Bio-PEPA Manual

Version March 2012

Authors: Allan Clark and Vashti Galpin

March 30, 2012

Contents

For up-to-date information about help and contacts, please see our website:

`http://www.biopepa.org`

Chapter 1

Introduction

This manual is intended as a quick-start guide to the Bio-PEPA Eclipse Plug-in. It is not intended as a reference manual to the Bio-PEPA language itself.

Chapter 2

The Graphical Views

2.1 The Editor View

The editor pane is the area where your model construction takes place. This is essentially an advanced text box which colourises your Bio-PEPA code appropriately. In addition, when the model is saved some validation, or static checking, of the model is performed. The results will be stored in the Problems View 2.3 but are also linked into the main editor pane, hence warnings and errors are underlined in yellow and red respectively. Additionally when the editor pane is in focus the items in the Bio-PEPA menu become available. The editor pane is shown in Figure 2.1. The table in Figure 2.2 shows the correspondence between Bio-PEPA formal operator syntax and the text representations of those operators which the Bio-PEPA software accepts and should be written in the editor pane.

2.2 The Outline View

Once the model is saved, if it is syntactically valid Bio-PEPA, then the outline view is automatically updated with the outline of your model. This describes each species and each reaction in a tree view. The outline view is depicted in Figure 2.3.

2.3 Problems View

Upon saving your Bio-PEPA model, the problems view is updated. The problems view is shown in Figure 2.4 and displays the results of static analysis on your model. The are two kinds of results, warnings and errors. Warnings are displayed in yellow and correspond to features of your model which the compiler believes may constitute a mistake. These warnings should be reviewed, however if you are sure that what you have written is indeed what you intended to write then the software is still able to execute or numerically evaluate your model.

Errors are displayed in red, these constitute serious problems with the model in that

goldbeter.biopepa | a-b-c.biopepa

```
omega   = 6022; // ~6.0221415E23 * 1E-14 * *1E-6
vi      = 0.025*omega;
kd      = 0.01;
kc      = 0.5*omega;
v1      = 12*omega;
v2      = 1.5;
v3      = 12;
v4      = 2;
k1      = 0.02*omega;
k2      = 0.02*omega;
k3      = 0.02*omega;
k4      = 0.02*omega;
vd2     = 0.0625;
kd2     = 0.02*omega;

a1  = [ fMA(vi) ] ;
a2  = [ fMA(kd) ] ;
a3  = [ ((v1*_C)/(kc + _C))*(inactive_M/(k1 + inactive_M)) ] ;
a4  = [ fMM(v2,k2) ] ;
a5  = [ fMM(v3,k3) ] ;
a6  = [ fMM(v4,k4) ] ;
a7  = [ (_C*vd2*active_X)/(_C + kd2) ] ;

_C          = a1 >> + a2 << + a3 (+) + a7 << ;
inactive_M  = a3 << + a4 >> ;
active_M    = a3 >> + a4 << + a5 (+);
inactive_X  = a5 << + a6 >> ;
active_X    = a5 >> + a6 << + a7 (+) ;
enzyme      = a4 (+) + a6 (+) ;

// _C[60] <*> active_M[60] <*> inactive_M[5962] <*> active_X[60]
//       <*> inactive_X[5962] <*> enzyme[6022]
_C[6] <*> active_M[5] <*> inactive_M[5] <*> active_X[5]
      <*> inactive_X[500] <*> enzyme[500]
```

Figure 2.1: The editor pane used to construct and modify the Bio-PEPA source of your models.

Operator	Text
$\downarrow$	<<
$\uparrow$	>>
$\odot$	(.)
$\oplus$	(+)
$\ominus$	(-)

Figure 2.2: Table showing the text representations of Bio-PEPA operators

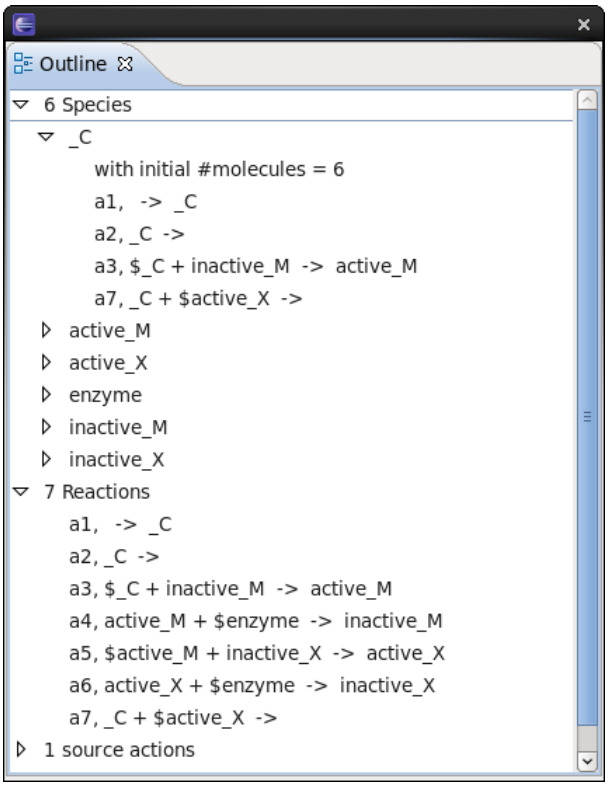

Figure 2.3: The outline view which is updated when your model is saved

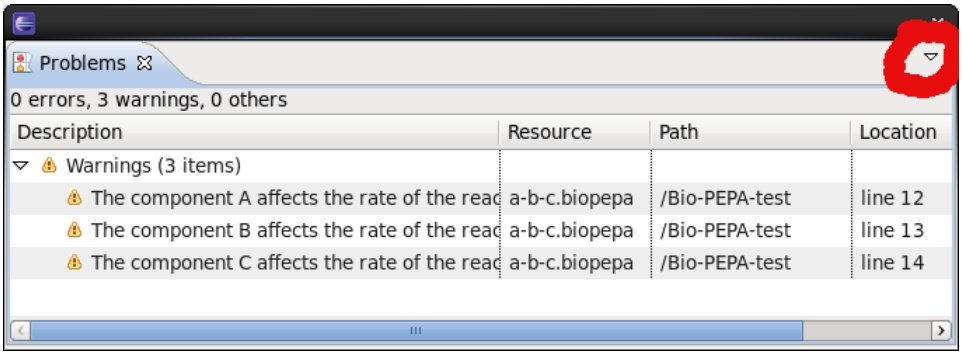

Figure 2.4: A screen shot of the problems view. Circled is the extra menu which allows one to limit the warnings and errors shown to the currently viewed model.

the compiler does not know how to numerically evaluate your model and hence such operations cannot be performed until the problem is fixed.

Clicking on warning or error descriptions in the problems view will cause the focus to shift to the editor pane and the problem part of the model will be highlighted with the cursor at the end.

A third class of problem is the set of messages labelled 'Info', these can be regarded as the same as a warning.

As an example, a warning might warn that the rate of a reaction which consumes a particular species 'A' does not depend upon the population of 'A'. This is a warning because it is rather unusual and means that population of 'A' may become negative (because the reaction rate may be higher than zero even when the population of 'A' is zero). However the model may still be evaluated. However, if you attempt to provide a behaviour for a particular reaction which does not have an associated rate law, then this constitutes an error since the compiler has no way to numerically evaluate this model.

By default all of the problems with models which are open in the editor pane, but not necessarily the current tab, are shown. This behaviour can be modified by selecting "Show → Errors/Warnings on Selection" from the menu accessed by the small arrow in the top right hand corner of the problems view and circled in Figure 2.4.

2.4 Graph View

Once time series analysis has been performed over your model, the results are displayed in line-graph form in the Graph View. From here these graphs may be saved as either an image (in portable network graphics format) or as the data, in a comma-separated-value text file. The graph view is shown in Figure 2.5.

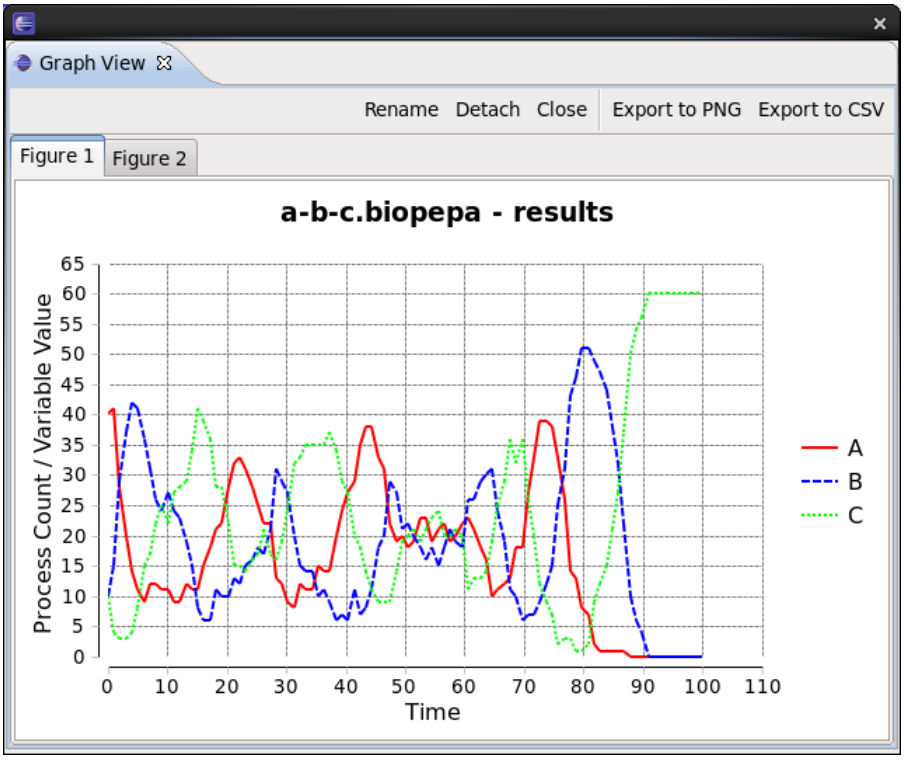

Figure 2.5: A screen shot of the graph view

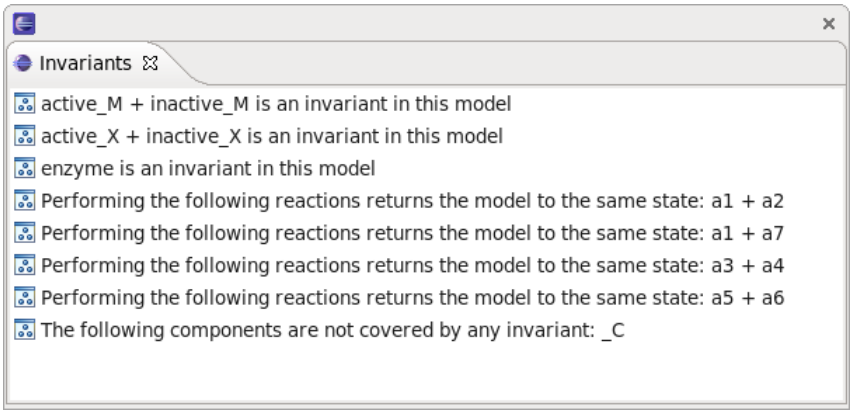

Figure 2.6: A screen shot of the invariants view

2.5 Invariants View

Performing invariant analysis over your Bio-PEPA model will cause the results to be displayed in the Invariants View. This is a simple view which describes each of the component and reaction invariants discovered for the model. It will also warn you of any components which are not involved in any invariant. The invariants view is shown in Figure 2.6.

Chapter 3

Available Analyses

3.1 Performing Timeseries Analysis

All the actions you will perform can be found under the 'Bio-PEPA' menu. The most important is the "Time series analysis ..." menu item. The wizard which appears is used to numerically solve the model using either a Stochastic Simulation Algorithm (SSA) or Ordinary Differential Equations (ODEs).

However time series analysis can be computationally expensive so one may wish to perform other analyses first to determine that your model is correct. In particular one might perform invariant analysis see section 3.2

The first screen in the Time Series Analysis Wizard allows you to select which species and variables in your model you are interested in tracking. The species you select here will be the species which are returned by the solver and are hence represented in the graph of results. Clicking next will take you to the algorithm specification page. Here the first thing you should decide upon is which algorithm you wish to use to perform the time series analysis. The first three are algorithms perform stochastic simulation, Gillespies, Gibson-Bruck and Gillespie's Tau-Leap. The final three algorithms perform ODE analysis, Adaptive step-size 5th order Dormand-Prince, Implicit-Explicit Runga Kutta and our own implementation of a Runga Kutta solver.

You will see that the remaining options on this page change depending upon which solver is selected. The number of independent replications is important for stochastic simulations but meaningless for an ODE analysis. In any case here is the point at which you need to set your start and stop times. The stop time is where the simulation ends while the start time is where the results start reporting, the simulation itself always begins at time t=0.

There are remaining pages to this wizard but they can be safely ignored and are there to allow for parameter scan experiments and comparison with experimental data.

3.2 Invariant Analysis

To perform invariant analysis on the model select "Infer Invariants ..." from the Bio-PEPA menu. The opening screen of the resulting wizard allows you to select the kind of invariants that you wish to infer. State Invariants refer to a sum of some or all of the populations of the species in the model. The activity invariants refer to a set of reactions which, if all performed exactly once, will return the model to the same state it was in prior to the activity invariant. The simplest activity invariant is a reversible reaction, if both directions of the reaction are performed once then the model is returned to its original state.

On the second page one is given the opportunity to ignore some of the reactions in the model for the purposes of invariant analysis. Clicking finish will perform invariant analysis and display the results in the Invariants View 2.5.

The algorithm used to perform invariant analysis is a greedy algorithm which performs well in most cases. The worst case performance however is exponential, meaning that it may take a long time to complete. It is also the case that it may take a long time to infer the state invariants but not the activity invariants or vice-versa. In general you can safely perform both sets of analysis on your model but if you find that invariant analysis appears to be stuck, try selecting only state or only activity invariants on the first page of the wizard.

3.2.1 Conservation of Mass

When performing state invariant analysis the user is warned about any species which are not involved in a single state invariant. If there are no such species then the model is entirely covered by state invariants and hence we could sum those individual invariants to obtain a single state invariant which covers all the species in the model. If this is the case, then it means that the amount of mass in the system as modelled remains constant, as we would expect it to.

However most models will have some source or sink reactions. These are reactions which consume or produce mass seemingly from/to nothing. In reality those source or sink reactions are really modelling abstractions for some process which consumes or produces some components which are outside of the realm of the model. So one way to make a consistency check on your model is to perform state invariant analysis over the model and use the second page of the wizard to ignore any source or sink reactions. If your model is still not completely covered by invariants then it suggests that there is an inconsistency within your model. Of course it might be that your model, as an abstraction, does not conserve mass and deliberately so. This would be the case in some model with births and/or deaths. However any such inconsistency in your model you should be able to explain and invariant analysis performed in this manner remains a useful model validation technique.

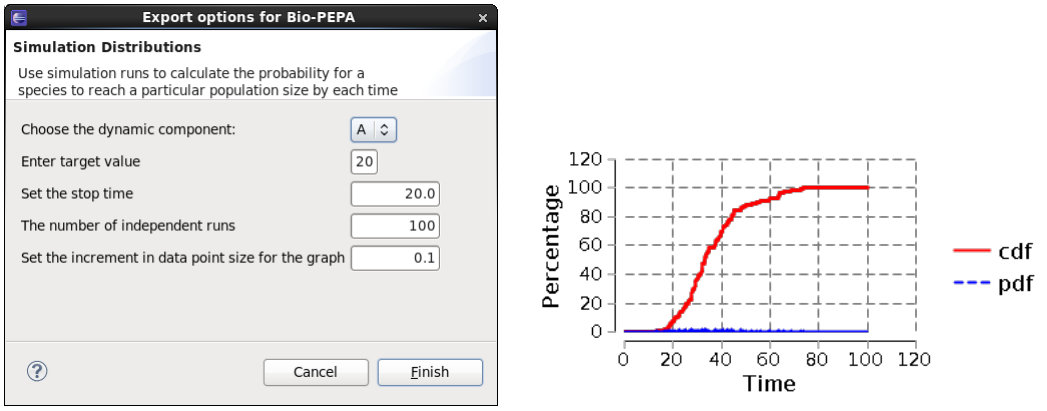

Figure 3.1: A screen shot of the Simulation Distributions wizard and a results graph.

3.3 Simulation Distributions

Performing many stochastic simulations provides an opportunity to obtain many more statistics about our model. One possibility is to obtain the percentage of simulations for which some property is true at or before a given time t. The menu item "Simulation Distributions" opens a wizard which uses stochastic simulation to perform this kind of analysis.

On the first page of the wizard the user is asked to select a component and a target population. The condition which must be met is that the chosen component reaches a prescribed population count. Once the simulation time is set by the user and 'Finish' clicked the software will simulate the model the defined number of times. The output is a graph which shows two lines, the most interesting of which is the red 'cdf' line. This plots at time t the percentage of simulations in which the condition was met at or before time t. Figure 3.1 shows the wizard for simulation distributions as well as an example results graph.

Chapter 4

Frequently Asked Questions

Q. All the items under the Bio-PEPA menu are greyed out and cannot be selected

A. Make sure the editor pane is selected

Q. I cannot see the Invariants View

A. From the "Window" menu select "Show View → Other" the Invariants view will be under the "Analysis" section

Q. I cannot see one of the views described here

A. From the "Window" menu select "Show View → Other" and use the "Search" bar to search for the name of the view you require.