

The Ipc Manual (developer version)

Allan Clark

October 10, 2007

1 Introduction

PEPA[1] is a popular stochastic process algebra.

About this document . . .

This document is a manual for users of the Imperial Pepa Compiler or `ipc`. Users are expected to already know how to interpret their models. A companion document, "The `ipc` tutorial" is available. The tutorial explains how to use `pepa` for stochastic modelling.

2 Usage

2.1 Trivial Invocation Options

These options only apply to the trivial invocations of `ipc`.

`-v, --version` The command-line:

```
$ ipc --version
```

will print the version of the `ipc` compiler and exit.

`-h, --help` The command-line

```
$ ipc -- help
```

will print an option and usage summary and exit.

2.2 Output Re-directing Options

By default the `ipc` command will generate output in the same directory and with a file name based on the input source name. The options in this section allow the user to redirect the output from the `ipc` compiler.

`--mod-file FILE` Sets the output `.mod` file which will be the input to the run of Hydra.

`--output FILE` A generic output flag, this works regardless of the kind of output that `ipc` is set to produce. For example instead of the default Hydra the compiler may have been set to produce a prism model file. The `--output` can be used to set the output file name.

`--stdout` When debugging the compiler it can often be useful for the output to be redirected to the terminal for immediate inspection by the programmer. The `--stdout` sets the output file to be the standard out. This may also prove useful for piping the output into further processing tools.

2.3 Static Analysis Options

These options control the way that static analysis is performed.

`--staunch` The flag `-- staunch` allows the compiler to ignore warnings. By default this flag is off and when `ipc` performs any static-analysis over the input PEPA model if there are any warnings it will cause the compiler to cancel compilation. This behaviour can be suppressed with the `-- staunch` flag. This will cause the warnings to still be emitted but compilation will proceed anyway.

`--no-static-analysis` The flag `-- no-static-analysis` causes `ipc` to avoid performing the static analysis over the PEPA model. It will therefore produce no warnings or errors. Because of this compilation may fail mysteriously and hence the user is advised only to use this flag if they know exactly what they are doing and expect their model to fail static-analysis for some reason but wish to proceed to compilation anyway.

2.4 Probe Specification Options

This section describes the options which control the performance measure specification probes which are added to the model.

- p, --probe PROBESPEC** The flag `-- probe` with the shortened version `p` is used to give a full probe specification.
- no-master** The flag `-- no-master` is used to specify that no master probe should be automatically added to the model.
- s, --source ACTIONS** The flags `-- source` and `-- target` with the short versions `s` and `t` respectively, set the state switching actions used in the master probe. If other probes are added using the `-- probe` option then they may perform immediate communication actions which are specified in the source and target action list. Both the `-- source` and `-- target` flags accept as argument a comma separated list of action names.
- t, --target ACTIONS** see `--source`

2.5 Running Hydra

This section details options used for running the hydra tool after processing of the model by `ipc`.

- run-hydra** The flag `--run-hydra` causes `ipc` to automatically run the hydra tool on the produced `.mod` file.
- hydra PATH** If the hydra tool is not installed in a standard location the path to the Hydra executable can be given as an argument to the `--hydra` option.

2.6 Miscellaneous Options

This section describes options which don't fit under any of the previous subsections.

- V, --verbose** No manual entry but the usage information states: logfile output to `STDOUT`
- passage** No manual entry but the usage information states: perform a passage-time performance measurement (default)
- steady** No manual entry but the usage information states: perform a steady-state measurement
- transient** No manual entry but the usage information states: perform a transient analysis measurement
- count-measure ACTIONS** No manual entry but the usage information states: perform a count measure over the given actions
- no-measurement** No manual entry but the usage information states: do not output a performance measurement specification

--steady-mean No manual entry but the usage information states: use the 'mean' estimator in a steady-state measure

--steady-variance No manual entry but the usage information states: use the 'variance' estimator in a steady-state measure

--steady-stddev No manual entry but the usage information states: use the 'stddev' estimator in a steady-state measure

--steady-distrib No manual entry but the usage information states: use the 'distribution' estimator in a steady-state measure

--start-time TIME No manual entry but the usage information states: specify a time at which to start a performance measure eg passage-time

--stop-time TIME No manual entry but the usage information states: specify a time at which to stop a performance measure eg passage-time

--time-step TIME No manual entry but the usage information states: specify a the time steps for a performance measurement

--solver SOLVER No manual entry but the usage information states: specify which solution method to use/specify to hydra

--rate DOUBLE No manual entry but the usage information states: Override/specify a rate value on the command-line

--rename-proc P=s No manual entry but the usage information states: cause a renaming on the given process within the model

--rename-rate r=s No manual entry but the usage information states: cause a renaming on the given rate within the model

--prioritise ACTIONS No manual entry but the usage information states: increase the priority of the given actions

--no-reduce-rate-exps No manual entry but the usage information states: do not reduce the rate expressions

--hide-non-coop No manual entry but the usage information states: hide any activities which a component performs but does not cooperate over

--process-num NUM No manual entry but the usage information states: provide a process number which is used to select rates and processes

--show-simplified No manual entry but the usage information states: Show the simplified model in the comments of the compiled file

--fsp **--fsp** this is an experimental option to produce an LTSA model. This should not be considered working

2.7 Using ipc to evaluate a single model

In this section we will learn how to invoke `ipc` to perform various analyses over simple models.

As a simple example consider the PEPA model shown here in the concrete syntax of `ipc`.

```
r1 = 1.0;
r2 = 2.0;
r3 = 2 * r1;

P1 = (start, r1).P2;
P2 = (run, r2).P3;
P3 = (stop, r3).P1;

P1 <> P1
```

2.8 Basic Invocation

Invoking the compiler to give basic information can be done by specifying either the `--version` or `--help` flags.

```
ipc --version

ipc --help
```

Many of the flags and options have a corresponding one letter alias, for example the version can be printed with the command:

```
ipc -v
```

For the remainder of this manual we will use the long versions for clarity. The short versions can be found by reading the output of the help flag.

Beyond the trivial invocations involving version and help information, `ipc` must be invoked with a PEPA model argument. The general form is

```
ipc [flags and options] tiny.pepa
```

The simplest command-line for our `tiny.pepa` file would be:

```
ipc tiny.pepa
```

Such a simple command will produce a `tiny.mod` file suitable for analysis with Hydra. However it is generally more useful to provide some kind of measurement specification. Generally of course the point of solving the model is to make some measurement of the model. The following section 2.8.1 begins with the simplest way in which to specify a measurement.

2.8.1 Specifying source and target options

The `ipc`/Hydra tool chain can be used to calculate three major types of performance measure:

1. Passage time quantiles and distributions

2. Transient measures
3. Steady state measures

There are two ways in which to specify the states in which the measurement is concerned. In this section simple measurements which can be done by specifying sets of source and target actions are considered. Section 3 is concerned with more complex measurements.

As the names suggest, the source activity is used to start the measurement and the target activity is used to end the measurement.

Thus if we were interested in computing the passage from an occurrence of the `start` activity to the occurrence of the `stop` activity by either of the copies of the process in the tiny model then we would use the following command

```
ipc --source start --target stop tiny.pepa
```

This will generate a Hydra model file (`tiny.mod`) to be analysed by a subsequent run of Hydra. With each of these arguments more than one action may be given as a comma separated list of actions. For example one might give a command such as:

```
ipc --source halfFull,halfEmpty --target full,empty optimism.pepa
```

Since Hydra will more often than not be invoked on the output `.mod` file one can specify the `--run-hydra` flag. By default the kind of measurement is a passage-time measurement and `ipc` knows this and will run the corresponding `hydra-uniform` command after the initial run of Hydra itself. So by issuing the command:

```
ipc tiny.pepa --source start --target stop --run-hydra
```

You should produce the file: `tiny.PT_RESULTS`. The `gnuplot` program can then be used to produce a graph of the results such as the one in Figure 1. This graph plots the cumulative probability of completion of the passage against time. That is the probability that t seconds after observing a `start` activity a `stop` activity is observed.

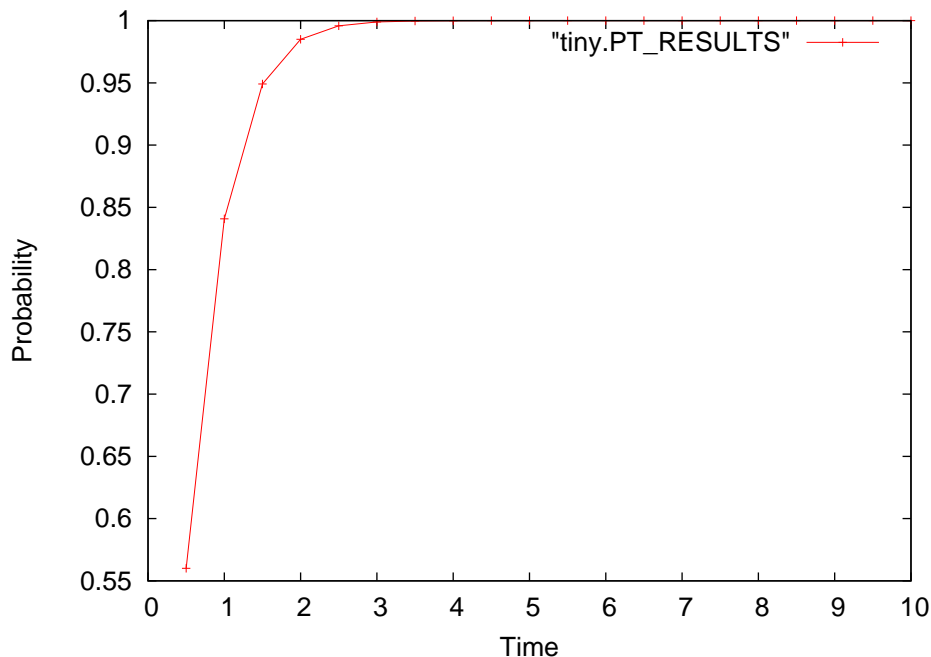
2.9 Changing the analysis type

We can specify the three different kinds of analysis with the three flags:

- `--steady`
- `--transient`
- `--passage`
- `--no-measurement`

A passage-time measure is the default unless the user does not specify either a measurement probe or a set of source/target actions. In this case `ipc` defaults to no measurement.

*(**TODO: Explain what each of the measurement kinds mean and in particular for the given example what it would compute ??**)* .



htb

Figure 1: Shows the cumulative probability of completion against time.

2.10 Re-directing the output

The output model used as input to the Hydra tool will by default be placed in the file `tiny.mod`, but this can be overridden with the command line option `--mod-file` as in:

```
ipc --source start --target stop --mod-file ipc.mod tiny.pepa
```

To aid debugging it's often useful to dump the output to standard out for inspection. The `--stdout` flag redirects the output to the terminal.

Note that as mentioned above we will be giving command line options in their long forms, but most of them have short forms. The `--start` and `--target` options have the short forms `-s` and `-t` respectively. Because start and stop actions must be given for almost all runs of `ipc` we will use the short options to reduce the length of command lines where appropriate in the remainder of this manual.

2.11 Specifying start and stop times

Via Hydra, `ipc` computes performance measures such as probability density functions (PDFs) and cumulative distribution functions (CDFs). Together the PDF and CDF give a complete description of the probability distribution of a random variable. Both are evaluated against increasing time in order to produce function plots. `ipc` has default values for the start time and stop time for the plot, as well as the time step which determines how many evaluations of these functions Hydra is to do. The command-line options `--start-time`, `--stop-time` and `--time-step` allow us to override these defaults.

To specify a start time of 50 seconds and an end time of 100 seconds with a timestep of 10 seconds we would pass the following command-line options to `ipc`.

```
ipc --starttime 50 --stoptime 100 --timestep 10 ...
```

3 Probe Specifications

4 Hydra Compilation Details

This section details the transformations that the input `pepa` model goes through in the translation to a Hydra model. We begin by listing the stages and then each stage will be discussed in greater detail.

- The input PEPA model is parsed.
- The syntactic sugar forms are removed: these are such constructs as:
 - Parallel Definitions eg $P = Q \langle \rangle R$.
 - Non-aliased sub-components, eg $(a, r).P + (b, r).Q$ is turned into $P1 + Q1$ with the appropriate definitions added.
- Qualify the model. Briefly this means that each defined process is used exactly once from the main cooperation.
- Remove the hiding operator

At this point the model is still in PEPA format and could be output suitably for other PEPA related tools.

- This model is then translated into a structured form of the Hydra transitions. Essentially this is sets of `dnam` transitions which may cooperate with each other.
- The structured sets of Hydra transitions are flattened into one list. At this point the apparent rates are calculated.
- The Hydra model is formatted appropriately for the Hydra[2] Markov chain analyser.

4.1 Removal of Syntactic Sugar

As we will see later PEPA models of a specific form are more convenient to convert into a model suitable for Hydra. The specific form of a PEPA model is a normal form, the grammar for this form of PEPA model is given in Figure 2. The interesting notions are the following:

- There are no parallel definitions eg($P \stackrel{def}{=} Q \bowtie R$) These are removed by substituting their definitions into the main system equation.
- There are no alias definitions, eg($P \stackrel{def}{=} Q$). Again these are removed by substitution.
- All component sums have only identifiers as their left and right choices.
- Prefix operations only have identifiers as their successor components.

$$\begin{aligned}
def & := P \stackrel{def}{=} sequential; \\
sequential & := (a, r).P \\
& \quad | P + Q \\
parallel & := P \\
& \quad | P < actions > Q \\
& \quad | P/actions \\
& \quad | P[N][actions] \\
model & := def^+ parallel
\end{aligned}$$

Figure 2: The grammar for a pepa model in normal form.

$$\begin{aligned}
P & \stackrel{def}{=} (a, r).Q \\
Q & \stackrel{def}{=} (b, r).P \\
& \quad P \underset{a}{\bowtie} P \\
\\
P_1 & \stackrel{def}{=} (a, r).Q_1 \\
Q_1 & \stackrel{def}{=} (b, r).P_1 \\
P_2 & \stackrel{def}{=} (a, r).Q_2 \\
Q_2 & \stackrel{def}{=} (b, r).P_2 \\
& \quad P_1 \underset{a}{\bowtie} P_2
\end{aligned}$$

Figure 3: An example model and the qualified version.

4.2 Qualifying the Model

Definitely rewrite this better. The model in normal form however is still not ready to be converted into a Hydra model. We must distinguish between each invocation of a given defined process. The transformation for this must be deep, in that it must see through definitions otherwise one process may become another one and hence violate the uniqueness property that we require.

The transformation provided adds a suffix number to each process name. Figure 3 gives an example of qualifying a very simple PEPA model. We start with the main composition and the list of original definitions are used as a dictionary but are not included in the final model. The transformation then must return a new main parallel composition plus a new list of process definitions.

The first part of the algorithm requires us to qualify a sequential process. We turn a given process identifier into a list of new definitions. We assume that $lookup(P)$ will return the sequential component which the process identifier P is bound to in the original list of process definitions (which will ultimately be discarded).

In order to do this we define two mutually recursive functions. Since most components will be cyclic we must end the recursion, we do this by maintaining a list of 'seen' process identifiers. Whenever we attempt to redefine a process in the 'seen' list then we return the empty set of new definitions. The T_{def} function checks if we have already seen the given process identifier and if not looks it up the initial list of process definitions. The T function takes in a process name to define and a sequential component to qualify. Notice that this algorithm is trivial because we already know that the model is in normal form.

$$\begin{aligned} T(i, \text{seen}, P, (a, r).Q) &= [P_i = (a, r) . Q_i] \text{ union} \\ &\quad T_{def}(i, \text{seen}, Q) \\ T(i, \text{seen}, P, Q + R) &= ([P_i = Q_i + R_i] \text{ union} \\ &\quad T_{def}(i, \text{seen}, Q) \text{ union} \\ &\quad T_{def}(i, \text{seen}, R) \\ T_{def} = (i, \text{seen}, P) &= \text{if } P \text{ is in seen then } [] \\ &\quad \text{else } T(i, P:\text{seen}, P, \text{lookup}(P)) \end{aligned}$$

The second part of the algorithm proceeds via recursion through the main system component. This simply updates a counter upon each process encountered. Because we know we have already removed parallel definitions and aliases each encountered process can be assumed to be a sequentially defined component. We must return three results; the list of newly defined components, the updated suffix number, and the qualified parallel process.

$$\begin{aligned} T(i, P) &= (i+1, T_{def}(i, [], P), P_i) \\ T(i, r \langle \text{acts} \rangle s) &= \text{let } (j, \text{defs}_r, r_i) = T(i, r) \\ &\quad (k, \text{defs}_s, s_i) = T(j, s) \\ &\quad (k, \text{defs}_r \text{ union } \text{defs}_s, r_i \langle \text{acts} \rangle s_i) \\ T(i, p/\{\text{actions}\}) &= \text{let } (j, \text{defs}, p_i) = T(i, p) \\ &\quad (j, \text{defs}, p_i/\{\text{actions}\}) \\ T(i, P[N][\text{actions}]) &= \text{let } (j, \text{defs}, p_i) = T_{def}(i, [], p) \\ &\quad (j, \text{defs}, p_i/\{\text{actions}\}) \end{aligned}$$

4.3 Removal of the Hiding Operator

The hiding operator can now be trivially removed. This is because the model is in qualified form. The hiding operator can be removed simply by renaming the hidden actions. This works because each occurrence of a process is unique therefore its definition is used only once and each action that it may perform is either hidden or not hidden. Figure 4 gives an example qualified model and the same model with the hiding operator removed. Notice that the processes P_1 and P_2 may still cooperate over the a action, though it is renamed to a_1 .

4.4 Structured Hydra Transitions

Because of the previous set of transformations the PEPA models which reach this stage of compilation are in a very strict form. Each sequential component can be easily transformed into a list of transitions which it may perform. Where a transition consists of a rate and the pre-conditions for the transition to be active, and the post-conditions or effects on the model's state that the transition has.

$$\begin{aligned}
P_1 &\stackrel{def}{=} (a, r).Q_1 \\
Q_1 &\stackrel{def}{=} (b, r).P_1 \\
P_2 &\stackrel{def}{=} (a, r).Q_2 \\
Q_2 &\stackrel{def}{=} (b, r).P_2 \\
P_3 &\stackrel{def}{=} (a, r).Q_3 \\
Q_3 &\stackrel{def}{=} (b, r).P_3 \\
&\quad (P_1 \bowtie_a P_2) / \{a\} \bowtie P_3
\end{aligned}$$

$$\begin{aligned}
P_1 &\stackrel{def}{=} (a_1, r).Q_1 \\
Q_1 &\stackrel{def}{=} (b, r).P_1 \\
P_2 &\stackrel{def}{=} (a_1, r).Q_2 \\
Q_2 &\stackrel{def}{=} (b, r).P_2 \\
P_3 &\stackrel{def}{=} (a, r).Q_3 \\
Q_3 &\stackrel{def}{=} (b, r).P_3 \\
&\quad (P_1 \bowtie_{a_1} P_2) \bowtie P_3
\end{aligned}$$

Figure 4: An example qualified model showing the removal of the hiding operator.

The pre-conditions at this stage is simply that there be at least one of the named process.

Sequential components then are turned into a list of transitions for the P_1 component of the example given in Figure 4 we can make the list of transitions:

```
{ kind = a
  rate = r
  preconditions = P_1 > 0
  postconditions = P_1 = P_1 - 1
                  Q_1 = Q_1 + 1
}
{ kind = b
  rate = r
  preconditions = Q_1 > 0
  postconditions = Q_1 = Q_1 - 1
                  P_1 = P_1 + 1
}
```

These sets are of transitions are stored within the original models structure for the main system equation. So that we have sets of transitions cooperating over action names.

4.5 Flattened List of Transitions

The structured sets of transitions are combined together into single list of transitions. Where the transitions are not cooperated over this is a simple union. Where the transitions cooperate then two transitions are combined into a single one. The apparent rate of the single transition is calculated at this stage.

5 Attaching Performance Measurement Probes

References

- [1] J. Hillston. *A Compositional Approach to Performance Modelling*. Cambridge University Press, 1996.
- [2] J.T. Bradley, N.J. Dingle, S.T. Gilmore, and W.J. Knottenbelt. Extracting passage times from PEPA models with the HYDRA tool: A case study. In S. Jarvis, editor, *Proceedings of the Nineteenth annual UK Performance Engineering Workshop*, pages 79–90, University of Warwick, July 2003.