

UPPAAL in a nutshell

Kim G. Larsen¹, Paul Pettersson², Wang Yi²

¹Department of Computer Science and Mathematics, Aalborg University, Denmark; E-mail: kgl@cs.auc.dk

²Department of Computer Systems, Uppsala University, Sweden; E-mail: {paupet,yi}@docs.uu.se

Abstract. This paper presents the overall structure, the design criteria, and the main features of the tool box UPPAAL. It gives a detailed user guide which describes how to use the various tools of UPPAAL version 2.02 to *construct* abstract models of a real-time system, to *simulate* its dynamical behavior, to *specify* and *verify* its safety and bounded liveness properties in terms of its model. In addition, the paper also provides a short review on case-studies where UPPAAL is applied, as well as references to its theoretical foundation.

Key words: Modeling real-time systems – Dynamic modeling – Modeling tools – UPPAAL

1 Introduction

UPPAAL is a tool box for modeling, simulation and verification of real-time systems, based on constraint-solving and on-the-fly techniques, developed jointly by Uppsala University and Aalborg University. It is appropriate for systems that can be modeled as a collection of non-deterministic processes with finite control structure and real-valued clocks, communicating through channels and (or) shared variables [26, 34]. Typical application areas include real-time controllers and communication protocols in particular, those where timing aspects are critical. It is designed mainly to check invariant and reachability properties by exploring the state-space of a system, i.e., reachability analysis in terms of symbolic states represented by constraints.

The two main design criteria for UPPAAL have been *efficiency* and *ease of usage*. The restriction to reachability analysis has been crucial to the efficiency of the UPPAAL model-checker. Another important key to efficiency

is the application of *on-the-fly* searching techniques combined with the *symbolic* technique that reduces verification problems to that of manipulating and solving simple constraints [26, 34]. To facilitate modeling and debugging, the UPPAAL model-checker may automatically generate a *diagnostic trace* that explains why a property is (or is not) satisfied by a system description. The diagnostic traces generated by the model-checker may be graphically visualized using the simulator.

Since its first release in 1995, UPPAAL has been applied in a number of case studies (see Sect. 6 for a summary). To meet requirements arising from the case studies, It has been extended with various features. The current version of UPPAAL is implemented in C++, XFORMS, and MOTIF. This paper is devoted to an informal presentation of UPPAAL. We present the semantics model implemented in UPPAAL, its various features, and review and provide references on the theoretical foundation and applications to case-studies. We also provide a detailed user guide. The paper is organized as follows:

Section 2 describes the overall structure and the main components of UPPAAL. Section 3 is an *informal* presentation of the syntax and semantics of the UPPAAL model. Section 4 presents the logic and the kernel of the model-checking algorithm of the UPPAAL model-checker. Section 5 serves as a user guide, describing in details how to use the various tools of UPPAAL. Section 7 concludes the paper with a brief description on recent and possible future development of UPPAAL.

2 Overview of UPPAAL

An overview of UPPAAL is shown in Fig. 1. In this section we briefly describe the main features of UPPAAL.

UPPAAL consists of three main parts: a description language, a simulator, and a model-checker. The descrip-

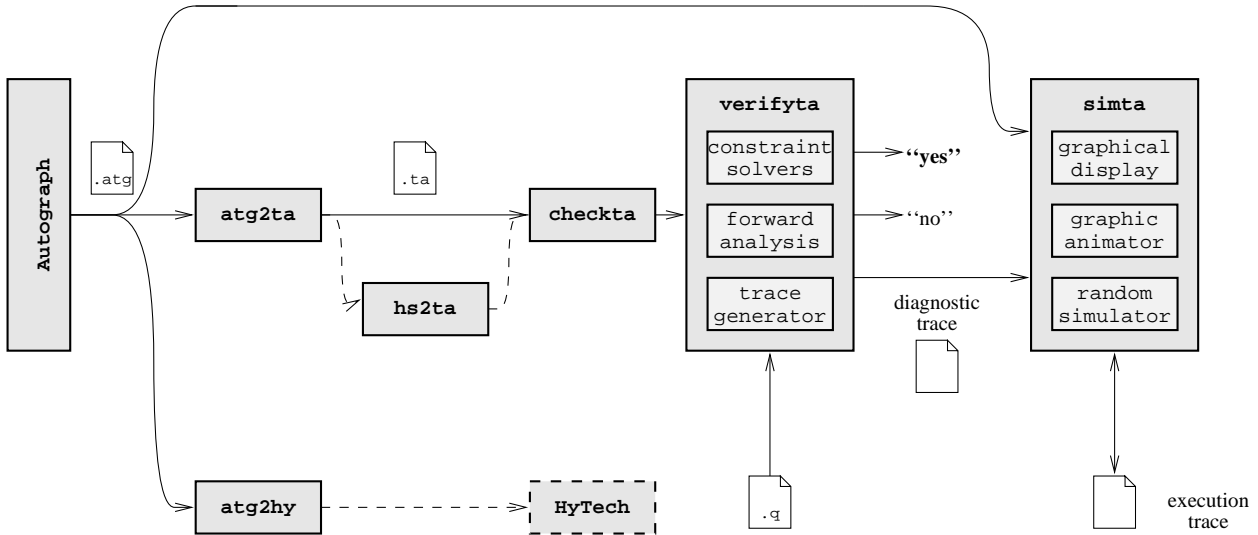


Fig. 1. Overview of UPPAAL

tion language is a non-deterministic guarded command language with data types¹. It serves as a modeling or design language to describe system behavior as networks of timed automata extended with data variables. The simulator and the model-checker are designed for interactive and automated analysis of system behavior by manipulating and solving constraints that represent the state-space of a system description. They have a common basis, i.e., constraint-solvers. The simulator enables examination of *possible* dynamic executions of a system during early modeling (or design) stages and thus provides an inexpensive mean of fault detection prior to verification by the model-checker which covers the *exhaustive* dynamic behavior of the system.

2.1 Modeling

To facilitate modeling, UPPAAL provides both graphical and textual formats for the description language. One can use either the textual format or the AUTOGRAPH-based graphical user interface [10] to define system descriptions, namely networks of timed automata. As an example, the textual representation of the graphical system description in Fig. 3 is shown in Fig. 2.

The textual format (i.e., `.ta`) provides a basic *programming language for timed automata*. In certain cases, the textual format can be more convenient (and faster) to work with than the graphical interface. The compiler `atg2ta` automatically transforms system description in the graphical `.atg`-format into the textual `.ta`-format, thus supporting the important principle WYSIWYV².

The UPPAAL description language also supports modeling of simple linear hybrid automata, that is, timed

```
//
// Global declaration section
//
clock x, y;
int n;
chan a;

//
// Component description section
//
process A {
  state A0 { y<=6 }, A1, A2, A3;
  init A0;
  trans A0 -> A1 {
    guard y>=3;
    sync a!;
    assign y:=0;
  },
  A1 -> A2 {
    guard y>=4;
  },
  A2 -> A3 {
    guard n==5;
  };
}

process B {
  state B0 { x<=4 }, B1, B2, B3;
  commit B1;
  init B0;
  trans B0 -> B1 {
    guard x>=2;
    sync a?;
    assign n:=5,x:=0;
  },
  B1 -> B2 {
    assign n:=n+1;
  },
  B2 -> B3 {
  };
}

//
// System description section
//
system A, B;
```

Fig. 2. Textual description

¹ Currently, only `integer` and `clock` with restricted forms of operations are implemented.

² What You See Is What You Verify.

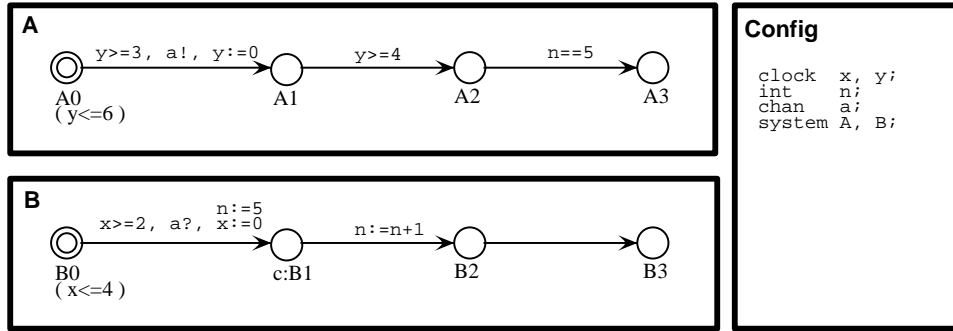


Fig. 3. An example UPPAAL model

automata with clocks whose rates may vary in a certain interval [31]. This extension of timed automata is useful for modeling of hybrid systems where the behavior of the system variables can be described or approximated using lower and upper bounds on their rates. Using abstraction techniques, this class of linear hybrid system can be transformed into timed automata and thus be verified using the techniques available for timed automata, implemented in UPPAAL. UPPAAL allows linear hybrid automata where the rates of clocks are given by an interval. Philips Audio-Control Protocol of [9] is an example of such a linear hybrid systems.

2.2 Analysis

Model-checking

The model-checker is designed to check for invariant and reachability properties, in particular whether certain combinations of control-nodes and constraints on clocks and integer variables are reachable from an initial configuration. Other properties such as bounded liveness properties can be checked by reasoning about the system in the context of testing automata or simply decorating the system description with debugging information and then checking reachability properties. Model-checking is performed by the module `verifyta` which takes as input a network of automata in the textual-format (i.e., `.ta`) and a formula. In checking a property, a *diagnostic trace* can be automatically reported by `verifyta` [25], that explains why the property is satisfied or not. Such a trace may be considered as diagnostic information of a system error, useful during the subsequent debugging of the system.

Simulation

The simulator allows the user to examine in an interactive and graphical fashion the dynamic behavior of a system. In contrast to the model-checker which explores the whole reachable state-space of a system - examining all the behavior of the system, the simulator explores only

a particular execution trace i.e., a sequence of states of the system. This will in early stages of modeling (or design) provide an inexpensive mean of fault detection. In comparison the model-checker is obviously more expensive as it amounts to an exhaustive simulation covering all behavior of the system. Another useful application of the simulator is to visualize a diagnostic trace generated by the model-checker; thus the user can in an interactive and graphical fashion examine the execution trace that may result in a system error.

3 The bolts of UPPAAL – modeling

In this section, we present the basic ingredients of the UPPAAL model based on small examples. For a precise semantical treatment we refer the reader to [4].

We assume that a typical real-time system is a network of non-deterministic sequential processes communicating with each other over channels. In UPPAAL we use finite-state automata extended with clock and data variables to describe processes and networks of such automata to describe real-time systems.

3.1 Syntax

The basis of the UPPAAL model is the notion of timed automata [2] developed by Alur and Dill as an extension of classical finite-state automata with clock variables. To provide a more expressive model and to ease the modeling task, we further extend timed automata with more general types of data variables such as boolean and integer variables. Our final goal is to develop a modeling (or design) language which is as close as possible to a high-level real-time programming language with various data types. Clearly, this will create problems for decidability of model-checking. However, we can always require that the value domains of the data variables should be finite in order to guarantee the termination of a verification procedure (as in $\text{Mur}\varphi$ [16]).

In the current implementation of UPPAAL a system description (or model) consists of a collection of timed automata extended with integer variables in addition to

clock variables. Consider the UPPAAL model in Fig. 3. The model consists of two components A and B with control nodes $\{A_0, A_1, A_2, A_3\}$ and $\{B_0, B_1, B_2, B_3\}$ respectively. In addition to these discrete control structures, the model uses two clocks x and y , one integer variable n and one channel a .

The edges of the automata are decorated with three types of labels: a *guard*, expressing a condition on the values of clocks and integer variables that must be satisfied in order for the edge to be taken; a synchronization *action* which is performed when the edge is taken and finally a number of *clock resets* and *assignments* to integer variables. All three types of labels are optional.

In addition, control nodes may be decorated with so-called *invariants*, which are conditions expressing constraints on the clock values in order for control to remain in a particular node.

3.1.1 Guards

Guards express conditions on the values of clocks and integer variables that must be satisfied in order for the edge to be taken. In Fig. 3 the edge between A_0 and A_1 can only be taken, when the value of the clock y is greater than or equal to 3. Similarly the edge between A_2 and A_3 can only be taken when the value of the integer variable n equals 5. Formally, guards are conjunctions of timing and data constraints; a timing constraint is of the form: $x \sim n$ or $x - y \sim n$, where n is a natural number and $\sim \in \{\leq, \geq, =, >, <\}$; a data constraint is of a similar form $i \sim j$ or $i - j \sim k$ but with k being an arbitrary integer. The default guard of an edge is *true*.

3.1.2 Reset-operations

When taking an edge clock and data variables may be subject to simple manipulations in terms of *resets* being assignments of the form $w := e$, where w is a clock or data variable and e is an expression. In the current version of UPPAAL reset-operations on clock variables must be of the simple form $x := n$, where n is a natural number, and reset-operations on integer variables should be in the form $i := c * i + c'$, where c, c' are integer constants (note that c, c' may be zero or negative). As examples reconsider Fig. 3. Here the clock y is reset to 0 when the edge between A_0 and A_1 is taken. Similarly the integer variable n is incremented when the edge from B_1 to B_2 is taken.

3.1.3 Channels, synchronization, and urgency

A UPPAAL model consists of a network of (extended) timed automata. Automata may communicate either via integer variables (which in UPPAAL are global), or using communication channels. As in CCS [30] communication on channels occur as two-process synchronizations. In Fig. 3 the two processes may communicate via the channel a .

To denote the actions that processes can perform when synchronizing with each other we use the notation $a!$ and $a?$ (denoting the complementary actions of sending and receiving on channel a). Absence of synchronization action indicates an internal (non-synchronizing) edge similar to τ -transitions in CCS. In Fig. 3, the edge between A_1 and A_2 is an example of an internal edge of the process A.

To prevent a network from delaying in a situation where two components are already able to synchronize, a channel may be declared as being *urgent*. For efficiency reasons edges labeled with synchronization actions on urgent channels may not have guards on clocks.

3.1.4 Committed locations

To introduce the notion of committed locations in timed automata, consider the scenario shown in Fig. 4. A sender S is to broadcast a message m to two receivers R_1 and R_2 . As this requires synchronization between *three* processes this cannot directly be expressed in UPPAAL where synchronization, as in CCS, is between two processes based on complementarity of actions. However, as an initial attempt we may model the broadcast as a sequence of two two-process synchronizations, where first S synchronizes with R_1 on m_1 and then with R_2 on m_2 . However, this is not an accurate modeling as the intended atomicity of the broadcast is not preserved (i.e., other processes may interfere during the “broadcast” sequence). To ensure atomicity, we mark the intermediate location S_2 of the sender S as a so-called *committed location* (indicated by the *c*-prefix). The atomicity of the action sequence $m_1!m_2!$ is now achieved by insisting that a committed location must be left immediately! This behavior is quite similar to what has been called “urgent transitions” [6, 14, 19] which insists that the next transition taken must be an action (and not a delay). The precise semantics of committed locations will be formalized in the transition rules for networks of timed automata with data variables in the following.




Fig. 4. Broadcasting communication and committed locations

3.1.5 Invariants

To enforce progress in a system, control nodes may be decorated with so-called *invariants*, which express constraints on the clock values in order for control to remain in a particular node. The default of a location invariant is *true*. Thus, in Fig. 3, control can only remain in A_0 as long as the value of y is no more than 6.

3.2 Semantics

Formally, states of a UPPAAL model are of the form (\bar{l}, v) , where \bar{l} is a *control vector* indicating the current control node for each component of the network and v is an *assignment* given the current value for each clock and integer variable. A UPPAAL model determines the following two types of *transitions* between states:

Delay transitions As long as none of the invariants of the control nodes in the current state are violated time may progress without affecting the control node vector and with all clock values incremented with the elapsed duration of time. In Fig. 3, from the initial state $((A_0, B_0), x = 0, y = 0, n = 0)$ time may elapse 3.5 time units leading to the state $((A_0, B_0), x = 3.5, y = 3.5, n = 0)$. However, time cannot elapse 5 time units as this would violate the invariant of B_0 .

Action transitions If two complementary labeled edges of two different components are enabled in a state then they can synchronize. Thus in state $((A_0, B_0), x = 3.5, y = 3.5, n = 0)$ the two components can synchronize on a leading to the new state $((A_1, B_1), x = 0, y = 0, n = 5)$ (note that x, y, n have been appropriately updated). If a component has an internal edge enabled, the edge can be taken without any synchronization. Thus in state $((A_1, B_1), x = 0, y = 0, n = 5)$, the B -component can perform without synchronizing with A , leading to the state $((A_1, B_2), x = 0, y = 0, n = 6)$.

The above two types of transitions may be overruled by the presence of urgent channels and committed locations in the following ways:

Urgent channels In a state where two components may synchronize on an urgent channel no further delay is allowed. Thus, in Fig. 3 if channel a is urgent³, time may *not* elapse 3.5 units from the initial state $((A_0, B_0), x = 0, y = 0, n = 0)$ as synchronization on a is already possible in the state $((A_0, B_0), x = 3, y = 3, n = 0)$.

Committed locations If in a state one of the components is in a control node labeled as being committed, no delay is allowed to occur and any action transition (synchronization or not) *must* involve the particular

component (the component is, so to speak, committed to continue). In the state $((A_1, B_1), x = 0, y = 0, n = 5)$ B_1 is committed; thus without any delay the next transition must involve the B -component; i.e., the next state of the network is $((A_1, B_2), x = 0, y = 0, n = 6)$. Hence the two first transitions of B are guaranteed to be performed atomically.

4 The nuts of UPPAAL – specifying

The UPPAAL model-checker is designed to check for simple invariant and reachability properties. A number of other properties, including bounded reachability properties, may be checked by reasoning about the system in the context of testing automata. We give an informal presentation of the UPPAAL logic and an example use of testing automata in the next two sections. Hereafter, we give a short review of the model-checking technique used in UPPAAL and point out some recent developed and implemented space-saving improvements.

4.1 UPPAAL specifications

In the current version, UPPAAL is able to check for reachability properties, in particular whether certain combinations of control-nodes and constraints on clock and data variables are reachable from an initial configuration. The properties that can be analyzed are of the forms:

$$\varphi ::= \forall \square \beta \mid \exists \diamond \beta \qquad \beta ::= a \mid \beta_1 \wedge \beta_2 \mid \neg \beta$$

Where a is an atomic formula being either an atomic clock (or data) constraint or a component location ($A_i \text{ at } l$). Atomic clock (data) constraints are either integer bounds on individual clock (data) variables (e.g., $1 \leq x \leq 5$) or integer bounds on differences of two clock (data) variables (e.g., $3 \leq x - y \leq 7$).

Intuitively, for $\forall \square \beta$ to be satisfied all reachable states must satisfy β . Dually, for $\exists \diamond \beta$ to be satisfied some reachable state must satisfy β . Formally let \rightsquigarrow denote the transitive closure of the delay- and action-transition relations between states. Then the satisfaction relation \models between states and formulas are defined as follows:

$$\begin{aligned} \langle \bar{l}, v \rangle \models \exists \diamond \beta &\iff \exists \langle \bar{l}', v' \rangle. \langle \bar{l}, v \rangle \rightsquigarrow \langle \bar{l}', v' \rangle \wedge \langle \bar{l}', v' \rangle \models \beta \\ \langle \bar{l}, v \rangle \models \forall \square \beta &\iff \forall \langle \bar{l}', v' \rangle. \langle \bar{l}, v \rangle \rightsquigarrow \langle \bar{l}', v' \rangle \implies \langle \bar{l}', v' \rangle \models \beta \end{aligned}$$

Satisfaction with respect to a boolean combination β of atomic formulas is defined inductively on the structure of β (behaving as usual with respect to the boolean connectives). Satisfaction with respect to an atomic formula is given by the following definitions:

$$\begin{aligned} \langle \bar{l}, v \rangle \models c &\iff v \in c \\ \langle \bar{l}, v \rangle \models A_i \text{ at } l &\iff l_i = l \end{aligned}$$

³ Note that, strictly speaking, this would violate the syntactic restriction that the guard of edges labeled with urgent actions must be empty.

4.2 Test automata – beyond reachability

Our (simple and efficient) model-checking technique extends to the logic presented in [25], which also allows for bounded liveness properties to be specified. Currently, bounded liveness properties must be obtained by reachability analysis of the system in the context of testing (and time-sensitive) automata.

Consider the following real-time property

$$\psi = \phi \text{ Until}_{<t} a$$

stating that the (atomic) property a must hold before t time units and that ϕ must hold until then. Here we assume that a is of the form $A_i \text{ at } l$. Now, to verify that a system S satisfies the formulae we extend it with the test automata T of Fig. 5 as a component.

Here T' is assumed to be an already constructed test automata for the sub-property ϕ , and 'a' is a (urgent) probe action inserting into the component A_i at location l . Now, it may be shown that our original system S will satisfy $\phi \text{ Until}_{<t} a$ if and only if $S | T$ satisfies the invariance property $\forall \square \neg (T \text{ at } \text{bad})$.

We conjecture that all bounded liveness properties of the logic in [25] can be translated into reachability problems in this manner, and in a forthcoming version of UPPAAL we intend to provide automatic support for generation of test automata from logical formulas (as has been done in the tool SPIN, where “never-claims” are directly generated from Linear Temporal Logic properties). For an initial investigation we refer the reader to [20].




Fig. 5. Test Automata for $\phi \text{ until}_t a$

4.3 Model-checking

The model-checking procedure implemented in UPPAAL is based on an interpretation using a finite-state symbolic semantics of networks. More precisely, we interpret the logic with respect to symbolic states of the form (\vec{l}, D) , where D is a constraint system (i.e., a conjunction of atomic clock and data constraints), and \vec{l} a control-vector. Thus, a symbolic state (\vec{l}, D) represents all the states (\vec{l}, v) where v satisfies the constraint D . Based on this notion of symbolic state, the heart of the UPPAAL model-checking procedure is the abstract reachability algorithm shown in Fig. 6,⁴ which reduces the reachability problem to that of solving simple constraint systems. The algorithm is to check whether a timed automaton may reach a state satisfying a given state formula β .

We observe that several operations of the algorithm are critical for efficient implementations. Firstly, the algorithm depends heavily on the test operations for checking the inclusion $D \subseteq D'$ (i.e., the inclusion between constraints D and D') and the emptiness of D_s in constructing the successor set SUCC of (\vec{l}, D) . Clearly, it is important to design efficient data structures and algorithms for the representation and manipulation of clock constraints. One such well-known data structure is that of Difference Bounded Matrices [3, 15, 33], DBM, which offers a canonical representation for constraint systems. It has been successfully employed by several real-time verification tools, e.g., UPPAAL [7] and KRONOS [14].

In the UPPAAL-implementation the reachability algorithm of Fig. 6 is extended so that a *diagnostic trace* is automatically generated as a side-effect in case reachability is established. If the symbolic state-space is examined in a breadth-first manner (corresponding to organizing

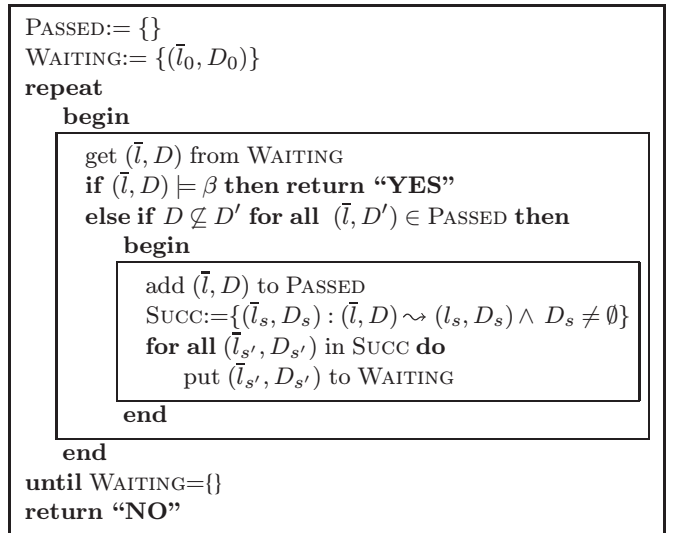


Fig. 6. An algorithm for symbolic reachability analysis

⁴ The relation \rightsquigarrow has been extended to symbolic states in the obvious fashion.

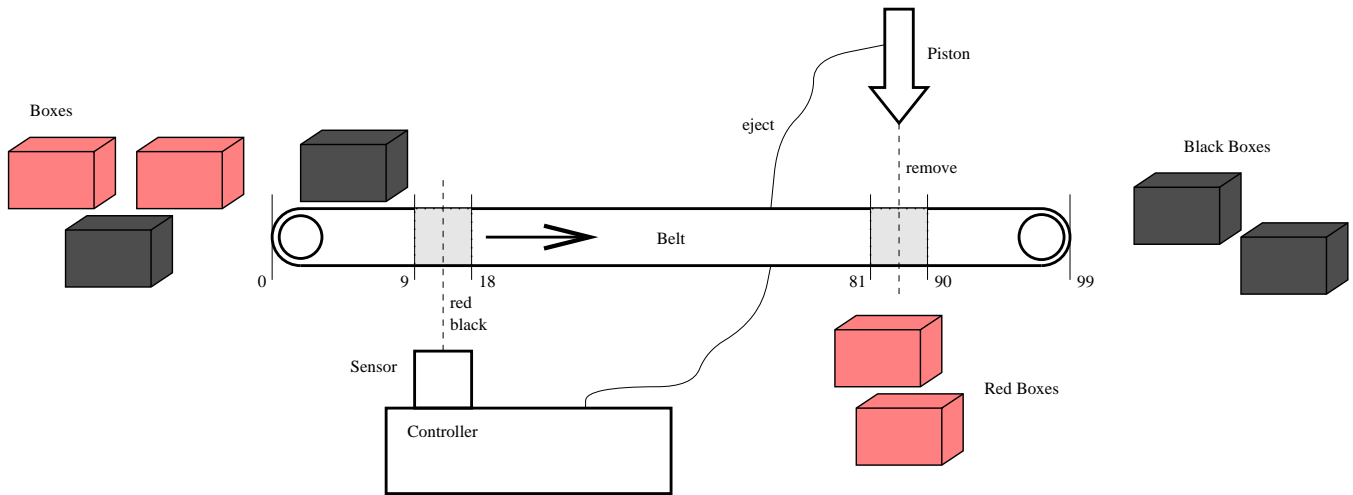


Fig. 7. The box-sorter unit

the waiting set as a queue), the trace is guaranteed to be the shortest possible.

4.4 Optimizations

A DBM representation is in fact a weighted directed graph where the vertices correspond to clocks (including a zero-clock) and the weights on the edges stand for the bounds on the differences between pairs of clocks [3, 15, 33]. As it gives an explicit bound for the difference between each pair of clocks, its space-usage is in the order of $\mathcal{O}(n^2)$ where n is the number of clocks. However, in practice it often turns out that most of these bounds are redundant. In [27], we have presented a new compact data structure for DBM, which provides *minimal* and *canonical* representations of clock constraints and also allows for efficient inclusion checks. The representation is obtained by a minimization of the weighted directed graph representing the constraint system, and our experimental results demonstrate truly significant space-savings as well as better time-performance.

In addition to the *local* reduction technique above, which is to minimize the space-usage of each individual symbolic state, we have developed a *global* reduction technique to reduce the total number of symbolic states to save in the global data structure, i.e., the passed list. It is completely orthogonal to the local technique and is based on static analysis of the control structure of the system. Again, our experimental results demonstrate significant space-savings and improved time-performance.

Finally, it is possible to save the symbolic state-space generated during checking of a property and re-use it in the checking of other properties. In cases where several correctness properties have to be examined this leads to significant time-savings.

5 How to use UPPAAL – a user guide

This section describes how to model, simulate and verify real-time systems using UPPAAL. We will focus on the graphical interface and the graphical modeling language of UPPAAL. However, it should be noticed that the tool also has a text-based interface and a textual representation of the modeling language.

As a running example throughout this section we will use the box-sorting unit of [23].

Example 1 (The box-sorter unit). The box-sorter shown in Fig. 7 is made in LEGO.⁵ It sorts red and black boxes in the unit, which consists of four components: a color-sensor, a piston, a controller, and an observer.

A box starts at the leftmost extreme of the belt, represented by position 0. At some position from 9 to 18 its color is sensed by the color-sensor which is attached to the controller. The controller reacts if the color is red by sending an eject request to the piston, after a certain delay. The piston ejects within one time unit after the arrival of a request. When the piston is ejected it is guaranteed to remove the box if it is positioned in the interval 81 to 90. If the box is not removed (i.e., if the box is black), it proceeds to position 99, representing the rightmost extreme of the belt, where it falls off the belt.

The observer is not participating in the sorting of boxes. Its only task is to observe that no red boxes appear on the rightmost extreme of the belt. As the observer is not part of the sorting mechanism, we have not shown him/her in Fig. 7 but he/she can be imagined to sit at the far right end.

⁵ Information about LEGO can be found on World Wide Web at: <http://www.lego.com/>.

For simplicity we regard the system as being correct if the observer sees no red boxes (i.e., only black boxes) at the rightmost extreme of the belt. □

5.1 XUPPAAL

XUPPAAL is to provide a user-friendly graphical interface to the tools in UPPAAL. It offers support to the user by managing the working file names, by providing an easy way to give optional setting, and to execute the various tools of the tool-box. The application is implemented in C and the Forms Library for X.⁶

In the following we give a short overview of the tool. The description continues in Sect. 5.4 which explains how to verify and generate diagnostic traces with XUPPAAL. A more detailed description of XUPPAAL can be found in [4].

5.1.1 The files

During an XUPPAAL session the user works with three different kind of files: a system description, a requirement specification, and a trace file.

The *system description* file contains the system description. It is assumed to be given in the textual format (called `.ta`) or on the graphical format (called `.atg`) used by Autograph and the simulator. Section 5.2 provides a guide for definition of system descriptions using Autograph. Once the file has been created its syntax can be checked by invoking **Syntax Check** in the **Run** menu of XUPPAAL.

The *requirement specification* file holds a set of formulae in a textual representation (called `.q`). The file can be created using a simple editor invoked in the **Run** (item **Req. Spec. Editor**) menu of XUPPAAL or using an external text editor.

The *trace* file is used to store information about diagnostic traces generated by the verifier. A trace is a sequence of (symbolic) states and transitions, that represents an (symbolic) execution of the system. It is often useful for discovering why a property is (or is not) satisfied by a system description. The trace file can be input to the simulator which is able to display traces graphically (see Sect. 5.3).

5.1.2 Getting started

XUPPAAL is activated from the command line with the command `xuppaal`. On startup, XUPPAAL shows the user with its main window as shown in Fig. 8. The main window consists of four parts: the menu bar, the two input

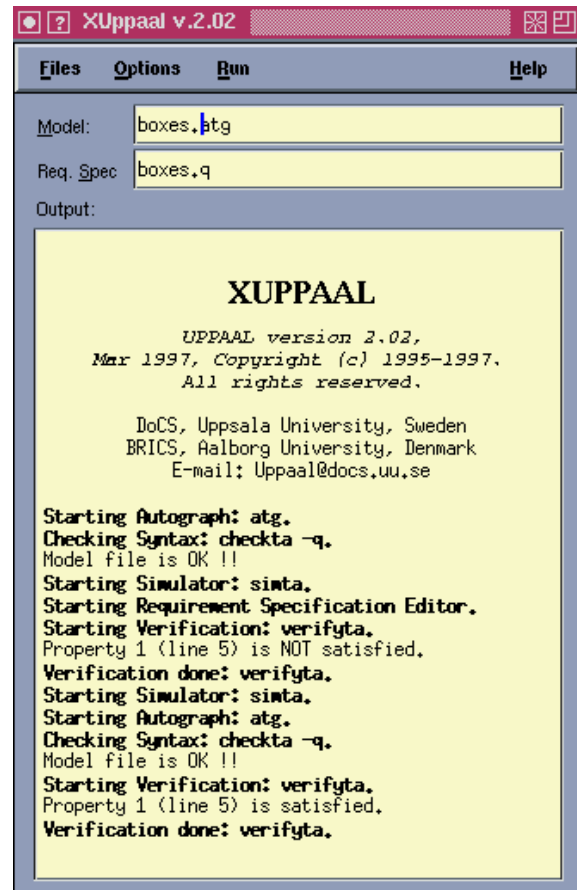


Fig. 8. The graphical interface of XUPPAAL

fields that displays the name of the currently specified system description file (labeled **Model**) and requirement specification file (labeled **Req. Spec.**), and the output browser (labeled **Output**).

5.1.3 The menu bar

The menu bar has four sub-menus: **Files**, **Options**, **Run**, and **Help**. Each sub-menu is invoked by clicking the menu label or using snap-keys. The snap-keys are given below enclosed in parenthesis (where **C** denotes the **Ctrl**-button).

The **File** menu (**C-x**) contains entries for selecting the files to work with, and for exiting XUPPAAL. The **Help** menu (**C-h**) provides online help on the six topics: **General** (**C-g**), **Files** (**C-f**), **Options** (**C-o**), **Run** (**C-r**), **Problems** (**C-p**), and **Version** (**C-v**). When an entry is selected, the help text is displayed in the output browser. The output generated by selecting the menu item **Version** is shown on the first seven lines of text in the browser in Fig. 8.

We proceed by giving a more detailed description of the two sub-menus: **Options** and **Run**.

The Options Menu (**C-o**) provides a list of choices that mainly affect the verification session. Each entry in the

⁶ Information about the Forms Library for X is available on World Wide Web: <http://www.uwm.edu/~zhao> and on anonymous ftp: <ftp://bloch.phys.uwm.edu/pub/xforms>.

menu can be toggled on or off (which is visualized in the menu by a filled or empty check-box to the left of each menu item):

Auto Check Syntax (C-a): automatically perform syntax check before simulation and verification sessions.

Diagnostic Info (C-i): generate diagnostic traces in textual format and present the result in the output browser.

Diagnostic File (C-f): produce diagnostic traces on the specified trace file.

Breadth-First (C-b): explore the state-space of the system by breadth-first search⁷.

Depth-First (C-d): explore the state-space of the system by depth-first search.

Local Reduction (C-l): use compact data-structures to represent constraints (instead of DBM, see Sect. 4.4).

Global Reduction (C-s): perform control structure analysis (see Sect. 4.4).

Re-use State-Space (C-r): re-use the generated portion of the state-space when verifying *several* reachability properties (see Sect. 4.4).

The Run Menu (C-c) contains commands that are associated with the various tool programs of UPPAAL. The outputs produced by the programs are always displayed in the output browser.

Autograph (C-a): start the graphical editor Autograph.
Syntax Check (C-c): perform syntactical check on the textual format of the system description.

Simulation (C-s): start the simulator with the specified system description file.

Req. Spec. Editor (C-e): open the requirement specification editor shown in Fig. 9.

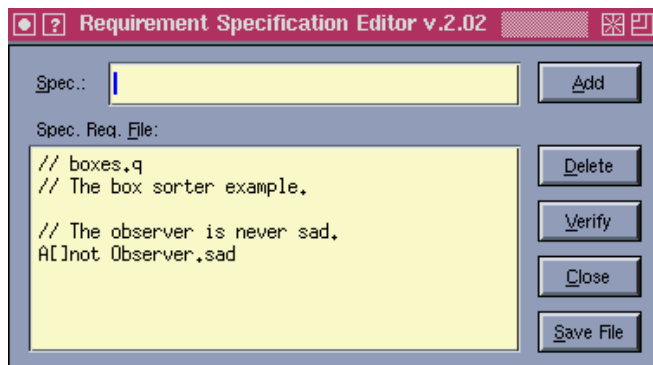


Fig. 9. The Requirement Specification Editor

Verification (C-v): check if the system description satisfies the requirement specification by model-checking.

Show Model (C-f): display the system description in textual representation.

Show Req. Spec. (C-r): show the contents of the requirement specification file.

Clear Output (C-o): clear the output browser (i.e., erase the contents of the browser).

Example 2 (Specifying the file names). As a preliminary step when working with the box-sorter example we define the file names for XUPPAAL to work with. We do this by typing in the system description file name “boxes.atg” in the field labeled **System** and the requirement specification file name “boxes.q” in the field labeled **Req. Spec.** The trace file is specified by selecting **Set Trace File** from the **Files** menu. We use the file name `boxes.tr`. The resulting XUPPAAL input fields are shown in Fig. 8. □

5.2 How to model

UPPAAL allows for systems descriptions to be defined textually or by drawing using Autograph. In this section we describe how to define a system using Autograph. For a description of the textual representation of systems used in UPPAAL we refer the reader to [4].

5.2.1 Systems descriptions in autograph

Autograph is a graphical tool for drawing automata-based system descriptions. It is very general though it was developed for the FC2TOOLS set [10]. This section will explain the subset of features of Autograph, that are needed to define UPPAAL system descriptions.

To define network of timed automata in Autograph it is necessary to define a mapping from the elements of timed automata to graphical objects in Autograph. Here we summarize the mapping of the components that a system description should consist of.

A location is denoted by a vertex labeled with its name (required) and its invariant (which is optional). By default, the invariant is *true*.

A transition is denoted by an edge connecting two vertices. The transition may be labeled with a guard which is *true* by default, a synchronization action (which is τ by default), and a list of assignments.

A timed automaton is denoted by a box containing the locations and transitions of the automaton with the initial location being denoted by an *initial vertex*. The box should be labeled with the name of the automaton.

A declaration for the variables, clocks and channels, both normal and urgent, used in the network of timed automata are placed in a box which must be labeled with **Config**.

A network of timed automata is represented by a set of set of automata boxes and the **Config** box.

If a graph is created according to these definitions and to the small set of rules that will be introduced in the

⁷ This option generates a shortest diagnostic trace when used in combination with one of the two trace options above.




Fig. 10. The menu bar with the ObjectEdit menu selected

remainder of this section, it can be used directly in the UPPAAL toolkit.

5.2.2 Getting started

Autograph can be activated from the Run menu of XUPPAAL. On startup the program displays a menu bar which is depicted in Fig. 10 with the ObjectsEdit menu selected. All menus mentioned in the remainder of this section are parts of this menu bar.

Note that all items inside a menu bar, can also be selected by a snap-key given below in parenthesis, and a selection will be active until another selection is made.

To get started it is necessary to open a window to draw in. This is done by selecting Create from the Window menu. It is possible to resize and reposition the window.

5.2.3 Creating locations

Locations of a timed automata are created by choosing Create (v) from the Vertice sub-menu in the ObjectsEdit menu. To instantiate a location just find the wanted position and click the left mouse button to place it.

If the vertex is supposed to model an initial location, select Initial (i) from the Vertice sub-menu of the ObjectsEdit menu and click at the chosen vertex.

To create the name and the invariant of the location, a label must be attached to the vertex. The actual creation of labels will be described in the next section. There is a set of rules concerning location labels that must be followed to avoid errors when using the drawing in UPPAAL:

- Every location must have a name associated.
- All names are specified by the regular expression $[a-zA-Z][a-zA-Z0-9_]*$. If the location is committed, its name must be prefixed with C: or c:.
- Location names are local to a process and can be reused in other processes.
- An invariant is a conjunction of inequalities. It is written as a list of inequalities separated by commas and surrounded by a pair of parentheses, e.g., ($clock1 \leq 42, clock2 \leq 117$).
- The invariant can be put either on the same line as the location name or on a line for itself.

5.2.4 Creating labels

In Autograph there are four kinds of labels: Struct, Behav, Logic, and Hook. The different kinds of labels may have different meanings depending on the kind of component they are attached to.

In UPPAAL the notion of different labels has been avoided as much as possible. Therefore, the default label (automatically selected by Autograph) is always used for components with only one label. To label a component with its default label select Create/Edit Default (a) from the Label menu, then click on the component to label. This opens an editor window for typing in text for the label.

The only component type with more than one label is the Config box described in Sect. 5.2.8.

Labels can be re-edited by selecting Reedit (R) from the Label menu. After the label has been clicked an editor for editing (the four kinds of) labels will appear.

5.2.5 Creating edges

The next step in defining a network of automata is to connect locations with transitions, denoted by edges. Edges can only be drawn between two vertices belonging to the same automaton. The start and end vertices can be the same.

To create an edge, first select Create (e) from the Edges sub-menu under ObjectsEdit and then select the start and end vertex. One can create curved edges. The simplest way is to drag the mouse from the start of the edge to a point and then continue to the end of the edge.

The optionals that can belong to a transition are put in a label attached to the edge. An edge can have more than one set of labels; whereas all other components can *only* have one.

As for location labels, there is a set of rules for transition labels:

- If a transition is synchronizing, its label should have a line containing the name of the channel it is synchronizing on.
- Guards are written as lists of comma-separated constraints.

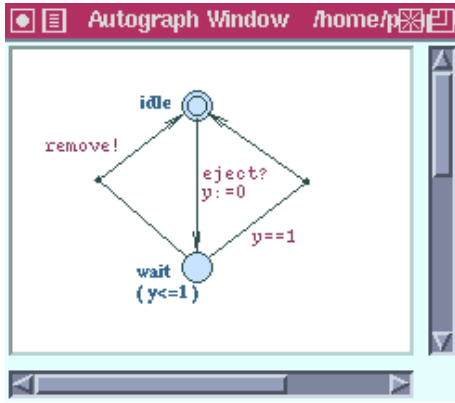


Fig. 11. Locations and transitions of the piston

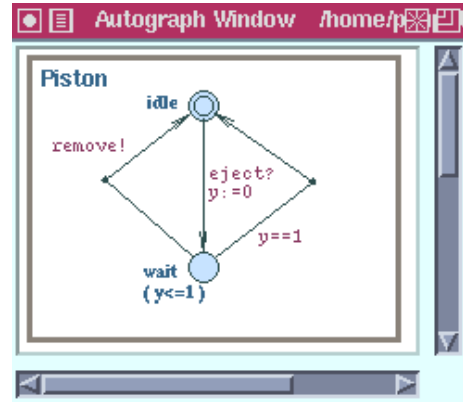


Fig. 12. An automaton defining the piston of the box-sorter unit

- Assignments are written as lists of comma-separated assignments on the form: `<name> := <expression>`, where expressions follow the syntax defined in Sect. 3.

Example 3. To start modeling the box-sorter unit we draw the locations and the transitions of the piston automaton. The piston waits for the controller to send an `eject`-signal. It then ejects within one second, possibly resulting in the removal of a box from the belt.

A model of the piston is shown in Fig. 11. The two locations `idle` and `wait` are used to model the two operational modes of the piston. Initially the piston is `idle`. The piston will enter location `wait` when the signal `eject` is received. In location `wait` the piston is ready to remove a box (by synchronizing on `remove`) for 1 time unit after which the piston will return to the location `idle`. The clock variable `y` is used to model the timing behavior of the piston. □

5.2.6 Creating automata

An automaton is represented by a box containing its locations and transitions. Boxes are created by choosing **Create (b)** from the **Boxes** sub-menu in the **ObjectsEdit** menu. The box is drawn by picking the position of its upper left corner and dragging the mouse until the box has the wanted size (i.e., the position of the lower right corner).

Every automaton must be labeled with a name that is written according to the regular expression `[a-zA-Z][a-zA-Z0-9_]*`.

Example 4. We now finish the piston automaton by adding a bounding box around its locations and transitions. We name the automaton `Piston` by adding a label to the box. The resulting automaton is shown in Fig. 12.

To create an automaton, one can also start with drawing the box and then the vertices, edges, etc. □

5.2.7 Creating a network of automata

A network of automata is denoted by a collection of automata boxes representing the component automata and a distinct box, labeled with `Config`, for declarations.

5.2.8 Creating declarations

Declarations of objects in the network, i.e., variables, clocks, and channels are placed in the `Config` box. In addition, comments for documentation can be put in this box. Declarations should be created according to the following rules.

- Declarations are written as lists of comma-separated object names preceded by the type of the entities declared and ended with a semi-colon.
- The type of a variable declaration is `int`.
- The type of a clock declaration is `clock`.
- The type of a declaration of normal channels is `chan` and for urgent channels it is `urgent chan`.
- There can be zero or more declarations of each type.
- Names of objects are specified by the regular expression `[a-zA-Z][a-zA-Z0-9_]*`.
- Comments are prepended with `//`. Except for this, there are no rules concerning the syntax of comments.

In addition to the object declarations the `Config` box also holds the definition of which automata the system consists of. This definition is written using the same syntax as above with the type of the declaration being `system`. Only the automata mentioned in the `system` definition are considered when the system is verified or simulated. There must be exactly one `system` definition.

There are also some rules concerning the `Config` box and other boxes with comments:

- There must be exactly one `Config` box.
- The `Config` box cannot contain any components. It is only allowed to have a label attached.
- Boxes other than the components and `Config` box are ignored by UPPAAL.

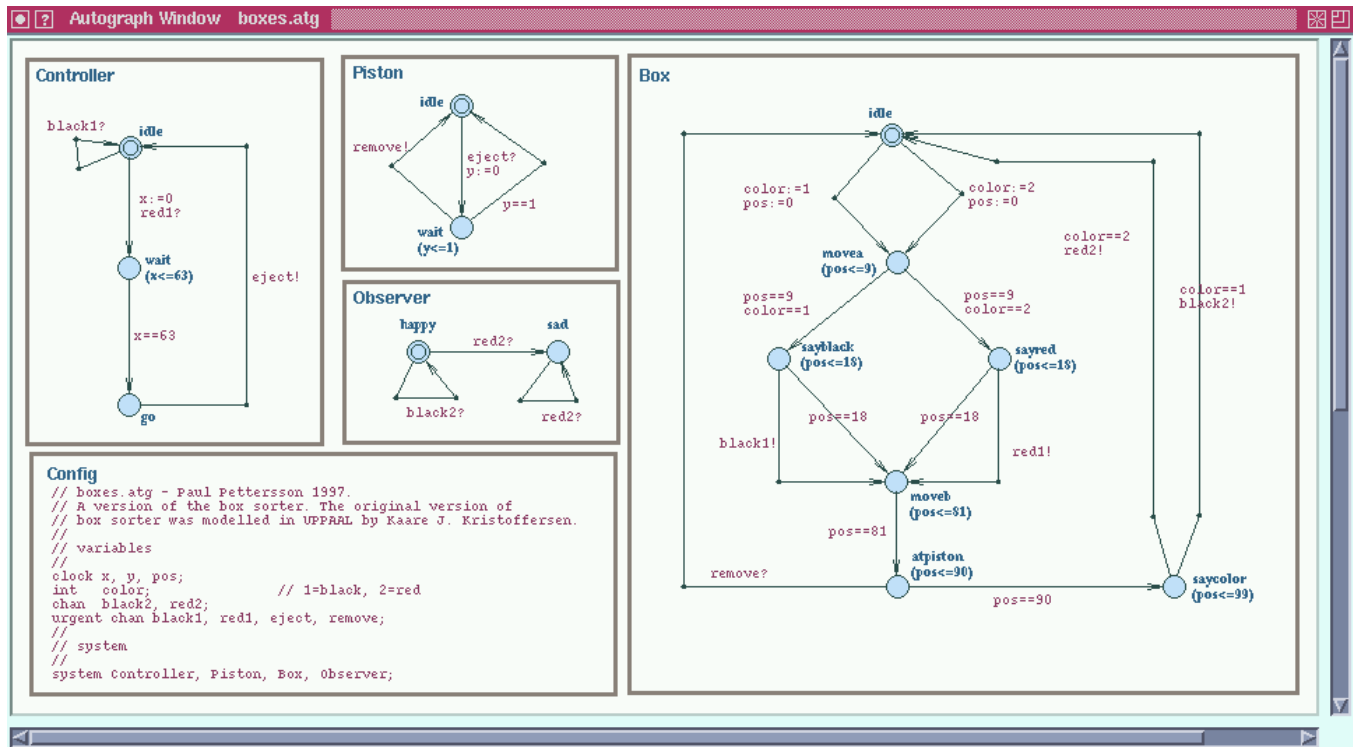


Fig. 13. The complete system description of the box sorter unit

Example 5. The complete system description of the box-sorter unit is shown in Fig. 13. It consists of four automata: Piston, Controller, Box, and Observer. In this example we describe the two automata Box and Controller. The Piston automaton was described in Example 0 and the Observer will be described in Example 0.

The Box automaton models the behavior of a box. It uses the integer variable `color` to represent the color of the box (where 1 is black and 2 is red), and the clock variable `pos` to represent its position on the belt. In the initial location `idle` is the box not yet placed on the belt (even if $pos \geq 0$). On the outgoing transitions from `idle` to `movea` the box is put on the belt and assigned a color. The two locations `sayblack` and `sayred` model the behavior of a box when it passes the color-sensor. It offers the signals `black1!` or `red1!` to the sensor until position 18 is reached, where it enters location `moveb`. The location `atpiston` models the situation when the box passes the piston, therefore the box offers synchronization on the channel `remove` to model the possibility of getting removed. If not removed, the box proceeds to location `saycolor` where it returns to its initial location by synchronizing on channel `black2` or `red2`, which are synchronized with the observer.

The Controller automaton models the controller and its integrated color-sensor. It stays in its initial location `idle` as long as only black boxes appear, modeled by a `black1`-synchronizing self-loop. When a red box appears, it delays for 63 time units in location `wait`

and there after requests a reject from the piston by synchronization on the `eject`-channel as soon as possible.

To model that the color-sensor and the piston react immediately on input, the four channels `black1`, `red1`, `eject`, and `remove` are declared as urgent. \square

5.2.9 Saving the system description file

To save a file, select either **Save atg (s)** or **Save atg as (S)** from the **Files** menu and click on the drawing to be saved. This leads to the opening of a file selector window, where the path and name of the file can be entered. When ready press the OK button.

The saved `.atg` files can then be used in UPPAAL.

5.3 How to simulate

The symbolic simulator enables the user to simulate and debug the dynamical executions of a network of timed automata given its statical structure i.e., a system description. This section provides a guide for the interface of the simulator and show how to use it.

5.3.1 Getting started

The simulator can be activated from the **Run** menu in XUPPAAL. The system description file (an `.atg`-file)

specified in the `Model` field of XUPPAAL will be loaded into the simulator on startup.

The graphical interface to the simulator consists of two windows: the `System` window to show the system description to simulate, and the `Main` window to control the simulator.

The interface uses the standard X notions for mouse and keyboard control meaning that, for example, double-clicking the mouse and using tab-key to switch between groups in windows can be used.

5.3.2 The System window

The `System` window holds the Autograph drawing of the system being simulated. Figure 14 shows an example of a `System` window containing the definition of the box sorter unit.

During a simulation, the current location vector and one of the currently available transitions are depicted in the window. For the location vector this is accomplished by marking the current location of each of the automata, e.g., in the figure, `Box` is in location `idle`. A possible transition is depicted as a highlighted arrow going from the source location of the transition to its destination. If the transition is a synchronization, an arrow will be shown for each of the two synchronizing transitions, otherwise only one. In the figure a transition from `idle` to `movea` in `Box` is shown.

It is not possible to manipulate the contents of the `System` window directly as its only purpose is to display the

structure of the current system and to show how simulations progress. All changes to the `System` window are performed indirectly through the `Main` window described below.

5.3.3 The Main window

The simulator is controlled from the `Main` window which is split up into two parts: one containing the basic simulation control, and the other containing the more advanced control mechanisms concerned with traces. Fig. 15 shows the `Main` window with the trace so far and the possible step for the system depicted in Fig. 14.

The upper part of the window has six buttons and a field holding a list of steps possible in the current state of the system. The lower part also has six buttons and a field containing the trace of the current simulation. When the simulator is activated only the upper half of this window is opened.

The upper six buttons are used for controlling the most basic functionality of the simulator. The semantics of the buttons is as follows:

- Take Step:** take the step selected in the list of possible steps. A step is selected if it is highlighted in the list.
- Run:** open the `Run` window which is used for controlling the automatic mode, where the simulator itself randomly selects transitions.
- Reset:** set the automata in the system to their initial locations and clears the trace.

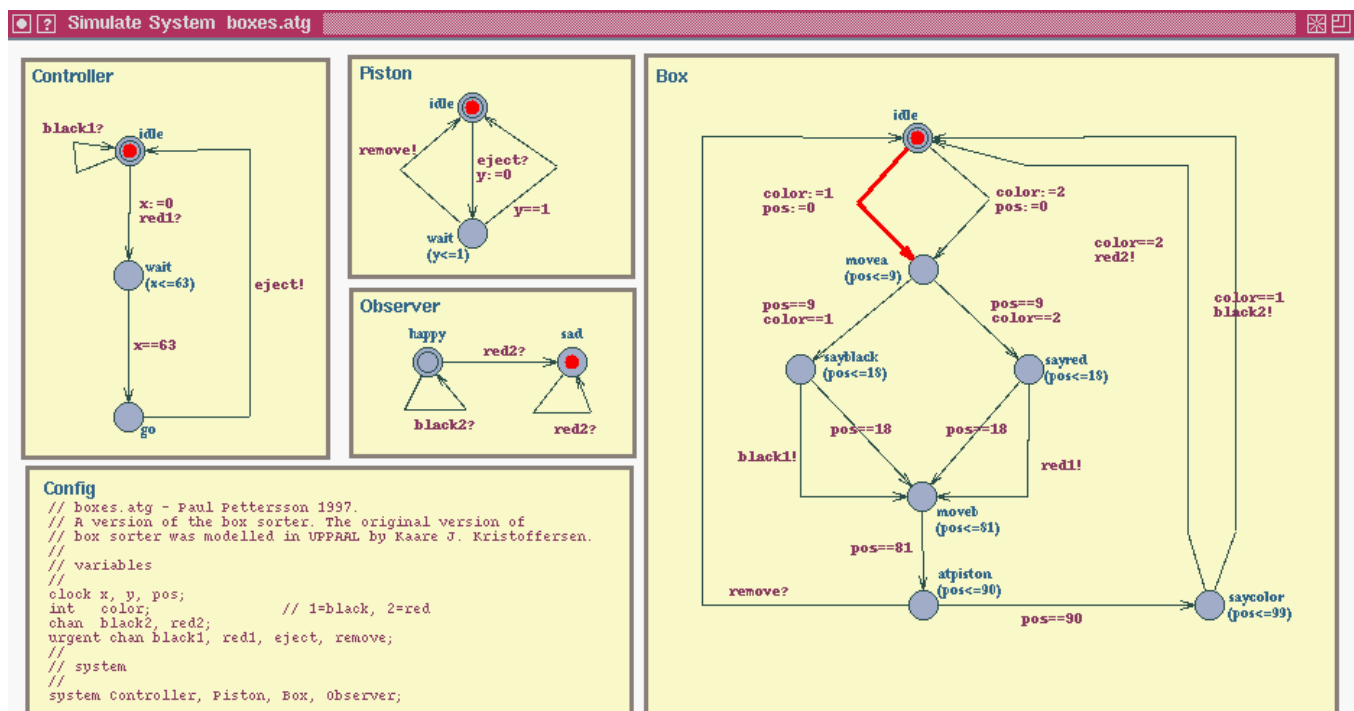


Fig. 14. A System window

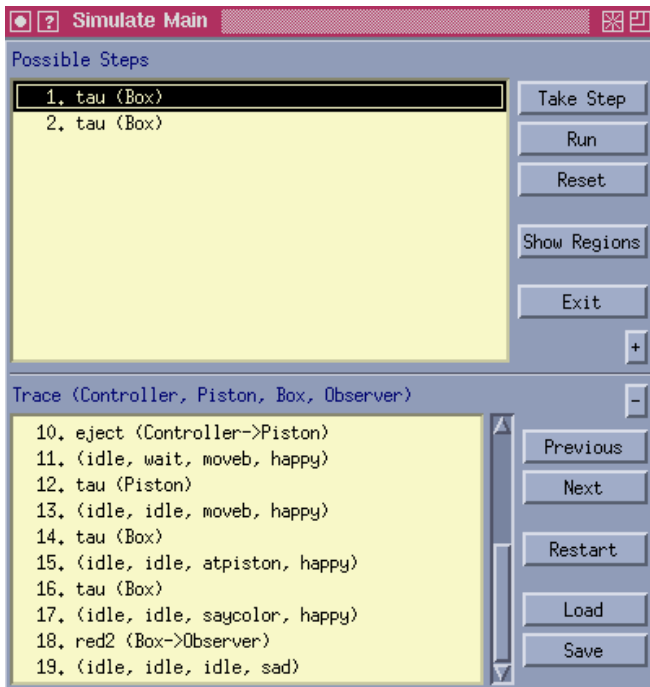


Fig. 15. The Main window

Show Regions: open the Regions window showing the regions valid for the current configuration of the simulated system.

Exit: exit the simulator.

+ (plus): open the lower half of the Simulate Main window, which displays the trace generated so far.

In the list of possible steps, the elements are written in one of the forms listed below. The first is used if the transition of the step is a synchronization between two processes and the second is used if the transition is non-synchronizing. It is possible to take a step directly from the list by double-clicking it.

- <channel> (<sender> -> <receiver>)
- tau (<process>)

A trace is a list whose elements alternate between location vectors and transitions. Location vectors are simply written as a list of location names where the names are ordered the same way as the list of automata given in the header of the lower half of the Simulate Main window shown in Fig. 15 is: Controller, Piston, Box, and Observer.

Transitions are written in the same form as the possible steps. Selecting an element in the trace changes the System window so that instead of reflecting the current state of the system it will show the selected transition or location.

The six buttons in the lower part of the Main window are all used for controlling the trace functionality. Their semantics is given in the following:

- (minus): close this half of the window.

Previous: highlight the element immediately preceding the current selection in the trace. If no element is selected, nothing happens.

Next: highlight the element immediately following the current selection in the trace. If no element is selected, nothing happens.

Restart: restart the simulation from the current selection in the trace.

Load: open a file selector window for loading a trace.

Save: open a file selector window for saving the current trace.

5.3.4 The Run window

Instead of creating traces manually, by taking one step at a time, it is possible to let the simulator proceed automatically by randomly selecting enabled steps and show the progress of the simulation in the System window. This can be done from the Run window as shown in Fig. 16.




Fig. 16. The Run window

The window shows two text input fields, **Number Of Steps** and **Speed**, plus three buttons. The field **Number Of Steps** decides how many steps the simulator should take when generating the random trace and the **Speed** field decides how fast the screen should be updated during the run, with 1 being the highest speed. The meaning of the three buttons is as follows:

Close: stop the run if it is currently executing, and close the Run window.

Stop: stop the run. If the simulator is not running, nothing happens.

Go: start or restart the execution of a run.

5.3.5 The Regions window

While doing a simulation the user can watch the changes of the regions in the Regions window. Fig. 17 shows the Regions window corresponding to the system state and transition selection of Figs. 14 and 15.

The window is split up into four parts each holding one particular kind of region.

The regions are represented as sets of equations and inequalities on the variables, clocks and differences between clocks in the system. Inequalities are written as intervals that the clock or clock difference must be in (see Fig. 17).

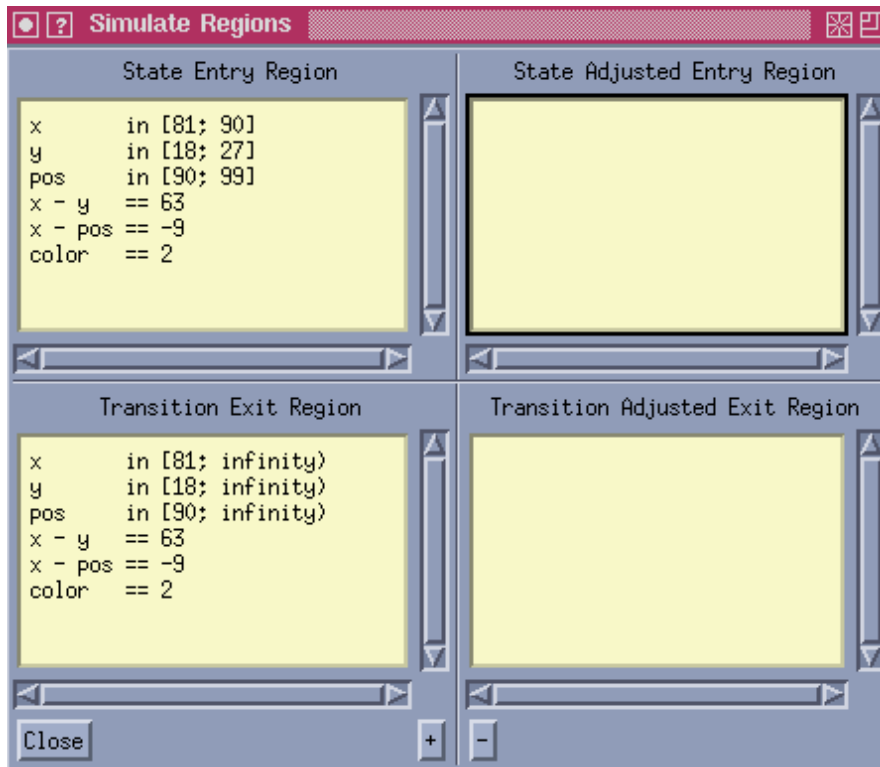


Fig. 17. The Regions window

The upper left quarter of the Regions window shows the *state entry region*. It is the region valid immediately after the last step, i.e., the step that took the system to the current state. The lower left quarter holds the *transition exit region*. For any of the possible steps that can take the system out of the current state, the transition exit region is the region where the system is able to take that particular step.

The *next state entry* region will then be the same as the *transition exit region* of the possible step which is taken next, except that it has been updated by the assignments of the transition.

The constraints of the latest step taken can put additional constraints on the regions for previous steps. This is not taken into account in the calculations of the regions in the left half of the window, but it is considered in the right where the adjusted counterparts of the two above mentioned types of regions are displayed.

5.4 How to verify

The verification of a system description with respect to its requirement specification is conducted entirely from the XUPPAAL window. It is often the case that the verification will *not* succeed immediately even if the system has been carefully validated in the simulator. More likely, a number of problems, usually in the system description, will have to be resolved before the verification finally succeeds.

In this section we describe how XUPPAAL supports the work by allowing diagnostic traces, generated by the verification procedure, to be loaded and graphically displayed in the simulator. We begin by describing how to verify in UPPAAL, using XUPPAAL.

5.4.1 Verification

To verify with XUPPAAL is very simple. When the file names have been specified, the next step is to select the optional settings that affect the verification method, and then it is ready for the start of a verification session. The menu of verification options has been described in Sect. 5.1.3.

The verifier (called `verifyta`) is activated by clicking **Verification** in the **Run** menu. XUPPAAL then first transforms the options, selected in the **Options** menu, to a list of flags to be accepted by the verifier. It then proceeds by compiling the graphical system description to its textual format.

Finally, it spawns a child process running the verifier with the required parameters, i.e., the list of flags, the system description file, the requirement specification, and (optionally) the specified trace file.

All output⁸ produced in the child process is displayed in the output browser of XUPPAAL. In particular, answers to whether each property of the requirement specification

⁸ The output which normally appears on standard output and standard error (i.e. `stdout` and `stderr`).

is satisfied or not are displayed. Each answer is numbered in order of appearance in the requirement specification file and with line numbers in parenthesis (see Fig. 8). If an answer line is double-clicked in the output browser of XUPPAAL, the requirement specification editor (if open) will display the corresponding line of the requirement specification file in its output browser (labeled **Req. Spec. File**). This is useful when working with large specifications.

Example 6 (Verifying the box-sorter). Now we show how to verify the correctness of the box-sorter unit. Since the sorter is considered correct if the observer sees no red boxes, we formalize the correctness in terms of reachable locations of the observer. We begin by describing the automaton **Observer** which models the observer.

The **Observer** automaton, which can be found in Fig. 13, has two locations: **happy** and **sad**. It is initially in location **happy** and stays there as long as it sees only black boxes at the rightmost extreme of the belt. If a red box appears, it enters the location **sad** where it stays forever.

To prove the box-sorter correct, it suffices to show that it is impossible for the observer to reach location **sad**, i.e., $\forall \square \neg \text{Observer.sad}$. We phrase the property in the UPPAAL syntax, $A[] \text{not Observer.sad}$, and put the resulting formula in the requirement specification file, named **boxes.q**. This can be done using the requirement specification editor which is activated in the **Run** menu of XUPPAAL.

Finally, we start the actual verification (**Verification** in the **Run** menu). The verifier outputs the text lines 13 and 14 in the output browser of Fig. 8, indicating that the property is not satisfied. \square

5.4.2 Visualizing diagnostic traces

The simulator generates traces and allows for traces to be displayed, reexamined, replayed, and reset from any intermediate point. Traces may also be reused using the save and load facilities. In particular, diagnostic traces generated by the verifier may be loaded for examination.

To use this facility, the first step is to specify the trace file name (see Sect. 5.1.1), and then instruct the verifier to generate traces on file. This is done by selecting item **Diagnostic File** in the **Options** menu of XUPPAAL. As a result, all succeeding verifications will output traces in the specified file until the **Diagnostic File** option is toggled off.

It should be noticed that it is possible to generate a *shortest* diagnostic trace, simply by instructing the verifier to explore the state space by breath-first search.

When the verification is finished, the trace file is loaded in the simulator by clicking the **Load** button in the **Main** window of the simulator (see Sect. 5.3.3).

Example 7 (Correcting the box-sorter). To find out why the requirement specification is not satisfied by the sys-

tem description we generate a diagnostic trace. As an exercise, we first produce a textual representation of a shortest diagnostic trace. To do so, we first select the options **Diagnostic Info** and **Breadth-First** in the **Options** menu and then initiate a new verification. The resulting trace is shown in Fig. 18.⁹

We can also try to debug graphically using the simulator. Recall that we have specified the trace file name in Example 0. A trace will be produced in that file if we select the option **Diagnostic File** and then verify the system again. Next, we load the trace into the simulator using button **Load** in the main window of the simulator. The windows of the simulator now look exactly the same as in Fig. 14, 15 and 17.

It can be deduced that there is an error in the **Controller** automaton. After the controller has sensed a red box, it will miss the box if it waits too short a time before ejecting the piston. If it waits another 9 time units before the **eject** signal is sent, the piston will hit the box.

The error can be corrected simply by modifying the time bound from 63 to 72 in the invariant and the guard on the edge outgoing from location **wait** in the **Controller**. It can be verified that the modified system is correct. \square

6 Case studies

UPPAAL has been applied in a number of (industrial) case studies. In this section we briefly review some of them.

Audio/Video Protocol: this is an audio control protocol highly dependent on real-time. The protocol was developed by Bang & Olufsen, to transmit messages between audio/video components over a single bus, and further studied in [18]. Though it was known to be faulty, the error was not found using conventional testing methods. Using UPPAAL, an error-trace is automatically produced, which revealed the error. Furthermore, a correction is suggested and automatically proved using UPPAAL.

Bounded Retransmission Protocol: the protocol was proposed and studied at *COST 247, International Workshop on Applied Formal Methods in System Design*. It is based on the alternating bit protocol over a lossy communication channel, but allows for a bounded number of retransmissions. In [12, 13], it is reported that a number of properties of the protocol is automatically checked with UPPAAL. In particular, it is shown that the correctness of the protocol is dependent on correctly chosen time-out values.

Collision Avoidance Protocol: the protocol in [20] is implemented on top of an Ethernet-like medium such as the CSMA/CD protocol. It is to ensure an upper bound on the communication delay between the network nodes. It was designed and proved correct using

⁹ Normally, the trace would have been displayed in the output browser of XUPPAAL.


```

(Controlller.idle Piston.idle Box.idle Observer.happy )
{ x=0 y=0 pos=0 color=0 }
(Controlller.idle Piston.idle Box.movea Observer.happy )
{ x=0 y=0 pos=0 color=2 }

delay(9)
(Controlller.idle Piston.idle Box.movea Observer.happy )
{x=9 y=9 pos=9 color=2 }
(Controlller.idle Piston.idle Box.sayred Observer.happy )
{ x=9 y=9 pos=9 color=2 }

Urgent sync: red1
(Controlller.wait Piston.idle Box.moveb Observer.happy )
{ x=0 y=9 pos=9 color=2 }
delay(63)
(Controlller.wait Piston.idle Box.moveb Observer.happy )
{x=63 y=72 pos=72 color=2 }
(Controlller.go Piston.idle Box.moveb Observer.happy )
{ x=63 y=72 pos=72 color=2 }
Urgent sync: eject

(Controlller.idle Piston.wait Box.moveb Observer.happy )
{ x=63 y=0 pos=72 color=2 }
delay(1)

(Controlller.idle Piston.wait Box.moveb Observer.happy )
{x=64 y=1 pos=73 color=2 }
(Controlller.idle Piston.idle Box.moveb Observer.happy )
{ x=64 y=1 pos=73 color=2 }
delay(8)

(Controlller.idle Piston.idle Box.moveb Observer.happy )
{x=72 y=9 pos=81 color=2 }
(Controlller.idle Piston.idle Box.atpiston Observer.happy )
{x=72 y=9 pos=81 color=2 }
delay(9)
(Controlller.idle Piston.idle Box.atpiston Observer.happy )
{x=81 y=18 pos=90 color=2 }
(Controlller.idle Piston.idle Box.saycolor Observer.happy )
{ x=81 y=18 pos=90 color=2 }

Sync: red2
(Controlller.idle Piston.idle Box.idle Observer.sad )
{ x=81 y=18 pos=90 color=2 }

```

Fig. 18. A diagnostic trace

UPPAAL. The two main properties established show that the protocol is collision-free, and it does ensure an upper bound on the user-to-user communication delay (assuming a perfect medium).

Gear-Box Controller: in this industrial case-study, UPPAAL is applied to the design and analysis of a prototype gear box controller for vehicles by Mecel AB¹⁰ [28]. The gear box controller is a component in the real-time distributed system that controls a modern car. The gear requests from the driver are delivered via the man/machine interface over a communication network to the gear box controller. The controller implements the actual gear change by actuating the lower level components of the system (such as the clutch, the engine, and the gear box).

In the design of the controller, the symbolic simulator of UPPAAL is applied to validate the system behavior. The correctness of the gear box controller design is established by automatic proofs of 46 properties derived from requirements specified by Mecel AB.

Philips Audio Protocol: the protocol was developed and implemented by Philips to exchange control information between components in audio equipment using Manchester encoding. The correctness of the encoding relies on timing delays between signals. It was first studied and manually verified in [9]. In [25] the protocol is modeled and verified using UPPAAL.

Philips Audio Protocol with Bus Collision: this is an extended variant of Philips audio control protocol with bus collision detection. It is significantly larger than

the version above since several new components (and variables) are introduced, and existing components are modified to deal with bus collisions. Its correctness was originally proved by hand in [17], and by model-checking for the first time using UPPAAL in [5].

Mutual Exclusion Protocol: the so-called Fischer protocol has been studied previously in many experiments, e.g., [1, 32]. The protocol is to ensure mutual exclusion among several processes competing for a critical section using timing constraints and a shared variable. Different versions of the protocol have been verified using UPPAAL, cf. [22, 24, 34].

TDMA Protocol Start-Up Mechanism: in [29], a formal verification of the start-up algorithm of a TDMA (Time Division Multiple Access) protocol is reported. It was proved using UPPAAL that an ensemble of three communicating stations becomes synchronized and operational within a bounded time from an arbitrary initial state, given a clock-drift corresponding to 10^{-3} . Furthermore, an upper time-bound for the start-up to complete was derived.

7 Conclusions

In this paper, we have presented the overall structure, the design criteria, and the main features of UPPAAL. It is intended to be an informal but reasonably “complete” description of the syntax and semantics of the UPPAAL modeling and specification languages as well as the kernel of its model-checking algorithm. It contains neither definitions nor theorems, but mainly examples. The intended readers are engineers working with development

¹⁰ Mecel AB is a Swedish company developing control systems for vehicle industries.

and analysis of embedded systems in particular, real-time protocols and controllers. A user guide is given, which describes in detail how to use the various tools of UPPAAL to *construct* abstract models of a system, to *simulate* its dynamical behavior, to *specify* and *verify* its safety and bounded liveness properties in terms of its model. In addition, we also provide a short review on case-studies where UPPAAL is applied, as well as references to its theoretical foundation.

However the document is by no means complete as UPPAAL is still under development. In fact, since this paper was written, a new version of UPPAAL has been implemented, which is currently under internal test and evaluation. Future versions of UPPAAL will extend the current description language with data types such as arrays, records, lists, etc., as well as various primitive operations on data types. In the near future, a new graphical interface will be available, which provides graphical support for not only system descriptions at process level as does the current version, but also hierarchical design and modularization. But that is all for now.

Acknowledgements. The authors wish to thank the people involved in the development of UPPAAL, including: Johan Bengtsson, Palle Christensen, Jesper Jensen, Per Jensen, Kåre Kristoffersen, Fredrik Larsson, and Thomas Sørensen.

References

- Martin Abadi, Leslie Lamport. An old-fashioned recipe for real time. In *Proc. of REX Workshop Real-Time: Theory in Practice*, Lecture Notes in Computer Science 600. Berlin Heidelberg New York: Springer Verlag, 1993
- R. Alur, D. Dill. Automata for modeling real-time systems. *Theoretical Computer Science* 126(2): pp183–236, April 1994
- Richard Bellman. *Dynamic Programming*. Princeton University Press, 1957
- Johan Bengtsson, Palle Christensen, Per Jensen, Kim G. Larsen, Fredrik Larsson, Paul Pettersson, Thomas Sørensen, Wang Yi. UPPAAL: a tool suite for validation and verification of real-time systems. <http://www.docs.uu.se/rtmv/uppaal/uppaal-guide.ps.gz>, 1996
- Johan Bengtsson, David Griffioen, Kåre Kristoffersen, Kim G. Larsen, Fredrik Larsson, Paul Pettersson, Wang Yi. Verification of an audio protocol with bus collision using UPPAAL. In Rajeev Alur, Thomas A. Henzinger, eds., *Proc. of 8th Int. Conf. on Computer Aided Verification*, in Lecture Notes in Computer Science 1102, pp. 244–256. Berlin Heidelberg New York: Springer-Verlag, July 1996
- Johan Bengtsson, Kim G. Larsen, Fredrik Larsson, Paul Pettersson, Wang Yi. UPPAAL – a tool suite for automatic verification of real-Time Systems. In *Proc. of Workshop on Verification and Control of Hybrid Systems III*, Lecture Notes in Computer Science 1066, pp. 232–243. Berlin Heidelberg New York: Springer-Verlag, October 1995
- Johan Bengtsson, Kim G. Larsen, Fredrik Larsson, Paul Pettersson, Wang Yi. UPPAAL in 1995. In *Proc. of the 2nd Workshop on Tools and Algorithms for the Construction and Analysis of Systems*, Lecture Notes in Computer Science 1055, pp. 431–434. Berlin Heidelberg New York: Springer-Verlag, March 1996
- Johan Bengtsson, Fredrik Larsson. UPPAAL a tool for automatic verification of real-time systems. Master's thesis, Uppsala University, 1996
- D. Bosscher, I. Polak, F. Vaandrager. Verification of an audio-control protocol. In *Proc. of Formal Techniques in Real-Time and Fault-Tolerant Systems*, Lecture Notes in Computer Science 863, 1994
- Amar Bouali, Annie Ressouche, Valérie Roy, Robert de Simone. The FC2TOOLS set. In Rajeev Alur, Thomas A. Henzinger, editors, *Proc. of 8th Int. Conf. on Computer Aided Verification*, Lecture Notes in Computer Science 1102, pp. 441–445. Berlin Heidelberg New York: Springer-Verlag, 1996
- Palle Christensen, Thomas Mark Sørensen. Master's thesis, Aalborg University, 1997
- P.R. D'Argenio, J.-P. Katoen, T. Ruys, J. Tretmans. Modeling and verifying a bounded retransmission protocol. In *Proc. of COST 247, International Workshop on Applied Formal Methods in System Design*, 1996. Also available as Technical Report CTIT 96-22, University of Twente, July 1996
- P.R. D'Argenio, J.-P. Katoen, T.C. Ruys, J. Tretmans. The bounded retransmission protocol must be on time! In *Proc. of the 3rd Workshop on Tools and Algorithms for the Construction and Analysis of Systems*, Lecture Notes in Computer Science 1217, pp. 416–431. Berlin Heidelberg New York: Springer-Verlag, 1997
- C. Daws, S. Yovine. Two examples of verification of multirate timed automata with KRONOS. In *Proc. of the 16th IEEE Real-Time Systems Symposium*, pp. 66–75, December 1995
- David Dill. Timing assumptions and verification of finite-state concurrent systems. In J. Sifakis, editor, *Proc. of Automatic Verification Methods for Finite State Systems*, in Lecture Notes in Computer Science 407, pp. 197–212. Berlin Heidelberg New York: Springer-Verlag, 1989
- David Dill. The Mur ϕ Verification System. In Rajeev Alur and Thomas A. Henzinger, eds., *Proc. of 8th Int. Conf. on Computer Aided Verification*, in Lecture Notes in Computer Science 1102, pp. 390–393. Berlin Heidelberg New York: Springer-Verlag, 1996
- W.O.D. Griffioen. Analysis of an Audio Control Protocol with Bus Collision. Master's thesis, University of Amsterdam, Programming Research Group, 1994
- Klaus Havelund, Arne Skou, Kim G. Larsen, Kristian Lund. Formal modeling and analysis of an audio/video protocol: an industrial case study using UPPAAL. Accepted for presentation at the *18th IEEE Real-Time Systems Symposium*, 1997
- Thomas A. Henzinger, Pei-Hsin Ho, Howard Wong-Toi. HYTECH: the next generation. In *Proc. of the 16th IEEE Real-Time Systems Symposium*, pp. 56–65, December 1995
- H.E. Jensen, K.G. Larsen, A. Skou. Modeling and Analysis of a Collision Avoidance Protocol Using SPIN and UPPAAL. In *Proc. of 2nd International Workshop on the SPIN Verification System*, pp. 1–20, August 1996
- Jesper Gravgaard Jensen, Per Stoffer Jensen. Design and Implementation of NewPaal – a modeling language for real-time systems. Master's thesis, Aalborg University, 1997
- Kåre J. Kristoffersen, Francois Larroussinie, Kim G. Larsen, Paul Pettersson, Wang Yi. A compositional proof of a real-time mutual exclusion protocol. In *Proc. of the 7th International Joint Conference on the Theory and Practice of Software Development*, April 1997
- Kim G. Larsen, Hans Hüttel. UPPAAL – an automatic tool for verification of real time and hybrid systems. Seminar slides from *Livslang Uddannelse 96*, 1996. Email: {kg1,hans}@cs.auc.dk.
- Kim G. Larsen, Paul Pettersson, Wang Yi. Compositional and Symbolic Model-Checking of Real-Time Systems. In *Proc. of the 16th IEEE Real-Time Systems Symposium*, p 76–87, December 1995
- Kim G. Larsen, Paul Pettersson, Wang Yi. Diagnostic Model-Checking for Real-Time Systems. In *Proc. of Workshop on Verification and Control of Hybrid Systems III*, Lecture Notes in Computer Science 1066, pp. 575–586. Berlin Heidelberg New York: Springer-Verlag, October 1995
- Kim G. Larsen, Paul Pettersson, Wang Yi. Model-Checking for Real-Time Systems. In *Proc. of Fundamentals of Computation Theory*, Lecture Notes in Computer Science 965, pp. 62–88. Berlin Heidelberg New York: Springer-Verlag, 1995
- Fredrik Larsson, Kim G. Larsen, Paul Pettersson, Wang Yi. Efficient verification of real-time systems: compact data structures and state-space reduction. Accepted for presentation at

- the *18th IEEE Real-Time Systems Symposium*, 1997
28. Magnus Lindahl, Paul Pettersson, Wang Yi. Formal design and analysis of a gear-box controller: an industrial case study using UPPAAL. Technical Report ASTEC 09/97, Dept. of Computer Systems, Uppsala University, August 1997. Available at <http://www.docs.uu.se/docs/rtmv/papers/lpw-astec-97.ps.gz>
 29. Henrik Lönn, Paul Pettersson. Formal Verification of a TDMA Protocol Startup Mechanism. Accepted for presentation at the *Pacific Rim International Symposium on Fault-Tolerant Systems*, December 1997
 30. R. Milner. *Communication and Concurrency*. Englewood Cliffs, NJ: Prentice Hall, 1989
 31. A. Olivero, J. Sifakis, S. Yovine. Using abstractions for the verification of linear hybrids systems. In *Proc. of 7th Int. Conf. on Computer Aided Verification*, Lecture Notes in Computer Science 818, Berlin Heidelberg New York: Springer-Verlag 1994
 32. N. Shankar. Verification of real-time systems using PVS. In *Proc. of 5th Int. Conf. on Computer Aided Verification*, Lecture Notes in Computer Science 697. Berlin Heidelberg New York: Springer-Verlag, 1993
 33. Mihalis Yannakakis, David Lee. An efficient algorithm for minimizing real-time transition systems. In *Proc. of 5th Int. Conf. on Computer Aided Verification*, Lecture Notes in Computer Science 697, pp. 210–224, Berlin Heidelberg New York: Springer-Verlag, 1993
 34. Wang Yi, Paul Pettersson, Mats Daniels. Automatic verification of real-time communicating systems by constraint-solving. In *Proc. of the 7th International Conference on Formal Description Techniques*, 1994

Appendix A: BNF for τ a-format

<i>Ita</i>	→	<i>VarList ProcList Globals</i>
<i>VarList</i>	→	ϵ <i>Channel VarList</i> <i>Var VarList</i>
<i>ProcList</i>	→	<i>Proc</i> <i>Proc ProcList</i>
<i>Globals</i>	→	system <i>IdList</i> ;
<i>Channel</i>	→	urgent chan <i>IdList</i> ; chan <i>IdList</i> ;
<i>Var</i>	→	<i>Type IdList</i> ;
<i>Proc</i>	→	process <i>Id</i> { <i>ProcBody</i> }
<i>IdList</i>	→	<i>Id</i> <i>Id</i> , <i>IdList</i>
<i>ProcBody</i>	→	<i>StateDecls TransDecls</i>
<i>StateDecls</i>	→	state <i>IdList</i> ; commit <i>IdList</i> ; init <i>Id</i> ; state <i>IdList</i> ; init <i>Id</i> ;
<i>Transdecls</i>	→	trans <i>TransList</i> ;
<i>TransList</i>	→	<i>Trans</i> <i>Trans</i> , <i>TransList</i>
<i>Trans</i>	→	<i>Id SInv</i> -> <i>Id</i> { <i>OpG OpS OpA</i> } -> <i>Id</i> { <i>OpG OpS OpA</i> }
<i>SInv</i>	→	ϵ { <i>InvList</i> }
<i>InvList</i>	→	<i>Inv</i> <i>Inv</i> , <i>InvList</i>
<i>Inv</i>	→	<i>Id</i> <= <i>Nat</i> <i>Id</i> < <i>Nat</i>
<i>OpG</i>	→	ϵ guard <i>GuardList</i> ;
<i>OpS</i>	→	ϵ <i>Id!</i> <i>Id?</i>
<i>OpA</i>	→	ϵ assign <i>AssignList</i> ;
<i>GuardList</i>	→	<i>Guard</i> <i>Guard</i> , <i>GuardList</i>
<i>AssignList</i>	→	<i>Assign</i> <i>Assign</i> , <i>AssignList</i>
<i>Type</i>	→	clock int
<i>Assign</i>	→	<i>ClockAssign</i> <i>IntAssign</i>
<i>Guard</i>	→	<i>Id RelOp Nat</i> <i>Id RelOp Id Op Nat</i>
<i>ClockAssign</i>	→	<i>Id</i> := <i>Nat</i>
<i>IntAssign</i>	→	<i>Id</i> := <i>IntExpr</i>
<i>IntExpr</i>	→	<i>Int</i> * <i>Id Op Nat</i> <i>Id Op Nat</i> <i>Id</i> <i>Int</i>
<i>RelOp</i>	→	< <= >= > ==
<i>Op</i>	→	+ -
<i>Id</i>	→	<i>Alpha</i> <i>Id AlphaNum</i>
<i>Nat</i>	→	<i>Num</i> <i>Num Nat</i>
<i>Int</i>	→	<i>Nat</i> - <i>Nat</i>
<i>Alpha</i>	→	A ... Z a ... z
<i>Num</i>	→	0 ... 9
<i>AlphaNum</i>	→	<i>Alpha</i> <i>Num</i> _

Appendix B: BNF for q-format

<i>Prop</i>	→	E <> <i>StateProp</i> A [] <i>StateProp</i>
<i>StateProp</i>	→	<i>AtomicProp</i> (<i>StateProp</i>) not <i>StateProp</i> <i>StateProp</i> or <i>StateProp</i> <i>StateProp</i> and <i>StateProp</i> <i>StateProp</i> imply <i>StateProp</i>
<i>AtomicProp</i>	→	<i>Id . Id</i> <i>Id RelOp Nat</i> <i>Id RelOp Id Op Nat</i>
<i>RelOp</i>	→	< <= >= > ==
<i>Op</i>	→	+ -
<i>Id</i>	→	<i>Alpha</i> <i>Id AlphaNum</i>
<i>Nat</i>	→	<i>Num</i> <i>Num Nat</i>
<i>Alpha</i>	→	A ... Z a ... z
<i>Num</i>	→	0 ... 9
<i>AlphaNum</i>	→	<i>Alpha</i> <i>Num</i> _