[ ddt(A,I) = Alpha * A + Beta + Gamma*B,
  ddt(B,I) = Delta*(B + 0.1*sin(B)),
  eval(A,0)=AO, eval(A,T)=A1,
  eval(B,0)=1,
  A in [-1.0E100,1.0E100],
  B in [-1.0E100,1.0E100],
  T=T1-TO, T in [0,I]
  ]}.

test(S,X,[(on,XO,Range)|R],D) :-
  % Step, Min, Max, state0 temp0 ODES for on/off states
  C=[ 2.0, 1.8, 2.3, 0.05, 0.1,
  [on=[-1,4,0,1],off=[-1,0,0,1],
  sw0=[-1,4,0,1],sw1=[-1,0,0,1],
  heating=[-1,4,-4,D],cooling=[-1,0,4,D]
  ]],
  in_trajectory((on,(0,2,0)),Range,C,XO),
  evolve((on,XO),C,(S,X),R).

```

Figure 5: More Complete Model of Thermostat

```

| ?- test2(S,(T,A,TO),H,-10),{A=2,T>0},S=on.
H = [
%state T A Z range of A in state
(on, (0.1625...,2.300*, 0.1625..., [2.0000*, 2.3000*])),
(sw0, (0.2125...,2.3829..., 0.050* , [2.2999*, 2.3829*])),
(cooling, (0.3125...,2.3858..., 0.100* , [2.3827*, 2.4904*])),
(off, (0.5942...,1.8000*, 0.2817..., [1.7999*, 2.3858*])),
(sw1, (0.6442...,1.7122..., 0.050* , [1.7122*, 1.8000*])),
(heating, (0.7442...,1.7002..., 0.100* , [1.6042*, 1.7123*])),
(on, (0.8839...,2.000, 0.1396 , [1.7002*, 2.0000*]))]
S = on
T = 0.8839...
A = 2
TO = 0.1396... ?

(25080 ms) yes
| ?-

```

Figure 6: Example of a more complex model



## 5.2 A More Realistic Model

In the example shown in Figure 5, we refine the previous model by using six states `on,sw0,cooling,off,sw1,heating` corresponding to the states in Henzinger’s model. The model also represents the continuous state as a triple  $T, A, Z$  where  $T$  is the total elapsed time,  $A$  is the temperature at time  $T$ , and  $Z$  is the time since the system entered the current state. The  $Z$  parameter is needed to implement the “switching” specification which states that the system waits some amount of time after the threshold is passed before switching on/off the heating element. Likewise, the time in which the system is heating/cooling before it “jumps” to the maximum/minimum value is given by a time unit. This represents a discontinuity in the model since the heating temperature is assumed to immediately rise to the maximum at the end of the element-heating period.

The `sw0,sw1` states are when the system is waiting before switching the heating element on or off. The `heating,cooling` states are when the element is warming up or cooling down. The `on,off` states are when the element is fully on or off. Observe that the ODEs for each state are specified in the variable `C` of the test procedure. Also, observe that the switching conditions are given declaratively in the `statechange` procedure. Finally, note that the system is assumed to be modeled by the following more complex non-linear family of ODEs, where the parameters  $(\alpha, \beta, \gamma, \delta)$  vary from state to state:

$$\begin{aligned} \forall t \in [0, I] \quad A'(t) &= \alpha A(t) + \beta + \gamma B(t) \\ \forall t \in [0, I] \quad B'(t) &= \delta(B(t) + 0.1 \sin(B(t))) \\ T &= T1 - T0, \quad 0 \leq T \leq I \\ A(0) &= A0, \quad A(T) = A1, \quad B(0) = 1 \\ A([0, I]), B([0, I]) &\subset [-10^{100}, 10^{100}] \end{aligned}$$

The variable  $B$  represents the temperature of the heating element and we assume that the rate at which it heats and cools depends on its temperature in a non-linear fashion. The test query in Figure 6 demonstrates the use of this model to determine the value of the length of one complete cycle.

The following code shows a more interesting example in which the modeler is used to find all values of the ODE parameter  $\delta$  in the range  $[-8, -7]$  for which the system evolves to the state with `S=off` and  $A = 2$  in exactly 0.5 seconds.

```
| ?- {D in [-8,-7]},
    test2(S,(T,A,Z),_H,D),{T=0.5},S=off,
        narrow_all(1000000).
A = 2
D = -7.6651...
S = off
T = 0.500*
Z = 0.1874810705022... ?;
(14280 ms) no
```

## 6 Conclusions

### 6.1 Strengths of CLP(F)

A novel aspect of the ACLP approach to hybrid system analysis is that it establishes a close correspondence between the semantics of a particular class of constraint programs and the behavior of hybrid systems providing several advantages: The simple mapping means that one needn’t worry about the “translation” from ODEs to CLP(F). Interval techniques guarantee that calculated safety properties are correct, and protect against round-off errors, while also providing a natural technique to handle error bars on physical measurements. While other hybrid system models can be extended to handle error bars in measurements, CLP(F) handles them naturally with no extra effort either in specifying them or in the calculation. A

further advantage is that CLP(F) handles non-linear ODEs directly and soundly. CLP(F) also allows one to incrementally refine a model as one learns more about the physics of a system. At areas near a transition, one need not understand the details of the transition, but can simply bound the behaviour in that area, and get sound results. As one learns more about the physics, one can refine the model.

In addition CLP(F) is a very expressive language. It is easy to state problems involving finding values of the control parameters which result in specified behavior. Similarly, finding the time of state transitions is a simple matter of describing the transition as a constraint. Once the problem is stated, the underlying Prolog interpreter automatically solves it. Other ODE approaches often require an explicit binary search to find such times.

## 6.2 Limitations and Future Work

Currently there are several limitations on using CLP(F) to model hybrid systems. Each of them is an obvious possibility for future work: It would be helpful to develop more efficient interpreters for CLP(F) to handle very large complex systems. So far we have put almost no work into the efficiency of CLP(F), so there is a great deal of room for improvement here. We would like to extend this work to hybrid systems where the sensors are governed by PDEs rather than ODEs. Further efficiency improvements would come from developing primitive implementations of the `ode` procedures so that one does not need to use the full power of ACLP (and its accompanying inefficiencies).

We would like to develop more realistic models of hybrid systems using this approach.

CLP is not complete. If the search is bounded (e.g. there is a limit on the time parameter), then the entire search space will be traversed via backtracking and the incompleteness of Prolog is not an issue. The incompleteness of the CLP solver is however an issue whose consequence is that we do not know for certain whether any particular answer constraint actually contains a solution.

There is room for improvement in the heuristics of the underlying constraint solver. Since the CLP system guarantees soundness, extra attempts to narrow cannot introduce errors, but they do take time. Better heuristics could improve performance, both by reducing the running time and by narrowing the resulting intervals.

This paper has demonstrated that Analytic Constraint Logic Programming provides a promising approach to modeling hybrid systems by providing a program whose semantics precisely match the behavior of the Hybrid System. Further research is needed to see if such an approach can be scaled up to real-life systems.

## References

- [ACH<sup>+</sup>95] Rajeev Alur, Costas Courcoubetis, Nicolas Halbwachs, Thomas A. Henzinger, Pei-Hsin Ho, Xavier Nicollin, Alfredo Olivero, Joseph Sifakis, and Sergio Yovine. The algorithmic analysis of hybrid systems. *Theoretical Computer Science*, 138:3–34, 1995.
- [Act96] Forman S. Acton. *Real computing made real: Preventing Errors in Scientific and Engineering calculations*. Princeton University Press, Princeton, New Jersey, 1996.
- [BO97] Frédéric Benhamou and William J. Older. Applying interval arithmetic to real, integer, and boolean constraints. *Journal of Logic Programming*, 32(1):1–24, Jul 1997.
- [CF00] Angelo E.M. Ciarlini and Thom Frühwirth. Automatic derivation of meaningful experiments for hybrid systems. In *ACM SIGSIM Conference on AI, Simulation and Planning (AIS '2000)*, Mar 2000.
- [DGS] Akash Deshpande, Aleks Göllü, and Luigi Semenzato. *The SHIFT Programming Language and Run-time System for Dynamic Networks of Hybrid Automata*. Department of Electrical Engineering and Computer Sciences; University of California at Berkeley, Berkeley, CA 94720. <http://www.path.berkeley.edu/shift/doc/ieeshift.ps.gz>.

- [DJVH02] Yves Deville, Micha Janssen, and Pascal Van Hentenryck. Consistency techniques in ordinary differential equations. *Constraints*, 7(3):289–315, Jul 2002.
- [GJS96] Vineet Gupta, Radha Jagadeesan, and Vijay Saraswat. Hybrid cc, hybrid automata and program verification. In Rajeev Alur, Thomas A. Henzinger, and Eduardo D. Sontag, editors, *Hybrid Systems III: Verification and Control*, volume 1066 of *LNCS*, pages 52–63. Springer Verlag, 1996.
- [HHMWT00] Thomas A. Henzinger, Benjamin Horowitz, Rupak Majumdar, and Howard Wong-Toi. Beyond HYTECH: Hybrid systems analysis using interval numerical methods. In Nancy Lynch and Bruce H. Krogh, editors, *Hybrid Systems: Computation and Control (HSCC 2000)*, volume 1790 of *LNCS*, pages 130–144. Springer Verlag, 2000.
- [HHWT98] Thomas A. Henzinger, Pei-Hsin Ho, and Howard Wong-Toi. Algorithmic analysis of nonlinear hybrid systems. *IEEE Transactions on Automatic Control*, 43:540–554, 1998.
- [Hic00] Timothy J. Hickey. Analytic constraint solving and interval arithmetic. In *POPL’00 ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 338–351, 2000. published as vol. 27 of SIGPLAN notices.
- [Hic01] Timothy J. Hickey. Metalevel interval arithmetic and verifiable constraint solving. *Journal of Functional and Logic Programming*, 2001(7), October 2001. <http://danae.uni-muenster.de/lehre/kuchen/JFLP/articles/2001/S01-02/JFLP-A01-07.pdf>.
- [JL87] J. Jaffar and J.L. Lassez. Constraint logic programming. In *Proceedings 14th ACM Symposium on the Principles of Programming Languages*, pages 111–119, 1987.
- [JM94] Joxan Jaffar and Michael J. Maher. Constraint logic programming: A survey. *Journal of Logic Programming*, 19/20:503–581, 1994.
- [Kui93] Benjamin J. Kuipers. Qualitative simulation: Then and now. *Artificial Intelligence*, 59:133–140, 1993.
- [LSV01] Nancy Lynch, Roberto Segala, and Frits Vaandrager. Hybrid I/O automata revisited. In Maria Domenica Di Benedetto and Alberto Sangiovanni-Vincentelli, editors, *HSCC*, number 2034 in *LNCS*, pages 403–417. Springer Verlag, 2001.
- [LSVW99] Nancy Lynch, Roberto Segala, Frits W. Vaandrager, and H.B. Weinberg. Hybrid I/O automata. Technical Report CSI-R9907, Computing Science Institute Nijmegen; Faculty of Mathematics and Informatics; Catholic University of Nijmegen, Toernooiveld 1; 6525 ED Nijmegen; The Netherlands, Apr 1999.
- [MMP91] Oded Maler, Zohar Manna, and Amir Pnueli. From timed to hybrid systems. In J.W. de Bakker, C. Huizing, W.P. de Roever, and G. Rozenberg, editors, *Real-Time: Theory in Practice*, volume 600 of *LNCS*, pages 447–484. Rex Workshop, Springer Verlag, 1991.
- [Moo66] Ramon E. Moore. *Interval Analysis*. Prentice-Hall, 1966.
- [Mos99] Pieter J. Mosterman. An overview of hybrid simulation phenomena and their support by simulation packages. In Frits W. Vaandrager and Jan H. van Schuppen, editors, *Hybrid Systems: Computation and Control*, volume 1569 of *LNCS*, pages 165–177. Springer Verlag, 1999.
- [Pod00] Andreas Podelski. Model checking as constraint solving. In Jens Palsberg, editor, *Proceedings of SAS’2000: Static Analysis Symposium*, Jun 2000.

- [Urb96] Luis Urbina. Analysis of hybrid systems in  $\text{CLP}(\mathcal{R})$ . In Eugene C. Freuder, editor, *Principles and Practice of Constraint Programming – CP96*, volume 1118 of *LNCS*, pages 451–467. Springer, Aug 1996.