Research Article

Lutin: A Language for Specifying and Executing Reactive Scenarios

Pascal Raymond, Yvan Roux, and Erwan Jahier

VERIMAG (CNRS, UJF, INPG), 2 avenue de Vignate, Gières 38610, France

Correspondence should be addressed to Pascal Raymond, pascal.raymond@imag.fr

Received 13 September 2007; Accepted 10 January 2008

Recommended by Michael Mendler

This paper presents the language Lutin and its operational semantics. This language specifically targets the domain of reactive systems, where an execution is a (virtually) infinite sequence of input/output reactions. More precisely, it is dedicated to the description and the execution of constrained random scenarios. Its first use is for test sequence specification and generation. It can also be useful for early simulation of huge systems, where Lutin programs can be used to describe and simulate modules that are not yet fully developed. Basic statements are input/output relations expressing constraints on a single reaction. Those constraints are then combined to describe non deterministic sequences of reactions. The language constructs are inspired by regular expressions and process algebra (sequence, choice, loop, concurrency). Moreover, the set of statements can be enriched with user-defined operators. A notion of stochastic directives is also provided in order to finely influence the selection of a particular class of scenarios.

Copyright © 2008 Pascal Raymond et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

1. INTRODUCTION

The targeted domain is the one of reactive systems, where an execution is a (virtually) infinite sequence of input/output reactions. Examples of such systems are control/command in industrial process, embedded computing systems in transportation.

Testing reactive software raises specific problems. First of all, a single execution may require thousands of atomic reactions, and thus as many input vector values. It is almost impossible to write input test sequences by hand; they must be automatically generated according to some concise description. More specifically, the relevance of input values may depend on the behavior of the program itself; the program influences the environment which in turn influences the program. As a matter of fact, the environment behaves itself as a reactive system, whose environment is the program under test. This feedback aspect makes offline test generation impossible; testing a reactive system requires to run it in a simulated environment.

All these remarks have led to the idea of defining a language for describing random reactive systems. Since testing is the main goal, the programming style should be close to the intuitive notion of test scenarios, which means that the language is mainly control-flow oriented.

The language can also be useful for early prototyping and simulation, where constrained random programs can implement missing modules.

1.1. Our proposal: Lutin

For programming random systems, one solution is to use a classical (deterministic) language together with a random procedure. In some sense, nondeterminism is achieved by relaxing deterministic behaviors. We have adopted an opposite solution, where nondeterminism is achieved by constraining chaotic behaviors; in other terms, the proposed language is mainly relational not functional.

In the language Lutin, nonpredictable atomic reactions are expressed as input/output relations. Those atomic reactions are combined using statements like sequence, loop, choice or parallel composition. Since simulation (execution) is the goal, the language also provides stochastic constructs to express that some scenarios are more interesting/realistic than others.

Since the first version [1], the language has evolved with the aim of being a user-friendly, powerful programming language. The basic statements (inspired by regular expressions) have been augmented with more sophisticated control structures (parallel composition, exceptions) and

a functional abstraction has been introduced in order to provide modularity and reusability.

1.2. Related works

This work is related to synchronous programming languages [2, 3]. Some constructs of the language (traps and parallel composition) are directly inspired by the imperative synchronous language Esterel [4], while the relational part (constraints) is inspired by declarative languages like Lustre [5] and Signal [6].

Related works are abundant in the domain of models for nondeterministic (or stochastic) concurrent systems: Input/Output automata [7], and their stochastic extension [8] (stochastic extension of process algebra [9, 10]). There are also relations with concurrent constraint programming [11], in particular, with works that adopt a synchronous approach of time and concurrency [12, 13]. However, the goals are rather different; our goal is to maintain an infinite interaction between constraints generators, while concurrent constraint programming aims at obtaining the solution of a complex problem in a (hopefully) finite number of interactions.

Moreover, a general characteristic of these models is that they are defined to perform analysis of stochastic dynamic systems (e.g., model checking, probabilistic analysis). On the contrary, Lutin is specifically designed for simulation rather than general analysis. On one hand, the language allows to concisely describe, and then execute a large class of scenarios. On the other hand, it is in general impossible to decide if a particular behavior can be generated and even less with which probability.

1.3. Plan

The article starts with an informal presentation of the language. Then, the operational semantics is formally defined in terms of constraints generator. Some important aspects, in particular constraints solving, are parameters of this formal semantics; they can be adapted to favor the efficiency or the expressive power. These aspects are presented in the implementation section. Finally, we conclude by giving some possible extensions of this work.

2. OVERVIEW OF THE LANGUAGE

2.1. Reactive, synchronous systems

The language is devoted to the description of nondeterministic reactive systems. Those systems have a cyclic behavior; they react to input values by producing output values and updating their internal state. We adopt the synchronous approach, which here simply means that the execution is viewed as a sequence of pairs “input values/output values.”

Such a system is declared with its input and output variables; they are called the *support variables* of the system.

Example 1. We illustrate the language with a simple “tracker” program that receives a boolean input (c) and a real input (t) and produces a real output (x). The high-level specification

of the tracker is that the output x should get closer to the input t when the command c is true or should tend to zero otherwise. The header of the program is

```
system tracker (c: bool; t: real) returns (x: real)
= statement. (1)
```

The core of the program consists of a *statement* describing the program behavior. The definition of *statement* is developed later.

During the execution, inputs are provided by the system environment; they are called *uncontrollable variables*. The program reacts by producing outputs; they are called *controllable variables*.

2.2. Variables, reactions, and traces

The core of the system is a statement describing a sequence of atomic reactions.

In Lutin, a reaction is not deterministic; it does not define uniquely the output values, but states some *constraints* on these values. For instance, the constraint $((x > 0.0) \text{ and } (x < 10.0))$ states that the current output should be some value comprised between 0 and 10.

Constraints may involve inputs, for instance, $((x > t - 2.0) \text{ and } (x < t))$. In this case, during the execution, the actual value of t is substituted, and the resulting constraint is solved.

In order to express temporal constraints, *previous values* can be used; $\text{pre } id$ denotes the value of the variable id at the previous reaction. For instance, $(x > \text{pre } x)$ states that x must increase in the current reaction. Like inputs, pre variables are uncontrollable; during the execution, their values are inherited from the past and cannot be changed—this is the *nonbacktracking principle*.

Performing a reaction consists in producing, if it exists, a particular solution of the constraint. Such a solution may not exist.

Example 2. Consider the constraint

$$(c \text{ and } (x > 0.0) \text{ and } (x < \text{pre } x + 10.0)), \quad (2)$$

where c (input) and $\text{pre } x$ (past value) are uncontrollable. During the execution, it may appear that c is false and/or that $\text{pre } x$ is less than -10.0 . In those cases, the constraint is unsatisfiable; we say that the constraint *deadlocks*.

Local variables may be useful auxiliaries for expressing complex constraints. They can be declared within a program:

```
local ident : type in statement. (3)
```

A local variable behaves as a hidden output; it is controllable and must be produced as long as the execution remains in its scope.

2.3. Composing reactions

A constraint (Boolean expression) represents an atomic reaction; it defines relations between the current values of

the variables. Scenarios are built by combining such atomic reactions with *temporal statements*. We introduce the type *trace* for typing expressions made of temporal statements. A single constraint obviously denotes a trace of length 1; in other terms, expressions of type *bool* are implicitly cast to type *trace* when combined with temporal operators.

The basic trace statements are inspired by regular expression, and have following signatures:

- (i) $\text{fby} : \text{trace} \times \text{trace} \rightarrow \text{trace}$ [sequence]
- (ii) $\text{loop} : \text{trace} \rightarrow \text{trace}$ [unbounded loop]
- (iii) $| : \text{trace} \times \text{trace} \rightarrow \text{trace}$ [nondeterministic choice]

Using regular expressions makes the notion of sequence quite different from the one of Esterel, which is certainly the reference in control-flow oriented synchronous language [4]. In Esterel, the sequence (semicolon) is instantaneous, while the Lutin construct *fby* “takes” one instant of time, just like in classical regular expressions.

Example 3. With those operators, we can propose a first version of our example. In this version, the output tends to 0 or according to a first-order filter. The nondeterminism resides in the initial value, and also in the fact that the system is subject to failure and may miss the *c* command.

```
((-100.0 < x) and (x < 100.0)) fby—initial constraint
loop{
  (c and (x = 0.9*(pre x) + 0.1*t))—x gets closer to t
  | ((x = 0.9*(pre x))—x gets closer to 0
}
```

(4)

Initially, the value of *x* is (randomly) chosen between -100 and $+100$, then forever, it may tend to *t* or to 0.

Note that, inside the loop, the first constraint (*x* tends to *t*) is not satisfiable unless *c* is true, while the second is always satisfiable. If *c* is false, the first constraint *deadlocks*. In this case, the second branch (*x* gets closer to 0) is necessarily taken. If *c* is true, both branches are feasible: one is randomly selected and the corresponding constraint is solved.

This illustrates an important principle of the language, the *reactivity*, principle, which states that a program may only deadlock, if all its possible behaviors deadlock.

2.4. Traces, termination, and deadlocks

Because of nondeterminism, a behavior has in general several possible first reactions (constraints). According to the reactivity principle, it deadlocks only if all those constraints are not satisfiable. If at least one reaction is satisfiable, it must “do something;” we say that it is *startable*.

Termination, startability, and deadlocks are important concepts of the language; here is a more precise definition of the basic statements according to these concepts.

- (i) A constraint *c*, if it is satisfiable, generates a particular solution and terminates, otherwise it deadlocks.

- (ii) $st1 \text{ fby } st2$ executes *st1*, and if and when it terminates, it executes *st2*. If *st1* deadlocks, the whole statement deadlocks.
- (iii) Loop *st*, if *st* is startable, behaves as $st \text{ fby loop } st$, otherwise it terminates. Indeed, once started, $st \text{ fby loop } st$ may deadlock if the first *st*, and so on. Intuitively, the meaning is “loop as long as starting a step is possible.”
- (iv) $\{st1\} \cdots \{stn\}$ randomly chooses one of the *startable* statements from $st1, \dots, stn$. If none of them are startable, the whole statement deadlocks.
- (v) The priority choice $\{st1\} > \cdots > \{stn\}$ behaves as *st1* if *st1* is startable, otherwise, behaves as *st2* if *st2* is startable and so on. If none of them are startable, the whole statement deadlocks.
- (vi) Try *st1* do *st2* catches any deadlock occurring *during the execution* of *st1* (not only at the first step). In case of deadlock, the control passes to *st2*.

2.5. Well-founded loops

Let us denote by ε the identity element for *fby* (i.e., the unique behavior such that $b \text{ fby } \varepsilon = \varepsilon \text{ fby } b = b$). Although this “empty” behavior is not provided by the language, it is helpful for illustrating a problem raised by nested loops.

As a matter of fact, the simplest way to define the loop is to state that “loop *c*” is equivalent to “ $c \text{ fby loop } c \mid > \varepsilon$ ”, that is, try in priority to perform one iteration and if it fails, stop. According to this definition, nested loops may generate infinite and instantaneous loops, as shown in the following example.

Example 4.

```
loop{loop c}. (5)
```

Performing an iteration of the outer loop consists in executing the inner $\text{loop } c$. If *c* is not currently satisfiable, $\text{loop } c$ terminates immediately and thus, the iteration is actually “empty”—it generates no reaction. However, since it is not a deadlock, this strange behavior is considered by the outer loop as a normal iteration. As a consequence, another iteration is performed, which is also empty, and so on. The outer loop keeps the control forever but does nothing.

One solution is to reject such programs. Statically checking whether a program will infinitely loop or not is undecidable, it may depend on arbitrarily complex conditions. Some over-approximation is necessary, which will (hopefully) reject all the incorrect programs, but also lots of correct ones. For instance, a program as simple as “ $\text{loop } \{\{\text{loop } a\} \text{ fby } \{\text{loop } b\}\}$ ” will certainly be rejected as potentially incorrect.

We think that such a solution is too much restrictive and tedious for the user and we prefer to slightly modify the semantics of the loop. The solution retained is to introduce the *well-founded loop principle*; a loop statement may stop or continue, but if it continues it must do something. In other terms, empty iterations are dynamically forbidden.

The simplest way to explain this principle is to introduce an auxiliary operator $st \setminus_\epsilon$. If st terminates immediately, $st \setminus_\epsilon$ deadlocks, otherwise it behaves as st . The correct definition of $\text{loop } st$ follows:

- (i) if $st \setminus_\epsilon$ is startable, it behaves as $st \setminus_\epsilon$ fby $\text{loop } st$,
- (ii) otherwise $\text{loop } st$ terminates.

2.6. Influencing non-determinism

When executing a nondeterministic statement, the problem of which choice should be preferred arises. The solution retained is that, if k out of the n choices are startable, each of them is chosen with a probability $1/k$.

In order to influence this choice, the language provides the concept of *relative weights*:

$$\{st1 \text{ weight } w1 \mid \dots \mid stn \text{ weight } wn\}. \quad (6)$$

Weights are basically integer constants and their interpretation is straightforward. A branch with a weight 2 has twice the chance to be tried than a branch with weight 1. More generally, a weight can depend on the environment and on the past; it is given as an integer expression depending on uncontrollable variables. In this case, weight expressions are evaluated at runtime before performing the choice.

Example 5. In a first version (Example 3), our example system may ignore the command c with a probability $1/2$. This case can be made less probable by using weights (when omitted, a weight is implicitly 1):

```
loop{
  (c and (x = 0.9*(pre x) + 0.1*t)) weight 9
  | ((x = 0.9*(pre x))
}
```

(7)

In this new version, a true occurrence of c is missed with the probability $1/10$.

Note that, weights are not only directives. Even with a big weight, a nonstartable branch has a null probability to be chosen, which is the case in the example when c is false.

2.7. Random loops

We want to define some loop structure, where the number of iterations is not fully determined by deadlocks. Such a construct can be based on weighted choices, since a loop is nothing but a binary choice between stopping and continuing. However, it seems more natural to define it in terms of expected number of iterations. Two loop “profiles” are provided as follows.

- (i) $\text{loop}[\text{min}, \text{max}]$: the number of iterations should be between the constants min and max .
- (ii) $\text{loop} \sim av : sd$: the average number of iteration should be av , with a standard deviation sd .

Note that random loops, just like other nondeterministic choices, follow the *reactivity principle*; depending on deadlocks, looping may sometimes be required or impossible. As

a consequence, during an execution, the actual number of iterations may significantly differ from the “expected” one (see Sections 4 and 5.3).

Moreover, just like the basic loop, they follow the *well-founded loop principle*, which means that, even if the core contains nested loops, it is impossible to perform “empty” iterations.

2.8. Parallel composition

The parallel composition of Lutin is synchronous; each branch produces, at the same time, its local constraints. The global reaction must satisfy the conjunction of all those local constraints. This approach is similar to the one of temporal concurrent constraint programming [12].

A parallel composition may deadlock for the following two reasons.

- (i) Obviously, if one or more branches deadlock, the whole statement aborts.
- (ii) It may also appear that each individual statement has one or more possible behaviours, but that none of the conjunctions are satisfiable, in which case the whole statement aborts.

If no deadlock occurs, the concurrent execution terminates, if and when all the branches have terminated (just like in the Esterel Language).

One can perform a parallel composition of several statements as follows:

$$\{st1 \ \& \ > \dots \& \ > \ stn\}. \quad (8)$$

The concrete syntax suggests a noncommutative operator; this choice is explained in the next section.

2.9. Parallel composition versus stochastic directives

It is impossible to define a parallel composition which is fair according to the stochastic directives (weights), as illustrated in the following example.

Example 6. Consider the statement

$$\{\{X \text{ weight } 1000 \mid Y\} \& \> \{\{A \text{ weight } 1000 \mid B\}\}, \quad (9)$$

where $X, A, X \wedge B, A \wedge Y$ are all startable, but not $X \wedge A$.

The higher priority can be given to

- (i) $X \wedge B$, but it would not respect the stochastic directive of the second branch;
- (ii) $A \wedge Y$, but it would not respect the stochastic directive of the first branch;

In order to deal with this issue, the stochastic directives are not treated in parallel, but in *sequence*, from left to right.

- (i) The first branch “plays” first, according to its local stochastic directives.
- (ii) The next ones make their choice according to what has been chosen for the previous ones.

In the example, the priority is then given to $X \wedge B$.

The concrete syntax ($\& \>$) has been chosen to reflect the fact that the operation is not commutative. The treatment is parallel for the constraints (conjunction), but sequential for stochastic directives (weights).

2.10. Exceptions

User-defined exceptions are mainly means for by-passing the normal control flow. They are inspired by exceptions in classical languages (Ocaml, Java, Ada) and also by the trap signals of Esterel.

Exceptions can be globally declared outside a system (exception *ident*) or locally within a statement, in which case the standard binding rules hold

exception *ident* in *st*. (10)

An existing exception *ident* can be raised with the statement:

raise *ident* (11)

and caught with the statement:

catch *ident* in *st1* do *st2*. (12)

If the exception is raised in *st1*, the control immediately passes to *st2*. The do part may be omitted, in which case the control passes in sequence.

2.11. Modularity

An important point is that the notion of system is not a sufficient modular abstraction. In some sense, systems are similar to main programs in classical languages. They are entry point for the execution but are not suitable for defining “pieces” of behaviors.

Data combinators

A good modular abstraction would be one that allows to enrich the set of combinators. Allowing the definition of data combinators is achieved by providing a functional-like level in the language. For instance, one can program the useful “within an interval” constraint;

let within (*x*, *min*, *max* : real) : bool
= (*x* >= *min*) and (*x* <= *max*). (13)

Once defined, this combinator can be instantiated, for instance,

within (*a*, 0.8, 0.9) (14)

or

within (*a* + *b*, *c* - 1.0, *c* + 1.0). (15)

Note that, such a combinator is definitively not a function in the sense of computer science—it actually computes nothing. It is rather a well-typed *macro* defining how to build a Boolean expression with three real expressions.

Reference arguments

Some combinators specifically require support variables as argument (input, output, local). This is the case for the operator *pre*, and as a consequence, for any combinator using a *pre*. This situation is very similar to the distinction between “by reference” and “by value” parameters in imperative languages. Therefore, we solve the problem in a similar manner by adding the flag *ref* to the type of such parameters.

Example 7. The following combinator defines the generic first-order filter constraint. The parameter *y* must be a real support variable (*real ref*) since its previous value is required. The other parameters can be any expressions of type *real*.

Let fof (*y* : real; *gain*, *x* : real) : bool =
(*y* = *gain* * (*pre y*) + (1.0 - *gain*) * *x*). (16)

Trace combinators

User-defined temporal combinators are simply macros of type *trace*.

Example 8. The following combinator is a binary parallel composition, where the termination is enforced when the second argument terminates.

Let, as long as, (*X*, *Y* : trace) : trace =
exception Stop in
catch Stop in {
 X & > {*Y* fby raise Stop}
}.

Local combinators

A macro can be declared within a statement, in which case the usual binding rules hold; in particular, a combinator may have no parameter at all;

Let *id* ([*params*]) : type = *statement* in *statement*. (18)

Example 9. We can now write more elaborated scenarios for the system of Example 3. For the very first reaction (line 2), the output is randomly chosen between -100 and +100, then the system enters its standard behavior (lines 3 to 14). A local variable *a* is declared, which will be used to store the current gain (line 3). An intermediate behavior (lines 4 to 6) is declared, which defines how the gain evolves; it is randomly chosen between 0.8 and 0.9, then it remains constant during 30 to 40 steps, and so on. Note that, this combinator has no parameter since it directly refers to the variable *a*. Lines 7 to 14 define the actual behavior; the user-defined combinator as long as runs in parallel the behavior *gen_gain* (line 8) with the normal behavior (9 to 11). In the normal behavior, the system works almost properly for about 1000 reactions; if *c* is true, *x* tends to *t* 9 times out of 10 (line 10), otherwise it tends to 0 (line 11). As soon as the normal behavior terminates, the whole parallel composition

```

(1) system tracker(c : bool; t : real) returns (x : real) =
(2)   within(x, -100.0, 100.0) fby
(3)   local a : real in
(4)   let gen_gain() : trace = loop{
(5)     within(a, 0.8, 0.9) fby loop[30, 40] (a = pre a)
(6)   } in
(7)   as_long_as(
(8)     gen_gain(),
(9)     loop~1000 : 100{
(10)      (c and fof(x, a, t)) weight 9
(11)      | fof(x, a, 0.0)
(12)    }
(13)   ) fby
(14)   loop fof(x, 0.7, 0.0)

```

FIGURE 1: A full example: the “tracker” program.

FIGURE 2: An execution of the tracker program.

terminates (definition of as long as). Then, the system breaks down and x quickly tends to 0 (line 14).

Figure 2 shows the timing diagram, a particular execution of this program. Input values are provided by the environment (i.e., us) according to the following specification, the input t remains constant (150) and the command c toggles each about 100 steps.

3. SYNTAX SUMMARY

Figure 3 summarizes the concrete syntax of Lutin. The detailed syntax for *expression* is omitted. They are made of classical algebraic expressions with numerical and logical operators, plus the special operator `pre`. The supported type identifiers are currently `bool`, `int`, and `real`.

We do not present the details of the type checking, which is classical and straightforward. The only original check concerns the arguments of the loop profiles and of the weight directive, that must be *uncontrollable* expressions (not depending on output or local variables).

4. OPERATIONAL SEMANTICS

4.1. Abstract syntax

We consider here a type checked Lutin program. For the sake of simplicity, the semantics is given on the flat language. User-defined macros are inlined, and local variables are made

global through some correct renaming of identifiers. As a consequence, an abstract system is simply a collection of variables (inputs, outputs, and locals) and a single abstract statement.

We use the following abstract syntax for statements, where the intuitive meaning of each construct is given between parenthesis:

$$\begin{aligned}
 t &::= c \text{ (constraint)}. | \varepsilon \text{ (empty behavior)}. \\
 &| t \setminus \varepsilon \text{ (empty filter)}. | t \cdot t' \text{ (sequence)}. \\
 &| t^* \text{ (priority loop)}. | t_k^{(w_c, w_s)} \text{ (random loop)}. \\
 &| \overset{x}{\hookrightarrow} \text{ (raise)}. | [t \overset{x}{\hookrightarrow} t'] \text{ (catch)}. \\
 &| \succ_{i=1}^n t_i \text{ (priority)}. | |_{i=1}^n t_i / w_i \text{ (choice)}. \\
 &| \&_{i=1}^n t_i \text{ (parallel)}.
 \end{aligned} \tag{19}$$

This abstract syntax slightly differs from the concrete one on the following points.

- (i) The empty behavior (ε) and the empty behavior filter ($t \setminus \varepsilon$) are internal constructs that will ease the definition of the semantics.
- (ii) Random loops are *normalized* by making explicit their weight functions:
 - (a) the stop function ω_s takes the number of already performed iterations and returns the relative weight of the “stop” choice;
 - (b) the continue function ω_c takes the number of already performed iterations and returns the relative weight of the “continue” choice.

These functions are completely determined by the loop profile in the concrete program (interval or average, together with the corresponding static arguments). See Section 5.3 for a precise definition of these weight functions.

- (iii) The number of already performed iterations (k) is syntactically attached to the loop; this is convenient to define the semantics in terms of rewriting (in the initial program, this number is obviously set to 0).

```

system ::= system ident([params]) returns (params) = statement
params ::= ident : type {; ident : type}
statement ::= expression | {statement} | statement fby statement
           | loop[loop-profile] statement
           | exception ident in statement | raise ident
           | try statement [do statement] | catch ident in statement
           | let ident ([params]) : type = statement in statement
           | local ident : type in statement
           | statement[weight expression] { | statement[weight expression] }
           | statement { | > statement }
           | statement { & > statement }
loop-profile ::= [expression, expression]
              | ^ expression : expression
type ::= ident[ref]

```

FIGURE 3: The concrete EBNF syntax of Lutin.

Definition 1. $\mathcal{T}$ denotes the set of trace expressions (as defined above) and $\mathcal{C}$ denotes the set of constraints.

4.2. The execution environment

The execution takes place within an environment which stores the variable values (inputs and memories). Constraint resolution, weight evaluation, and random selection are also performed by the environment. We keep this environment abstract. As a matter of fact, resolution capabilities and (pseudo)random generation may vary from one implementation to another and they are not part of the reference semantics.

The semantics is given in term of constraints generator. In order to generate constraints, the environment should provide the two following procedures.

Satisfiability

the predicate $e \models c$ is true if and only if the constraint c is satisfiable in the environment e .

Priority sort

Executing choices first requires to evaluate the weights in the environment. This is possible (and straightforward) because weights may dynamically depends on uncontrollable variables (memories, inputs), but not on controllable variables (outputs, locals). Some weights may be evaluated to 0, in which case the corresponding choice is forbidden. Then a random selection is made, according to the actual weights, to determine a total order between the choices.

For instance, consider the following list of pairs (trace/weight), where x and y are uncontrollable variables,

$$(t_1/x + y), (t_2/1), (t_3/y), (t_4/2). \quad (20)$$

In an environment, where $x = 3$ and $y = 0$, weights are evaluated to

$$(t_1/3), (t_2/1), (t_3/0), (t_4/2). \quad (21)$$

The choice t_3 is erased and the remaining choices are randomly sorted according to their weights. The resulting (total) order may be

- (i) t_1, t_2, t_4 with a probability $3/6 \times 1/3 = 1/6$,
- (ii) t_1, t_4, t_2 with a probability $3/6 \times 2/3 = 1/3$,
- (iii) t_4, t_1, t_2 with a probability $2/6 \times 3/4 = 1/4$,
- (iv) so on.

All these treatments are “hidden” within the function $Sort_e$ which takes a list of pairs (choice/weights) and returns a totally ordered list of choices.

4.3. The step function

An execution step is performed by the function $Step(e, t)$ taking an environment e and a trace expression t . It returns an *action* which is either

- (i) a transition $\xrightarrow{c} n$, which means that t produces a *satisfiable* constraint c and rewrite itself in the (next) trace n ,
- (ii) a termination $\xrightarrow{x}$, where x is a termination flag which is either ε (normal termination), δ (deadlock) or some user-defined exception.

Definition 2. $\mathcal{A}$ denotes the set of actions and $\mathcal{X}$ denotes the set of termination flags.

4.4. The recursive step function

The step function $Step(e, t)$ is defined via a recursive function $S_e(t, g, s)$, where the parameters g and s are continuation functions returning actions.

- (i) $g : \mathcal{C} \times \mathcal{T} \mapsto \mathcal{A}$ is the *goto* function defining how a local transition should be treated according to the calling context.
- (ii) $s : \mathcal{X} \mapsto \mathcal{A}$ is the *stop* function defining how a local termination should be treated according to the calling context.

At the top-level, $\mathcal{S}_e$ is called with the trivial continuations,

$$Step(e, t) = \mathcal{S}_e(t, g, s) \text{ with } g(c, v) = \xrightarrow{c} v, \quad s(x) = \xrightarrow{x} \quad (22)$$

Basic traces

The empty behavior raises the termination flag in the current context. A raise statement terminates with the corresponding flag. At last, a constraint generates a goto or raises a deadlock depending on its satisfiability;

$$\begin{aligned} \mathcal{S}_e(\varepsilon, g, s) &= s(\varepsilon) \\ \mathcal{S}_e(\overset{x}{\hookrightarrow}, g, s) &= s(x) \\ \mathcal{S}_e(c, g, s) &= \text{if } (e \models c) \text{ then } g(c, \varepsilon) \text{ else } s(\delta). \end{aligned} \quad (23)$$

Sequence

The rule is straightforward;

$$\begin{aligned} \mathcal{S}_e(t \cdot t', g, s) &= \mathcal{S}_e(t, g', s'), \\ \text{where } g'(c, n) &= g(c, n \cdot t') \\ s'(x) &= \text{if } x = \varepsilon \text{ then } \mathcal{S}_e(t', g, s) \text{ else } s(x). \end{aligned} \quad (24)$$

Priority choice

We only give the definition of the binary choice, since the operator is right-associative. This rule formalizes the reactivity principle. All possibilities in t must have failed before t' is taken into account,

$$\begin{aligned} \mathcal{S}_e(t > t', g, s) &= \text{if } (r \neq \overset{\delta}{\hookrightarrow}) \text{ then } r \text{ else } \mathcal{S}_e(t', g, s), \\ \text{where } r &= \mathcal{S}_e(t, g, s). \end{aligned} \quad (25)$$

Empty filter and priority loop

The empty filter intercepts the termination of t and replaces by a deadlock,

$$\begin{aligned} \mathcal{S}_e(t \setminus \varepsilon, g, s) &= \mathcal{S}_e(t, g, s'), s \quad \text{where } s'(x) = \text{if } (x = \varepsilon) \\ &\quad \text{then } s(\delta) \text{ else } s(x). \end{aligned} \quad (26)$$

The semantics of the loop results from the equivalence

$$t^* \iff (t \setminus \varepsilon) \cdot t^* > \varepsilon. \quad (27)$$

Catch

This internal operator ($[n \overset{z}{\hookrightarrow} t]$) covers the cases of `try` ($z = \delta$) and `catch` (z is a user-defined exception)

$$\begin{aligned} \mathcal{S}_e([t \overset{z}{\hookrightarrow} t'], g, s) &= \mathcal{S}_e(t, g', s'), \\ \text{where } g'(c, n) &= g(c, [n \overset{z}{\hookrightarrow} t']) \\ s'(x) &= \text{if } (x = z) \text{ then } \mathcal{S}_e(t', g, s) \text{ else } s(x). \end{aligned} \quad (28)$$

Parallel composition

We only give the definition of the binary case, since the operator is right-associative;

$$\begin{aligned} \mathcal{S}_e(t \& t', g, s) &= \mathcal{S}_e(t, g', s'), \\ \text{where } s'(x) &= \text{if } (x = \varepsilon) \text{ then } \mathcal{S}_e(t', g, s) \text{ else } s(x), \\ g'(c, n) &= \mathcal{S}_e(t', g'', s''), \\ \text{where } s''(x) &= \text{if } (x = \varepsilon) \text{ then } g(c, n) \text{ else } s(x) \\ g''(c', n') &= g(c \wedge c', n \& n'). \end{aligned} \quad (29)$$

Weighted choice

The evaluation of the weights and the (random) total ordering of the branches are both performed by the function Sort_e (cf., Section 4.2).

$$\begin{aligned} \text{If } \text{Sort}_e(\{t_i/w_i\}_{i=1\dots n}) &= \emptyset : \mathcal{S}_e(\big|_{i=1}^n t_i/w_i, g, s) = s(\delta) \\ \text{otherwise, } \mathcal{S}_e(\big|_{i=1}^n t_i/w_i, g, s) &= \mathcal{S}_e(> \text{Sort}_e(t_1/w_1, \dots, \\ &\quad t_n/w_n), g, s). \end{aligned} \quad (30)$$

Random loop

We recall that this construct is labelled by two weight functions (ω_c for continue, ω_s for stop) and by the current number of already performed iterations (i). The weight functions are evaluated for i and the statement is then equivalent to a binary weighted choice,

$$t_i^{(\omega_c, \omega_s)} \iff (t \setminus \varepsilon) \cdot t_{i+1}^{(\omega_c, \omega_s)} / \omega_c(i) \mid \varepsilon / \omega_s(i), \quad (31)$$

Note that, the semantics follows the well-founded loop principle.

4.5. A complete execution

Solving a constraint

The main role of the environment is to store the values of uncontrollable variables; it is a pair of stores “past values, input values.” For such an environment $e = (\rho, \iota)$ and a satisfiable constraint c , we suppose given a procedure able to produce a particular solution of $c : \text{Solve}_{\rho, \iota}(c) = \gamma$ (where γ is a store of controllable variables). We keep this Solve function abstract, since it may vary from one implementation to another (see Section 5).

Execution algorithm

A complete run is defined according to

- (i) a given sequence of input stores $t_0, t_1, \dots, t_n$,
- (ii) an initial (main) trace t_0 ,
- (iii) an initial previous store (values of pre variables) ρ_0 .

A run produces a sequence of (controllable variables) stores $\gamma_1, \gamma_2, \dots, \gamma_k$, where $k \leq n$. For defining this output sequence, we use intermediate sequences of traces $(t_1, \dots, t_{k+1})$, previous stores $(\rho_1, \dots, \rho_k)$, environments $(e_0, \dots, e_k)$, and constraints $(c_0, \dots, c_k)$. The relation between those sequences are listed below, for all step $j = 0 \dots k$:

- (i) the current environment is made of previous and input values, $e_j = (\rho_j, \iota_j)$,
- (ii) the current trace makes a transition, $e_j : t_j \xrightarrow{c_j} t_{j+1}$,
- (iii) a solution of the constraint is elected, $\gamma_j = \text{Solve}_{e_j}(c_j)$,
- (iv) the previous store for the next step is the union of current inputs/outputs: $\rho_{j+1} = (\iota_j \oplus \gamma_j)$.

At the end, we have

- (i) either $k = n$, which means that the execution has run to completion,
- (ii) or $(\rho_{k+1}, \iota_{k+1}) : t_{k+1} \xrightarrow{x}$ which means that it has been aborted.

5. IMPLEMENTATION

A prototype has been developed in Ocaml. The constraint generator strictly implements the operational semantics presented in the previous section. The tool can do the following.

- (i) Interpret/simulate Lutin programs in a file-to-file (or pipe-to-pipe) manner. This tool serves for simulation/prototyping; several Lutin simulation sessions can be combined with other reactive process in order to animate a complex system.
- (ii) Compile Lutin programs into the internal format of the testing tool Lurette. This format, called Lucky, is based on flat, explicit automata [14]. In this case, Lutin serves as a high-level language for designing test scenarios.

5.1. Notes on constraint solvers

The core semantics only defines how constraints are generated, but not how they are solved. This choice is motivated by the fact that there is no “ideal” solver.

A required characteristic of such a solver is that it must provide a constructive, complete decision procedure; methods that can fail and/or that are not able to exhibit a particular solution are clearly not suitable. Basically, a constraint solver should provide the following.

- (i) A syntactic analyzer for checking if the constraints are supported by the solver (e.g., linear arithmetics); this is necessary because the language syntax allows to write arbitrary constraints.
- (ii) A decision procedure for the class of constraints accepted by the checker.
- (iii) A precise definition of the election procedure which selects a particular solution (e.g., in terms of fairness).

Even with those restrictions, there is no obvious best solver as follows.

FIGURE 4: A BDD containing 10 solutions (ade , $\bar{a}bc\bar{e}$, and $\bar{a}\bar{b}d$).

- (i) It may be efficient, but limited in terms of capabilities.
- (ii) It may be powerful, but likely to be very costly in terms of time and memory.

The idea is that the user may choose between several solvers (or several options of a same solver) the one which best fits his needs.

The solver that is currently used is presented in the next section.

5.2. The Boolean/linear constraint solver

Actually, we use the solver [15] that have been developed for the testing tool Lurette [16, 17]. This solver is quite powerful, since it covers Boolean algebra and linear arithmetics. Concretely, constraints are solved by generating a normalized representation mixing binary decision diagrams and convex polyhedra. This constraint solver is sketched below and fully described in [15]

First of all, each atomic numeric constraint (e.g., $x + y > 1$) is replaced by a fresh Boolean variable. Then, the resulting constraint is translated into a BDD. Figure 4 shows a graphical representation of a BDD; *then* (resp., *else*) branches are represented at the left-hand-side (resp., right-hand-side) of the tree. This BDD contains 3 paths to the true leaf: ade , $\bar{a}bc\bar{e}$, and $\bar{a}\bar{b}d$. When we say that the monomial (conjunction of literals) $\bar{a}bc\bar{e}$ is a solution of the formula. It means that variables a and e should be false; variables b and c should be true; and variable d can be either true or false. The monomial $\bar{a}bc\bar{e}$, therefore, represents two solutions, whereas ade and $\bar{a}\bar{b}d$ represents 4 solutions each, since 2 variables are left unconstrained.

In Figure 4 and in the following, for the sake of simplicity, we draw trees instead of DAGs. The key reason why BDDs work well in practice is that in their implementations, common subtrees are shared. For example, only one node “true” would be necessary in that graph. Anyway, the algorithms work on DAGs the same way as they work on trees.

Random choice of Boolean values

The first step consists in selecting a Boolean solution. Once the constraint has been translated into a BDD, we have a (hopefully compact) representation of the set of solutions.

We first need to randomly choose a path into the BDD that leads to a true leaf. But if we naively performed a fair toss at each branch of the BDD during this traversal, we would be very unfair. Indeed, consider the BDD of Figure 4; the monomial ade has 50% of chances to be tried, whereas $\bar{a}bc\bar{e}$ and $\bar{a}\bar{b}d$ have 25% each. One can easily imagine situation, where the situation is even worse. This is the reason why counting the solutions before drawing them is necessary.

Once each branch of the BDD is decorated with its solution number performing a fair choice among Boolean solutions is straightforward.

Random choice of numeric values

From the BDD point of view, numeric constraints are just Boolean variables. Therefore, we have to know if the obtained set of atomic numeric constraints is satisfiable. For that purpose, we use a convex polyhedron library [18].

However, a solution from the logical variables point of view may lead to an empty set of solutions for numeric variables. In order to chose a Boolean monomial that is valid with respect to numerics, a (inefficient) method would be to select at random a path in the BDD until that selection corresponds to a satisfiable problem for the numeric constraints. The actual algorithm is more sophisticated [15], but the resulting solution is the same.

When there are solutions to the set of numeric constraints, the convex polyhedron library returns a set of generators (the vertices of the polyhedron representing the set of solutions). Using those generators, it is quite easy to choose point inside (or more interestingly, at edges or at vertices) the polyhedron.

Using polyhedra is very powerful, but also very costly. However the solver benefits from several years of experimentation and optimizations (partitioning, switch from polyhedra to intervals, whenever it is possible).

5.3. Notes on predefined loop profiles

In the operational semantics, loops with iteration profile are translated into binary weighted choices. Those weights are dynamic; they depend on the number of (already) performed iterations k .

Interval loops

For the “interval” profile, those weights functions are formally defined and thus, they could take place in the reference semantics of the language. For a given pair of integers (min, max) such that $0 \leq \min \leq \max$ and a number k of already performed iterations, we have the following:

- (i) if $k < \min$, then $\omega_s(k) = 0$ and $\omega_c(k) = 1$ (loop is mandatory);
- (ii) if $k \geq \max$, then $\omega_s(k) = 1$ and $\omega_c(k) = 0$ (stop is mandatory);
- (iii) if $\min \leq k < \max$, then $\omega_s(k) = 1$ and $\omega_c(k) = 1 + \max - k$.

Average loops

There is no obvious solution for implementing the “average” profile in terms of weightss. A more or less sophisticated (and accurate) solution should be retained, depending on the expected precision.

In the actual implementation, for an average value av and a standard variation sv , we use a relatively simple approximation as follows.

- (i) First of all, the underlying discrete repartition law is approximated by a continuous (Gaussian) law. As a consequence, the result will not be accurate if av is too close to 0 and/or if sv is too big comparing to av . Concretely, we must have $10 < 4 * sv < av$.
- (ii) It is well known that no algebraic definition for the Gaussian repartition function exists. This function is then classically approximated by using an interpolation table (512 samples with a fixed precision of 4 digits).

6. CONCLUSION

We propose a language for describing constrained-random reactive systems. Its first purpose is to describe test scenarios, but it may also be useful for prototyping and simulation.

We have developed a compiler/interpreter which strictly implements the operational semantics presented here. Thanks to this tool, the language is integrated into the framework of the Lurette tool, where it is used to describe test scenarios. Further works concerns the integration of the language within a more general prototyping framework.

Other works concern the evolution of the language. We plan to introduce a notion of signal (i.e., event) which is useful for describing values that are not always available (this is related to the notion of clocks in synchronous languages). We also plan to allow the definition of (mutually) tail-recursive traces. Concretely, that means that a new programming style would be allowed, based on explicit concurrent, hierarchic automata.

REFERENCES

- [1] P. Raymond and Y. Roux, “Describing non-deterministic reactive systems by means of regular expressions,” in *Proceedings of the 1st Workshop on Synchronous Languages, Applications and Programming (SLAP '02)*, Grenoble, France, April 2002.
- [2] A. Benveniste and G. Berry, “The synchronous approach to reactive and real-time systems,” *Proceedings of the IEEE*, vol. 79, no. 9, pp. 1270–1282, 1991.
- [3] N. Halbwachs, *Synchronous Programming of Reactive Systems*, Kluwer Academic Publishers, Dordrecht, The Netherlands, 1993.
- [4] G. Berry and G. Gonthier, “The Esterel synchronous programming language: design, semantics, implementation,” *Science of Computer Programming*, vol. 19, no. 2, pp. 87–152, 1992.
- [5] N. Halbwachs, P. Caspi, P. Raymond, and D. Pilaud, “The synchronous data flow programming language LUSTRE,” *Proceedings of the IEEE*, vol. 79, no. 9, pp. 1305–1320, 1991.
- [6] T. Gauthier, P. Le Guernic, and L. Besnard, “Signal: a declarative language for synchronous programming of real-time systems,” in *Proceedings of the 3rd Conference on Functional*

- Programming Languages and Computer Architecture*, vol. 274 of *Lecture Notes in Computer Science*, pp. 257–277, Springer, Portland, Ore, USA, September 1987.
- [7] N. A. Lynch and M. R. Tuttle, “An introduction to Input/Output automata,” *CWI-Quarterly*, vol. 2, no. 3, pp. 219–246, 1989.
 - [8] S.-H. Wu, S. A. Smolka, and E. W. Stark, “Composition and behaviors of probabilistic I/O automata,” *Theoretical Computer Science*, vol. 176, no. 1-2, pp. 1–38, 1997.
 - [9] M. Bernardo, L. Donatiello, and P. Ciancarini, “Stochastic process algebra: from an algebraic formalism to an architectural description language,” in *Performance Evaluation of Complex Systems: Techniques and Tools*, vol. 2459 of *Lecture Notes in Computer Science*, pp. 173–182, Springer, Berlin, Germany, 2002.
 - [10] B. Jonsson, K. G. Larsen, and W. Yi, “Probabilistic extensions of process algebras,” in *Handbook of Process Algebra*, North-Holland, Amsterdam, The Netherlands, 2001.
 - [11] V. A. Saraswat, Ed., *Concurrent Constraint Programming*, MIT Press, Cambridge, Mass, USA, 1993.
 - [12] M. Nielsen, C. Palamidessi, and F. D. Valencia, “Temporal concurrent constraint programming: denotation, logic and applications,” *Nordic Journal of Computing*, vol. 9, no. 1, pp. 145–188, 2002.
 - [13] V. A. Saraswat, R. Jagadeesan, and V. Gupta, “Foundations of timed concurrent constraint programming,” in *Proceedings of the 9th Annual IEEE Symposium on Logic in Computer Science (LICS '94)*, pp. 71–80, Paris, France, July 1994.
 - [14] P. Raymond, E. Jahier, and Y. Roux, “Describing and executing random reactive systems,” in *Proceedings of the 4th IEEE International Conference on Software Engineering and Formal Methods (SEFM '06)*, pp. 216–225, Pune, India, September 2006.
 - [15] E. Jahier and P. Raymond, “Generating random values using Binary decision diagrams and convex polyhedra,” in *Proceedings of the Workshop on Constraints in Software Testing, Verification and Analysis (CSTVA '06)*, Nantes, France, September 2006.
 - [16] E. Jahier, P. Raymond, and P. Baufreton, “Case studies with Lurette V2,” *International Journal on Software Tools for Technology Transfer*, vol. 8, no. 6, pp. 517–530, 2006.
 - [17] P. Raymond, D. Weber, X. Nicollin, and N. Halbwachs, “Automatic testing of reactive systems,” in *Proceedings of the 19th IEEE Real-Time Systems Symposium*, pp. 200–209, Madrid, Spain, December 1998.
 - [18] B. Jeannet, “The Polka Convex Polyhedra library Edition 2.0,” May 2002, www.irisa.fr/prive/bjeannet/newpolka.html.

Special Issue on Selected Papers from SPL 2009 Programmable Logic and Applications

Call for Papers

Field programmable logic devices, field-programmable gate arrays (FPGAs), are nowadays one of the most important alternatives to construct high-speed digital systems. This technology was marketed in the mid eighties with a simple but strong argument: its capability to be in-house erased and reconfigured in a few milliseconds. Designers soon discovered that reconfigurability facilitates the mapping of particular algorithms in hardware. This idea led to new lines of research such as FPGA-based custom computer machines (FCCMs) or custom digital signal processors (CDSPs).

The purpose of this special issue is to provide an opportunity for researchers, engineers, and designers to report recent advances in this important area of reconfigurable logic and applications using FPGAs.

Topics of interest for this special issue cover, among others, the following:

- Design methodology (low-power design, high-speed techniques, physical design, synchronization and soft-timed systems, dynamic and run-time reconfiguration, reconfigurable embedded systems, field programmable analogue arrays, interconnects, and NoCs)
- FPGA in education (roadmap of programmable logic, teaching reconfigurable systems, history and surveys of programmable logic, emerging device technologies, tutorials)
- Platform-based design (embedded processors, custom computers, IP cores, Java, Handel-C, and system-C)
- Applications (robotics, artificial vision, communications/networking, cryptography, bioinformatics, application acceleration, evolvable and bioinspired applications, and rapid prototyping)
- Custom computers and DSPs (computer arithmetics, digital signal processing, software radio, and FCCMs)
- EDA tools (logic and architectural, synthesis, modelling and simulation, emulation, formal methods in system design, CAD for reconfigurable architectures, System-level design methods, testing, verification and benchmarking, and hardware/software codesign)

- Aerospace applications (design verification and validation, reliability and fault tolerance, FIT rates analysis, qualification process, device obsolescence, high-reliability processor cores, noise, radiation effects and EMC, and experiences and lessons learned)

Submission to this special issue is limited to the participants of the SPL 09 conference who have been invited to submit to this issue.

The criteria for invitations include: (1) the evaluation of SPL09 program committee (only papers very good rated in SPL evaluation system), (2) the relevance of the topics, and (3) a positive impression and comments in the presentation during the conference. The invitation to submit an extended version to the special issue does not guarantee that the paper will be published. In particular, submissions will be reviewed by at least by 3 external referees to ensure that they meet the standard required to be published in the journal.

Before submission authors should carefully read over the journal's Author Guidelines, which are located at <http://www.hindawi.com/journals/ijrc/guidelines.html>. Prospective authors should submit an electronic copy of their complete manuscript through the journal Manuscript Tracking System at <http://mts.hindawi.com/> according to the following timetable:

Manuscript Due	July 1, 2009
First Round of Reviews	October 1, 2009
Publication Date	February 1, 2010

Lead Guest Editor

Elías Todorovich, School of Engineering, Universidad Autonoma de Madrid, Spain; etodorov@uam.es

Guest Editor

Valentin Obac Roda, Universidade de São Paulo/EESC, Brazil; valentin@sel.eesc.usp.br

Special Issue on Network Coding for Wireless Networks

Call for Papers

The main idea in network coding was introduced in 2000 by Ahlswede et al. With network coding, an intermediate node in a network cannot only forward its incoming packets but also encode them. It has been shown that the use of network coding can enhance the performance of wired networks significantly. Recent work indicates that network coding can also offer significant benefits for wireless networks.

Communications over wireless channels are error-prone and unpredictable due to fading, mobility, and intermittent connectivity. Moreover, in wireless networks, transmissions are broadcast and can be overheard by neighbors, which is treated in current systems as interference. Finally, security poses new challenges in wireless networks, where both passive and active attacks have quite different premises than in wired networks. Ideas in network coding promise to help toward all these issues, allowing to gracefully add redundancy to combat errors, take advantage of the broadcast nature of the wireless medium and achieve opportunistic diversity, exploit interference rather than be limited by it, and provide secure network communication against adversarial attacks.

In this special issue, we are interested in original research articles which can carry the momentum further and take the wireless network coding research to the next level. The areas of interest include novel network code designs and algorithms, new applications of wireless network coding, network coding capacity, and performance analysis. In addition to original research articles, we are open to review articles. The following list indicates topics of interest which is by no means exhaustive:

- Network codes and algorithms for wireless networks
- Physical layer network coding
- Joint source coding and network coding
- Graph codes and network coding
- Reduced complexity decoding for network coding
- Secure network coding
- Capacity and fundamental bounds on network coding performance
- Cross-layer optimization and network coding
- Energy-efficient network coding
- TCP, routing, MAC, or scheduling algorithms for network codes

- Wireless network coding for multimedia application
- Wireless network coding for bio-medical application

Before submission authors should carefully read over the journal's Author Guidelines, which are located at <http://www.hindawi.com/journals/wcn/guidelines.html>. Prospective authors should submit an electronic copy of their complete manuscript through the journal Manuscript Tracking System at <http://mts.hindawi.com/> according to the following timetable:

Manuscript Due	October 1, 2009
First Round of Reviews	January 1, 2010
Publication Date	April 1, 2010

Lead Guest Editor

Heung-No Lee, Department of Information and Communications, Gwangju Institute of Science and Technology (GIST), Gwangju, South Korea; heungno@gist.ac.kr

Guest Editors

Sae-Young Chung, School of EECS, KAIST, Daejeon, South Korea; sychung@ee.kaist.ac.kr

Christina Fragouli, School of Computer and Communication Sciences, EPFL, Lausanne, Switzerland; christina.fragouli@epfl.ch

Zhi-Hong Mao, ECE/Bio-Medical Dept., the University of Pittsburgh, Pittsburgh, PA, USA; zhm4@pitt.edu

Special Issue on Simulators and Experimental Testbeds Design and Development for Wireless Networks

Call for Papers

In the context of wireless networking, performance evaluation of protocols and distributed applications is generally conducted through simulation or experimentation campaigns. An efficient and accurate simulation of wireless networks raises various issues which generally need to be addressed from several research domains simultaneously. As examples, we can consider the wireless physical layer modeling and simulation, the support of large-scale networks, the simulation of complex RF systems such as MIMO ones, the emulation of wireless nodes or the interconnection of simulators, experimental testbeds, and so forth.

The aim of this Special Issue is to bring together academic and industry researchers and practitioners from both the wireless networking and the simulation communities to discuss current and future trends in simulation or experimentation techniques, models, and practices for the future communication system and to foster interdisciplinary collaborative research in this area. The guest editors seek high-quality papers on aspects of wireless network simulation, and value both theoretical and practical research contributions. Topics of interest include, but are not limited to:

- Radio medium modeling and cross-layer simulation
- Scalability, large-scale networks support
- Validation of simulators and simulation results
- Simulators benchmarking and comparisons
- Fluid-flow simulation for assessing QoS in large-scale networks
- Support of new emerging technologies (WiMax, 3.5G, Wireless Mesh Networks, 802.11x, etc.) in simulators
- Support of advanced RF systems (Multi-carrier schemes, MIMO, smart-antenna) in simulators
- Wireless node simulation or emulation
- Interoperability of simulators, emulators, and experiments
- Support of distributed physical layer schemes (distributed signal processing; cooperative schemes)
- Distributed simulation, and scalability of simulators
- Implementation of simulators
- Experimental testbeds for wireless networks

- Methodology for protocol and distributed application performance evaluation
- SDR techniques, cognitive radio approaches, dynamic spectrum access testbeds, and simulators as well as modeling
- Simulation and testbeds for cooperative communication protocols

Before submission, authors should carefully read over the journal's Author Guidelines, which are located at <http://www.hindawi.com/journals/wcn/guidelines.html>. Prospective authors should submit an electronic copy of their complete manuscript through the journal Manuscript Tracking System at <http://mts.hindawi.com/>, according to the following timetable:

Manuscript Due	June 1, 2009
First Round of Reviews	September 1, 2009
Publication Date	December 1, 2009

Lead Guest Editor

Faouzi Bader, Centre Tecnológico de Telecomunicaciones de Catalunya (CTTC), PMT, Spain; faouzi.bader@cttc.es

Guest Editors

Guillaume Chelius, Institut National de Recherche en Informatique et en Automatique (INRIA), France; guillaume.chelius@inria.fr

Christian Ibarras, Centre Tecnológico de Telecomunicaciones de Catalunya (CTTC), PMT, Spain; cibarras@cttc.cat

Mohamed Ibnkahla, Department of Electrical and Computer Engineering, Queen's University, Canada; ibnkahla@post.queensu.ca

Nikos Passas, Department of Informatics and Telecommunications, University of Athens, Greece; passas@di.uoa.gr

Arnd-Ragnar Rhiemeier, Defence Electronics, EADS Deutschland GmbH, Germany; arnd-ragnar.rhiemeier@eads.com