A Modular Representation of Fluid Stochastic Petri Nets

Fernando J. Barros
Departamento de Engenharia Informática
Universidade de Coimbra
3030 Coimbra, Portugal

Abstract

In this paper we develop a modular representation of Fluid Stochastic Petri Nets (FSPNs) using the Hybrid Flow Systems Specification (HFSS, a formalism that combines the concepts of sampling and discrete events to describe hybrid systems. We show that HFSS provides a sound representation of FSPNs supporting a direct mapping between FSPNs elements and HFSS components. FSPNs can be modeled by a composition of HFSS components preserving the structure of the original FSPNs, removing the need for a model transformation layer to simulate FSPNs, or making it easy to develop such a mapping mechanism. We show that the continuous flow representation used by HFSS enables an efficient simulation of FSPNs. Simulation results are presented for a simple manufacturing system with machines subjected to breakdowns.

1 Introduction

Petri Nets (PNs) have widely been used in modeling and analysis of systems. Since their creation many extensions have been developed, including, for example, Timed Petri Nets (TPNs), and Fluid Stochastic Petri Nets (FSPNs) aimed to model timed hybrid systems exhibiting both discrete and continuous elements [9]. In this paper we develop a modular representation of FSPNs using the Hybrid Flow System Specification Formalism (HFSS) [2]. HFSS combines both continuous [1], and discrete (event) flows [10], to represent hybrid systems. We develop a library of HFSS components to represent the elements of a FSPN. These elements are described using the HFSS-Groovy toolkit and include discrete places, transitions and continuous places. We introduce a conflict manager component to explicitly represent transition tie-breaking rules, enabling the use of application dependent algorithms to choose among conflicting transitions. We use infinite server semantics enabling an arbitrary number of transitions to fire simultaneously. The library of HFSS components permits FSPNs to be represented by a structural equivalent HFSS network, that can be obtained through composition using simple transformation rules, in a direct mapping. Structure preserving makes this conversion a very simple process that can be performed manually, also making it easier to define conversion tools, avoiding namely the (costly) compilers that are common in Domain Specific Languages (DSLs) approaches [6]. We present simulation results for a simple manufacturing system with machines subjected to breakdowns. Our results show that HFSS representation of continuous systems by HFSS continuous flows enable an efficient simulation of FSPNs.

The paper is organized as follows. Section 2 presents the semantics of FSPNs in an informal manner. In Section 3 we present a library of HFSS-Groovy components that provides a representation of basic FSPNs elements. This section describes also FSPNs as a composition of HFSS components. Simulation results are presented in Section 4. Related work is discussed in Section 5.

2 Fluid Stochastic Petri Nets

FSPNs introduce continuous marking for supporting a representation of hybrid systems [9]. The continuous places of PNs are described by constant rate differential equations enabling a fluid approximation of systems with a large number of tokens. We present next an informal description of FSPN semantics.

2.1 Discrete TPNs

Time was introduced in Petri Nets to model system delays. The time to complete a task or a delay in the system can usually be modeled by a stochastic distribution. Timed Petri Nets (TPNs) define a set of transitions, places and arcs. A transition checks its preconditions that depend on the marking of its input places. If preconditions are satisfied, a set of tokens is removed from the input places. After the transition, tokens are added to transition output places. In this paper we assume that time elapses inside transitions [8], departing from the more common TPNs where time elapses in places or arcs. In this paper we also assume infinite server semantics, that allows many transitions to start simultaneously as long as their preconditions are satisfied. The semantics of TPNs can be described by Figure 1 that depicts the behavior of a TPN with transition t_0 and places $p_0, ..., p_3$. The transition precondition requires two units of p_0, one unit of p_1 and it imposes an inhibitor arc of two units of p_2. When the precondition is satisfied a token representing an activity (transition instance) is created within the transition to model a time advance (delay). After this time interval, the transition finishes, and new tokens are created in the corresponding output places. Given the initial marking of Figure 1a, t_0 can start two activities, Figure 1b. Transition t_0 removes all tokens from t_0 and t_1 and creates two time events to signal the end of the scheduled activities.

Although transition marking is commonly omitted in Petri net analysis, a PN simulator needs to consider it. A standard marking representation could also be used, since a time-in-transition PN can be mapped into a time-in-place PN [8]. The non-standard representation introduced here simplifies simulator description.

When an instance of t_0 finishes the execution it creates new tokens in its output places, Figure 1c, in this Petri net, 1 unit of p_2 and 3 units of p_4. The final marking after the firing of the second instance of t_0 is depicted in Figure 1d. We have assumed that each transitions instances were assigned to a different time duration. Was the t_0 associated with a fixed processing time and the two instance would have finished at the same time, jumping the intermediate step of Figure 1c.

2.2 Continuous Flows

Fluid Stochastic Petri Nets (FSPNs) were introduced to enable the description of systems requiring a large number of tokens, since an explicit representation of each token would make the PN difficult to analyze and also time consuming to simulate. A fluid approximation is used, instead, becoming tokens represent by a real number whose value is governed by a piecewise constant rate.

Before detailing the semantics of FSPNs we define $|t_k|$ as the number of instances of transition t_k currently active. Similarly the quantity (integer or real) of tokens in place p_k if given by $|p_k|$.

In FSPNs, a transition t_k is considered active iff $|t_k| > 0$. When a transition is active the corresponding flow is enabled and place content is influenced by that flow. On the contrary, when not active the corresponding flow is zero. Another constraint imposes that places can only contain positive values, i.e., $|p_k| \geq 0$.

Figure 2 represents a FSPN with transitions t_0, t_1, t_2, place p_0, constant flows a and b, and a variable flow controlled by the number of tokens in transition t_2.

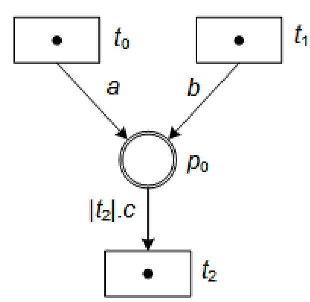

Figure 2: Continuous flow Petri net with variable rate $|t_2| \cdot c$.

When the all transitions are enabled the content of place p_0 is described by:

$$\frac{\mathrm{d}|p_0|}{\mathrm{d}t} = a + b - |t_2| \cdot c$$

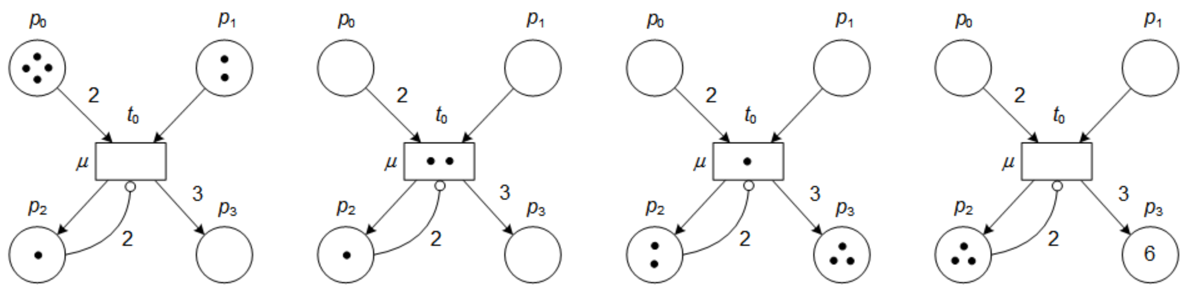

(a) initial marking (b) fire of transition t_0 (c) end of processing of a token in t_0 (d) final marking

Figure 1: Discrete Petri net evolution.

When a transition is disable, the corresponding flow is zero. For example, when $|t_0| = 0$, then $\frac{\mathrm{d}|p_0|}{\mathrm{d}t} = b - |t_2| \cdot c$

2.3 Hybrid Flows

FSPNs enable the representation of systems with both discrete and continuous semantics. The FSPN of Figure 3 models a manufacturing system with N machines that process at an (exponential) rate μ and breakdown at an (exponential) rate λ. Entities enter the system at rate a and are processed at rate $|t_1| \cdot d$. The initial number of entities to be produced is given by L and all machines are initially available, $|t_1| = N$. Given the semantics defined before, $|t_1|$ represents the number of machines available for production and $|t_2|$ is the number of machines being repaired (not working).

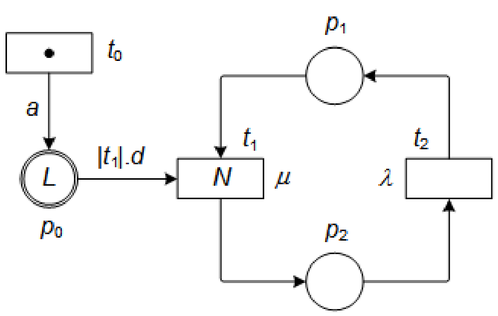

Figure 3: Petri net with hybrid flows (FSPN).

We have described the main elements of FSPNs. In the next section we provide their representation in the HFSS formalism.

3 Modular Representation of FSPNs

The mapping of FSPNs into a deterministic modeling and simulation formalisms has several advantages. It establishes FSPNs semantics, since modeling formalism like HFSS have deterministic semantics [3]. Modularity enable also the composition of systems from simple elements. Given a FSPN model library developed in HFSS, we can create complex FSPNs by simple composition of basic elements without the need to develop a compiler to generate new HFSS models from a FSPN specification, a requirement usual in Domain Specific Languages for representing (non-timed) PNs [6]. Additionally, mapping FSPNs into a modeling formalism can also exploit the advantages of existing and efficient simulation kernels without the need to create a specific solution for FSPNs.

The HFSS formalism combines several abstractions, including adaptive sampling, continuous flows, and discrete events. HFSS models are modular communicating through a well defined interface. A HFSS model can read and produce continuous and discrete flows (events), offering a framework for defining hybrid models [3]. The HFSS-Groovy toolkit is a Groovy language implementation of the HFSS formalism and it is used in the next sections to describe the HFSS components required to represent FSPNs.

3.1 Discrete Places

We start by describing a HFSS model of FSPNs discrete places. This model requires the ability to change its content supporting the basic discrete event operations of adding and removing tokens, and to communicate changes in this value. Since a key operation in a FSPN is to test preconditions we provide the access to the place current number of tokens through the component continuous output flow function. The HFSS `Place` is represented in Figure 4

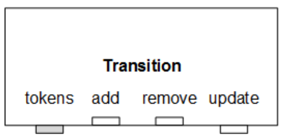

Figure 4: HFSS discrete `Place` model.

HFSS-Groovy definition of the `Place` model is given in Listing 2. Class `Model` provides the basic support to HFSS and it defines variable `alpha` and `beta` to set the time to read (sample), and the time to write (produce a discrete flow). A `Place` is created with an initial number of `tokens` and it becomes passive (`alpha = beta` = ∞), waiting for an input.

```
1  public class Place extends Model {
2     private int tokens;
3     public Place(String name, int tokens) {
4        super(name);
5        this.tokens = tokens;
6        alpha = Double.POSITIVE_INFINITY;
7        beta = Double.POSITIVE_INFINITY;
8     }
9     public void transition(double e, def xc, def xd) {
10       beta = Double.POSITIVE_INFINITY;
11       if (xd == null) return;
12       int prev = tokens;
13       xd.at("add").each {Port p-> tokens += p.value()}
14       xd.at("remove").each {Port p-> tokens -= p.value()}
15       if (prev != tokens) beta = 0;
16    }
17    public def outputC(double e) {return new Port("tokens", tokens)}
18    public def outputD(double e) {return new Port("update", tokens)}
19 }
```

Listing 1: HFSS-Groovy `Place` model.

A `Place` receives commands to change its content in ports `add` and `remove`. Since HFSS uses a parallel semantics, an input is usually a list of pairs in the form (port name, value). The transition function (line 9), specifies the behavior of `Place` input arrival. Each `add` increases the number of tokens, and each `remove` decreases this value. When the number of tokens is modified the updated value is sent through the discrete port `update`, as defined by function `outputD` (discrete output). The current number of tokens is always available through the continuous output flow port `tokens`, as specified by the function `outputC` (continuous output). This continuous value plays a key role in simplifying the definition of the HFSS model of a FSPN as we show in the next sections. Typical `Place` trajectories are shown in Figure 5. While the input is discrete, the output has a (piecewise constant) continuous flow with the current number of tokens, and a discrete flow trajectory, signaling a change in this number.

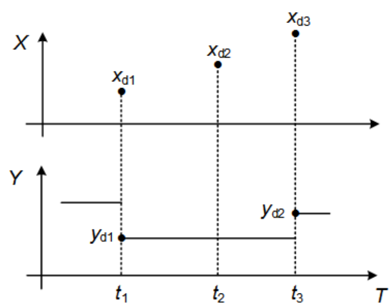

Figure 5: HFSS `Place` trajectory.

3.2 Continuous/Fluid Places(Reservoirs)

When the number of tokens in a FSPN is very large a fluid approximation simplifies the analyses and enables a more efficient simulation. In this approximation transitions are continuous and become characterized by the rate they modify reservoir (fluid places) contents. The HFSS model of a reservoir in given in Figure 6.

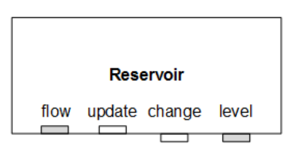

Figure 6: HFSS `Reservoir` model.

The `Reservoir` samples the input flow and integrates this value, that is constrained to be positive. Given a negative input rate, reservoir level decreases until it reaches zero, and keeps this value irrespective to a negative input rate. Since, FSPNs constrain input rates to be piecewise constant, fluid integration involves only transition when these rates changes. Input port `update` receives a signal when the rate is modified. The current rate is available at input port `flow`. The reservoir defines also the continuous output port `level`, to provide access to the current reservoir level, and the discrete port `change` to signal a modification in the reservoir flow rate. The HFSS-Groovy implementation is described in Listing 2. Where method `rate`, line 10, computes the effective input rate, and method `level`, line 34, computes reservoir contents. When the rate is zero, line 28, the model becomes passive.

```
1  class Reservoir extends Model {
2     private double level;
3     private double rate = 0.0;
4     public boolean isZERO(double x) {return Math.abs(x) <= 1.0e-12}
5     public Tank(String name, double level) {
6        super(name);
7        this.level = level;
8        alpha = 0.0;
9     }
10    private double rate(double r) {
11       if (isZERO(r)) return 0;
12       if (isZERO(level) && r < 0) {
13          level = 0;
14          return 0;
15       }
16       return r;
17    }
18    public void transition(double e, def xc, def xd) {
19       this.passivate();
20       level = level(e);
21       double nextRate = rate(xc.value());
22       if (! isZERO(nextRate - rate)) {
23          rate = nextRate;
24          beta = 0.0;
25          return;
26       }
27       rate = nextRate;
28       if (isZERO(rate)) return;
29       if (rate < 0) {
30          beta = -level / rate;
31          return;
32       }
33    }
34    private double level(double e) {return level + rate * e}
35    public def outputC(double e) {return new Port("level", level(e))}
36    public def outputD(double e) {return new Port("change",
             level(e))}
37 }
```

Listing 2: HFSS-Groovy `Reservoir` model.

Typical reservoir input and output trajectories are depicted in Figure 7. The input flow changes at instants t_1, t_2 and t_4. This value is sampled from the continuous input trajectory and imposed by the arrival of discrete flows. Tank flow level is continuous, a piecewise linear flow, as a consequence of the piecewise constant flow rate constraint. A discrete flow signal is produced each time the flow changes. In the interval $[t_3, t_4]$ the reservoir content is zero due to the negative input rate.

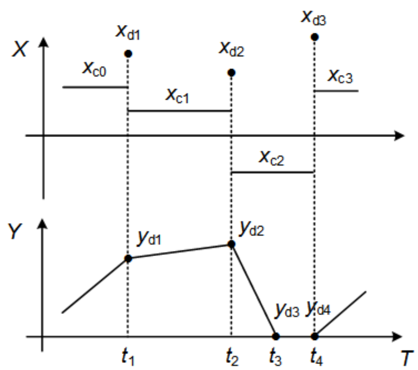

Figure 7: HFSS reservoir trajectories.

3.3 Transitions

HFSS transition model is represented in Figure 8. A transition samples the precondition from the continuous ports `tokens[n]`, that are connected to the input places the transition depends upon. If the precondition is true the transition tries to seize the tokens from the `Conflict Manager` component that manages place access conflicts. When an acknowledged is received the transition executive $\texttt{Transition}_\eta$ creates a `Delay` to signal transition end. We assume the infinite server semantics and thus multiple copies of a transitions (instances) can execute in parallel. When a `Delay` finishes it signals the executive that removes it from the network and releases the corresponding tokens through the output ports `add[n]`. A transition also checks for its precondition when receives an `update` message sent by a place that has changed its number of tokens. The current number of `Delay` instances can be sampled at the continuous output port `number`. When this value changes a signal is sent through the discrete output port `number`.

3.4 Conflict Manager

Given the infinite server assumption, and since HFSS has a parallel semantics, implying that all transitions scheduled to the same time must be fired simultaneously, conflicts can arise. Conflicts need to be solved in a deterministic manner, and under the control of the modeler. These requirements impose a centralized controller to decide what transitions can be fired and those that need to hold. The `ConflictManager` model is represented in Figure 9. It receives requests in port `seize` from transitions whose preconditions are satisfied. The information about tokens availability is sampled at input ports `tokens[n]`. At this point the conflict manager has full control on what transitions can proceed. Many strategies to break ties are possible, including a probabilistic rule, fairness considerations, waiting time, priorities, etc. Since we have an explicit representation of conflicts it becomes possible to use any algorithm suitable for a particular goal.

When the manager decides to enabling a transition it removes the tokens from the corresponding places using output port `remove[n]`. Transition acknowledgment is made through output port `ack[n]`.

We have described the basic HFSS components that can be used to represent FSPNs elements. In the next sections we show how these HFSS components can be combined to create arbitrary FSPNs.

3.5 Continuous Flow HFSS Model

We start the HFSS representation of FSPNs by considering the continuous FSPNs of Figure 2. The mapping from the FSPN to the corresponding HFSS network model is straightforward and is depicted in Figure 10.

Reservoir P0 samples the input rate from transitions T1, T2 and T3, and any change in the number of tokens, forces a sampling operation in P0. Thus P0 input rate is updated whenever there is a change in any input transition.

3.6 Discrete Flow HFSS Model

The mapping of the FSPN of Figure 1 into a HFSS model is represented in Figure 11. FSPN

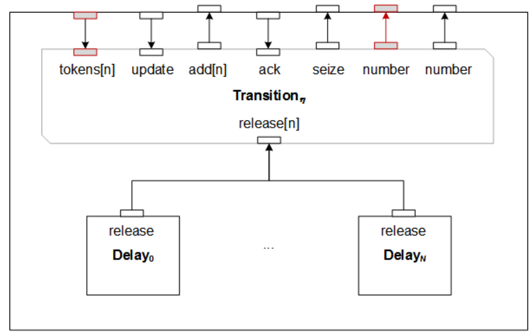

Figure 8: HFSS transition model.

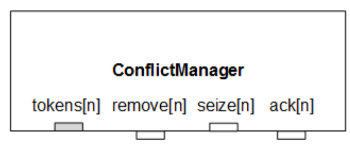

Figure 9: HFSS conflict manager model.

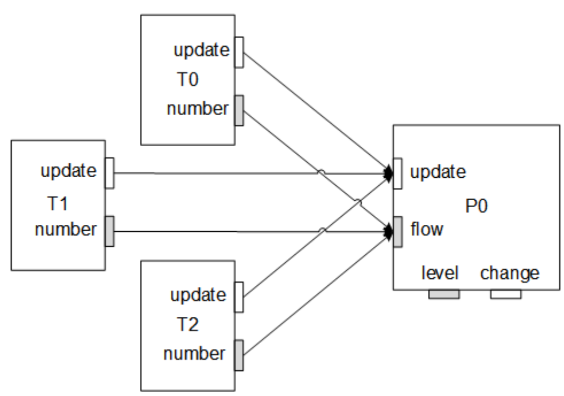

Figure 10: HFSS continuous flow network model.

places and transitions are mapped into the corresponding HFSS components described before. The conflict manager completes the HFSS model since it is required to handle the access to (common) places.

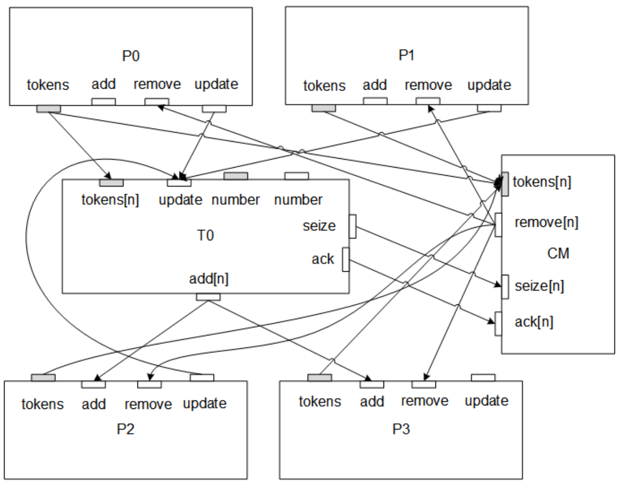

Figure 11: HFSS discrete flow network model.

We can observe that the structure of the original PN is mainly preserved. Transition T0 reads from places P0 and P1 and it produces tokens to places P2 and P3. However, the arity of the arcs is stored in ports input function and it is not explicitly represented in the Figure 11. The conflict manager is introduced to explicitly represent the algorithm for solving collisions when transitions access to the same resources. This element is not used in PN diagrams and need to be specified by some textual annotation. Our approach makes it possible to define different strategies and making them reusable HFSS components that can be chosen according to the system requirements.

4 Simulation Results

For the validation of the HFSS representation of FSPNs we use the petri net of Figure 3, with parameters $L = 90$, $N = 5$, $\mu = 1/4.0$, $\lambda = 1/1.5$, $a = 50$, and $d = 14$. The input rate of p_0 is set by HFSS-Groovy method `influencers` defined by Listing 3.

```
influencers("P0", ["T0", "T1"],
  {List xc->
    def res = xc.filterByPort(["number"]);
    return res[0].value() * 50 - res[1].value() * 14;
  }, {List xd->
    def res = xd.filterByPort(["number"]);
    return res;
});
```

Listing 3: P0 input function.

Simulation results for the contents of reservoir p_0 are depicted in Figure 12, where discrete flows are represented by squares and the continuous flow is piecewise linear.

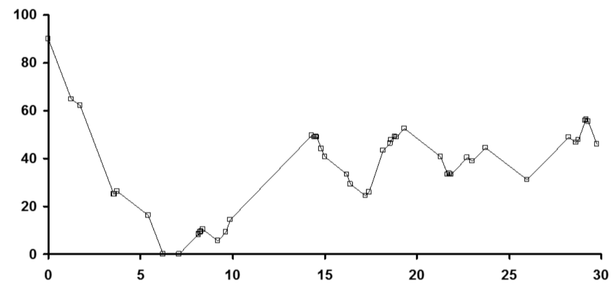

Figure 12: Marking of place p_0.

As can be observed from the graphic, the number of transitions is very small since the HFSS representation exploits the piecewise constraint of input flows to achieve an efficient simulation. Given that HFSS models can produce continuous output flows, reservoir continuous trajectory can be calculated without involving any transition, reducing the computation cost.

5 Related Work

To the best of our knowledge we have developed the first modular description of FSPNs. The representation of Petri Nets in discrete event formalisms, like DEVS, have been described [7]. However, these approaches relate only to non-timed PNs with single server semantics having simpler requirements when compared to the FSPNs modeled in this paper. The discrete event simulation of FSPNs has been described in [4]. This work, however, provides a description of an *ad hoc* implementation not supported by a modular representation. Implementations based on fixed sampling rates have been developed [5], but results are dependent on the sampling rate, and loosing accuracy when compared with the *exact* results achieved by HFSS models. The use of HFSS sampling provides a simplification in the messages required to retrieve information. Taken for example the conflict manager component, it requires one sampling operation to retrieve token information from all connected places. The alternative discrete event representation would require sending a message to each place and then to wait for every answer to arrive. HFSS sampling enables all this information to be exchanged in a single atomic operation, being easier to represent, and levering model reuse. Domain Specif languages (DSLs) have been developed to create simulations from (non-timed) PNs. However, these approaches require the development of specific compilers in order to be applied [6]. On the contrary, our approach based on a domain specific (FSPN) library makes the mapping between FSPNs and simulations very simple, removing the need for expensive compilers.

6 Conclusion and Future Work

We have developed a modular representation of FSPNs based on the HFSS formalism that enables a direct map of net elements, like places and transitions, into a HFSS components. The HFSS representation preserves the structure of the original FSPN, becoming very easy to develop and amenable to a manual translation. Our approach makes it also possible FSPNs to be defined directly from a library of HFSS components without requiring the support for translation/generation tools. As future work we plan to model Hybrid Petri Nets [5], an alternative to FSPNs with more powerful semantics, namely the ability to describe inhibitor arcs associated with continuous places.

References

[1] F. Barros. Towards a theory of continuous flow models. *International Journal of General Systems*, 31(1):29–39, 2002.

[2] F. Barros. Dynamic structure multi-paradigm modeling and simulation. *ACM Transactions on Modeling and Computer Simulation*, 13(3):259–275, 2003.

[3] F. Barros. Semantics of discrete event systems. In *Distributed Event-Based Systems*, pages 252–258, 2008.

[4] G. Ciardo, D. Nicol, and K. Trivedi. Discrete-event simulation of fluid stochastic Petri Nets. *IEEE Transactions on Software Engineering*, 25(2):207–217, 1999.

[5] R. David and H. Alla. *Discrete, Continuous, and Hybrid Petri Nets*. Springer, 2005.

[6] R. Esser and J. Janneck. A framework for defining domain-specific visual languages. In *OOPSLA Workshop on Domain-Specific Visual Languages*, 2001.

[7] C. Jacques and G. Wainer. Using the C++ DEVS toolkit to develop Petri Nets. In *Summer Computer Simulation Conference*, pages 252–258, 2002.

[8] C. Ramchandani. *Analysys of Asynchronous Concurrent Systems by Timed Petri Nets*. PhD thesis, Massachusetts Institute of Tecnology, 1974.

[9] K. Trivedi and V. Kulkarni. FSPNs: Fluid stochastic Petri Net. In *Application and Theory of Petri Nets, Lecture Notes in Computer Science*, volume 691, pages 24–31, 1993.

[10] B. Zeigler, H. Praehofer, and T. Kim. *Theory of Modeling and Simulation: Integrating Discrete Event and Continuous Complex Dynamic Systems*. Academic Press, 2000.