

Contents

Introduction

Installing PRISM

- [Instructions](#)
- [Common Problems And Questions](#)

The PRISM Language

- [Introduction](#)
- [Example 1](#)
- [Model Type](#)
- [Modules And Variables](#)
- [Initial States](#)
- [Commands](#)
- [Parallel Composition](#)
- [Local Nondeterminism](#)
- [Module Renaming](#)
- [Example 2](#)
- [Constants](#)
- [Expressions](#)
- [Synchronisation](#)
- [Process Algebra Operators](#)
- [Global Variables](#)
- [Formulas And Labels](#)
- [PRISM Language Files](#)
- [Costs And Rewards](#)

Property Specification

- [Introduction](#)
- [Identifying A Set Of States](#)
- [The P Operator](#)
- [Quantitative Properties](#)
- [The S Operator](#)
- [Reward-based Properties](#)
- [Syntax And Semantics](#)
- [Properties Files](#)

Running PRISM

- [Starting PRISM](#)
- [Loading And Building a Model](#)
- [Debugging Models With The Simulator](#)
- [Exporting The Model](#)
- [Model Checking](#)
- [Approximate Computation Of Quantitative Properties](#)
- [Computing Steady-state And Transient Probabilities](#)

- [Experiments](#)
- [Support For PEPA Models](#)
- [Support For SBML](#)
- [Explicit Model Import](#)

Configuring PRISM

- [Introduction](#)
- [Computation Engines](#)
- [Iterative Numerical Methods](#)
- [Other Options](#)

References

Introduction

PRISM is a *probabilistic model checker*, a tool for the modelling and analysis of systems which exhibit probabilistic behaviour. Probabilistic model checking is a *formal verification* technique. It is based on the construction of a precise mathematical model of a system which is to be analysed. Properties of this system are then expressed formally in temporal logic and automatically analysed against the constructed model.

PRISM has direct support for three types of probabilistic models:

- discrete-time Markov chains (DTMCs),
- Markov decision processes (MDPs),
- continuous-time Markov chains (CTMCs).

Models are supplied to the tool by writing descriptions in the [PRISM language](#), a simple, high-level modelling language.

Properties of these models are written in the PRISM [property specification language](#) which is based on temporal logic. It incorporates several well-known probabilistic temporal logics:

- PCTL (probabilistic computation tree logic),
- CSL (continuous stochastic logic),
- LTL (linear time logic),
- PCTL* (which subsumes both PCTL and LTL).

plus support for [costs/rewards](#) and several other custom features and extensions.

PRISM performs probabilistic model checking to automatically analyse such properties. It also contains a [discrete-event simulation](#) engine for approximate verification.

Installing PRISM

Instructions

Prerequisites

PRISM is known to run on Linux, Windows, Mac OS X and Solaris, including 64-bit variants of several of these operating systems.

You will need Java, more specifically [Sun's distribution of Java](#). To run binary versions of PRISM, you only need the Java Runtime Environment (JRE), not the full Java Development Kit (JDK). Binaries are currently built using Java 6. If you are compiling from source, version 5.0 (sometimes also called version 1.5) or above is sufficient. If you don't know what version of Java you have, type `java -version` from a command prompt.

To compile PRISM from source, you need the Java Development Kit (JDK), GNU make and a C/C++ compiler (e.g. `gcc/g++`). For compilation under Windows, you will need Cygwin. See below for more information.

Installation on Windows

To install the latest versions of PRISM on Windows, just run the self-extracting installer which you [downloaded](#). For older versions (3.0 and earlier) you need to unpack the zip file by hand. In either case, you do **not** need administrator privileges to install PRISM - you just need to have write access to the directory in which you plan to install it.

If requested, the installer will place shortcuts to run PRISM on the desktop and/or start menu. If not, you can run by PRISM double-clicking the file `xprism.bat` (which may just appear as `xprism`) in the `bin` folder of your PRISM folder. If nothing happens, the most likely explanation is that Java is not installed or not in your path. To check, open a command prompt window, navigate to the PRISM directory, type `cd bin`, then `xprism.bat` and examine the resulting error.

If you want to create shortcuts to `xprism.bat` manually, you will find some PRISM icons in the `etc` folder. It is also best to change the "Run" option in the shortcut properties window from "Normal window" to "Minimized".

If you wish to use the command-line version of PRISM on Windows, open a command prompt window and type for example:

```
cd "c:\Program Files\prism-3.0.beta1-win\bin"
prism ../examples/dice/dice.pm
```

You can also edit the file `bin\prism.bat` to allow it to be run from any location. See the instructions within the file for further details.

PRISM is also known to run on 64-bit versions of Windows. Since we currently only provide 32-bit binary versions, this will only work if you are also running a 32-bit version of Java. If this is not possible for some reason, you will have to build PRISM from source using Cygwin (See below).

Problems? See the section "[Common Problems And Questions](#)".

Installation of (non-Windows) binary versions

To ensure compatibility, we recommend that you compile PRISM from source on non-Windows platforms. See below for instructions. However, we do [provide](#) pre-compiled binary distributions for Linux (32-bit) and, depending on availability, other operating systems such as Mac OS X and Solaris.

To install a binary distribution, unpack the tarred/zipped PRISM distribution into a suitable location, enter the directory and run the `install.sh` script, e.g.:

```
gunzip prism-3.2-linux.tar.gz
tar xf prism-3.2-linux.tar
cd prism-3.2-linux
./install.sh
```

You do not need to be root to install PRISM. The install script simply makes some small customisations to the scripts used to launch PRISM. The PRISM distribution is self-contained and can be freely moved/renamed, however if you do so you will need to re-run `./install.sh` afterwards.

To run PRISM, execute either the `xprism` or `prism` script (for the graphical user interface or command-line version, respectively). These can be found in the `bin` directory. These scripts are designed to be run from anywhere and you can

easily create symbolic links or aliases to them. If you want icons to create desktop shortcuts to PRISM, you can find some in the etc directory.

Problems? See the section "[Common Problems And Questions](#)".

Building PRISM from source

To compile PRISM from source code, you will need:

- GNU make (sometimes called gmake)
- a C/C++ compiler (e.g. gcc/g++)
- Sun's [Java Development Kit](#) (JDK), version 5.0 (sometimes called 1.5) or higher

To check what version of Java you have, type `java -version`. To check that you have the development kit, type `javac`. If you get an error message that `javac` cannot be found, you probably do not have the JDK installed (or your path is not set up correctly).

Hopefully, you can build PRISM simply by entering the PRISM directory and running make, e.g.:

```
gunzip prism-3.2-src.tar.gz
tar xf prism-3.2-src.tar
cd prism-3.2-src
make
```

For this process to complete correctly, PRISM needs to be able to determine both the operating system you are using and the location of your Java distribution. If there is a problem with either of these, you will see an error message and will need to specify one or both of these manually, such as in these examples:

```
make OSTYPE=linux
make JAVA_DIR=/usr/java/jdk1.5.0
make OSTYPE=cygwin JAVA_DIR="/cygdrive/c/Program Files/Java/jdk1.5.0"
```

Note the use of double quotes for the case where the directory contains a space. If you don't know the location of your Java installation, try typing `which javac`. If the result is e.g. `/usr/java/jdk1.5.0/bin/javac` then your Java directory is `/usr/java/jdk1.5.0`. Sometimes `javac` will be a symbolic link, in which case use `"ls -l"` to determine the actual location.

It is also possible to set the environment variables `OSTYPE` and `JAVA_DIR` directly or edit their values in the Makefile directly. Note that even when you specify `JAVA_DIR` explicitly (in either way), PRISM still uses the versions of `javac` (and `javah`) that are in your path so make sure this is set up correctly.

Problems? See the section "[Common Problems And Questions](#)".

Building PRISM on 64-bit machines

PRISM has also been successfully compiled under Linux on a variety of 64-bit architectures. In general this should work exactly as described above. The Makefile will try to detect that you are compiling on a 64-bit machine but, if this does not work for some reason, you can override detection by setting `ARCH` to either `amd64` (for AMD/Intel 64) or `ia64` (for Itanium). For example:

```
make ARCH=amd64
```

For 64-bit Macs, see below.

Problems? See the section "[Common Problems And Questions](#)".

Building PRISM on Macs

In most cases, you should be able to build PRISM from source under Mac OS X using the instructions given above. There are, however, a few known issues to be aware of.

The most common problem is when building on recent 64-bit Macs, under which PRISM has some issues with Java. The best advice, currently, is to compile and run a 32-bit version of PRISM. To do this:

1. Make sure you are using a 32-bit version of Java (e.g. through "Java Preferences" under /Applications/Utilities).
2. Make sure you have the very latest version of PRISM [available](#) (for version 3.3, get the final release, not beta1 or beta2).
3. Build PRISM (as above).

If this still does not work, it may be that PRISM is detecting automatically that you are on a 64-bit machine. You can override this as follows:

```
make clean_all
make ARCH=
```

Problems? See the section "[Common Problems And Questions](#)".

Building PRISM from source on Windows using Cygwin

The compilation of PRISM currently relies on a Unix-like environment. On Windows, this can be achieved using the [Cygwin](#) development environment (or alternatively using [MSYS](#)). Once Cygwin is installed, you can proceed as described in the previous section. Note that the PRISM compilation process uses the [MinGW](#) libraries so that the final result is independent of Cygwin at run-time.

One thing to note: please unzip the PRISM distribution from within Cygwin (e.g. using `tar xzf prism-XXX-src.tar.gz`). Don't use a Windows program (Winzip, etc.) since this can cause problems.

Problems? See the section "[Common Problems And Questions](#)".

Building PRISM from source on Windows using MSYS

Compiling from source in [MSYS](#) is less obvious as this environment is currently not directly supported in the makefile. Additionally, MSYS does not handle symlinks in the same way as cygwin does. The first problem is fixed by providing a `OSTYPE` variable to the makefile, whereas the second problem currently has to be solved manually.

```
make OSTYPE=cygwin
```

At some point it will fail, saying that it cannot find the CUDD library, this is due to the failing symlinks. We can solve this by means of a few commands:

```
cd cudd/
rmdir lib/
./setup.sh
cd ..
make OSTYPE=cygwin
./install.sh
```

Problems? See the section "[Common Problems And Questions](#)".

Common Problems And Questions

This section describes some of the most common problems and questions related to the installation and running of PRISM. These are grouped into the following categories:

- Running PRISM on Windows
 - Running PRISM on non-Windows platforms
 - Compiling PRISM
 - Other issues
-

Running PRISM on Windows

When I try to run PRISM on Windows, I double-click the PRISM shortcut but nothing happens.

The most common cause of this is that you either do not have Java installed or the java executable is not in your path. In any case, to determine the exact problem, launch a command shell and navigate to the bin directory inside the directory where you installed PRISM (on recent versions the "PRISM (console)" shortcut installed in the start menu does this for you). Then, type `xprism.bat` and see what error message is displayed.

Running PRISM on non-Windows platforms

When I try to run PRISM, I get an error of the form:

Exception in thread "main" java.lang.NoClassDefFoundError: ...

Check:

- Did you run `install.sh` from the PRISM directory? (non-Windows platforms)
- If you compiled PRISM from source code, are you sure no errors occurred during the process? To check, go into the PRISM directory, type `make clean_all` and then re-compile, checking the output (especially at the end) carefully for any error messages.

When I try to run PRISM, I get an error of the form:

java.lang.UnsatisfiedLinkError: no prism in java.library.path

Check:

- Did you run `install.sh` from the PRISM directory? (non-Windows platforms)
- If you compiled PRISM from source code, are you sure no errors occurred during the process? To check, go into the PRISM directory, type `make clean_all` and then re-compile, checking the output (especially at the end) carefully for any error messages.

Are you on a 64-bit machine? If so, make sure that you are running 64-bit versions of java and javac. (Look for "64-Bit Server VM" in the output of `java -version`).

When I try to run PRISM, I get an error of the form:

Exception in thread "main" java.lang.UnsupportedClassVersionError: Bad version number in .class file

Your version of Java is too old. Install a newer version of Sun Java and then try again.

When I try to run a (Linux) binary version of PRISM, I get an error saying that `libstdc++.so.5` cannot be found or `libstdc++.so.6` cannot be found.

This is usually due to a discrepancy between the version of Linux that was used to build the binary distribution and the version that you are using to run it.

If the error message is about `libstdc++.so.5`, you will just need to install an old version of the `libstdc++` library. This should be quite easy to find for most Linux distributions. On Fedora Core, for example, just type: `yum install compat-libstdc++-33` as root.

If the error message is about `libstdc++.so.6`, you are running an older version of Linux than the binary release was compiled for. The easiest solution is to compile PRISM yourself from the source code version instead.

When I try to run PRISM, I get an error of the form:

```
Exception in thread "main" java.lang.ExceptionInInitializerError
at java.lang.Class.initializeClass(libgcj.so.7)
at prism.PrismSettings.<init>(PrismSettings.java:297)
```

You are not running the [Sun](#) version of Java. You will need to install it.

When I try to run PRISM, I get an error of the form:

```
java.lang.UnsatisfiedLinkError: libprism: ... cannot restore segment prot after reloc:
Permission denied
```

This is likely to be caused by the default settings of [SELinux](#) on newer versions of Linux. Open up the "Security Level Configuration" (on Fedora, for example, this is found under "Administration | Security Level and Firewall" under the main menu or by running `system-config-securitylevel`). Look in the "Compatibility" section of the SELinux Policy settings and make sure "Allow the use of shared libraries with Text Relocation" is ticked. You may need to reboot for changes to take effect.

Do I have to use Sun's version of Java to build/run PRISM?

Currently, this seems to be the case. We will aim to address this in the future.

Compiling PRISM

I can't compile PRISM on Fedora 9. Why not?

This is because Fedora 9 comes with GCC 4.3 by default, which causes some problems. The easiest fix is to first install an older version of GCC (`yum install compat-gcc-34 compat-gcc-34-c++`) and then recompile using this version (`cd prism; make clean; make CPP=g++34`).

When I try to compile PRISM, make seems to get stuck in an infinite loop

This is probably due to the detection of Java failing. Specify the location of your Java directory by hand, e.g. `make JAVA_DIR=/usr/java/jdk1.5.0`. See the [Instructions](#) page for more on this.

When I try to compile PRISM, I get errors of the form:

```
/usr/bin/libtool: for architecture: cputype (16777234) cpusubtype (0) file: -lSystem is not an
object file (not allowed in a library)
```

Are you compiling PRISM on Mac OS X? If so, the likely explanation is that you have upgraded to a new version of Mac OS X but have not upgraded the developer tools (eg. XCode). Upgrade and try again.

When I try to compile PRISM, nothing seems to happen

Perhaps you are not using the GNU version of make. Try typing `make -v` to find out. On some systems, GNU make is called `gmake`.

When I try to compile PRISM, I get errors of the form:

Unexpected end of line seen...

or:

make: Fatal error in reader: Makefile, line 58: Unexpected end of line seen...

Perhaps you are not using the GNU version of make. Try typing `make -v` to find out. On some systems, GNU make is called `gmake`.

When I try to compile PRISM, I get an error of the form:

./setup.sh: line 33: syntax error: unexpected end of file

Are you building on Cygwin? And did you unpack PRISM using WinZip? If so, unpack from Cygwin, using `tar xzf` (or similar) instead.

When I try to compile PRISM, I get an error of the form:

Assembler messages: Fatal error: can't create ../../obj/dd/dd_abstr.o: No such file or directory

Did you unpack PRISM using a graphical tool or file manager? If so, unpack using `tar xzf` (or similar) instead.

When I try to compile PRISM, I get errors of the form:

dirname: extra operand `Files/Java/jdk1.5.0_09/bin/javac' Try `dirname --help' for more information.

This error occurs if the path to your Java distribution contains a space (a common example is when it is somewhere in "Program Files" on Windows). Hopefully, this will be fixed soon. A workaround is to move the java installation to e.g. `C:\java`.

When I try to compile PRISM, I get an error of the form:

/bin/sh: line 43: [: :/cygdrive/c/Program: binary operator expected...

See answer to previous question.

Do I have to use GNU make to build PRISM?

Strictly speaking, no, but you will have to modify the various PRISM Makefiles manually to overcome this.

Can I build PRISM on operating systems other than those currently supported?

PRISM should be suitable for any Unix/Linux variant.

The first thing you will need to do is compile CUDD (the BDD library used by and included in PRISM) on that platform. Fortunately, CUDD has already been successfully built on a large number of operating systems. Have a look at the sample Makefiles we provide (i.e. the files `cudd/Makefile.*`) which are slight variants of the original Makefile provided with CUDD (found here: `cudd/modified/orig/Makefile`). They contain instructions on how to modify it for various platforms. You can then call your new modified makefile something appropriate (`cudd/Makefile.$OSTYPE`) and proceed to build PRISM as usual. To just build CUDD, not PRISM, type `make cuddpackage` instead of `make`.

Next, look at the main PRISM Makefile, in particular, each place where the variable `$OSTYPE` is referred to. Most lines include comments and further instructions. Once you have done this, proceed as usual.

If you do successfully build PRISM on other platforms, please let us know so we can include this information in future releases. Thanks.

Other issues

How do I uninstall PRISM?

If you installed PRISM on Windows using the self-extracting installer, you can uninstall it using the option on the start menu. If you didn't add these shortcuts, just run `uninstall.exe` from the directory where you installed PRISM.

For older versions of PRISM on Windows or on any other platform, simply delete the directory containing it.

The only thing that is not removed via either of these methods is the `.prism` file containing your PRISM settings which is in your home directory (see the section "[Configuring PRISM](#)"). You may wish to retain this when upgrading.

I still have a problem installing/running PRISM. What can I do?

Please post a message in the discussion group (see the [support](#) section of the PRISM website).

The PRISM Language

Introduction

In order to construct and analyse a model with PRISM, it must be specified in the PRISM language, a simple, state-based language, based on the Reactive Modules formalism of Alur and Henzinger [[AH99](#)]. In this section, we describe the PRISM language and present a number of small illustrative examples. A precise definition of the semantics of the language is available from the "[Documentation](#)" section of the PRISM web site. One of the best ways to learn what can be done with the PRISM language is to look at some existing examples. A number of these are included with the tool distribution in the `examples` directory. Many additional examples can be found on the "[Case Studies](#)" section of the [PRISM website](#).

The fundamental components of the PRISM language are *modules* and *variables*. A model is composed of a number of *modules* which can interact with each other. A module contains a number of local *variables*. The values of these variables at any given time constitute the state of the module. The *global state* of the whole model is determined by the *local state* of all modules. The behaviour of each module is described by a set of *commands*. A command takes the form:

```
[ ] guard -> prob_1 : update_1 + ... + prob_n : update_n;
```

The *guard* is a predicate over all the variables in the model (including those belonging to other modules). Each *update* describes a transition which the module can make if the guard is true. A transition is specified by giving the new values of the variables in the module, possibly as a function of other variables. Each update is also assigned a probability (or in some cases a rate) which will be assigned to the corresponding transition.

Example 1

We will use the following simple example to illustrate the basic concepts of the PRISM language. Consider a system comprising two identical processes which must operate under mutual exclusion. Each process can be in one of 3 states: {0,1,2}. From state 0, a process will move to state 1 with probability 0.2 and remain in the same state with probability 0.8. From state 1, it tries to move to the critical section: state 2. This can only occur if the other process is not in its critical section. Finally, from state 2, a process will either remain there or move back to state 0 with equal probability. The PRISM code to describe this system can be seen below. In the next sections, we explain each aspect of the code in turn.

```
// Example 1
// Two process mutual exclusion

mdp

module M1

    x : [0..2] init 0;

    [] x=0 -> 0.8:(x'=0) + 0.2:(x'=1);
    [] x=1 & y!=2 -> (x'=2);
    [] x=2 -> 0.5:(x'=2) + 0.5:(x'=0);

endmodule

module M2

    y : [0..2] init 0;

    [] y=0 -> 0.8:(y'=0) + 0.2:(y'=1);
    [] y=1 & x!=2 -> (y'=2);
    [] y=2 -> 0.5:(y'=2) + 0.5:(y'=0);

endmodule
```

The PRISM Language: Example 1

Model Type

The PRISM language can be used to describe three types of probabilistic models: discrete-time Markov chains (DTMCs), continuous-time Markov chains (CTMCs) and Markov decision processes (MDPs). To indicate which type is being described, a PRISM model includes one of the keywords **dtmc**, **ctmc** or **mdp**. This is typically at the very start of the file (as in [Example 1](#)), but can actually occur anywhere in the file (except inside modules and other declarations). The keywords **probabilistic**, **stochastic** and **nondeterministic** can be used as alternatives for **dtmc**, **ctmc** and **mdp**, respectively. If no such model type declaration is included, the model is by default assumed to be an MDP.

Modules And Variables

The [previous example](#) uses two modules, M1 and M2, one representing each process. A module is specified as:

```
module name ... endmodule
```

The definition of a module contains two parts: its variables and its guards. In this example, each module has one integer variable with range [0..2]. A variable declaration looks like:

```
x : [0..2] init 0;
```

Notice that the initial value of the variable is also specified. See the next subsection for more information about this subject. Boolean variables can also be used, e.g.:

```
b : bool init false;
```

The names given to modules and variables are referred to as *identifiers*. Identifiers can be made up of letters, digits and the underscore character, but cannot begin with a digit, i.e. they must satisfy the regular expression `[A-Za-z_][A-Za-z0-9_]*`, and are case-sensitive. Furthermore, identifiers cannot be any of the following, which are all reserved keywords in PRISM: **A, bool, const, ctmc, C, double, dtmc, E, endinit, endmodule, endrewards, endsystem, false, formula, func, F, global, G, init, I, int, label, max, mdp, min, module, X, nondeterministic, Pmax, Pmin, P, probabilistic, prob, rate, rewards, Rmax, Rmin, R, S, stochastic, system, true, U, W,**

Initial States

The state space of a probabilistic model described in the PRISM language is the set of all possible valuations of its variables. The set of initial states of the model can be specified in one of two ways. The most common approach is simply to specify an initial value for each variable in the model. This is done in the declaration of each variable, as illustrated in [Example 1](#) and as discussed in the previous section. If the initial value of a variable is omitted from its declaration, it is taken to be the minimum value of the variable's range.

This approach results in a single initial state. It is also possible to construct a model with multiple initial states. This is done using the **init...endinit** construct, which can be placed anywhere in the file except within a module definition, and removing any initial values from variable declarations. Between the **init** and **endinit** keywords, there should be a predicate over all the variables of the model. Any state which satisfies this predicate is an initial state.

Consider [Example 1](#). As it stands, there is a single initial state $(0, 0)$ (i.e. $x=0$ and $y=0$). If we remove the **init 0** part of both variable declarations and add the following to the end of the file:

```
init x=0 endinit
```

there will be three initial states: $(0, 0)$, $(0, 1)$ and $(0, 2)$. Similarly, we could instead add:

```
init x+y=1 endinit
```

in which case there would be two initial states: $(0, 1)$ and $(1, 0)$.

Commands

The behaviour of each module is described by commands, comprising a guard and one or more updates. The first command of module M1 in our example is:

```
[] x=0 -> 0.8:(x'=0) + 0.2:(x'=1);
```

The guard $x=0$ indicates that this describes the behaviour of the module when the variable x has value 0. The updates $(x'=0)$ and $(x'=1)$ and their associated probabilities state that the value of x will remain at 0 with probability 0.8 and change to 1 with probability 0.2. For DTMCs and MDPs, we require that the probabilities on the right hand side of a command sum to one. Note that the inclusion of updates in parentheses, e.g. $(x'=1)$, is essential. While older versions of PRISM did not report the absence of parentheses as an error, newer versions do.

The second command:

```
[] x=1 & y!=2 -> (x'=2);
```

illustrates that guards can contain constraints on any variable, not just the ones in that module, i.e. the behaviour of one module can depend on the state of another. Updates, however, can only specify values for variables belonging to the module. The command above also shows that, when there is a single update with probability 1, the $1.0:$ can be omitted.

If a module has more than one variable, updates describe the new value for each of them. For example, if it had two variables x_1 and x_2 , a possible command would be:

```
[ ]  $x_1=0 \ \& \ x_2>0 \rightarrow 0.5:(x_1'=1)\&(x_2'=x_2+1) + 0.5:(x_1'=2)\&(x_2'=x_2-1);$ 
```

Notice that elements of the updates are concatenated with $\&$ and that each element must be bracketed individually. If an update does not give a new value for a local variable, it is assumed not to change. As a special case, the keyword **true** can be used to denote an update where no variable's value changes.

Parallel Composition

The probabilistic model corresponding to a PRISM language description is constructed as the parallel composition of its modules. In every state of the model, there is a set of commands (belonging to any of the modules) which are enabled, i.e. whose guards are satisfied in that state. The choice between which command is performed (i.e. the scheduling) depends on the model type. For an MDP, as in Example 1, the choice is *nondeterministic*. By way of example, consider state $(0,0)$ (i.e. $x=0$ and $y=0$). In the MDP, there would be a nondeterministic choice between two probability distributions:

- $0.8:(0,0) + 0.2:(1,0)$ (module M1 moves)
- $0.8:(0,0) + 0.2:(0,1)$ (module M2 moves)

For a DTMC, the choice is *probabilistic*: each enabled command is selected with equal probability. If Example 1 was a DTMC, then in state $(0,0)$ of the model the following probability distribution would result:

- $0.8:(0,0) + 0.1:(1,0) + 0.1:(0,1)$

For a CTMC, the choice is modelled as a "race" between transitions.

Local Nondeterminism

PRISM models which are MDPs can also exhibit *local nondeterminism*, which allows the modules themselves to make nondeterministic choices. In Example 1, we can make the probabilistic choice in the first state of module M1 nondeterministic by replacing the command:

```
[ ]  $x=0 \rightarrow 0.8:(x'=0) + 0.2:(x'=1);$ 
```

with the commands:

```
[ ]  $x=0 \rightarrow (x'=0);$   
[ ]  $x=0 \rightarrow (x'=1);$ 
```

Assuming we do the same for module M2, in state $(0,0)$ of the MDP there will be a nondeterministic choice between three (trivial) probability distributions:

- $1.0:(0,0)$
- $1.0:(1,0)$
- $1.0:(0,1)$

PRISM also permits local nondeterminism in models which are DTMCs, although the nondeterministic choice is randomised when the parallel composition of the modules occurs. Since the appearance of nondeterminism in a DTMC is often the result of a user error in the model specification, PRISM displays a warning when local nondeterminism is detected in a DTMC.

Module Renaming

PRISM also supports *module renaming*, which allows duplication of modules. In Example 1, module M2 is identical to module

M1 so we can in fact replace its entire definition with:

```
module M2 = M1 [ x=y, y=x ] endmodule
```

Note that this renaming is done at a textual level, so any identifiers (including [action labels](#), [constants](#) and [functions](#)) used in the module definition can be changed in this way. All of the variables in the module being renamed (in this case, just *x*) must be renamed to new, unused names.

Example 2

We now introduce a second example, which is a model of an *N*-place queue of jobs and a server which removes jobs from the queue and processes them. The PRISM code can be found below.

```
// Example 2
// N-place queue + server

ctmc

const int N = 10;
const double mu = 1/10;
const double lambda = 1/2;
const double gamma = 1/3;

module queue
  q : [0..N];

  [] q<N -> mu:(q'=q+1);
  [] q=N -> mu:(q'=q);
  [serve] q>0 -> lambda:(q'=q-1);
endmodule

module server
  s : [0..1];

  [serve] s=0 -> 1:(s'=1);
  [] s=1 -> gamma:(s'=0);
endmodule
```

The PRISM Language: Example 2

As can be seen from the start of the file, the model is a continuous-time Markov chain (CTMC). One of the main differences, therefore, is that transitions of modules are labelled with (positive-valued) *rates*, rather than probabilities. The notation used in commands, however, to associate rates to transitions is identical to that used to assign probabilities in DTMCs and MDPs:

```
rate_1:update_1 + rate_2:update_2 + ...
```

Example 2 also introduces a number of other PRISM language concepts: constants, action labels and synchronisation. These are described in the following sections.

Constants

PRISM also supports the use of *constants*, as seen in [Example 2](#). Constants can be integers, doubles or Booleans and can be defined using literal values or as constant expressions (including in terms of each other) using the **const** keyword. For example:

```
const int radius = 12;
const double pi = 3.141592;
const double area = pi * radius * radius;
const bool yes = true;
```

The identifiers used for their names are subject to the same rules as [variables](#).

Constants can be used anywhere that a constant value would be expected, such as the lower or upper range of a variable (e.g. N in Example 2), the probability or rate associated with an update (mu in Example 2), or anywhere in a guard or update. As will be described later constants can also be left undefined and specified later, either to a single value or a range of values, using [experiments](#).

Note: For the sake of backward-compatibility, the notation used in earlier versions of PRISM (**const** for **const int** and **rate** or **prob** for **const double**) is still supported.

Expressions

The definition of the area constant, in the example above, uses an *expression*. We now define more precisely what types of expression are supported by PRISM. Expressions can contain literal values (12, 3.141592, **true**, **false**, etc.), identifiers (corresponding to variables, constants, etc.) and operators from the following list:

- - (unary minus)
- *, / (multiplication, division)
- +, - (addition, subtraction)
- <, <=, >=, > (relational operators)
- =, != (equality operators)
- ! (negation)
- & (conjunction)
- | (disjunction)
- => (implication)
- ? (condition evaluation: condition ? a : b means "if condition is true then a else b")

All of these operators except ? are left associative (i.e. they are evaluated from left to right). The precedence of the operators is as found in the list above, most strongly binding operators first. Operators on the same line (e.g. + and -) are of equal precedence.

The notation for expressions is hence essentially equivalent to that of C/C++ or Java. One notable exception to this is that the division operator / always performs floating point, not integer, division, i.e. the result of 22/7 is 3.142857... not 3. All expressions must evaluate correctly in terms of type (integer, double or Boolean).

Built-in Functions

Expressions can make use of several built-in functions:

- min(...) and max(...), which select the minimum and maximum value, respectively, of two or more numbers;
- floor(x) and ceil(x), which round x down and up, respectively, to the nearest integer;
- pow(x,y) which computes x to the power of y;
- mod(i,n) for integer modulo operations;
- log(x,b), which computes the logarithm of x to base b.

Examples of their usage are:

```

min(x+1, x_max)
max(a,b,c)
floor(13.5)
ceil(13.5)
pow(2, 8)
pow(9.0, 0.5)
mod(1977, 100)
log(123, 2.71828183)

```

For compatability with older versions of PRISM, all functions can also be expressed via the **func** keyword, e.g. `func(floor, 13.5)`.

Use of Expressions

Expressions can be used in a wide range of places in a PRISM language description, e.g.:

- constant definitions
- lower/upper bounds and initial values for variables
- guards
- probabilities/rates
- updates

This allows, for example, the probability in a command to be dependent on the current state:

```
[ ] (x>=1 & x<=10) -> x/10 : (x'=max(1,x-1)) + 1-x/10 : (x'=min(10,x+1))
```

Synchronisation

Another feature of PRISM introduced in Example 2 is *synchronisation*. In the style of many process algebras, we allow commands to be labelled with *actions*. These are placed inside the square brackets which mark the start of the command, for example serve in this command from [Example 2](#):

```
[serve] q>0 -> lambda:(q'=q-1);
```

These actions can be used to force two or more modules to make transitions simultaneously (i.e. to *synchronise*). For example, in state (3,0) (i.e. q=3 and s=0), the composed model can move to state (2,1), synchronising over the serve action. The rate of this transition is equal to the product of the two individual rates (in this case, $\lambda * 1 = \lambda$). The product of two rates does not always meaningfully represent the rate of a synchronised transition. A common technique, as seen here, is to make one action *passive*, with rate 1 and one action *active*, which actually defines the rate for the synchronised transition. By default, all modules are combined using the standard CSP parallel composition (i.e. modules synchronise over all their common actions).

Process Algebra Operators

To make the concept of synchronisation described above more powerful, PRISM allows you to define precisely the way in which the set of modules are composed in parallel. This is specified using the **system** ... **endsystem** construct, placed at the end of the model description, which should contain a process-algebraic expression. This expression should feature each module exactly once, and can use the following (CSP-based) operators:

- M1 || M2 : alphabetised parallel composition of modules M1 and M2 (synchronising on only actions appearing in both M1 and M2)
- M1 ||| M2 : asynchronous parallel composition of M1 and M2 (fully interleaved, no synchronisation)
- M1 |[a,b,...]| M2 : restricted parallel composition of modules M1 and M2 (synchronising only on actions from the set {a, b,...})

- $M / \{a, b, \dots\}$: hiding of actions $\{a, b, \dots\}$ in module M
- $M \{a \leftarrow b, c \leftarrow d, \dots\}$: renaming of actions a to b , c to d , etc. in module M .

The first two types of parallel composition ($||$ and $|||$) are associative and can be applied to more than two modules at once. When evaluating the expression, the hiding and renaming operators bind more tightly than the three parallel composition operators. No other rules of precedence are defined and parentheses should be used to specify the order in which modules are composed.

Some examples of expressions which could be included in the **system ... endsystem** construct are as follows:

- $(\text{station1} ||| \text{station2} ||| \text{station3}) \text{ } |[\text{serve}] | \text{server}$
- $((P1 \text{ } |[a]| P2) / \{a\}) || Q$
- $((P1 \text{ } |[a]| P2) \{a \leftarrow b\}) |[b]| Q$

When no parallel composition is specified by the user, PRISM implicitly assumes an expression of the form $M1 || M2 || \dots$ containing all of the modules in the model. For a more formal definition of the process algebra operators described above, check the semantics of the PRISM language, available from the "[Documentation](#)" section of the PRISM web site.

PRISM is also able to [import](#) model descriptions written in (a subset of) the stochastic process algebra [PEPA](#) [[Hil96](#)].

Global Variables

In addition to the local variables belonging to each module, a PRISM model can also include *global variables*, which can be written to, as well as read, by all modules. Like local variables, these can be integers or Booleans. Global variables are declared in identical fashion to a module's local variables, except that the declaration must not be inside the definition of any module. Some example declarations are as follows:

```
global g : [1..10];
global b : bool init true;
```

A global variable can be modified by any module and provides another way for modules to interact. An important restriction on the use of global variables is the fact that commands which synchronise with other modules (i.e. those with an action label attached; see the section "[Synchronisation](#)") cannot modify global variables. PRISM will detect this and report an error.

Formulas And Labels

PRISM models can include *formulas* which are used to avoid duplication of code. A formula comprises a name (an identifier) and an [expression](#). The formula name can then be used as shorthand for the expression anywhere an expression might usually be accepted. A formula is defined as follows:

```
formula lfree = p2=0..4,6,10;
```

It can then be used anywhere within that file, as for example in this command:

```
[ ] p1=2 & lfree -> (p1'=4);
```

The effect is exactly as if the following had been typed:

```
[ ] p1=2 & (p2=0..4,6,10) -> (p1'=4);
```

Formulas defined in a model can also be used when specifying its properties.

Formulas and renaming

During parsing of the model, expansion of formulas is done before module renaming so, if a module which uses formulas is renamed to another module, it is the contents of the formula which will be renamed, not the formula itself.

Labels

PRISM models can also contain *labels*. These are a way of identifying sets of states that are of particular interest. Labels can only be used when specifying [properties](#) but, for convenience, can be defined in model files as well as property files.

Labels differ from formulas in two other ways: firstly, they must be of Boolean type; secondly, they are written using quotation marks ("..."), as illustrated in the following example:

```
label "safe" = temp<=100 | alarm=true;
label "fail" = temp>100 & alarm=false;
```

PRISM Language Files

Files containing model descriptions written in the PRISM language can contain any amount of white space (spaces, tabs, new lines, etc.), all of which is ignored when the file is parsed by the tool. Comments can also be used included in files in the style of the C programming language, by preceding them with the characters `//`. This is illustrated by the PRISM language examples from earlier in this section. By convention, PRISM model files which describe MDPs, DTMCs and CTMCs are given the extensions `.nm`, `.pm` and `.sm`, respectively.

Costs And Rewards

PRISM now includes support for the specification and analysis of properties based on *costs* and *rewards*. This means that PRISM can be used to reason, not just about the probability that a model behaves in a certain fashion, but about a wider range of quantitative measures relating to model behaviour. For example, PRISM can be used to compute properties such as "expected time", "expected number of lost messages" or "expected power consumption". The implementation of cost- and reward-based techniques in the tool is only partially completed and is still ongoing. If you have questions, comments or feature-requests relating to this functionality, please feel free to contact the PRISM team about this.

The basic idea is that probabilistic models (of all three types) developed in PRISM can be augmented with costs or rewards: real values associated with certain states or transitions of the model. In fact, since there is no practical distinction between costs and rewards (except that costs are generally perceived to be "bad" and rewards to be "good"), PRISM only supports rewards. The user is, however, free to interpret the values however they choose.

In this section, we describe how models described in the PRISM language can be augmented with rewards. Later, we will discuss how to express properties that relate to these rewards. Rewards are associated with models using **rewards** ... **endrewards** constructs, which can appear anywhere in a model file except within a module definition. These constructs contains one or more *reward items*. Consider the following simple example:

```
rewards
  true : 1;
endrewards
```

This assigns a reward of 1 to every state of the model. It comprises a single reward item, the left part of which (**true**) is a guard and the right part of which (1) is a reward. States of the model which satisfy the predicate in the guard are assigned the corresponding reward. More generally, state rewards can be specified using multiple reward items, each of the form `guard : reward;`, where `guard` is a predicate (over all the variables of the model) and `reward` is an expression (containing any variables, constants, etc. from the model). For example:

```
rewards
  x=0 : 100;
  x>0 & x<10 : 2*x;
  x=10 : 100;
endrewards
```

assigns a reward of 100 to states satisfying $x=0$ or $x=10$ and a reward of $2*x$ to states satisfying $x>0$ & $x<10$. Note that a single reward item can assign different rewards to different states, depending on the values of model variables in each one. Any states which do not satisfy the guard of any reward item will have no reward assigned to them. For states which satisfy multiple guards, the reward assigned to the state is the sum of the rewards for all the corresponding reward items.

Rewards can also be assigned to transitions of a model. These are specified in a similar fashion to state rewards, within the **rewards ... endrewards** construct. Reward items describing transition rewards are of the form `[action] guard : reward;`, the interpretation being that transitions from states which satisfy the guard `guard` and are labelled with the action `action` acquire the reward `reward`. For example:

```
rewards
  [] true : 1;
  [a] true : x;
  [b] true : 2*x;
endrewards
```

assigns a reward of 1 to all transitions in the model with no action label, and rewards of x and $2*x$ to all transitions labelled with actions `a` and `b`, respectively.

As is the case for states, multiple reward items can specify rewards for a single transition, in which case the resulting reward is the sum of all the individual rewards. A model description can specify rewards for both states and transitions. These are all placed together in a single **rewards ... endrewards** construct.

A PRISM model can have multiple reward structures. Optionally, these can be given labels such as in the following example:

```
rewards "total_time"
  true : 1;
endrewards

rewards "num_failures"
  [fail] true : 1;
endrewards
```

Property Specification

Introduction

In order to analyse a probabilistic model which has been specified and constructed in PRISM, it is necessary to identify one or more *properties* of the model which can be evaluated by the tool. PRISM's *property specification language* subsumes several well-known probabilistic temporal logics, including PCTL, CSL, LTL and PCTL*. PCTL is used for specifying properties of DTMCs and MDPs; CSL is an extension of PCTL for CTMCs; LTL and PCTL* can be used to specify properties of DTMCs and MDPs (or untimed properties of CTMCs).

In fact, PRISM also supports numerous additional customisations and extensions of these two logics. Full details of the property specifications permitted in PRISM are provided in the following sections. The presentation given here is relatively informal. For the precise syntax and semantics of the various logics, see [HJ94],[BdA95] for PCTL, [ASSB96],[BKH99] for CSL and, for example, [Bai98] for LTL and PCTL*. You can also find various pointers to useful papers in the [About](#) and [Documentation](#) sections of the PRISM website.

Before discussing property specifications in more detail, it is perhaps instructive to first illustrate some typical examples of properties which PRISM can handle. The following are a selection of such properties. In each case, we give both the PRISM syntax and a natural language translation:

```
P>=1 [ F "terminate" ]
```

"the algorithm eventually terminates successfully with probability 1"

```
"init" => P<0.1 [ F<=100 num_errors > 5 ]
```

"from any initial state, the probability that more than 5 errors occur within the first 100 time units is less than 0.1"

```
"down" => P>0.75 [ !"fail" U[1,2] "up" ]
```

"when a shutdown occurs, the probability of system recovery being completed in between 1 and 2 hours without further failures occurring is greater than 0.75"

```
S<0.01 [ num_sensors < min_sensors ]
```

"in the long-run, the probability that an inadequate number of sensors are operational is less than 0.01"

Note that the above properties are all assertions, i.e. ones to which we would expect a "yes" or "no" answer. This is because all references to probabilities are associated with an upper or lower bound which can be checked to be either true or false. In PRISM, we can also directly specify properties which evaluate to a numerical value, e.g.:

```
P=? [ !proc2_terminate U proc1_terminate ]
```

"the probability that process 1 terminates before process 2 does"

```
Pmax=? [ F<=T messages_lost > 10 ]
```

"the maximum probability that more than 10 messages have been lost by time T"

```
S=? [ queue_size / max_size > 0.75 ]
```

"the long-run probability that the queue is more than 75% full"

Furthermore, PRISM makes it easy to combine such properties into more complex expressions, compute their values for a range of parameters and plot graphs of the results using [experiments](#). This is often a very useful way of identifying interesting patterns or trends in the behaviour of a system. See the [Case Studies](#) section of the [PRISM website](#) for many examples of this kind of analysis.

Identifying A Set Of States

One of the most fundamental tasks when specifying properties of a model is to identify particular sets or classes of states of the model. For example, to verify a property such as "the algorithm eventually terminates successfully with probability 1", it is first necessary to identify the states of the model which correspond to situations where "the algorithm has terminated successfully". In terms of the way temporal logics are usually presented, these correspond to *atomic propositions*.

In PRISM, this is achieved simply by writing an [expression](#) in the PRISM language which evaluates to a Boolean value. This expression will typically contain references to variables (and constants) from the model to which it relates. The set of states corresponding to this expression is those for which it evaluates to true. We say that the expression is "satisfied" in these states.

For example, in the property given above:

```
"init" => P<0.1 [ F<=100 num_errors > 5 ]
```

the expression `num_errors > 5` is used to identify states of the model where more than 5 errors have occurred.

It is also common to use [labels](#) to identify states in this way, like "terminate" in the example:

```
P>=1 [ F "terminate" ]
```

Properties can refer to [labels](#) either from the model to which the property relates, or included in the same [properties file](#).

The P Operator

One of the most important operators in the PRISM property specification language is the **P** operator, which is used to reason about the probability of an event's occurrence. This operator was originally proposed in the logic PCTL but also features in the other logics supported by PRISM, such as CSL. The **P** operator is applicable to all three model types: DTMCs, MDPs and CTMCs.

Informally, the property:

```
P bound [ pathprop ]
```

is true in a state s of a DTMC, MDP or CTMC if "the probability that path property `pathprop` is satisfied by the paths from state s meets the bound `bound`". A typical example of a bound would be:

```
P>0.98 [ pathprop ]
```

which means: "the probability that `pathprop` is satisfied by the paths from state s is greater than 0.98". More precisely, `bound` can be any of `>=p`, `>p`, `<=p` or `<p`, where p is a PRISM language expression evaluating to a double in the range $[0,1]$.

The types of path property supported by PRISM and their semantics are discussed below. PRISM also supports *quantitative* properties of the form `P=?` which are covered [later](#) in this section.

MDPs

For MDPs, some additional clarifications are necessary. Whereas for DTMCs and CTMCs, probability measures over paths are well defined (see e.g. [KSK66] and [BKH99], respectively), for MDPs a probability measure can only be feasibly defined once all nondeterminism has been removed. Hence, the actual meaning of the property `P bound [pathprop]` for an MDP is taken to be "the probability that `pathprop` is satisfied by the paths from state s meets the bound `bound` for *all possible resolutions of nondeterminism*". This means that, for an MDP, properties using the **P** operator actually reason about the *minimum* or *maximum* probability, over all possible resolutions of nondeterminism, that a certain type of behaviour is observed. This depends on the bound attached to the **P** operator: a lower bound (`>` or `>=`) relates to minimum probabilities and an upper bound (`<` or `<=`) to maximum probabilities.

Path properties

PRISM supports a wide range of path properties that can be used with the **P** operator. A path property is a formula that evaluates to either true or false for a single path in a model. Here, we review some of the simpler properties that feature a single *temporal operator*, as used for example in the logics PCTL and CSL. Later, we briefly describe how PRISM also supports more complex LTL-style path properties.

The basic different types of path property that can be used inside the **P** operator are:

- **X** : "next"
- **U** : "until"
- **F** : "eventually" (sometimes called "future")

- **G** : "always" (sometimes called "globally")
- **W** : "weak until"
- **R** : "release"

In the following sections, we describe each of these *temporal operators*. We then discuss the (optional) use of time bounds with these operators. Finally, we also discuss LTL-style path properties.

"Next" path properties

The property **X prop** is true for a path if prop is true in its second state, An example of this type of property, used inside a **P** operator, is:

```
P<0.01 [ X y=1 ]
```

which is true in a state if "the probability of the expression $y=1$ being true in the next state is less than 0.01".

"Until" path properties

The property **prop1 U prop2** is true for a path if prop2 is true in some state of the path and prop1 is true in all preceding states. A simple example of this would be:

```
P>0.5 [ z<2 U z=2 ]
```

which is true in a state if "the probability that z is eventually equal to 2, and that z remains less than 2 up until that point, is greater than 0.5".

"Eventually" path properties

The property **F prop** is true for a path if prop eventually becomes true at some point along the path. The **F** operator is in fact a special case of the **U** operator (you will often see **F prop** written as **true U prop**). A simple example is:

```
P<0.1 [ F z>2 ]
```

which is true in a state if "the probability that z is eventually greater than 2 is less than 0.1".

"Globally" path properties

Whereas the **F** operator is used for "reachability" properties, **G** represents "invariance". The property **G prop** is true of a path if prop remains true at all states along the path. Thus, for example:

```
P>=0.99 [ G z<10 ]
```

states that, with probability at least 0.99, z never exceeds 10.

"Weak until" and "release" path properties

Like **F** and **G**, the operators **W** and **R** are derivable from other temporal operators.

Weak until (a **W b**), which is equivalent to $(a \text{ U } b) \mid G a$, requires that a remains true until b becomes true, but does not require that b ever does become true (i.e. a remains true forever). For example, a weak form of the until example used above is:

```
P>0.5 [ z<2 W z=2 ]
```

which states that, with probability greater than 0.5, either z is always less than 2, or it is less than 2 until the point where z is 2.

Release $(a \text{ R } b)$, which is equivalent to $!(\neg a \text{ U } \neg b)$, informally means that b is true until a becomes true, or b is true forever.

"Bounded" variants of path properties

All of the temporal operators given above, with the exception of X , have "bounded" variants, where an additional time bound is imposed on the property being satisfied. Since in DTMCs and MDPs, time progresses in discrete steps, whereas CTMCs model real (continuous) time, we treat these two cases separately for this operator.

For a DTMC or MDP, the time interval specification must be of the form " $\leq t$ " where t is a PRISM expression evaluating to a non-negative integer. A bounded until property $\text{prop1 U} \leq t \text{ prop2}$, for example, is satisfied along a path if prop2 becomes true within t steps and prop1 is true in all states before that point. A typical example of this would be:

$P \geq 0.98 \ [\text{ } y < 4 \text{ U} \leq 7 \text{ } y = 4 \]$

which is true in a state if "the probability of y first exceeding 3 within 7 time steps is greater than or equal to 0.98". Similarly:

$P \geq 0.98 \ [\text{ } F \leq 7 \text{ } y = 4 \]$

is true in a state if "the probability of y being equal to 4 within 7 time steps is greater than or equal to 0.98" and:

$P \geq 0.98 \ [\text{ } G \leq 7 \text{ } y = 4 \]$

is true if the probability of y staying equal to 4 for 7 time steps is at least 0.98.

The time bound can be an arbitrary (constant) expression, but note that you may need to bracket it, as in the following example:

$P \geq 0.98 \ [\text{ } G \leq (2 * k + 1) \text{ } y = 4 \]$

In the context of a CTMC, the time interval specification time can take any of the three forms: $\geq t$, $\leq t$ or $[t1, t2]$, where t , $t1$ and $t2$ are PRISM expressions evaluating to a non-negative doubles, and $t1$ is no greater than $t2$. In each case, time defines an interval of real values in which the path property must be true.

For example:

$P \geq 0.25 \ [\text{ } y < 1 \text{ U} \leq 6.5 \text{ } y > 1 \]$

means that the probability of y being greater than 1 within 6.5 time-units (and remaining less than or equal to 1 at all preceding time-points) is at least 0.25.

$P < 0.4 \ [\text{ } F \geq 5.5 \text{ } y > 1 \]$

states that the probability of y exceeding 1 at some point after 5.5 time-units have elapsed is less than 0.4, and:

$P > 0 \ [\text{ } G[5.5, 6.5] \text{ } y > 1 \]$

says that the probability that y exceeds 1 for the whole time interval $[5.5, 6.5]$ is greater than zero".

Transient probabilities

We can also use the bounded **F** operator to refer to a single time instant, e.g.:

```
P<0.01 [ F[10,10] y=6 ]
```

refers to the probability of y being 6 at time instant 10.

LTL-style path properties

PRISM also supports probabilistic model checking of the temporal logic LTL (and, in fact, PCTL*). LTL provides a richer set of path properties for use with the **P** operator, by permitting temporal operators to be combined. Here are a few examples of properties expressible using this functionality:

```
P>0.99 [ F ( "request" & (X "ack") ) ]
```

"with probability greater than 0.99, a request is eventually received, followed immediately by an acknowledgement"

```
P>=1 [ G F "send" ]
```

"a message is sent infinitely often with probability 1"

```
P<0.01 [ F G ("error" & !"repair") ]
```

"the probability of an error occurring that is never repaired is less than 0.01" occurs?"

Quantitative Properties

It is very often useful to take a *quantitative* approach to probabilistic model checking, computing the actual probability that some behaviour of a model is observed, rather than just verifying whether or not the probability is above or below a given bound. Hence, PRISM allows the **P** operator to take the following form:

```
P=? [ pathprop ]
```

These properties return a numerical rather than a Boolean value. The **S** and **R** operators, discussed later, can also be used in this way.

Note that the probability bound on a **P** operator can only be replaced with **=?** if it is the outermost operator of the property in which it appears; otherwise the semantics are not well defined. Note also that, for MDPs, the situation is slightly more complex. Probabilities for an MDP can only be computed once the nondeterminism has been resolved. Hence, PRISM actually computes either the *minimum* or *maximum* probability of a path property being satisfied, quantifying over all possible resolutions (i.e. the best and worst cases). Therefore, for MDPs we have two possible types of property:

```
Pmin=? [ pathprop ]  
Pmax=? [ pathprop ]
```

which return the minimum and maximum probabilities, respectively.

All of these operators return a single numerical value. In the simplest case, where the model being verified has a single initial state, the numerical value returned is the one corresponding to that state. Hence, for example:

```
P=? [ F x=5&y=5 ]
```

returns the probability of, from the initial state, reaching a state satisfying $x=5 \wedge y=5$. It is also possible, however, to obtain the probability for an arbitrary state, by specifying an atomic proposition, true only in that state, inside braces $\{ \dots \}$ before the closing $]$ bracket. This is known as a *filter*. For example, if the model contains two variables, X and y , then:

```
P=? [ F x=5&y=5 {x=1&y=2} ]
```

returns the probability of, from the state $(1, 2)$ (i.e. $x=1$ and $y=2$), reaching a state satisfying $x=5 \wedge y=5$. It is possible of course that the expression in the filter satisfies more than one state. If this the case, by default the first such state (lexicographically) is chosen. In this situation, PRISM will display a warning to notify you that the expression in the filter matched multiple states and will display what the first matching state is. The same approach is taken in the situation where no filter is given but the model contains multiple initial states. In this case, the numerical result obtained is for the first (lexicographically speaking) of the set of initial states. Note that if the expression in the filter is not satisfied by any states of the model, an error is reported.

It is also possible to request either the minimum and maximum value from a set of values. For example:

```
P=? [ F x=5&y=5 {y=2}{min} ]  
P=? [ F x=5&y=5 {y=2}{max} ]
```

return the minimum and maximum probability, respectively, of reaching a state satisfying $x=5 \wedge y=5$ from all the states satisfying $y=2$. In addition, PRISM will report the states of the model in which the minimum or maximum probability is attained. You can in fact also request that both the minimum and maximum value are computed simultaneously using, for example, $\{y=2\}\{\min\}\{\max\}$. In this case, both values will be reported, but the actual return value of the property will be the minimum value.

Finally, note that filters can also be used for classical (bounded) **P** properties. This has no bearing on the result of model checking a property; it simply causes the probabilities (where non-zero) for all states satisfying the expression in the filter to be printed during model checking (in the log, if using the GUI, or to the screen, if using the command-line). For example:

```
P>0.5 [ F x=5&y=5 {y=2} ]
```

will return a Boolean value, depending whether the property is true in all states of the model or not, but the probability (where non-zero) of path property $F \ x=5 \wedge y=5$ being satisfied will be displayed for all states where $y=2$.

The S Operator

The **S** operator is used to reason about the *steady-state* behaviour of a model, i.e. its behaviour in the *long-run* or *equilibrium*. Although this could in principle relate to all three model types, PRISM currently only provides support for DTMCs and CTMCs. The definition of steady-state (long-run) probabilities for finite DTMCs and CTMCs is well defined (see e.g. [Ste94]). Informally, the property:

```
S bound [ prop ]
```

is true in a state s of a DTMC or CTMC if "starting from s , the steady-state (long-run) probability of being in a state which satisfies the (Boolean-valued) PRISM property $prop$, meets the bound $bound$ ". A typical example of this type of property would be:

```
S<0.05 [ queue_size / max_size > 0.75 ]
```


which means: "the long-run probability of the queue being more than 75% full is less than 0.05".

Like the **P** operator, the **S** operator can be used in a [quantitative](#) form, e.g.:

```
S=? [ queue_size / max_size > 0.75 ]
```

including the use of [filters](#).

Reward-based Properties

PRISM models can be augmented with information about [rewards](#) (or, equivalently, costs). The tool can analyse properties which relate to the *expected values* of these rewards. This is achieved using the **R** operator, which works in a similar fashion to the **P** and **S** operators, and can be used either in a Boolean-valued query, e.g.:

```
R bound [ rewardprop ]
```

where *bound* takes the form $<r$, $\leq r$, $>r$ or $\geq r$ for an expression r evaluating to a non-negative double, or a real-valued query, e.g.:

```
R query [ rewardprop ]
```

where *query* is $=?$ for a DTMC or CTMC, and $\min=?$ or $\max=?$ for MDPs. In the latter case, [filters](#) can be used, as for the **P** and **S** operators.

Informally, "**R bound [rewardprop]**" is true in a state of a model if "the expected reward associated with *rewardprop* of the model when starting from that state" meets the bound *bound* and "**R query [rewardprop]**" returns the actual expected reward value.

There are four different types of reward properties, namely:

- "reachability reward": **F prop**
- "cumulative reward": **C \leq t**
- "instantaneous reward": **I=t**
- "steady-state reward": **S**.

Below, we consider each of these cases in turn. The descriptions here are kept relatively informal. For precise definitions (in the case of DTMCs and CTMCs), see [\[KNP07a\]](#).

"Reachability reward" properties

"Reachability reward" properties associate a reward with each path of a model. More specifically, they refer to the reward accumulated along a path until a certain point is reached. The manner in which rewards are accumulated depends on the model type. For DTMCs and MDPs, the total reward for a path is the sum of the state rewards for each state along the path plus the sum of the transition rewards for each transition between these states. The situation for CTMCs is similar, except that the state reward assigned to each state of the model is interpreted as the *rate* at which rewards are accumulated in that state, i.e. if t time units are spent in a state with state reward r , the reward accumulated in that state is $r \times t$. Hence, the total reward for a path in a CTMC is the sum of these products for each state along the path plus the sum of the transition rewards for each transition between these states.

The reward property "**F prop**" corresponds to the reward cumulated along a path until a state satisfying property *prop* is reached, where rewards are cumulated as described above. State rewards for the *prop*-satisfying state reached are not included in the cumulated value. In the case where the probability of reaching a state satisfying *prop* is less than 1, the reward is equal to infinity.

A common application of this type of property is the case when the rewards associated with the model correspond to time. One can then state, for example:

$R \leq 9.5$ [F $z=2$]

which is true in a state s if "the expected time taken to reach, from s , a state where z equals 2 is less than or equal to 9.5".

"Cumulative reward" properties

"Cumulative reward" properties also associate a reward with each path of a model, but only up to a given time bound. The property $C \leq t$ corresponds to the reward cumulated along a path until t time units have elapsed. For DTMCs and MDPs, the bound t must evaluate to an integer; for CTMCs, it can evaluate to double. State and transition rewards along a path are cumulated exactly as described in the previous section.

A typical application of this type of property is the following. Consider a model of a disk-drive controller which includes a queue of incoming disk requests. If we assign a reward of 1 to each transition of the model corresponding to the situation where an incoming request is lost because the queue is full, then the property:

$R=?$ [$C \leq 15.5$]

would return, for a given state of the model, "the expected number of lost requests within 15.5 time units of operation".

"Instantaneous reward" properties

"Instantaneous reward" properties refer to the reward of a model at a particular instant in time. The reward property $I=t$ associates with a path the reward in the state of that path when exactly t time units have elapsed. For DTMCs and MDPs, the bound t must evaluate to an integer; for CTMCs, it can evaluate to double.

Returning to our example from the previous section of a model for a disk-request queue in a disk-drive controller, consider the case where the rewards assigned to each state of the model give the current size of the queue in that state. Then, the following property:

$R < 4.4$ [$I=100$]

would be true in a state s of the model if "starting from s , the expected queue size exactly 100 time units later is less than 4.4". Note that, for this type of reward property, state rewards for CTMCs do not have to refer to rates; they can refer to any instantaneous measure of interest for a state.

"Steady-state reward" properties

Unlike the previous three types of property, "steady-state reward" properties relate not to paths, but rather to the reward in the long-run. A typical application of this type of property would be, in the case where the rewards associated with the model correspond to power consumption, the property:

$R \leq 0.7$ [S]

which is true in a state s if "starting from s , the long-run average power consumption is less than 0.7".

Which reward structure?

In the case where a PRISM model has multiple [reward structures](#) you may need to specify which reward structure your property refers to. This is done by placing the information in braces $\{\}$ after the R operator. You can do so either using the name assigned to a reward structure (if any) or using the index (where 1 means the first rewards structure in the PRISM model file, 2 the second, etc.). Examples are:

```
R{"num_failures"}=? [ C<=10.0 ]
R{"time"}=? [ F step=final ]
R{2}=? [ F step=final ]
```

Note that when using an index to specify the reward structure, you can actually put any expression that evaluates to an integer. This allows you to, for example, write a property of the form $R\{c\}=?[...]$ where c is an undefined integer constant. You can then vary the value of c in an experiment and compute values for several different reward structures at once.

If you don't specify a reward structure to the R operator, by default, the first one in the model file is used.

Availability

There are currently a few restrictions on the model checking [engines](#) that can be used for some reward properties. The following table summarises the currently availability, where S, M and H denote the "sparse", "MTBDD" and "hybrid" engines, respectively.

	F	C<=t	I=t	S
DTMCs	SMH	SMH	SMH	SMH
CTMCs	SMH	SMH	SMH	SMH
MDPs	SM-	---	SM-	n/a

Syntax And Semantics

The three principal operators of the PRISM property specification language, namely P , S and R , can be combined into any valid, well-typed PRISM [expression](#). This means that any of the following operators can be used:

- $-$ (unary minus)
- $*$, $/$ (multiplication, division)
- $+$, $-$ (addition, subtraction)
- $<$, $<=$, $>=$, $>$ (relational operators)
- $=$, $!=$ (equality operators)
- $!$ (negation)
- $\&$ (conjunction)
- $|$ (disjunction)
- \Rightarrow (implication)
- $?$ (condition evaluation: condition $? a : b$ means "if condition is true then a else b ")

Logical operators can be used, for example, to specify any property expressible in logics such as PCTL and CSL, e.g.:

```
"no_error" => P<0.1 [ F ("error1" | "error2") ]
```

Other operators can also be used to derive various arithmetic expressions such as:

```
P=? [ F[3600,7200] oper ]
```

"the probability that the system is **not** operational at any point during the second hour of operation"

```
R{"oper"}=? [ C<=t ] / t
```

"the expected fraction of time that the system is available (i.e. the expected interval availability) in the time interval $[0, t]$ "

```
P=? [ F fail_A ] / P=? [ F any_fail ]
```

"the (conditional) probability that component A eventually fails, given that at least one component fails"

Semantics

We omit a formal presentation of the semantics of the PRISM property language. The semantics of the probabilistic temporal logics that the language incorporates can be found from a variety of sources. See for example the pointers given in the [About](#) and [Documentation](#) sections of the PRISM website.

It is worth, however, clarifying a few points specific to PRISM. A property is evaluated with respect to a particular state of a model. Depending on the type of the property, this value may either be a Boolean, an integer or a double. When analysing a property of a model, though, PRISM actually determines the value for all states of the model. Furthermore, the final value that PRISM returns as the result of model checking depends on the model type.

For Boolean properties, such as those conforming to the classical syntax of PCTL, CSL, etc., PRISM returns a value of `true` if the property is true for **all** states of the model. To check if a property is satisfied in some subset of states, for example just the initial states, properties can be prefixed with an implication, e.g.:

```
"init" => P=>1 [ F leader_elected=true ]
```

For numeric properties, typically a single value (for a single state) is required. The default behaviour of PRISM is to return the value for the initial state of the model. This issue was discussed in more depth in the [Quantitative Properties](#) section.

Properties Files

Files containing properties to be analysed by PRISM can also contain constants, as is the case for model files. These are defined in identical fashion, for example:

```
const int k = 7;  
const double T = 9.5;  
const double p = 0.01;  
  
P<p [ true U<=T x=k ]
```

As before, these constants can actually be left undefined and then later assigned either a single value or a range of values using [experiments](#).

In fact, probability bounds for the **P** or **S** operators (like **P** above) and upper or lower bounds for the **U** operator (like **T** above) can be arbitrary expressions, provided they are constant. Furthermore, expressions in the properties file can also refer to constants previous defined in the model file.

Another feature of properties files is *labels*. These are a way of defining sets of states that will be referred to in properties (they correspond to *atomic propositions* in a temporal logic setting). As described [earlier](#), labels can be defined in either model files or property files.

Labels are defined using the keyword **label**, followed by a name (identifier) in double quotes, and then an expression which evaluates to a Boolean. Definition and usage of formulas are illustrated in the following example:

```
label "safe" = temp<=100 | alarm=true;  
label "fail" = temp>100 & alarm=false;  
  
P>=0.99 [ "safe" U "fail" ]
```

Two special cases are the "init" and "deadlock" labels which are always defined. These are true in initial states of the model and states where deadlocks were found (and fixed by adding self-loops), respectively. The former is useful if you only wish to check whether a property is true in the initial states, e.g.:

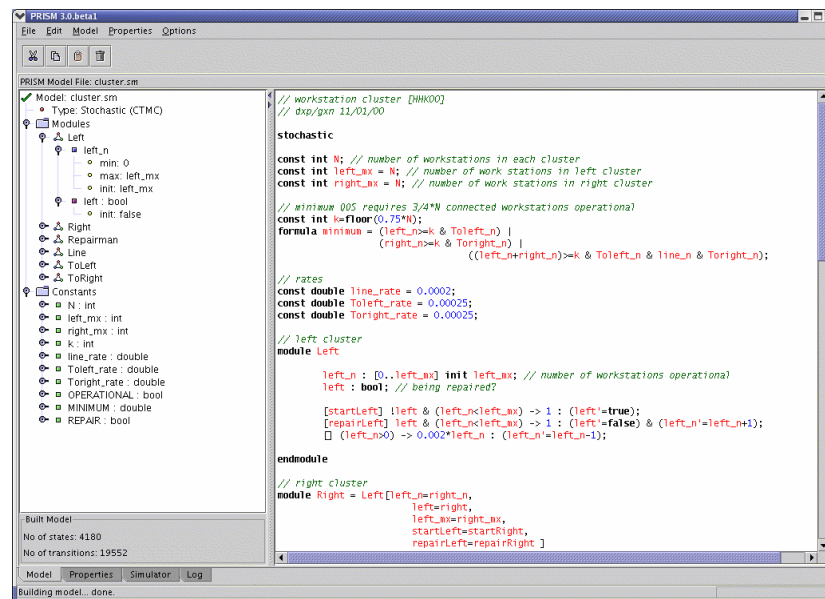
```
"init" => P>=0.99 [ "safe" U "fail" ]
```

A PRISM properties file can contain any number of properties. Like model files, they can also include any amount of white space (spaces, tabs, new lines, etc.) and C-style comments, which are both ignored. By convention, files are given the extension .pctl for properties of DTMCs and MDPs and extension .csl for properties of CTMCs.

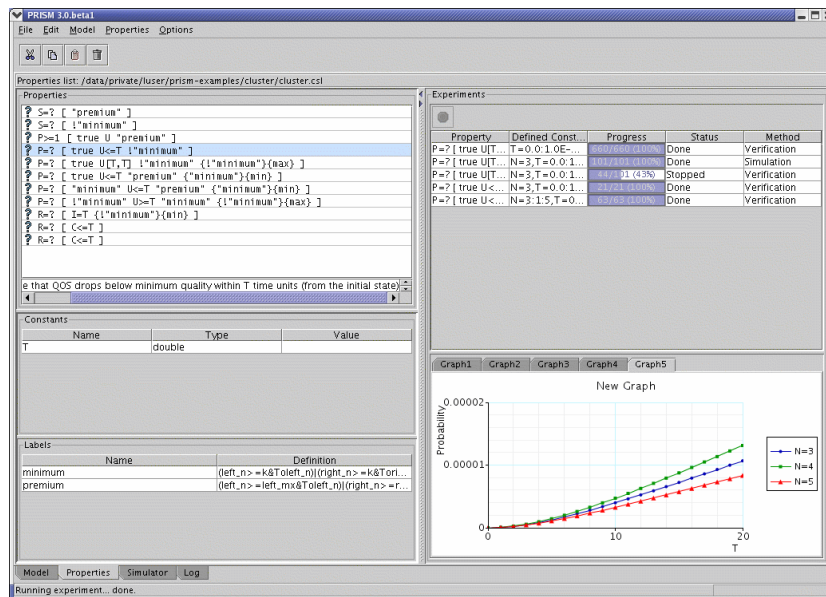
Running PRISM

Starting PRISM

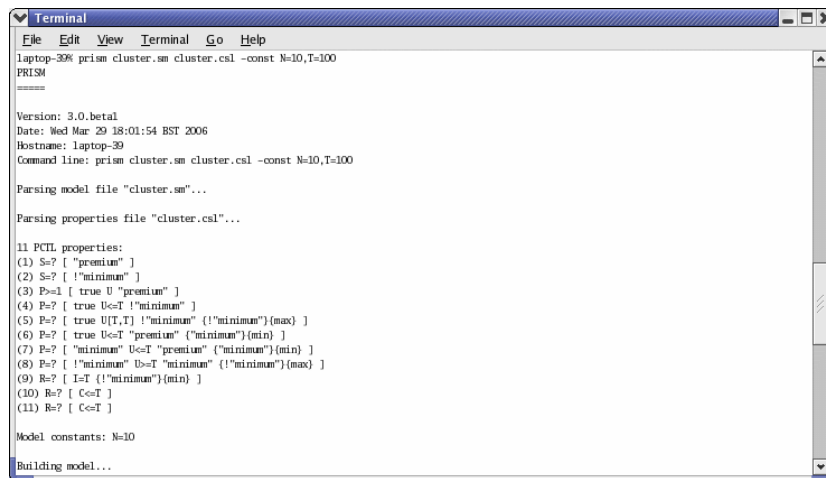
There are two versions of PRISM, one based on a graphical user interface (GUI), the other based on a command line interface. Both use the same underlying model checker. The latter is useful for running large batches of jobs, leaving long-running model checking tasks in the background, or simply for running the tool quickly and easily once you are familiar with its operation. Assuming that the bin directory of your PRISM distribution is in your path, the GUI version can be started by typing xprism and the command-line version by typing prism (typically with several arguments and switches). Below are screenshots of the PRISM GUI and the command-line version running.



The PRISM GUI



The PRISM GUI



The Command-Line

Loading And Building a Model

Use of PRISM typically starts with the construction of a probabilistic model, i.e. the conversion of a description in the PRISM language to an MDP, DTMC or CTMC, as appropriate. During this process, PRISM computes the set of states in the model which are reachable from the initial states and the transition matrix which represents the model. There are a selection of sample PRISM model files in the examples directory of the distribution. From the command-line, simply type:

```
prism model
```

where model is the name of the file containing the model description. In the GUI, select menu option "Model | Open Model" and choose a file. The model will be displayed in the editor in the "Model" tab of the GUI window. The file is parsed upon loading. If there are no errors, information about the modules, variables, and other components of the model is displayed in the panel to the left and a green tick will be visible. If there are errors in the file, a red cross will appear instead and the errors will be highlighted in the model editor. To view details of the error, position the mouse pointer over the source of the error (or over the red cross). Alternatively, select menu option "Model | Parse Model" and the error message will be displayed in a message box. Model descriptions can, of course, also be typed from scratch into the GUI's editor.

Once the file has been parsed successfully, the model can be built by selecting "Model | Build Model". If there are no

errors during model construction, the number of states and transitions in the model will be displayed in the bottom left corner of the window.

Deadlocks

The presence of *deadlock states* in the model, i.e. states which are reachable but from which there are no outgoing transitions, constitutes an error. From the GUI, you are offered the opportunity to automatically add self-loops to these states to resolve the situation. The same can be achieved from the command-line by using the `-fixdl` switch. Otherwise, the deadlock states are displayed (in the GUI, they appear in the "Log" tab of the main window). Note that if you choose to fix the deadlocks by adding self-loops, you can determine in which states this occurred by model checking the property "deadlock". Another useful technique in this situation is to generate a random path in the [simulator](#) in the hope of reaching a deadlock state. This will give you useful information about how the state can be reached.

Debugging Models With The Simulator

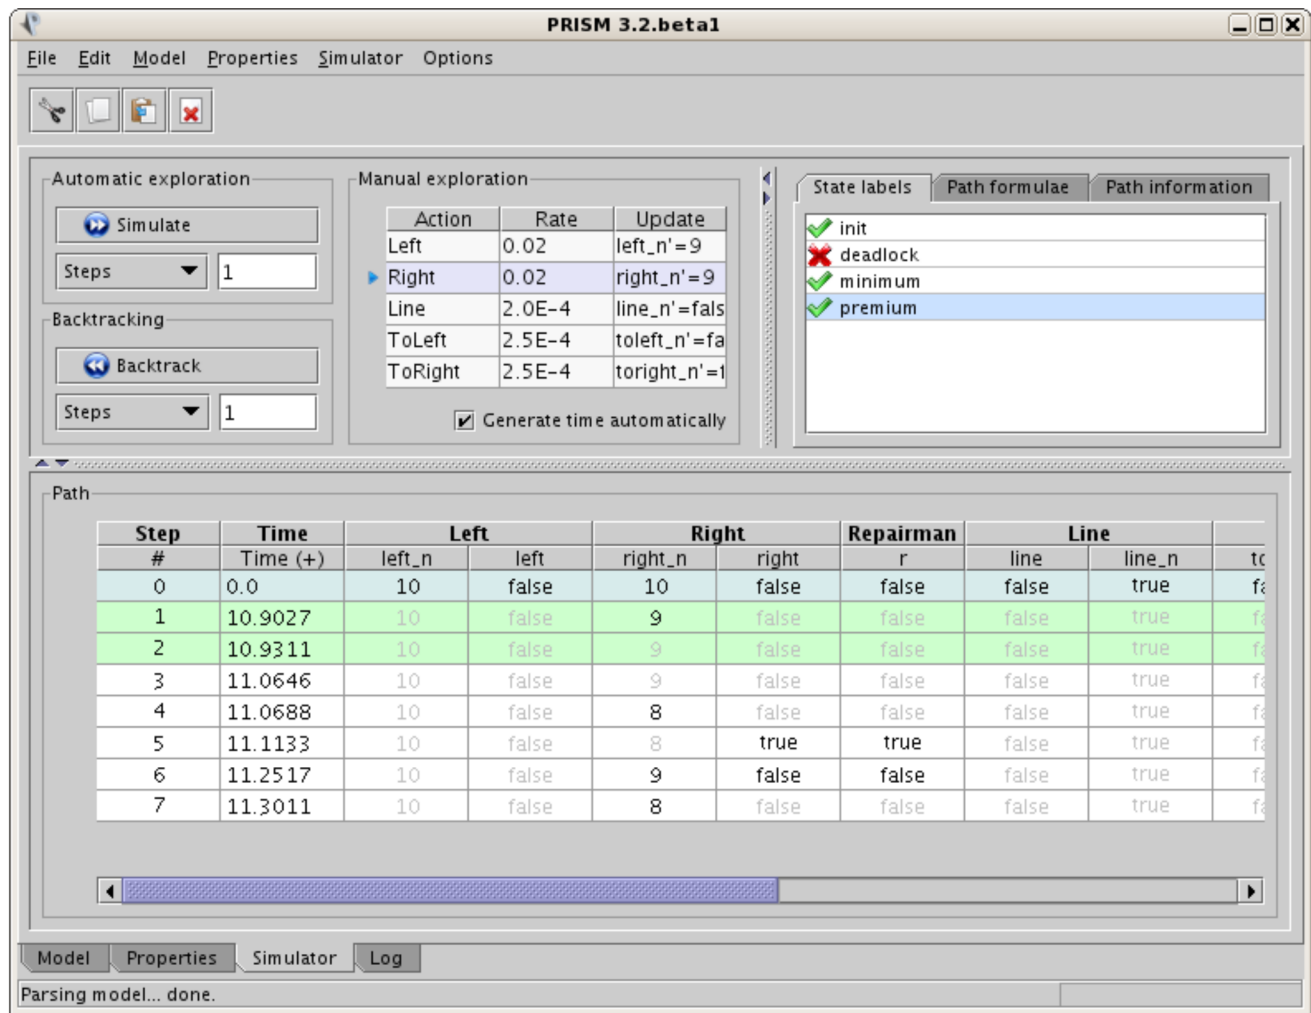
PRISM includes a *simulator*, a tool which can be used to generate sample paths (executions) through a PRISM model. From the GUI, the simulator allows you to explore a model by interactively generating such paths. This is particularly useful for debugging models during development and for running sanity checks on completed models. Paths can also be generated from the command-line.

Generating a path in the GUI

Once you have loaded a model into the PRISM GUI (note that it is not necessary to build the model), select the "Simulator" tab at the bottom of the main window. You can now start a new path by double-clicking in the bottom half of the window (or right-clicking and selecting "New path"). If there are undefined constants in the model (or in any currently loaded properties files) you will be prompted to give values for these. You can also specify the state from which you wish to generate a path. By default, this is the initial state of the model.

The main portion of the user interface (the bottom part) displays a path through the currently loaded model. Initially, this will comprise just a single state. The table above shows the list of available transitions from this state. Double-click one of these to extend the path with this transition. The process can be repeated to extend the path in an interactive fashion. Clicking on any state in the current path shows the transition which was taken at this stage. Click on the final state in the path to continue extending the path. Alternatively, clicking the "Simulate" button will select a transition randomly (according to the probabilities/rates of the available transitions). By changing the number in the box below this button, you can easily generate random paths of a given length with a single click. There are also options (in the accompanying drop-down menu) to allow generation of paths up until a particular length or, for CTMCs, in terms of the time taken.

The figure shows the simulator in action.



The PRISM GUI: exploring a model using the simulator

It is also possible to:

- backtrack to an earlier point in a path
- remove all of the states before some point in a path
- restart a path from its first state
- export a path to a text file

Notice that the table containing the path displays not just the value of each variable in each state but also the time spent in that state and any rewards accumulated there. You can configure exactly which columns appear by right-clicking on the path and selecting "Configure view". For rewards (and for CTMC models, for the time-values), you can opt to display the reward/time for each individual state and/or the cumulative total up until each point in the path.

At the top-right of the interface, any labels contained in the currently loaded model/properties file are displayed, along with their value in the currently selected state of the path. In addition, the built-in labels "init" and "deadlock" are also included. Selecting a label from the list highlights all states in the current path which satisfy it.

The other tabs in this panel allow the value of path operators (taken from properties in the current file) to be viewed for the current path, as well as various other statistics.

Path generation from the command-line

It is also possible to generate random paths through a model using the command-line version of PRISM. This is achieved using the `-simp` switch, which requires two arguments, the first describing the path to be generated and the second specifying the file to which the path should be output (as usual, specifying `stdout` sends output to the terminal). The

following examples illustrate the various ways of generating paths in this way:

- `prism model.pm -simpath 10 path.txt` (a path of 10 steps)
- `prism model.pm -simpath time=7.5 path.txt` (a path of at most 7.5 time units)
- `prism model.pm -simpath deadlock path.txt` (a path ending in deadlock)

Further options can also be appended to the first parameter. If you are only interested in the changes to certain variables of your model, use the `vars=(...)` option. For example:

```
prism model.sm -simpath 10 stdout
...
step a b c time
0 0 0 0 0.0
1 0 1 0 0.016569539535505946
2 0 1 1 0.04999280708731619
3 0 2 1 0.0637472535911344
4 0 2 2 0.243645533565261
5 0 2 3 0.5359625623773467
6 0 3 3 0.7862449420673264
7 1 3 3 0.8749262111456289
8 1 3 4 0.9472785807052686
9 1 3 5 1.040096742715008
10 1 3 6 1.2801655430222152
```

```
prism model.sm -simpath '10,vars=(a,b)' stdout
...
step a b time
0 0 0 0.0
1 0 1 0.20115547684708998
5 0 2 0.5822925951221433
7 1 2 0.9559600285257709
8 1 3 1.3395175958850654
10 1 4 1.869013176198441
```

Note the use of single quotes around the path description argument to prevent the shell from misinterpreting special characters such as "`(`".

You can also use the `sep=...` option to specify the column separator. Possible values are space (the default), tab and comma. For example:

```
prism model.sm -simpath '10,vars=(a,b),sep=comma' stdout
...
step,a,b,time
0,0,0,0.0
2,1,0,0.058443536856580006
3,1,1,0.09281024515535738
6,1,2,0.2556555786269585
7,1,3,0.284062896359802
8,1,4,1.1792064236954896
```

When generating paths to a deadlock state, additional `repeat=...` option is available which will construct multiple paths until a deadlock is found. For example:

```
prism model.sm -simpath 'deadlock,repeat=100' stdout
```

By default, the simulator detects deterministic loops in paths (e.g. if a path reaches a state from which there is a just a single self-loop leaving that state) and stops generating the path any further. You can disable this behaviour with the `loopcheck=false` option. For example:

```
prism dice.pm -simpath 10 stdout
...
Warning: Deterministic loop detected after 6 steps (use loopcheck=false option to extend path).
step s d state_reward transition_reward
0 0 0 0.0 1.0
1 1 0 0.0 1.0
2 3 0 0.0 1.0
3 1 0 0.0 1.0
4 4 0 0.0 1.0
5 7 3 0.0 0.0
6 7 3 0.0 0.0
```

```
prism dice.pm -simpath 10,loopcheck=false stdout
...
step s d state_reward transition_reward
0 0 0 0.0 1.0
1 2 0 0.0 1.0
2 6 0 0.0 1.0
3 7 6 0.0 0.0
4 7 6 0.0 0.0
5 7 6 0.0 0.0
6 7 6 0.0 0.0
7 7 6 0.0 0.0
8 7 6 0.0 0.0
9 7 6 0.0 0.0
10 7 6 0.0 0.0
```

One final note: the `-simpath` switch only generates paths up to the maximum path length setting of the simulator (the default is 10,000). If you want to generate longer paths, either change the [default setting](#) or override it temporarily from the command-line using the `-simpathlen` switch. You might also use the latter to decrease the setting, e.g. to look for a path leading to a deadlock state, but only within 100 steps:

```
prism model.sm -simpath deadlock stdout -simpathlen 100
```

Exporting The Model

If required, once the model has been constructed it can be exported, either for manual examination or for use in another tool. In the GUI, the "Model | Export" menu provides options to export various components of the model: the set of reachable states, the transition matrix, the state rewards vector and the transition rewards matrix. In all cases, it is possible to export this information either in plain text format or as [Matlab](#) code. Additionally, matrices and vectors can be exported in a format suitable for import into the [MRMC](#) tool and the transition matrix of the model can be exported in [Dot](#) format which allows easy graphical visualisation of the graph structure of the model. The GUI also provides options in the "Model | View" menu, which allow exports (in textual format) directly to the log. This is convenient for quick examination of small models.

All of this export functionality is available from the command-line version of PRISM, using the following switches:

- `-exportstates <file>`
- `-exporttrans <file>`
- `-exportstaterewards <file>`
- `-exporttransrewards <file>`
- `-exporttransdot <file>`

In each case, using `stdout` instead of a file name causes the information to be displayed directly to the screen. To change the export format from the default (plain text) to Matlab or MRMC, use the `-exportmatlab` and `-exportmrmc` switches. The export command-line switches can be used in combination. For example:

```
prism poll2.sm -exportstates poll2.sta -exporttrans poll2.tra -exportmatlab
```

exports both the state space and transition matrix in Matlab format.

Finally, there is some additional export functionality available only from the command-line.

Firstly, in the case where both a model file and properties file have been specified, it is possible to export the set of states satisfied by each label in the properties file, including the built-in labels "init" and "deadlock". This is done with the `-exportlabels` switch and the information can be output either in plain text format (default) or for use with Matlab or MPMC (see switches above).

Secondly, when outputting matrices for DTMCs or CTMCs, it is possible to request that PRISM does not sort the rows of the matrix, as is normally the case. This is achieved with the `-exportunordered` switch. The reason for this is that in this case PRISM does not need to construct an explicit version of the model in memory and the process can thus be performed with reduced memory consumption.

Finally, the `-exportrows` switch provides an alternative output format for transition matrices where the elements of each row of the matrix are grouped on the same line. This can be particularly helpful for viewing the matrix for MDPs. By way of example:

```
0 0 13 0.5
0 0 26 0.5
0 1 1 0.5
0 1 2 0.5
1 0 14 0.5
1 0 27 0.5
1 1 3 0.5
1 1 4 0.5
```

becomes:

```
0 0.5:13 0.5:26
0 0.5:1 0.5:2
1 0.5:14 0.5:27
1 0.5:3 0.5:4
```

Model Checking

Typically, once a model has been constructed, it is analysed through model checking. Properties are specified as described in the ["Property Specification"](#) section, and are usually kept in files with extension `.pctl` or `.csl`. There are properties files accompanying most of the sample PRISM models in the `examples` directory.

To load a file containing properties into the GUI, select menu option "Properties | Open properties list". The file can only be loaded if there are no errors, otherwise an error is displayed. Note that it may be necessary to have loaded the corresponding model first, since the properties will probably make reference to variables (and perhaps constants) declared in the model file. Once loaded, the properties contained in the file are displayed in a list in the "Properties" tab of the GUI. Constants and labels are displayed in separate lists below. You can modify or create new properties, constants and labels from the GUI, by right-clicking on the appropriate list and selecting from the pop-up menu which appears. Properties with errors are shaded red and marked with a warning sign. Positioning the mouse pointer over the property displays the corresponding error message.

The pop-up menu for the properties list also contains a "Verify" option, which allows you instruct PRISM to model check the currently selected properties (hold down Ctrl to select more than one property simultaneously). All properties can be model checked at once by selecting "Verify all". PRISM verifies each property individually. Upon completion, the icon next to the property changes according to the result of model checking. A green tick indicates that the property is satisfied in all states of the model; a red cross that this is not the case. As mentioned earlier, if you just want to know whether or not the property is satisfied in the initial states of the model, prefix the property with "init" \Rightarrow . For properties which have a numerical result (e.g. $P=? [\dots]$), position the mouse pointer over the property to view the result. In addition, this and further information about model checking is displayed in the log in the "Log" tab.

From the command-line, model checking is achieved by passing both a model file and a properties file as arguments, e.g.:

```
prism poll2.sm poll.csl
```

By default, all properties in the file are checked. To model check only a single property, use the `-prop` switch. For example, to check only the fourth property in the file:

```
prism poll2.sm poll.csl -prop 4
```

Alternatively, a single property can be specified directly from the command-line, using the (equivalent) `-pctl` and `-csl` switches:

```
prism poll2.sm -csl 'P>=0.5 [ true U<=5 (s=1 & a=0) ]'
```

Note the use of single quotes ('...') to avoid characters such as (and > being interpreted by the command-line shell. The results of model checking are sent to the display and are as described above for the GUI version.

Approximate Computation Of Quantitative Properties

The discrete-event simulator built in to PRISM (see the section "[Debugging Models With The Simulator](#)") can also be used to generate approximate values for PRISM properties which give quantitative results, namely properties of the form **P=?[...]** and **R=?[...]**. This is achieved by *sampling*: generating a large number of random runs through the model, evaluating the result of the given properties on each run, and then averaging the results. This approach is particularly useful on very large models when normal model checking is infeasible. This is because discrete-event simulation is performed using the PRISM language model description, without actually constructing the probabilistic model.

To use this functionality, load a model and some properties into PRISM, as described in the previous sections. To generate an approximate value for one or more properties, select them in the list, right-click and select "Simulate" (as opposed to "Verify"). As usual, it is first necessary to provide values for any undefined constants. Subsequently, a dialog appears. Here, the state for which values are to be computed (i.e. from which the paths will be generated) can be selected. By default, this is initial state of the model. The number of paths (samples) and the maximum length of each path which will be generated by PRISM to compute approximate values can be modified as required. The related parameters of *confidence* and *approximation* can also be viewed or edited. Informally, for a given number of paths, the probability that the difference between computed and actual answers is bigger than the approximation parameter is at most equal to the confidence parameter. See for example [\[HLMPO4\]](#) for a more in-depth discussion of these concepts. By default in PRISM the approximation parameter is 0.01 and the confidence parameter is 1e-10. In this case, the number of paths/samples required is 402,412.

Finally, approximate generation of results can also be activated from the command-line using the `-sim` switch. To change the approximation parameter, confidence parameter, number of samples and maximum path length, respectively, use the switches `-simapprox`, `-simconf`, `-simsamples` and `-simpathlen`. Here is an example of activating the simulator from the command-line:

```
prism dice.pm dice.pctl -sim -simconf 1e-1
```

Currently, the simulator does not support every part of the PRISM modelling and property languages. For example, it does not handle models with multiple initial states or with **system...endsystem** definitions and it does not handle properties containing LTL-style path properties or arithmetic combinations of **P=?[...]** and **R=?[...]** properties.

Computing Steady-state And Transient Probabilities

If the model is a CTMC or DTMC, it is possible to compute corresponding vectors of steady-state or transient probabilities directly (rather than indirectly by analysing a property which requires their computation). From the GUI, select an option from the "Model | Compute" menu. For transient probabilities, you will be asked to supply the time value for which you wish to compute probabilities. From the command-line, add the `-steadystate` (or `-ss`) switch:

```
prism poll2.sm -ss
```

for steady-state probabilities or the `-transient` (or `-tr`) switch:

```
prism poll2.sm -tr 100.0
```

for transient probabilities, again specifying a time value in the latter case. The probabilities are computed for all states of the model and displayed, either on the screen (from the command-line) or in the log (from the GUI). For both steady-state and transient probability computation, PRISM assumes that the initial probability distribution of the CTMC is an equiprobable choice over the set of initial states.

Experiments

PRISM supports *experiments*, which is a way of automating multiple instances of model checking. This is done by leaving one or more [constants](#) undefined, e.g.:

```
const int N;  
const double T;
```

This can be done for constants in the model file, the properties file, or both. Before any verification can be performed, values must be provided for any such constants. In the GUI, a dialog appears in which the user is required to enter values. From the command line, the `-const` switch must be used, e.g.:

```
prism cluster.sm cluster.csl -const N=4,T=85.9
```

You can also provide ranges of variables for one or more of the constants, e.g.:

```
prism cluster.sm cluster.csl -const N=4:6,T=60:10:100
```

where `N=4:6` means that values of 4, 5 and 6 are used for `N`, and `T=60:10:100` means that values of 60, 70, 80, 90 and 100 (i.e. steps of 10) are used for `T`. Model checking is performed for all combinations of the constant values provided. Using the switch `-exportresults file`, you can send a tabulated copy of all the model checking results generated to the file `file`. For convenience, constant specifications can split across separate instances of the `-const` switch, if desired.

From the GUI, the same thing can be achieved by selecting a single property, right clicking on it and selecting "New experiment" (or alternatively using the popup menu in the "Experiments" panel). Values or ranges for each undefined constant can then be supplied in the resulting dialog. Details of the new experiment and its progress are shown in the panel. To stop the experiment before it has completed, click the red "Stop" button and it will halt after finishing the current iteration of model checking. Once the experiment has finished, right clicking on the experiment produces a pop-up menu, from which you can view the results of the experiment or export them to a file.

For experiments based on properties which return numerical results, you can also use the GUI to plot graphs of the results. This can be done either before the experiment starts, by selecting the "Create graph" tick-box in the dialog used to create the experiment (in fact this box is ticked by default), or after the experiment's completion, by choosing "Plot results" from the pop-up menu on the experiment. A dialog appears, where you can choose which constant (if there are more than one) to use for the x-axis of the graph, and for which values of any other constants the results should be plotted. The graph will appear in the panel below the list of experiments. Right clicking on a graph and selecting "Graph options" brings up a dialog from which many properties of the graph can be configured. From the pop-up menu of a graph, you can also choose to print the graph (to a printer or Postscript file) or export it in a variety of formats: as an image (PNG or JPEG), as an encapsulated Postscript file (EPS), in an XML-based format (for reloading back into PRISM), or as code which can be used to generate the graph in Matlab.

Approximate computation of quantitative results obtained with the [simulator](#) can also be used on experiments. In the GUI, select the "Use Simulation" option when defining the parameters for the experiment. From the command-line, just add the `-sim` switch as usual.

Support For PEPA Models

For CTMCs, PRISM also accepts model descriptions in the stochastic process algebra [PEPA](#) [Hil96]. The tool compiles such descriptions into the PRISM language and then constructs the model as normal. The language accepted by the PEPA to PRISM

compiler is actually a subset of PEPA. The restrictions applied to the language are firstly that component identifiers can only be bound to sequential components (formed using prefix and choice and references to other sequential components only). Secondly, each local state of a sequential component must be named. For example, we would rewrite:

- $P = (a,r).(b,s).P;$

as:

- $P = (a,r).P';$
- $P' = (b,s).P;$

Finally, active/active synchronisations are not allowed since the PRISM definition of these differs from the PEPA definition. Every PEPA synchronisation must have exactly one active component. Some examples of PEPA model descriptions which can be imported into PRISM can be found in the `examples/pepa` directory.

From the command-line version of PRISM, add the `-importpepa` switch and the model will be treated as a PEPA description. From the GUI, select "Model | Open model" and then choose "PEPA models" on the "Files of type" drop-down menu. Alternatively, select "Model | New | PEPA model" and either type a description from scratch or paste in an existing one from elsewhere. Once the PEPA model has been successfully parsed by PRISM, you can view the corresponding PRISM code (as generated by the PEPA-to-PRISM compiler) by selecting menu option "Model | View | Parsed PRISM model".

Support For SBML

PRISM includes a (prototype) tool to translate specifications in [SBML](#) (Systems Biology Markup Language) to model descriptions in the PRISM language. SBML is an XML-based format for representing models of biochemical reaction networks. The translator currently works with Level 2 Version 1 of the SBML specification, details of which can be found [here](#). Level 1 SBML files will first need to be translated into equivalent Level 2 files, for example using [this on-line converter](#).

Since PRISM is a tool for analysing discrete-state systems, the translator is designed for SBML files intended for discrete stochastic simulation. A useful set of such files can be found in the CaliBayes [Discrete Stochastic Model Test Suite](#). There are also many more SBML files available in the [BioModels Database](#).

We first give a simple example of an SBML file and its PRISM translation. We then give some more precise details of the translation process.

Example

An SBML file comprises a set of *species* and a set of *reactions* which they undergo. Below is the SBML file for the simple reversible reaction: $\text{Na} + \text{Cl} \leftrightarrow \text{Na}^+ + \text{Cl}^-$, where there are initially 100 Na and Cl atoms and no ions, and the base rates for the forwards and backwards reactions are 100 and 10, respectively.

```

<?xml version="1.0" encoding="UTF-8"?>
<sbml xmlns="http://www.sbml.org/sbml/level2" metaid="_000000" level="2" version="1">
  <model id="nacl" name="Na+Cl">

    <listOfCompartments>
      <compartment id="compartment"/>
    </listOfCompartments>

    <listOfSpecies>
      <species id="na" initialAmount="100" hasOnlySubstanceUnits="true"/>
      <species id="cl" initialAmount="100" hasOnlySubstanceUnits="true"/>
      <species id="na_plus" initialAmount="0" hasOnlySubstanceUnits="true"/>
      <species id="cl_minus" initialAmount="0" hasOnlySubstanceUnits="true"/>
    </listOfSpecies>

    <listOfReactions>
      <reaction id="forwards" reversible="false">
        <listOfReactants>
          <speciesReference species="na"/>
          <speciesReference species="cl"/>
        </listOfReactants>
        <listOfProducts>
          <speciesReference species="na_plus"/>
          <speciesReference species="cl_minus"/>
        </listOfProducts>
        <kineticLaw>
          <math xmlns="http://www.w3.org/1998/Math/MathML">
            <apply><times/><ci>forwards_rate</ci>
            <apply><times/><ci>na</ci><ci>cl</ci></apply></apply>
          </math>
          <listOfParameters>
            <parameter id="forwards_rate" value="100"/>
          </listOfParameters>
        </kineticLaw>
      </reaction>

      <reaction id="backwards" reversible="false">
        <listOfReactants>
          <speciesReference species="na_plus"/>
          <speciesReference species="cl_minus"/>
        </listOfReactants>
        <listOfProducts>
          <speciesReference species="na"/>
          <speciesReference species="cl"/>
        </listOfProducts>
        <kineticLaw>
          <math xmlns="http://www.w3.org/1998/Math/MathML">
            <apply><times/><ci>backwards_rate</ci>
            <apply><times/><ci>na_plus</ci><ci>cl_minus</ci></apply></apply>
          </math>
          <listOfParameters>
            <parameter id="backwards_rate" value="10"/>
          </listOfParameters>
        </kineticLaw>
      </reaction>
    </listOfReactions>

  </model>
</sbml>

```

And here is the resulting PRISM code:

```

// File generated by automatic SBML-to-PRISM conversion
// Original SBML file: nacl.xml

ctmc

const int MAX_AMOUNT = 100;

// Parameters for reaction forwards
const double forwards_rate = 100; // forwards_rate

// Parameters for reaction backwards
const double backwards_rate = 10; // backwards_rate

// Species na
const int na_MAX = MAX_AMOUNT;
module na

na : [0..na_MAX] init 100; // Initial amount 100

// forwards
[forwards] na > 0 -> (na'=na-1);
// backwards
[backwards] na <= na_MAX-1 -> (na'=na+1);

endmodule

// Species cl
const int cl_MAX = MAX_AMOUNT;
module cl

cl : [0..cl_MAX] init 100; // Initial amount 100

// forwards
[forwards] cl > 0 -> (cl'=cl-1);
// backwards
[backwards] cl <= cl_MAX-1 -> (cl'=cl+1);

endmodule

// Species na_plus
const int na_plus_MAX = MAX_AMOUNT;
module na_plus

na_plus : [0..na_plus_MAX] init 0; // Initial amount 0

// forwards
[forwards] na_plus <= na_plus_MAX-1 -> (na_plus'=na_plus+1);
// backwards
[backwards] na_plus > 0 -> (na_plus'=na_plus-1);

endmodule

// Species cl_minus
const int cl_minus_MAX = MAX_AMOUNT;
module cl_minus

cl_minus : [0..cl_minus_MAX] init 0; // Initial amount 0

// forwards
[forwards] cl_minus <= cl_minus_MAX-1 -> (cl_minus'=cl_minus+1);
// backwards
[backwards] cl_minus > 0 -> (cl_minus'=cl_minus-1);

endmodule


// Reaction rates
module reaction_rates

// forwards
[forwards] (forwards_rate*(na*cl)) > 0 -> (forwards_rate*(na*cl)) : true;
// backwards
[backwards] (backwards_rate*(na_plus*cl_minus)) > 0 -> (backwards_rate*(na_plus*cl_minus)) : true;

. . .

```


From the latter, we can use PRISM to generate a simple plot of the expected amount of Na and Na+ over time (using both model checking and a single random trace from the simulator):



Expected amount of Na/Na+ at time T

Using the translator

At present, the SBML-to-PRISM translator is included in the PRISM code-base, but not integrated into the application itself. To perform a conversion:

```
cd prism
java -cp classes prism.SBML2Prism sbml_file.xml > prism_file.sm
```

Alternatively (on Linux or Mac OS X), ensure prism is in your path and then save the script below as an executable file called sbml2prism:

```
#!/bin/sh

# Startup script for SBML-to-PRISM translator

# Launch using main PRISM script
PRISM_MAINCLASS="prism.SBML2Prism"
export PRISM_MAINCLASS
prism "$@"
```

Then use:

```
sbml2prism sbml_file.xml > prism_file.sm
```

The following PRISM properties file will also be useful:

```
const double T;
const int c;

R{c}=? [I=T]
```

This contains a single property which, based on the reward structures in the PRISM model generated by the translator, means "the expected amount of species c at time T ". The constant c is an integer index which can range between 1 and N , where N is the number of species in the model. To view the expected amount of each species over time, create an [experiment](#) in PRISM which varies c from 1 to N and T over the desired time range.

Details of the translation

The basic structure of the translation process is as follows:

- Each *species* in the SBML file is represented by a [module](#) in the resulting PRISM file. This module, which (where possible) retains the SBML species id as its name, contains a single [variable](#) whose value represents the amount of the species present. A corresponding [reward structure](#) for computing the expected amount of the species at a given time instant is also created. Species for which the `boundaryCondition` flag is set to `true` in the SBML file do not have a corresponding module.
- Each *reaction* in the SBML file is associated with a unique [synchronisation action label](#). The module for each species which takes part in the reaction will include a synchronous [command](#) to represent this. An additional PRISM module called `reaction_rates` stores the expression representing the rate of each reaction (from the corresponding `kineticLaw` section in the SBML file). Reaction stoichiometry information is respected but must be provided in the scalar stoichiometry field of a `speciesReference` element, not in a separate `StoichiometryMath` element.
- Each *parameter* in the SBML file, either global to the file or specific to a reaction, becomes a [constant](#) in the PRISM file. If a value for this parameter is given, it is used. If not, the constant is left as undefined.

As described above, this translation process is designed for discrete systems and so the amount of each species in the model is represented by an integer variable. It is therefore assumed that the initial amount for each species specified in the SBML file is also given as an integer. If this is not the case, then the values will need to be scaled accordingly first.

Furthermore, since PRISM is primarily a model checking (rather than simulation) tool, it is important that the amount of each species also has an upper bound (to ensure a finite state space). When model checking, the efficiency (or even feasibility) of the process is likely to be very sensitive to the upper bound(s) chosen. When using the discrete-event simulation functionality of PRISM, this is not the case and the bounds can be set much higher. By default the translator uses an upper bound of 100 (which is increased if the initial amount exceeds this). A different value can be specified through a second command-line argument as follows:

```
cd prism
java -cp classes prism.SBML2Prism sbml_file.xml 1000 > prism_file.sm
```

Alternatively, upper bounds can be modified manually after the translation process.

Finally, The following aspects of SBML files are not currently supported and are ignored during the translation process:

- compartments
 - events/triggers
-

Explicit Model Import

It is also now possible to construct models in PRISM through direct specification of their transition matrix. The format in which this information is input to the tool is exactly the same as is currently output (see the section [Exporting The Model](#)). Presently, this functionality is only supported in the command-line version of the tool, using the `-importtrans` switch. At the moment, PRISM makes no attempt to discern the model type from the input file. By default it assumes that the model is an MDP. If this is not the case, the model type can be overwritten using the `-dtmc`, `-ctmc` and `-mdp` switches. For example:

```
prism -importtrans poll2.tra -ctmc
```

Please note that this method of constructing models in PRISM is typically less efficient than using the PRISM language. This is because PRISM is a symbolic model checker and the underlying data structures used to represent the model function better when there is high-level structure and regularity to exploit. This situation can be alleviated to a certain extent by importing not just a transition matrix, but also a definition of each state of the model in terms of a set of variables. The format of this information is again identical to PRISM's current output format, using the `-exportstates` switch. The following example shows how PRISM could be used to build, export and then re-import a model (not a good strategy in general):

```
prism poll2.sm -exporttrans poll2.tra -exportstates poll2.sta
prism -importtrans poll2.tra -importstates poll2.sta -ctmc
```

Lastly, since details about the initial state of a model is not preserved in the files output from `-exportstates` and `-exporttrans`, this information needs to be specified separately. If not, the default is to assume a single initial state, in which all variables take their minimum value (if `-importstates` is not used, the model has a single zero-indexed variable x , and the initial state is $x=0$).

To specify an alternative set of initial states, use the switch `-importlabels`, e.g.:

```
prism -importtrans poll2.tra -importstates poll2.sta -importlabels poll2.lab -ctmc
```

where the labels file (`poll2.lab` above) is in the format generated by the `-exportlabels` [export option](#) of PRISM. Currently, the `-importlabels` switch does not import any other label information - just the set of initial states.

Configuring PRISM

Introduction

The operation of PRISM can be configured in a number of ways. From the GUI, select "Options" from the main menu to bring up the "Options" dialog. The settings are grouped under several tabs. Those which affect the basic model checking functionality of the tool are under the heading "PRISM". Separate settings are available for the simulator and various aspects of the GUI (the model editor, the property editor and the log).

From the command-line version, options are controlled by switches. A full list can be displayed by typing:

```
prism -help
```

In the following sections, we give a brief description of the most important configuration options available.

User options and settings are saved locally and reused. They are stored in the file `.prism`, which can be found in your home directory (`~`) under Unix, Linux or Mac OS X and in e.g. `C:\Documents and Settings\username` under Windows. The settings are used by both the command-line and GUI versions of PRISM. Currently the "Options" dialog in the GUI represents the easiest way to modify the settings, but the `.prism` file is in a simple textual format and can be edited by hand.

To restore the default options for PRISM, click "Load Defaults" and then "Save Options" from the "Options" dialog in the GUI. Alternatively, delete the `.prism` file and then run PRISM.

Computation Engines

Computation engines

An important feature of the tool is its *engines*. PRISM is a symbolic model checker: its basic underlying data structures are

BDDs and MTBDDs. When performing numerical computation, however, the tool can use one of three engines. The first is implemented purely in *MTBDDs* (and BDDs); the second uses *sparse matrices*; and the third is a *hybrid*, using a combination of the other two.

The choice of engine ("MTBDD", "sparse" or "hybrid") will not affect the results of model checking - all engines perform the same calculations (although in a few cases, some types of computation can only be performed with certain engines). Performance (time and space), however, may vary significantly and if you are using too much time/memory with one engine, it may be worth experimenting. Below, we briefly summarise the key characteristics of each engine.

- **The hybrid engine** is enabled by default in PRISM. It uses a combination of *symbolic* and *explicit* data structures (as used in the MTBDD and sparse engines, respectively). In general it provides the best compromise between time and memory usage: it (almost) always uses less memory than the sparse engine, but is typically slightly slower. The size of model which can be handled with this engine is quite predictable. The limiting factor in terms of memory usage comes from the storage of 2-4 (depending on the computation being performed) arrays of 8-byte values, one for each state in the model. So, a typical PC can handle models with between 10^7 and 10^8 states (one vector for 10^7 states uses approximately 75 MB).
- **The sparse engine** can be a good option for smaller models where model checking takes a long time. For larger models, however, memory usage quickly becomes prohibitive. As a rule of thumb, the upper limit for this engine, in terms of model sizes which can be handled, is about a factor of 10 less than the hybrid engine.
- **The MTBDD engine** is much more unpredictable in terms of performance but, when a model exhibits a lot of structure and regularity, can be very effective. This engine has been successfully applied to extremely large structured (but non-trivial) models, in cases where the other two engines cannot be applied. The MTBDD engine often performs poorly when the model (or solutions computed from it) contain lots of distinct probabilities/rates; it performs best when there are few such values. For this reason the engine is often successfully applied to MDP models, but much less frequently to CTMCs. When using the MTBDD engine, the *variable ordering* of your model is especially important. This topic is covered in the [FAQ](#) section.

When using the PRISM GUI, the engine to be used for model checking can be selected from the "Engine" option under the "PRISM" tab of the "Options" dialog. From the command-line, engines are activated using the `-mtbdd`, `-sparse` and `-hybrid` (or `-m`, `-s` and `-h`, respectively) switches, e.g.:

```
prism poll2.sm -tr 1000 -s
```

For further information and technical details about PRISM's implementation and engines, see: [\[Par02\]](#), [\[KNP04b\]](#).

Note also that precise details regarding the memory usage of the current engine are displayed during model checking (from the GUI, check the "Log" tab). This can provide valuable feedback when experimenting with different engines.

Iterative Numerical Methods

For performing the computation of probabilities and expected costs/rewards during verification, PRISM uses iterative numerical methods. The methods used vary depending on the type of verification being performed. In some cases, PRISM solves systems of linear equation systems (e.g. "until" properties for DTMCs and CTMCs, steady-state properties for CTMCs and "reachability reward" properties for DTMCs). For this, a range of methods are available (see below). For "until" and "reachability reward" properties of MDPs, PRISM uses a method called "value iteration". For computations involving transient probabilities of CTMCs (e.g. "bounded until" and "cumulative reward" properties), it uses a method called "uniformisation".

Convergence

Common to all of these methods is the way that PRISM checks convergence, i.e. decides when to terminate the iterative methods because the answers have converged sufficiently. This is done by checking when the maximum difference between elements in the solution vectors from successive iterations drops below a given threshold. The default value for this threshold is 10^{-6} but it can be altered with the "Termination epsilon" option (switch `-epsilon <val>`). The way that the maximum difference is computed can also be varied: either "relative" or "absolute" (the default is "relative"). This can be changed using the "Termination criteria" option (command-line switches `-relative` and `-absolute`, or `-rel` and `-abs` for short).

Also, the maximum number of iterations performed is given an upper limit in order to trap the cases when computation will not converge. The default limit is 10,000 but can be changed with the "Termination max. iterations" option (switch `-maxiters <val>`). Computations that reach this upper limit will trigger an error during model checking to alert the user to this fact.

For the specific case of "steady-state convergence checking" during uniformisation (which is an optimisation for uniformisation), convergence checking can be disabled with the "Use steady-state detection" option (command-line switch `-nossdetect`).

Linear equation systems

For instances where PRISM has to solve a linear equation system (see above), the numerical method used can be selected by the user. Below is a list of the alternatives available and the switches used to select them from the command-line. The corresponding GUI option is "iterative method".

- Power method: `-power` (or `-pow`, `-pwr`)
- Jacobi method: `-jacobi` (or `-jac`)
- Gauss-Seidel method: `-gaussseidel` (or `-gs`)
- Backwards Gauss-Seidel method: `-bgaussseidel` (or `-bgs`)
- JOR method (Jacobi with over-relaxation): `-jor`
- SOR method: `-sor`
- Backwards SOR method: `-bsor`

When using the MTBDD engine, Gauss-Seidel/SOR based methods are not available. When using the hybrid engine, *pseudo* variants of Gauss-Seidel/SOR based method can also be used [Par02] (type `prism -help` at the command-line for details of the corresponding switches). For methods which use over-relaxation (JOR/SOR), the over-relaxation parameter (between 0.0 and 2.0) can also be specified with option "Over-relaxation parameter" (switch `-omega <val>`).

More information

For more information about all the issues discussed in this section, see e.g. [Ste94].

Other Options

Output options

To increase the amount of information displayed by PRISM (in particular, to display lists of states and probability vectors), you can use the "Verbose output" option (activated with command-line switch `-verbose` or `-v`). To display additional statistics about MTBDDs after model construction, use the "Extra MTBDD information" option (switch `-extraddin`) and, to view MTBDD sizes during the process of reachability, use option "Extra reachability information" (switch `-extrareachinfo`).

Fairness

Sometimes, model checking of properties for MDPs requires fairness constraints to be taken into account. See e.g. [BK98], [Bai98] for more information. To enable the use of fairness constraints (for **P** operator properties), use the `-fair` switch.

Probability/rate checks

By default, when constructing a model, PRISM checks that all probabilities and rates are within acceptable ranges (i.e. are between 0 and 1, or are non-negative, respectively). For DTMCs and MDPs, it also checks that the probabilities sum up to one for each command. These checks are often very useful for highlighting user modelling errors and it is strongly recommended that you keep them enabled, however if you need to disable them you can do so via option "do prob checks?" in the GUI or command-line switch `-noprobschecks`.

CUDD memory

CUDD, the underlying BDD and MTBDD library used in PRISM has an upper memory limit. By default, this limit is 204800 KB (200 MB). If you are working on a machine with significantly more memory this and PRISM runs out of memory when model checking, this may help. To set the limit (in KB) from the command-line, use the `-cuddmaxmem val` switch. You can also

change this setting in the GUI, but you will need to close and restart the GUI (saving the settings as you do) for this option to take effect.

Java memory

The Java virtual machine (JVM) used to execute PRISM also has an upper memory limit. Sometimes (for example if you are trying to load a model for which the PRISM code is extremely lengthy), this limit will be exceeded and you will see an error of the form `java.lang.OutOfMemory`. The current default limit for the Sun JVM is 64MB. To resolve this problem, you can increase this memory limit. On Unix, Linux or Mac OS X platforms, this is done by setting the environment variable `PRISM_JAVAMAXMEM`, for example (under a `tcsh` shell):

```
setenv PRISM_JAVAMAXMEM 512m
prism big_model.pm
```

or (under a `bash` shell):

```
PRISM_JAVAMAXMEM=512m
export PRISM_JAVAMAXMEM
prism big_model.pm
```

The lower case `m` in `512m` denotes MB. If you are running PRISM on Windows you will have to do this manually by modifying the `prism.bat` or `xprism.bat` scripts. To set the memory to 512MB for example, replace `java` at the start of the last line in the appropriate script with `java -Xmx512m`.

Precomputation

By default, PRISM's probabilistic model checking algorithms use an initial *precomputation* step which uses graph-based techniques to efficiently detect trivial cases where probabilities are 0 or 1. Occasionally, you may want to disable this step for efficiency (e.g. if you know that there are no/few such states and the precomputation process is slow). This can be done with the `-nopre` switch. Please note that it is not always safe to disable precomputation: unless you are checking until properties on MDPs or on DTMCs using the Power method (which is not the default option), then you should only disable precomputation if you know that it will not identify any states which (non-trivially) have probability 0 or 1.

References

AH99

R. Alur and T. Henzinger.
Reactive modules.
Formal Methods in System Design, 15(1):7-48, 1999.

ASSB96

A. Aziz, K. Sanwal, V. Singhal, and R. Brayton.
Verifying continuous time Markov chains.
In R. Alur and T. Henzinger, editors, *Proc. 8th International Conference on Computer Aided Verification (CAV'96)*, volume 1102 of *LNCS*, pages 269-276. Springer, 1996.

Bai98

C. Baier.
On algorithmic verification methods for probabilistic systems, 1998.
Habilitation thesis, Fakultät für Mathematik & Informatik, Universität Mannheim.

BKH99

C. Baier, J.-P. Katoen, and H. Hermanns.
Approximate symbolic model checking of continuous-time Markov chains.
In J. Baeten and S. Mauw, editors, *Proc. 10th International Conference on Concurrency Theory (CONCUR'99)*, volume 1664 of *LNCS*, pages 146-161. Springer, 1999.

BK98

C. Baier and M. Kwiatkowska.
Model checking for a probabilistic branching time logic with fairness.

Distributed Computing, 11(3):125-155, 1998.

BdA95

A. Bianco and L. de Alfaro.

Model checking of probabilistic and nondeterministic systems.

In P. Thiagarajan, editor, *Proc. 15th Conference on Foundations of Software Technology and Theoretical Computer Science*, volume 1026 of *LNCS*, pages 499-513. Springer, 1995.

HJ94

H. Hansson and B. Jonsson.

A logic for reasoning about time and reliability.

Formal Aspects of Computing, 6(5):512-535, 1994.

HLMP04

T. Hérault, R. Lassaigne, F. Magniette, and S. Peyronnet.

Approximate probabilistic model checking.

In *Proc. 5th International Conference on Verification, Model Checking and Abstract Interpretation (VMCAI'04)*, volume 2937 of *LNCS*. Springer, 2004.

HKN+03

H. Hermanns, M. Kwiatkowska, G. Norman, D. Parker, and M. Siegle.

On the use of MTBDDs for performability analysis and verification of stochastic systems.

Journal of Logic and Algebraic Programming: Special Issue on Probabilistic Techniques for the Design and Analysis of Systems, 56(1-2):23-67, 2003.

Hil96

J. Hillston.

A Compositional Approach to Performance Modelling.

Cambridge University Press, 1996.

KSK66

J. Kemeny, J. Snell, and A. Knapp.

Denumerable Markov Chains.

D. Van Nostrand Company, 1966.

KNP04b

M. Kwiatkowska, G. Norman, and D. Parker.

Probabilistic symbolic model checking with PRISM: A hybrid approach.

International Journal on Software Tools for Technology Transfer (STTT), 6(2):128-142, 2004.

KNP07a

M. Kwiatkowska, G. Norman, and D. Parker.

Stochastic model checking.

In M. Bernardo and J. Hillston, editors, *Formal Methods for the Design of Computer, Communication and Software Systems: Performance Evaluation (SFM'07)*, volume 4486 of *LNCS (Tutorial Volume)*, pages 220-270. Springer, 2007.

Par02

D. Parker.

Implementation of Symbolic Model Checking for Probabilistic Systems.

Ph.D. thesis, University of Birmingham, 2002.

Ste94

W. J. Stewart.

Introduction to the Numerical Solution of Markov Chains.

Princeton, 1994.