The PEPA Eclipse Plug-in

A modelling, analysis and verification platform for PEPA

Adam Duguid, Stephen Gilmore, Michael Smith and Mirco Tribastone

Wednesday 01 December 2010

Abstract:

This user manual describes the PEPA Eclipse Plug-in, a modelling environment for the high-level modelling language Performance Evaluation Process Algebra, PEPA [http://www.dcs.ed.ac.uk/pepa]. The PEPA Eclipse Plug-in allows modellers to analyse their PEPA models using Markov chain theory, model-checking, and discrete or continuous simulation. The conversion from models in the PEPA process algebra to Markov chain models, simulation models, or differential equation models is entirely automatic. Analysis results are returned to the user in the form of graphs, charts and time series plots.

A version of this document suitable for printing is available online [http://www.dcs.ed.ac.uk/pepa/documentation/plugin.pdf].

Contents

1	Introduction	2
	1.1 Getting Started	2
	1.2 The PEPA Perspective	2
2	Editing PEPA models	3
	2.1 The PEPA Editor	3
	2.2 Basic Syntax	4
3	Views	6
	3.1 The Abstract Syntax Tree View	6
	3.2 Graph View	6
4	Markovian analysis	8
	4.1 State Space Derivation	8
	4.2 State Space View	8
	4.3 Filtering the State Space	9
	4.4 Single Step Navigator 12	2
	4.5 Exporting the State Space	4
	4.6 Steady-State Analysis	5
5	Performance analysis 18	8
	5.1 Performance Evaluation View	8
	5.2 Experimentation $\ldots \ldots 1$	9
6	Discrete and continuous simulation 24	4
	6.1 Time Series Analysis	4
7	Importing UML models 23	5
	7.1 Importing from UML	5

Introduction

1.1 Getting Started

This guide provides a step by step description of the PEPA plugin.

It shows how you can edit PEPA models, carry out static analysis, obtain the underlying Continuous Time Markov Chain for steady-state analysis or perform time series analysis using Stochastic Simulation Algorithms or Ordinary Differential Equations. This guide will also describe how to use import/export tools for data exchange with third-party applications.

Familiarity with PEPA is assumed. To learn more about PEPA visit the PEPA web site [http://www.dcs.ed.ac.uk/pepa/].

1.2 The PEPA Perspective

The plugin provides an editor for PEPA model input files. The editor contributes a top-level menu item to the workbench bar through which all the PEPA-related tools are accessible. Additional information is provided via workbench views. The PEPA perspective provides a customisable layout for those views in the workbench.

To open the PEPA Perspective, select **Window > Open Perspective > Other**. Then select **PEPA** and press **OK** to finish.

Editing PEPA models

2.1 The PEPA Editor

The PEPA Editor lets you edit PEPA input model files. It is associated to open files with the .pepa extension in the workbench. In order for to open a PEPA input file in the Eclipse workbench, a project container has to be chosen.

To create a new project, select **File > New > Project...** . Then select **General > Project** . Choose a name and press **Finish** to complete.

To create a new file, right click the project and select File > New > File. Give the new file the extension .pepa

000	PEPA – paper.pepa – Eclipse SD	K
] 🖆 • 🔒 🖆] 💁 •] 🖋] 🕫	•]9]+8]+80+0+0+	🖹 🔀 PEPA 🖏 Java
Image: Second	<pre>M paper.pepa % r = 1.0; P1 = (a,r).P2; P2 = (b, 1).P1; Q1 = (a, infty).Q2; Q2 =(c,1).Q1; P1[1] <a> Q1[1]</pre>	Cutilisation Throughput
	Problems AST View State Space Vie O errors, 0 warnings, 0 infos Description	ew Graph View Properties Error Log

The input is parsed when the file is saved. If the parsing is successful static analysis is then carried out. Warning and errors generated by both the parser and the static analysis tool are shown in the Eclipse Properties View.



2.2 Basic Syntax

Rate Assignments must precede any process definition. A rate identifier is valid if it is a valid Java identifier starting with a lowercase letter. Rate assignments support expressions. They must end with a semicolon

> r1 = 1.0; r2 = 2.0 * r1 - 0.05;

Process Assignments A process identifier is valid if it is a valid Java identifier starting with an uppercase letter. As usual an identifier can represent a Choice or a Prefix. It can also identify subparts of the system when it is assigned a cooperation. Process assignments must end with a semicolon

$$P1 = (a, r1).P1 + (b, r2).P1;$$

Passive Rates can be specified as either infty or T. Weights can be assigned as integer numbers. Weights however cannot be rate identifiers. If no weight is specified, the default value 1 is assumed.

```
(a, T).P
(b, infty).P
(a, 3 * infty).P
```

Cooperation Sets are enclosed by . Cooperations with empty action set can be alternately specified with the parallel operation ||.

Aggregation is partially supported by the plugin. If a cooperation on an empty action set between n is declared as P[n] then a canonical form representation for the process is considered during state space derivation of the underlying Markov chain.

Views

3.1 The Abstract Syntax Tree View

The AST View is a read-only Eclipse view for navigating the Abstract Syntax Tree of the model. The view is always associated to the currently active editor in the workbench.



3.2 Graph View

The Graph View is a tab-based view responsible for the management of all the graphs produced by the PEPA plugin. It is automatically opened when a new graph is requested. If you want to open it explicitly, select **Window > Show View > Other...** and then select **PEPA > Graph View**. The view support tools for editing graph properties such as background color, line thickness and legends. In addition, it provides a number of exporters to save the image to graphic formats such as PNG or serialise the data to file.



To edit the graph, right click it and select **Properties...** . To serialise the data, click the drop-down button in the toolbar menu of the view in the top-right corner of the view and select the desired format. A wizard will guide you through the process of setting exporter options and selecting target resources.

Markovian analysis

4.1 State Space Derivation

Once a model is successfully parsed, Markovian analysis can begin via state space derivation of the underlying Markov chain.

To derive the state space of the model in the active editor in the work bench, select ${\bf PEPA} > {\bf CTMC} > {\bf Derive}$.

A tabular representation of the derived state space is shown in the State Space View

4.2 State Space View

The State Space View allows you to navigate the underlying CTMC of a PEPA model. After a model is derived, the view presents a tabular representation of the state space. The first column shows the state number, then there are as many columns as the number of top-level components. An additional column shows the steady-state probability of the state after the model is solved.

000	PEPA – paper.pepa – Eclipse SDK	
] 📬 • 🔛 👜] 💁] 🛷] 🖉 •	•]∲-§-*\$-	🗈 🔀 PEPA 🖏 Java
Image: Second		Performance 23 Outline P Utilisation Throughput aph View Properties Error Log
] =]	1

4.3 Filtering the State Space

The view provides filters for the state space. Here are the filter rules available. States with unnamed sequential components. States whose steady-state probability is in a given range. States matching a given pattern . States which have a top-level component in a given state.

A number of rules can be grouped into a collection. The collection is assigned a name which is unique within the model's rule sets. You can specify any number of rule sets. However, only one rule set can be applied at a time. The rule setting are persistent across workbench sessions.

To open the filter rule editor, click the drop-down arrow in the State Space View toolbar.



The State Space Filters dialog box will show the currently available filter sets for the model.

PEPA - Blank Model, pepa - Eclipse SDK File Edit Navigate Search Project Run PEPA Wit			
	₩ - ♥ ♦ • + +		E M PEPA
Image: Navigator 12 Image: BankScenario Image: BankSc	Image:		del.pepa 23 "45 C).Informed; anRequest).Waiti: eHessage,infty). le; est,infty).Reque: est,infty).Reque: et,infty).WaitingFo: ty).EvaluatingFo: ty).EvaluatingFo: ty).EvaluatingFo: liableMessage,gov
	Run	Cancel	
	1 Idle 2 Informed 3 WaitingBankResponse 4 WaitingBankResponse 5 WaitingBankResponse 6 OfferReceived 7 WaitingBankResponse	WaitingForCustomer WaitingForCustomer RequestReceived CustomerReliable WaitingManagerResponse WaitingForCustomer CustomerNotReliable	WaitingForEmployee WaitingForEmployee WaitingForEmployee WaitingForEmployee EvaluatingForEmployee WaitingForEmployee WaitingForEmployee
: 📭 BankScenario		:	

Click **New...** to edit a new set of rules. You are required to give it a name. The two radio buttons let you select between matching any or all of the defined rules. The rules are shown in the list box.

🖯 🔿 🖯 Fi	lter Rules
lame cannot be empty	
Filter name:	
• Match all of the following	O Match any of the following
	(-)
	Cancel OK
	Cancel OK

To manage the rule list, right click the list box. Click \mathbf{New} to show the supported rules.



4.4 Single Step Navigator

The Single Step Navigator is tool for walking the state space graphically. It is particularly useful for debugging purposes. It consists of two tables containing the list of incoming and outgoing states. The sequential components which cause the transition to be performed are highlighted and an option allows you to make filtered states not walkable.

To open the navigator, right click a state in the state space view and select **Single Step Navigator**. The tool will be shown in a separate dialog box.



PEPA - File Edit	Bla Nav	a nk Model.pepa - Eclip igate Search Project Ri	se SDK un PEPA Win	dow Help					
1.	1 4	👌 i 💁 🔹 🛷 i 🖗	• 1 2 - 1	- ₹	• 🔿 •			E M PEPA	»
🖏 Naviga	tor	8		paper2new	v.pepa 🛛 🔀 p	aper2.pepa	🖂 🖾 Blank Mod	del.pepa 🛿 🦯 4	s
	_			//Rates	; 	.			
🖶 🗁 Ba	٢	Single Step Navigator						_ 🗆 🔀	
🕀 🧀 Pi		Only states unaffected by fi	lters can be rea	ched					
🕒 🧁 Te		Enable colouring							
🗄 🗁 te	Inc	coming States							
	2	Informed	WaitingForO	ustomer	WaitingForEmp	loyee			
	5	waitingBankResponse	WaitingMana	agerResponse	EvaluatingOffe	r			=
	-								
	-								1+1
									y).
	2 ir	ncoming states							
	Гu	rrent State							que:
	3	WaitingBankResponse	RequestRec	eived	WaitingForEmp	loyee			Not:
									iti:
	Ou	tgoing States							gFo:
	4	WaitingBankResponse	CustomerRe	liable	WaitingForEmp	loyee			
	-	waitingbankResponse	CustomerNo	tReliable	WaitingForEmp	loyee			
	-								
									8
	20	utgoing states							8
									8
							Go to	Finish	
			1	7 waiungbai	nikkesponse	customenvi	лкенаре	waicing=orempioya	re re
0				<u> </u>					

Double click a state to set the new current state in the navigator.

4.5 Exporting the State Space

The State Space View provides a wizard to export the state space into a csv (comma-separated values) files. To open the wizard, click the export button in the top right corner of the state space view.

🛢 PEPA - Blank Model, pep	~ E 11 - 0014	
File Edit Navigate Search R	÷	
i 📬 ▼ 📄 👜 i 💁 ፣ i ,	Exporter Options Export the state space data model using the following settings	del.pepa 2 *45
BankScenario	Common Options Separator , State Space Options	
Deployment.profile.u Deployment.profile.um PEPA.profile.um ProfileBase.profile.u Deployment.profile.u ProfileBase.profile.u Deployment.profile.u ProfileBase.profile.u Deployment.profile.u ProfileBase.profile.u Deployment.profile.u	Include state number Complete state Separated top level components	≣
🖲 🗃 test	Include current steady-state solution).Informed; anRequest.Waiti: eMessage,infty). le; est,infty).Reque; rateReliabilityC: Woorcom pateWat
		skManager).WaitingFo
	Image: Idea Idea Finish Cancel 1 Idea WaitingForCustomer WaitingForCustomer 3 WaitingBanResponse ResuestReceived	WaitingForEmployee WaitingForEmployee WaitingForEmployee

The wizard will produce two files, the generator matrix and the state space in the form of an ordered list of strings. The separator character will be used to delimit the values in the generator matrix files. If the **Include state number** option is checked, the state number will be prepended to state. The state can be represented fully or as an array of separated top level components. Moreover, if the steady-state solution is available, the steady state probability can be included in the state space file. The values of the state space files will be delimited using the separator character of the generator matrix.

Click **Next** to select the destination for the target files. The wizard doesn't allow to overwrite existing files.

4.6 Steady-State Analysis

To obtain the steady-state probability distribution of the underlying Markov chain of a PEPA model, click **PEPA > CTMC > Steady State Analysis**. The item is not be enabled if the state space is not derived.

$\Theta \cap \Theta$	Markovian Analysis	
Steady State	Analysis	
Select solver	for calculating the steady state probability distribution	
Choose solv	er	
Direct solv	ver 🛟	
Output optic	ons	
E	Save solution to disk	
Location: F	rst/paper.steadystate	Browse
?	< Back Next > Cancel	Finish

The wizard will guide you through the process of selecting a solver and setting its parameters. You can choose between a direct solver and a range of iterative solvers for which a number of preconditioners is available.

Once the model is solved, the State Space View will be updated with a column showing the steady-state probability of each column.

€ € €	PEPA – paper.pepa – Eclipse SDK	\bigcirc
] 🗂 • 🔛 🗁] 💁 -] 🖋] 🗠	·]∲ •§ • % • % • ◆ •◆•	FEPA 🐉 Java
Se Navigator 22 □ □	Paper.pepa X Performance r = 1.0; r1 = 1.0; r3 = (notused, 1).03; P1 p1 = (a, r).P2; P2 = (b, 1).P1; p1 = (a, infty).02; Q2 = (c, 1).01; p1[1] <a> Q1[1]	Cutline Cutli
	4 states	
	1 P1 Q1 0.4 2 P2 Q2 0.2 3 P2 Q1 0.2 4 P1 Q2 0.2	
[] □ ◆	1	

In addition, throughput and utilisation analysis will be carried out automatically and results will be available in the Performance Evaluation View.

00	PEPA – paper.pepa – Eclipse SDK	0
] 📬 • 🔛 👜] 🏊 •] 🖋] 🔤	3•]∯₩-Φ↔-	🖹 🔀 PEPA 🖏 Java
San Navigator S		Performance 2 Outline
] = +		

Performance analysis

5.1 Performance Evaluation View

The Performance Evaluation View shows information about throughput and utilisation analysis. The view is automatically updated when the model in the active editor is solved.

The **Throughput Analysis** tab lists the rate at which the actions of the PEPA model are performed at steady-state. A bar plot can be obtained by right-click the list and selecting **Show graph**. The corresponding graph will be shown in the Graph View.

00	PEPA – paper.pepa – Eclipse SDK	\bigcirc
] 📬 • 🔛 👜] 💁 -] 🖋] 📾	·]♀!•§!•\$\$•\$\$•\$•	PEPA 🐉 Java
🔁 Navigator 🛿 📃 🗆	🗖 paper.pepa 🛿 🦳 🎴 🗖 Performance	🛛 🛛 Outline 🗖 🗖
C Q E S ♥ S First Image: project M paper.pepa ♥ S UML	r = 1.0; r1 = 1.0; 03 = (notused, 1).03; 03 = (notused, 1).03; p1 = (a, r).P2; p2 = (b, 1).P1; Q2 = (c, 1).01; 0.4 p1[1] <a> Q1[1] Problems AST View State Space View 🖾 Graph View Properties 4 states 1 P1 0.4	© + 😵 Throughput s Error Log □ □ ⊽
	2 P2 Q2 0.2	
	3 P2 Q1 0.2	
	4 P1 Q2 0.2	
] =>]	

The **Utilisation Analysis** tab is a tree-based view showing the long-run utilisation of each top-level component of the model. For each component, it shows the percentage of time it is in a particular local state. A pie chart of a component can be obtained by right-clicking the desired component in the tree and selecting **Show graph**. The corresponding graph will be shown in the Graph View.

00	PEPA – paper.pepa – Eclipse S	DK 🗆
] 🗗 • 🔚 📥] 🏊 •] 🖋] 🗠	•]⊈+8]•♥+♥+	🖹 🔀 PEPA 🖏 Java
Ravigator 23 Control Control Control	M paper.pepa 33 r = 1.0; r1 = (a, r).P2; P2 = (b, 1).P1; Q1 = (a, infty).02; Q2 = (c, 1).01; P1[1] <a> Q1[1]	Performance 23 Outline □ ↓ Utilisation Throughput ↓ P1 P2 = 0.4 P1 = 0.60000000000000000000000000000000000
] 0*][]	

5.2 Experimentation

The plugin features support for experimentation, i.e. running a model with values for its parameters varying across desired ranges. The results of the experiment are plotted in the Graph View and a number of exporters is available for data serialisation.

To launch the experimentation wizard on the model of the active PEPA editor, select **PEPA > CTMC > Experimentation...**

The first page of the wizard lists the model parameters on which experimentation can be run. This will include the rate definitions as well as all the aggregation operators included in the model system equation.

000		
Experimentation		
Select the element	ts to include in experiments	
r r1		
□ P1		
🗌 Q1		
	Cancel Next > Cancel	Finish
O	Cance	rinish

After **Next** is pressed, a wizard page for each selected parameter will let the user set the desidered range to be used in the experiments. Range can be specified in two different ways. For a uniform sampling over a given interval you may want to click the **Specify an interval** button. Otherwise, a comma separated list of values can be directly specified if the **Specify comma-separated list** is pressed.

periment Settings onfigure parameters fo	or rate r			
Specify an interval Interval Settings				
From:				
To:				
Step:				
Specify comma-sep List Settings Insert list:	arated list	_		

The last page of the wizard lets you create the experiments with the specified parameters. The plugin currently supports two types of experiments. A basic 2D graph plotting a performance metric against a parameter. A parametrised

periments		
ll Experiments		
2	<back next=""> (</back>	Cancel Finish

2D graph plots a number of lines a performance metric against a parameter, a line for each value in the desired range of a second parameter.

The performance metrics available are: action throughput, sequential state utilisation, or steady-state probability of a subset of the state space as specified by a filter in the filter rule editor of the model.

To create a new experiment for a basic 2D graph, right click the list-box on the left of the wizard page and select **New > XYPlot**. The right hand side of the page will be dinamically updated with the widgetry for defining the experiment.

ll Experiments	Experiment Settings	
XY Plot	Name	
	XY Plot	
	X Axis	
		;
	Performance metric	

The name of the experiment will be shown as the title of the generated graph. The X Axis combo box lists the available model parameters and selection on the performance metric combo box will update the user interface with an additional section related to that performance metric. In the case of throughput, for example, a combo box lets the user select the action type.

l Experiments	Experiment Settings	
XY Plot	Name	
	XY Plot	
	X Axis	
	Rate r	÷
	Performance metric	
	Throughput	
	Action	
	a	;

To create a new experiment for a parametrised 2D graph, right click the listbox on the left of the wizard page and select **New > Parametrised XYPlot**. The user interface is very similar to that of the basic 2D graph, except for an additional combo box requiring the user to specify the second parameter.

ll Experiments	Experiment Settings	
XY Plot	Name	
Parametrised XY Plot	Parametrised XY Plot	
	X Axis	
	Rate r	\$
	Parameter	
	Aggregation Q1	÷
	Performance metric	
	Utilisation	÷
	Local state	
	P1	;

To delete an experiment, right-click the experiment on the list box and select ${\bf Delete}$

When you are done setting up the experiments, you can run them by pressing the **Finish** button of the wizard. The experiments will be run in the background and the results will be shown in the Graph View.

Discrete and continuous simulation

6.1 Time Series Analysis

Time series analysis can be carried out on a PEPA model via two different techniques: ODE and SSA. The former is the result of the mapping of a PEPA model to a set of Ordinary Differential Equation, the latter runs Stochastic Simulation Algorithms on a PEPA model. These functionalities are available under the **PEPA > Time Series Analysis** menu and they share parts of the user interface.

To perform ODE analysis, click **PEPA** > **Time Series Analysis** > **ODE**... . You will be provided with a dialog box for configuring the ODE solver. On the left hand side you are required to check the components you want to investigate. The time series of each component will be shown as a different line on the 2D graph for this analysis. On the right hand side you are asked to select the ODE solver and set its main parameters (start time, stop time, step and tolerance). Output options are available to export the data model to disk. The wizard also let you save the intermediate CMDL file, which is the format used by the underlying library featuring the ODE solvers.

To perform SSA analysis, click **PEPA > Time Series Analysis > SSA...** . You will be provided with a similar dialog box to the ODE one. An important SSA-related parameter to set here is the number of replications used during the simulation.

Importing UML models

7.1 Importing from UML

To import a PEPA model from a UML model, right click the .uml file in the Navigator view and click **Generate PEPA Model...** . You will be provided with a wizard dialog asking you which collaboration line stereotype application you want to import and the destination file.