



ePNK: A generic PNML tool - Users' and Developers' Guide

version 0.9.1

Kindler, Ekkart

Publication date:
2011

Document Version
Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

Citation (APA):

Kindler, E. (2011). *ePNK: A generic PNML tool - Users' and Developers' Guide: version 0.9.1*. Technical University of Denmark, DTU Informatics, Building 321. IMM-Technical Report-2011-03
<http://orbit.dtu.dk/en/publications/the-epnk-a-generic-pnml-tool-users-and-developers-guide-for-version-100%2860c6f267-41d1-4273-8873-bc4c64a73f9b%29.html>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

ePNK: A generic PNML tool Users' and Developers' Guide

Ekkart Kindler
Denmark's Technical University
DTU Informatics
DK-2800 Kgs. Lyngby
Denmark
`eki@imm.dtu.dk`

February 17, 2011
version 0.9.1

IMM-Technical Report-2011-03

DTU Informatics
Department of Informatics and Mathematical Modeling
Technical University of Denmark

Building 321, DK-2800 Lyngby, Denmark
Phone +45 45253351, Fax +45 45882673
reception@imm.dtu.dk
www.imm.dtu.dk

Technical Report: ISSN 1601-2321

Abstract

The *Petri Net Markup Language (PNML)* is an XML based interchange format for all kinds of Petri nets, which was published as International Standard ISO/IEC 15909-2:2011 in February 2011. Technically, ISO/IEC 15909-2 is defining an interchange format for three different kinds of high-level Petri nets and a simple version of Place/Transition systems only. But, one of the objectives of PNML was to provide a means for exchanging any kind of Petri net [8, 15, 1]. To this end, the concept of a *Petri Net Type Definition (PNTD)* was introduced, which is subject of a newly issued standardisation project ISO/IEC 15909-3.

There are many tools supporting one form of PNML or the other, and, in particular, there is the PNML Framework [6], which helps tool developers to ease the implementation of PNML by providing a framework and an API for loading and saving Petri net documents in PNML. This framework is based on the *Eclipse Modeling Framework (EMF)* [2] and has the focus on the underlying meta-models of Petri nets. The PNML Framework, however, is not generic in the following sense: Whenever a new Petri net type is created, the code for the complete tool needs to be regenerated. Moreover, the PNML Framework does not come with a graphical editor for Petri nets.

The ePNK overcomes these limitations: It provides an extension-point, so that new Petri net types can be plugged in to the existing tool without touching the code of the ePNK. For defining a new Petri net type, the developer, basically, needs to give a class diagram (actually an ecore-diagram) defining the concepts of the new Petri net type, along with a mapping of these concepts to XML syntax. This type can then be plugged into the ePNK, and the graphical editor of the ePNK will be able to edit nets of this new type with all its features.

Actually, this was the idea when we started the development of the *Petri Net Kernel (PNK)* about 15 years ago [12, 10, 13]. At that time, however, we had to implement all of the IDE functionality of such a tool ourselves. Today, we can make use of the eclipse platform [14], which helps us focusing on the Petri net specific parts; we get all the other functionality of a nice IDE, basically, for free. This is why we named the tool *ePNK*: it can

be considered to be an eclipse based Petri Net Kernel. But, it is just the spirit and idea that the PNK and the ePNK have in common; technically, there is not a single line of code from the PNK in the ePNK, and they are not compatible.

What is more, we use the nice features of EMF, GMF, and Xtext for developing the ePNK in a model-based way. In this way, the complete development process of the ePNK, is a case study in model-based software engineering using EMF and related technologies. This, actually, was the driving force behind this project. The evaluation and the lessons learned during this project, however, will be reported at an other occasion and to a different audience. This manual will focus on how to use the ePNK as an end user, and it will show how a developer can use the extension mechanisms of the ePNK for providing new Petri net types along with their XML syntax, and how to add new functionality to the ePNK.

Contents

1	Installation	1
1.1	Prerequisites	1
1.2	Installing the ePNK in Eclipse	2
2	Introduction	5
2.1	Motivation	5
2.2	The Petri Net Markup Language	6
2.2.1	PNML core model	6
2.2.2	Petri net type definitions	7
2.2.3	Mapping to XML	8
2.3	ePNK: Objective	9
2.4	How to read this manual	9
3	Users' guide	13
3.1	Eclipse as an IDE	13
3.2	Creating Petri net files	16
3.3	The tree editor	17
3.3.1	The tree editor: overview	18
3.3.2	Creating elements	19
3.3.3	Saving the document	20
3.3.4	Validating and correcting the document	20
3.3.5	Other Petri net information	22
3.4	The graphical editor	23
3.4.1	Overview of the graphical editor	23
3.4.2	Labels	23
3.4.3	Sub-pages	26
3.5	Petri net types	27
3.5.1	PTNet	27
3.5.2	HLPNG	28

3.6	Limitations and pitfalls	35
3.6.1	Dirty-flag and saving	35
3.6.2	Graphical features	36
3.6.3	Petri net types	37
3.6.4	Wrapping labels	37
3.6.5	Graceful PNML interpretation	37
3.6.6	Deviation from PNML	37
4	Developers' guide	39
4.1	Eclipse: a development platform for the ePNK	40
4.1.1	Installing the ePNK source projects	40
4.1.2	Installing Ecore Tools	41
4.2	Adding functions	41
4.2.1	Accessing a PNML file and its contents: A file overview	42
4.2.2	Writing PNML files: Generating multi-agent mutex .	47
4.2.3	Long-running functions: A model checker	51
4.3	Adding Petri net types	62
4.3.1	Simple Petri net type definitions: PTNet	63
4.3.2	Petri net type definitions in general: HLPNG	70
4.4	Adding tool specific information	87
4.5	Overview of the ePNK and its API	90
4.6	Deploying extensions	94
4.7	Future plans	95
5	PNML: Observations and suggestions	97
5.1	Toolspecific extensions: Type attribute	97
5.2	Relaxing requirement for ids everywhere	97
5.3	Naming in PNML standard	98
5.4	Mapping of models to XML	98
5.5	Explicit result type for NamedOperators	99
5.6	Opposite references for source and target	99
5.7	Name labels for arcs	99
5.8	Remove the possibility of net labels	99
5.9	Attributes with default values	99
5.10	Recursive sort and operator definitions	100
5.11	Names for ArbitraryDeclarations	100
5.12	Incorrect types in the HLPNG models	100
5.13	Reference to Sort in FiniteIntRangeConstants	101
5.14	Omitting redundant opposit references	101
5.15	Useless operator: Append	101

5.16	Incorrect constraints for List Append	101
5.17	Implicit and constant parameteres	102
5.18	Mismatch between tables and grammars	102

Chapter 1

Installation

This chapter discusses the installation of the ePNK (version 0.9.1). Readers who are interested in getting an idea of what the ePNK is and who do not want to work with it right away, can skip this chapter.

1.1 Prerequisites

In order to install the ePNK, you need to have Java 1.6 (or higher¹) and eclipse 3.5 (Galileo) or eclipse 3.6 (Helios) installed on your computer. In this version of the manual, we discuss the installation of the ePNK on eclipse Helios only. You will find information on how to install the ePNK on earlier versions of eclipse on the *ePNK installation page*².

For the installation of Java, please refer to <http://www.java.com/>.

If you are new to eclipse, it is recommended that you install the *Eclipse Classic* version of Helios. Download this eclipse version for your operating system from <http://www.eclipse.org/downloads/> and extract the downloaded file to some directory; after the extraction, you will find a folder named “eclipse” in this directory and in this folder, there will be an executable file also called “eclipse” (e.g. “eclipse.exe” on the Windows platform). Executing this file will start eclipse.

If you are new to eclipse, you can get a quick start to the Eclipse Integrated Development Environment (IDE) at <http://www.vogella.de/>

¹Actually, there seems to be a problem with the Oracle/Sun VM 1.6.0.21 on Windows and eclipse (see http://wiki.eclipse.org/FAQ_How_do_I_run_Eclipse%3F). You should use an older or a newer version. The notes on the bug-report on this problem suggest also that the eclipse versions 3.3-3.6 might not work together with Java 1.7 anymore.

²<http://www2.imm.dtu.dk/~eki/projects/ePNK/install-details.html>

`articles/Eclipse/article.html`. Once you have installed and started eclipse, you can also find much more information on eclipse in the “Workbench User Guide” in the eclipse help: You will get access to it in the “Help” menu in the eclipse toolbar under “Help Contents”.

1.2 Installing the ePNK in Eclipse

Once you have installed eclipse Helios, you can install the ePNK from the eclipse workbench. To this end, the ePNK is made available via an eclipse update site at <http://www2.imm.dtu.dk/~eki/projects/ePNK/update/>.

To install the ePNK from there in your eclipse version, you should proceed as follows (after you have started it and selected a workspace):

1. In the eclipse tool bar, select “Help” → “Install New Software...”, which will open an install dialog.
2. In the install dialog, press the “Add...” button to add a new update site. In the “Add Site” dialog, enter some name (e. g. “ePNK Update Site”), enter the URL

`http://www2.imm.dtu.dk/~eki/projects/ePNK/update/`

as location, and then press okay.

3. Now, select the newly created ePNK update site in the still open install dialog. After some time, some ePNK items should pop up in the dialog. Select all of them and press okay. For eclipse Helios, select the following plugins:

- ePNK Basic Extensions 0.9.1
- ePNK Core 0.9.1
- ePNK HLPNGs (Xtext 1.0.1) 0.9.1
- ePNK Tutorial 0.9.0

Note: Make sure that the box “Contact all update sites during install to find required software” is checked; this will guarantee that all additional plugins from eclipse that the ePNK requires, will be installed automatically (EMF, GMF, Xtext, etc.)

4. Follow through the installation process (don’t forget to accept the license agreement).

Note: If you get an error of the kind

```
Cannot complete the install because one
or more required items could not be found.
...
```

you have probably selected the wrong combination of features of the ePNK – probably “ePNK HLPNGs 0.9.1” (which works for the eclipse Galiloe version) instead of “ePNK HLPNGs (Xtext 1.0.1) 0.9.1”. In that case, go back and select the right combination as explained above.

An other possibility for such an error is that you have forgotten to tick the checkbox “Contact all update sites during install to find required software”. In that case, go back in the installation process and make sure to check this box.

5. Then, the ePNK and all other required plugins will be installed; it is a good idea to restart eclipse after that (eclipse will ask you to do that anyway).

Chapter 2

Introduction

This chapter gives a brief overview of the *Petri Net Markup Language* (PNML) as well as on the concepts and ideas of the ePNK and its main features. In the end, of this chapter, there is some information for different kinds of readers on what to read and on how to read this manual.

2.1 Motivation

The PNML is an XML-based interchange format for all kinds of Petri nets, that allows different tools to exchange Petri net models among each other. One of its main features is that it is generic, which means that it provides a mechanism to define own types of Petri nets, which are called *Petri net type definitions* (PNTD). These Petri net type definitions define the additional concepts of the new Petri net type, as well as the representation of these new concepts in XML syntax. It is also possible that different tools include their *tool specific information* to PNML documents, which is information that can be ignored by other tools.

Up to now, there is no tool that fully supports these ideas, so that Petri net types, and tool specific extensions can be easily defined and plugged in to the tool. And there is no generic editor supporting all Petri net types, once they are plugged in.

The lack of such a generic tool support was the starting point for developing the *ePNK*.

2.2 The Petri Net Markup Language

In order to better understand the ideas behind the ePNK, we briefly discuss the main concepts and ideas of the PNML here. For more information on PNML and ISO/IEC 15909-2, we refer to [11, 5] or the recently published ISO/IEC 15909-2 itself [7].

2.2.1 PNML core model

As stated above, the extensibility and genericity were two of the main objectives behind the PNML [9]. This is achieved by identifying the concepts that are common to all Petri nets in the so-called *PNML core model*. These common concepts are mainly the *places*, *transitions* and *arcs*, and that these *objects* can have some kind of *label*. The PNML core model also takes into account that larger Petri nets can be split up into *pages* and that connections between the nodes on different pages can be established by *reference places* or *reference transitions*. And there are all kinds of graphical information that can be attached to the different elements, such as position, size, font-type, and font-size.

In addition, the PNML core model defines the possible relation between these elements. In particular, it defines that places and transitions, which are generalized as *nodes*, are contained in pages and that arcs may connect these nodes. Figure 2.1 shows the PNML core model as a UML diagram.

Note that there is only one concrete type of label in the PNML core model, which is *name*. All the other possible labels need to be defined by a Petri net type definition, which will be discussed in Sec. 2.2.2.

In addition to the concepts and relations between them, the PNML core model also states some restrictions. For example, there is an OCL constraint stating that arcs can only be between nodes that are on the same page. But, there is no constraint yet that arcs can only run between a place and a transition or the other way round. The reason for that is that there are some kinds of Petri nets that would allow arcs between places or between transitions. This is why these restrictions would be part of a Petri net type definition.

Note also that the PNML core model does not specify concrete tool specific extensions. It is up to a tool to define what it needs. But, any tool must be able to read – and later write – any tool specific extension; their contents however, can be ignored.




Figure 2.1: The PNML core model

2.2.2 Petri net type definitions

As stated above, it is the purpose of the Petri net type definitions to define which labels are possible in a specific kind of Petri net, and also to define some additional restrictions on the legal connections. Here we explain this idea by the help of a simple example: Place/Transition- Systems (P/T-Systems).

The two additional kinds of labels for Place/Transition-Systems are the initial marking for places, and the inscription for arcs. The initial marking can be any natural number (including 0) and the inscription for arcs can be any positive number. Figure 2.2 shows the UML model for these new concepts and how they are related to the concepts of the PNML core model.

In Fig. 2.2, there is also one additional OCL constraint. Without going into the details of OCL, this constraint states that an arc must run from a place to a transition or from a transition to a place. So, for P/T-Systems, it is no longer possible to connect places with places or transitions with




Figure 2.2: The PNTD for PT-Nets

transitions.

A Petri net type definition, would typically also define how the new concepts from Fig. 2.2 would be mapped to XML. But, for this simple kind of nets, this is actually done in a standard way.

2.2.3 Mapping to XML

As mentioned above, the PNML core model together with the model for a Petri net type definition, define the concepts of a specific kind of Petri net and how they can be connected. Therefore, these models are the centerpiece of PNML. Still, PNML is a XML transfer format for Petri nets. So, PNML also needs to define how these concepts will be saved or represented in XML. This is achieved by mapping every concept or feature of the models to some XML construct.

Here, we do not give these mappings, but rather show an example (for a




Figure 2.3: A simple P/T-System

detailed discussion of the mappings, see [5]). Figure 2.3 shows a simple example of a P/T-System and Listing 2.1 shows its representation in PNML’s XML-syntax¹.

Note that this also shows an example of a tool specific extension: the positions of the individual tokens in the place.

2.3 ePNK: Objective

The main objective of the ePNK is to fully support the concepts of PNML, so that new Petri net types along with the mapping to XML syntax can be easily plugged into this tool – and to provide all the Petri net type definitions for supporting ISO/IEC 15909-2.

As soon as such a new Petri net type definition is plugged in, it should be possible to load and save nets of this type (and also documents containing different net of different net types that are plugged in). And there should be a graphical editor that allows us to edit Petri nets of any plugged in net type and that is fully aware of all the features and the additional restrictions of the plugged-in net types.

For tool developers, the ePNK should provide an API to easily load and access Petri nets from PNML files, to manipulate them, and save them. Moreover, it should be easy to plug in new functionality for analysing and manipulating Petri nets.

2.4 How to read this manual

In this manual, we will explain the features of the ePNK in more detail.

This will cover the parts relevant for the “end user” who just wants to load, save and edit Petri nets of existing types and use some existing or plugged in functionality of the ePNK. In the rest of this manual, we will call

¹We deleted some line-breaks to make this listing fit to single page

Listing 2.1: PNML code of the example net in Fig. 2.3

```

1 <pnml xmlns="http://www.pnml.org/version-2009/grammar/pnml">
  <net id="n1" type="http://www.pnml.org/version-2009/grammar/ptnet">
    <page id="top-level">
      <name><text>An example P/T-net</text></name>
      <place id="p1">
6      <graphics><position x="20" y="20"/></graphics>
      <name>
        <text>ready</text>
        <graphics>
          <offset x="0" y="-10"/>
11      </graphics>
      </name>
      <initialMarking>
        <text>3</text>
        <toolspecific tool="org.pnml.tool" version="1.0">
16      <tokengraphics>
          <tokenposition x="-2" y="-2" />
          <tokenposition x="2" y="0" />
          <tokenposition x="-2" y="2" />
        </tokengraphics>
21      </toolspecific>
      </initialMarking>
    </place>
    <transition id="t1">
      <graphics><position x="60" y="20"/></graphics>
26 </transition>
    <arc id="a1" source="p1" target="t1">
      <graphics>
        <position x="30" y="5"/>
        <position x="60" y="5"/>
31 </graphics>
      <inscription>
        <text>2</text>
        <graphics>
          <offset x="0" y="5"/>
36 </graphics>
      </inscription>
    </arc>
  </page>
</net>
41 </pnml>

```

“end users” just *users*. All the information for these users can be found in Chapter 3.

This manual will also cover the parts relevant for *developers*, who are interested in using the ePNK for their purposes and extending it by defining new Petri net types, by defining new tool specific extensions, or by new functionality². Chapter 4 will provide this information for developers.

In some future version of this manual, there will (eventually) also be a part that discusses the architecture, the design, and some of the used technologies (and its problems). These parts might be interesting for developers, but actually addresses people interested in model-based software development technologies that are used (and extended) in this project: EMF, GMF, Xtext, ExtendedMetaData, EMF Validation, and OCL.

During the development of the ePNK, we came across some problems and issues of PNML or ISO/IEC 15909-2. Chapter 5 will list and discuss these issues and make some suggestions for improving future versions of PNML or the definition of some specific Petri net types. Chapter 5 is mostly addressed to people interested in the standardisation process of ISO/IEC 15909 – mostly concerning Parts 2 and 3.

²Note that up to now, this needs to be done by the standard eclipse mechanisms. But, it is planned to eventually provide some functionality extension mechanisms for the ePNK that is tuned to Petri nets and allows plugging in ePNK functionality in a uniform way.

Chapter 3

Users' guide

This chapter explains how to use the ePNK for creating, loading, saving, and editing Petri nets, and also how to use some of its functions. Since new Petri net types can be plugged in, we try to point out the general principles of these editors and how to use them. But for the particular syntax of some labels, it would be necessary to refer to the documentation of the specific extensions of that type. We will discuss these principles by some of the Petri net types that come with the basic version of the ePNK; and we use high-level nets (in terms of the ISO/IEC 15909-2 High-level Petri Net Graphs, or HLPNGs for short) to point out for which parts you would need to refer to the specific documentation of the specific Petri net type.

3.1 Eclipse as an IDE

For users who are new to eclipse and its IDE (Integrated Development Environment), we start with a brief overview of eclipse's workbench. Users who are familiar to eclipse already can directly read on in Sect. 3.2.

Once you installed and started eclipse (see Chapter 1), you will see the eclipse *workbench*. Depending on the chosen *perspective*, the different parts can be arranged in different ways. But, the principle behind is always the same. Figure 3.1 shows an example of the eclipse workbench, with some numbers marking some parts, which we will discuss next.

At the top (1), you can see the *menu bar* and the *tool bar*. Here, you will find the menus and tools for all the standard functionality, such as loading and saving files, and for standard editing operations. The menus that are shown in the menu bar depend on you installation and also on the editor that is currently active. For many operations, there are also the standard

shortcuts, like CNTRL-S (on the windows platform) for saving the contents of an editor to a file. For getting more information, on that you could chose the “Help Contents” in the menu “Help” in the menu bar, and read the “Workbench User Guide”.

Note: In automatically generated editors, such as the graphical editor of the ePNK, the copy/paste functionality with CNTRL-C, CNTRL-V, and CNTRL-X does not work properly without taking some extra measures, which we did not take yet. Therefore, you should not use the copy/paste shortcuts in these editors for now.




Figure 3.1: The eclipse workbench

On the left-hand side (2), you can see the *package explorer*, which gives you access to all the files in your workbench, and you can use it for browsing through the existing files and also to manipulate them, rename, copy, move, and delete them. This is very much like the file explorer of your operating system. Eclipse actually has different kinds of explorers, depending on the perspective and the user’s preferences. The package explorer is made for java development projects. For our purposes, any of these would do, like for example the “navigator” or the simple “project explorer”. To find and open

one of these other explorers, you can use the menu “Show View” in the menu bar menu “Window”. All these explorers have some important concept in common, which concerns the organisation of files in the workbench: the top-level “folders” are actually not folders, but they are *projects*. This is only relevant when creating these projects. You can create a folder or file in the workbench only after you have created some project; this can be done via the “File” menu or by a right-click in a resource browser and then selecting “New” → “Project”. Note that in the dialog, you can create many different kinds of projects; for us the project “Project” in category “General” will do. Then files can be created within this project.

In the center (3), you can see the *editor area* of eclipse. This is where all the editors that are started in eclipse will be opened. Note that there can be many editors open at the same time (in our example, there are four editors open). Typically, you can only see one at a time and the others are hidden below it. But, you can move the editor tab to some border of the editor area, so that you can see the contents of two or more editors at the same time. In our example, there is a tree editor of a complete Petri net document open on the left-hand side, and on the right-hand side you can see one specific page open in a graphical editor (the graphical editors for some other pages are hidden beneath). Note that, even though there can be many editors open and even visible at the same time, there will always be only one editor that is active. This editor and what is selected in it determines what you see in some other views. For example you can see the *overview* (4) of the page or you can see the property of the currently selected element in the *properties view* (5) at the bottom. In order to open an editor on some resource in the explorer, you would normally double click on the resource you want to open. This will open the *default editor* on the selected resource. You can use the right mouse button on a resource to open a pop-up menu and then select “Open with” to select a specific editor for this purpose. The way of how the content of a resource is edited very much depends on the kind of editor (mostly, it is straight forward). Saving the file can typically done by a shortcut (like CNTRL-S) in all editors (or via the “File” menu in the menu bar). An editor can be terminated (closed) either by clicking the close symbol on the tab of the editor or via the “File” menu in the menu bar.

Most editors support undo and redo of the latest changes, which you can access via the “Edit” menu or via the CNTRL-Z and CNTRL-Y shortcuts.

Note that the graphical editor for pages of a Petri net cannot be initiated directly from the explorer, since the resource could have many pages. The graphical editors for pages of a Petri net can be opened with a pop-up menu

(right mouse button) on pages in the tree editor for Petri nets (see Sect. 3.4 for details).

All the other areas of the workbench are *views*¹. In eclipse, views are used for many different purposes. The views that are most relevant for us, are the *outline* (4), the *properties* (5), and the *problems view* (not visible in Fig. 3.1). The outline gives an overview over the content of the currently active editor and in case of a graphical editor allows us to quickly move around in the visible area of this editor. The properties view shows some details of the currently selected element in the editor; in many cases, the properties view also allows us to edit some properties. Note that, initially, the properties view might not be open. You can, typically, open it from the active editor via a context menu on the right mouse button: In the pop-up menu that opens, there will be a menu “Show Properties View”, which opens the properties view. You can also open the properties view via “Window” → “Show View” → “Others ...” and then selecting “Properties” in the category “General”.

We mentioned already that the eclipse workbench can appear in different ways, which is defined by the chosen *perspective* (and some user-specific settings). The perspective can be changed via the tools at the top-right of the workbench, which is marked with (6) in our example. We do not need to change it; but if, for whatever reason, you end up in a wrong perspective, by clicking on the left symbol, you can open the “Open Perspective” dialog. There, you can select the perspective “Resource” or, if you like, “Java” (which is the default perspective).

If you are interested in more details in the eclipse workbench, you can have a look into the eclipse help (“Help” → “Help Contents”) or at one of the many books or online articles; <http://www.vogella.de/articles/Eclipse/article.html> could be a start.

3.2 Creating Petri net files

This section explains how to create new ePNK files. Note that there are two formats in which the ePNK can save a Petri net. The first and recommended format is PNML. The second format is the XMI-serialisation of the PNML models, which we call PNX. Note that PNX, is part of the ePNK since XMI is the standard serialisation mechanism of the technology (EMF and ecore), and came for free. Whether PNX really should be a part of the ePNK

¹Actually, also the resource browsers are views.

distribution is yet to be seen. Therefore, the focus of this users' guide is on PNML.

The easiest way of getting started with the ePNK is obtaining existing PNML files from somewhere else and just copy them to the workbench. For example, you could get some examples from the ePNK home page: <http://www2.imm.dtu.dk/~eki/projects/ePNK/>. You can also use a text editor and create a simple text file with file extensions “.pnml” and insert the single line

```
<pnml xmlns="http://www.pnml.org/version-2009/grammar/pnml"/>
```

to this file, which is an empty Petri net document without any nets in it.

The ePNK also provides you with a wizard for creating a PNML document. Like all eclipse creation wizards, this wizard will be started via the “New” menu, which can be either accessed by the “File” menu from the menu bar or via the popup menu that opens on a right-mouse click in the explorer. Then, select “Other...” (the short-cut to that would be pressing CNTRL-N in the explorer) and in the newly opened “Select a wizard” dialog chose “ePNK (PNML)” from the “ePNK” category and press “Next”. In the next dialog, you must chose a name and, if you want, you can chose a different folder to which this file should be added. Pressing “Finish” will create the file and also open the tree editor on it (see Sect. 3.3); note that you also can continue the creation process by pressing “Next”, which will allow you to chose an XML-Encoding. Note that in the dialog with the encoding, there is also a field asking for the “Model Object”; but you cannot chose anything here since PNML, in contrast to other formats, has a fixed root object that cannot be changed: “PetriNetDoc”.

Note that in the same wizard category “ePNK” there is another wizard called “Pnmlcoremodel Model”. When you use this wizard, a PNX file will be created. And in this wizard, you can select a root element different from the PetriNetDoc – but this would be reasonable only in very special cases (and when you know what you are doing).

3.3 The tree editor

As mentioned earlier already, the ePNK provides two kinds of editors for Petri nets. The *tree editor*, which allows us to create, modify, and delete all parts of the Petri net in a tree like structure. And there is a *graphical editor* in which a page of a Petri net with its places, transitions, and arcs can be edited in a graphical way. Clearly, the graphical editor is more convenient

for editing pages than the tree editor. But, other parts like for example the page structure and the complete Petri net document are more convenient to edit in the tree editor. This is why there are two different editors in the ePNK. The graphical editor for pages will always be started from a selected page – either in the tree editor or in the graphical editor. When opening a PNML document from the resource explorer, it will always be the tree editor that opens. The graphical editor can be accessed from pages in that tree editor.

3.3.1 The tree editor: overview

We have a closer look at the tree editors first. Figure 3.2 shows the Eclipse workbench with two PNML documents open in tree editors. The right one shows the tree editor opened with the PNML file (“test.pnml”) with the single line as discussed in Sect. 3.2. Therefore, it contains only the Petri net document element without any contents. The other PNML document, which is open on the left-hand side (“hlpng-gmf.pnml”), shows a Petri net document with three nets that have different types.




Figure 3.2: Two Petri net documents in tree editors

These documents were opened from the explorer by a double click² on the respective file in the workbench explorer. Let us briefly go through what

²Remember, that you can use the pop-up menu to make an explicit choice by which editor you want to open the file. This way, you can open the file with a text editor, so

you see in the Petri net document “hlpng-gmf.pnml”. The top-line shows the actual *resource* or file in which this document is stored; the second line is the symbol for the Petri net document itself – all documents will follow this structure in the tree editor. Then you can see that there are three Petri nets contained in this document (with *ids* n1, n2, and n3); the last one is actually not folded out because its was not fitting to the screen. The first line below the Petri net is the type of the Petri net. The first net is a high-level net, which is named HLPNG according to ISO/IEC 15909-2; the second net is a Place/Transition-System, called PTNet according to the standard. You can also see that the nets contain places, transitions, and arcs, which are indicated by corresponding icons. You can also see some pages and sub-pages. Note that the icons for the places, transitions, and arcs are different for the two different Petri net types, so that it easier to distinguish them on first glance. In the properties view at the bottom, you can see the properties of the currently selected element, which is Petri net n1, and the only property is its id. Note that this net also has a name; this, however, is not shown as a property, but as a *child element* of the net, which is true for all labels of Petri nets. In this case, the name is “A high-level next example”.

3.3.2 Creating elements

You can unfold all the sub nodes (children) of the net and this way inspect the complete document in all details. More importantly, however, you can create the basic elements of the Petri net document. You can create new nets, their type, and their pages. And from there, you would use the graphical editor to draw the rest. This basically works by inserting *child elements*. Inserting a child element is done by right-clicking on the element to which we want to add a child, then selecting “New Child” in the dialog that pops up, and then selecting the appropriate element. Figure 3.3 shows the pop-up dialog when inserting a new Petri net to the Petri net document.

The type of a Petri net, its name, and its pages can be inserted in the same way. Note that it is important that every Petri net is assigned a type right after it was created; and after its creation, the type should never be changed again. Otherwise, it could happen that a Petri net contains elements that do not make sense in a specific type. The tree editor does not enforce this yet³. There will be some more information on the Petri net

that you can see the PNML it produces. On a double click, the file will always be opened with the editor you had selected last time.

³Note that it would be easy to enforce this, and we might eventually implement this.




Figure 3.3: Pop-up menu when inserting a new Petri net

types that are deployed as part of the ePNK in Sect. 3.5.

3.3.3 Saving the document

As discussed in Sect. 3.1, you can save the net via the “File” menu or with the CNTRL-S shortcut. Note that saving the net should always be done in the tree-editor, for now; saving from the graphical editors works, but the dirty bit is not reliable. So, after finishing the editing of a net, close all the graphical editors and explicitly save the net from the tree editor again.

3.3.4 Validating and correcting the document

Before saving the document, it would be a good idea to validate it, which checks whether all the constraints that PNML imposes on Petri nets are met. It is possible to save a document that does not properly validate, and you would be able to load the file again. But, if you save a file that does not properly validate, you cannot be sure that the saved document is ISO/IEC 15909-2 conformant PNML.

In the current experimentation phase, we leave that in as a feature to play around and to experiment with this. But, you should really know what you are doing, when changing the type later.

There are many things that could be wrong and need validation on a Petri net document. Most of them are type specific (such as requiring that an arc must run from places to transitions or the other way round only); these Petri net type specific constraints will be discussed in Sect. 3.5. But, there are also some general constraints:

- Every Petri net object must have an id and this id must be unique in the scope of this document.
- Arcs may only connect nodes which are on the same page (as long as you are using the graphical editor, this constraint cannot be violated; but if you do changes in the tree editor, this could be violated).
- There must not be cycles on the references between reference nodes, and all reference nodes must refer to a node.
- A reference node must refer to a node within the same net.

In order to identify which constraints are violated, you can use the validation feature. Click the right mouse button on the Petri net document and, in the pop-up menu, select “Validate”. The result will be shown in a dialog; and the results will be also visible in the *problem view* later as shown in Fig. 3.4. Most of these errors are actually coming from high-level nets. But, there is also a general constraint violated in this example: some IDs collide (line 5 and 6 in the problems view). Remember that, if for some reason, the problems view is not open in your workbench, you can open it with “Windows” → “Show View” → “Problems”.

In order to reduce the number of errors, you can also do a validation of sub-elements of the Petri net document, which could be a net or even a single page. Ultimately, however, you must validate the complete Petri net document.

In our example, there are some problems that could be fixed automatically. For example, unique ids could be set automatically. The ePNK provides an action for this. To this end, select the Petri net document, click the right mouse-button, and then select “ePNK” → “Add missing IDs” in the pop-up menu. This will fix all problems with the ids within a Petri net document.

There is another tool, which will fix some of the type specific problems. In high-level nets, many of the labels of the net are related to each other. For example, there are variable, function, and sort declarations on pages. And these variables, functions, and sorts could be used at other places, i. e. in other labels. To establish the connection between the use and the




Figure 3.4: The problems view with many constraint violations

declaration of these symbols, the labels need to be linked with each other. Since this is part of the Petri net type definition interface, this can be done in a generic way. To properly link the *symbols* of any such type, you can select the Petri net in the tree editor, click the right mouse button, and select “ePNK” → “Link Labels (globally)”. We will discuss more details in Sect. 3.5.

3.3.5 Other Petri net information

In principle, you can inspect and edit all the information of a Petri net document in the tree editor. In our example from Fig. 3.3, you can also see some labels (declarations of high-level nets in this case, or a marking of a P/T-System) or graphical information. If you have a closer look at these examples, you will also find some other types of elements such as tool specific information – and once graphical editors are started some auxiliary data. But, it is strongly recommended not to change any of this information in the tree editor⁴.

⁴Once the features of the ePNK that are really needed are fixed, the parts that should not be edited in the tree editor will probably removed or at least made read only.

3.4 The graphical editor

For editing the contents of pages, the graphical editor should be used. This graphical editor can be opened by right-clicking on the respective page in the tree-editor and, in the pop-up menu, selecting “ePNK” → “Start GMF Editor on Page”.

Figure 3.4 shows the graphical editor with an opened page from a high-level Petri net (in this case one with several errors in it). Normally, this new editor shows on top of the tree editor, but it can be moved to the right side (click in the tab at the top of the editor window and move it while keeping the mouse pressed), so that the tree editor as well as the graphical editor are visible at the same time.

3.4.1 Overview of the graphical editor

On the left-hand side of the graphical editor for the page, you see the *canvas* with all the Petri net objects on that page represented in a graphical way. This includes also the *labels*, which are either attached to an object by a dashed line or attached to the page itself, in which case it is called a *page label*.

At the top, you see the *tab* of this page, which also shows the page name (if the page has a name label assigned) or its id, or the path to this page (if the page has neither a name nor an id).

On the right-hand side, you see the *palette* or tool bar of the graphical editor. These tools allow you to create all the Petri net objects. Note that you can also create sub pages.

There are two different tools for labels. The tool “Label” is for creating labels that are attached to objects, the tool “Page label” is for creating labels that are directly attached to the page that is shown in this editor.

For creating objects and labels, you first select the tool by clicking on it, and then clicking somewhere into the canvas. For creating an arc, you select the arc tool and then click on the source object, and keeping the mouse pressed and move the mouse to the target object. Note that the arc is not added between two objects, if the Petri net type you are editing does not allow this.

3.4.2 Labels

When creating a new page label on a page, the graphical editor will show you all the possible options of legal labels for that net via a pop-up menu.

Figure 3.5 shows the pop-up menu during the creation of a page label for high-level nets, where the only option “Declaration” is shown here. You can select an option, after which a label of that kind will be created. You can also abort by either pressing the “ESC” button or clicking somewhere outside the menu.




Figure 3.5: Pop-up menu during page label creation

The process for creating a label is a bit different. First you can create a label, which however will not be attached to any object yet. This will be indicated by the text “<not connected label>”. Such a label can be connected to some Petri net object (also to a sub-page) by choosing the tool “Link Label”, clicking on an object, and without releasing the mouse button moving it over the not yet connected label. Then, a pop-up menu will be opened showing you the possible options of labels for the chosen object. This is shown in Figure 3.6 for a label that is attached to a place. The possible options are “Name” (which is a legal option for any object, but only if there is no name attached yet) or “PTMarking 0”, which is the initial marking for P/T-Systems (where “0” is the default value). After the selection, the label of the chosen type will be attached to the object. Again, attaching the label can be aborted by pressing “ESC” or by clicking somewhere outside the pop-up menu.

After a label was attached to an object it can be edited “in place”, by clicking into it and pressing the ENTER key in the end. The legal syntax




Figure 3.6: Pop-up menu during attaching a unconnected label to an object

of the label depends on the Petri net type and which kind of label it is. In general, editing of labels as well as of page labels can be distinguished into two cases. The first case are *simple labels*, which typically are numbers or simple values like “true” or “false”. If such a label is typed in syntactically incorrectly, the new value will be rejected, and the value of that label will be reverted to the value it had before editing. The other case, are *structural labels*. These are typically, labels with a complex syntax, as for example the declarations of a high-level net (actually all labels of high-level nets except the names are structural). These labels will also be parsed and checked for syntactical correctness; but the entered text will be stored in all cases. If the text is incorrect, however, the structure is not set and this will very likely result in some validation error later (see Sect. 3.3.4). So this error needs to be fixed, by editing the label again. In case of such an error, the text of the label will be shown between “<! ” and “ !>”. If, for example, we delete the comma that separates the two sort declarations in the label “sorts B = (A*INT), C = (B*B);” this will be shown as “<! sorts B = (A*INT) C = (B*B); !>”. And upon validation, a validation error message will be given and shown in the problems view.

The documentation of the legal syntax of all type specific labels, in particular the one of the structural labels, is part of the documentation of the Petri net type definition. For the types deployed together with the

ePNK, this information can be found in Sect. 3.5.

Note that labels, in principle⁵, can have line-breaks. Since pressing the ENTER button will finish the editing of a label, however, a line-break is inserted to a label by pressing CNTRL-ENTER while editing the label.

3.4.3 Sub-pages

The graphical editor allows you also to create other pages on that page, which we call *sub-pages* in order to avoid confusion. This can be done with the “Page” tool, in the very same way places or transitions are created. In the graphical editor, sub-pages are graphically represented as rounded rectangles (or squares).

It is also possible to open a graphical editor on a sub-page from the graphical editor via a pop-up menu on the right mouse button: “ePNK” → “Start GMF Editor on Page” (as we have seen it for the tree-editor). Therefore, the tree-editor could be used for creating the top-level pages of the net only; all the sub-pages could be created by the graphical editor. But navigation is much easier in the tree-editor; this is why you would probably want to use the tree-editor for navigating and opening sub-pages further down in the tree-hierarchy. It is recommended not to create sub-pages in the tree editor, since they would not have a position in the graphical editor. Still, it is possible and the graphical editor would show these pages (as well as other objects created in the tree-editor) in the top-left corner, when it is opened with the graphical editor for the first time. Then, you could move it to a better position.

What is more important about pages is how to deal with their labels. All the type-specific labels will be shown as page labels within the sub-page. The name, however, will be shown as a label attached to that page on the super-page.

Note: There might be some Petri net types that have some labels for pages that, like the name, should rather be shown as a label attached to the page on the super-page than as a page label within that page. The mechanism for defining Petri net types could provide a way of identifying which is which. But, right now, we do not have any example for such labels. Therefore, all label of a page except for the name are considered to be page labels in any Petri net type. Once we learn about a reasonable example, it would be easy to extend the type definition

⁵That is, if the legal syntax of a specific Petri net type does allow it.

mechanism in such a way. Please, contact us, if you have a good example.

3.5 Petri net types

In this section, we give an overview of the Petri net types that are deployed together with the ePNK. Currently, these are P/T-Systems (PTNet) and high-level Petri nets (HLPNG). And there is the empty type (Empty), which, however, does not contain any concepts in addition to the PNML core model; therefore, we do not need to discuss this here. The empty type was introduced to explicitly indicate, that there are no Petri net type specific extensions.

Actually, HLPNGs come in different levels or kinds: “dot nets”, which is a representation of P/T-Nets in high-level nets; basically, “dot nets” are restricted to the sort “DOT” and a minimal version of operators on them; “symmetric nets”, which is a restriction to some special finite sorts and a limited set of operations; and the full version of high-level nets. The kind of a HLPNG can be changed by selecting the HLPNG type in the tree editor and selecting the kind in the properties view (identified by their ISO/IEC 15909-2 URI). For a detailed discussion of the legal constructs of the different kinds of HLPNG, we refer to [5]. Note that, in contrast to the Petri net type, the kind of a HLPNG can be changed anytime, since the kind of HLPNG concerns the validation only. The PNML syntax is the same.

3.5.1 PTNet

We start with explaining PTNets. In Sect. 2.2.2 we have already seen the additional features of PTNets, which are the initial marking for places and the inscription for arcs. Both labels are simple labels, which means that it will be checked right after editing whether the label is syntactically correct (see Sect. 3.4.2); if it is not correct, the value will be reverted to the value it had before.

The marking of a place must be a non-negative integer in any reasonable representation⁶. The arc-inscription is similar, just that it must represent a positive integer (i. e. must be greater than 0).

Moreover, PTNets have the restriction that arcs may only run from a place to a transition or from a transition to a place, which will be enforced in

⁶For those who want to bother with the technical details, it is any String that would be accepted by the Java `Integer.parseInt()` method as a number and evaluates to a number greater or equal than 0.

the graphical editor⁷. Actually, the constraint is slightly more complicated due to reference nodes: We can connect place-like nodes (PlaceNodes) with transition-like nodes (TransitionNodes) and vice versa; but semantically, i. e. when flattening reference nodes, this amounts to the above condition.

3.5.2 HLPNG

HLPNGs are much more involved than P/T-Nets and we cannot explain them in all details here. For a detailed motivation and full account on what HLPNGs are, we refer to ISO/IEC 15909-2 [7] or [5].

For HLPNGs, there are the following labels (in addition to names):

Declaration A *declaration* is a page label, which is used to define *variables*, *sorts*, and *operators*, which can then be used in the other labels. Every page can have any number of declarations and, within a single declaration, different kinds of declarations can be mixed. Note that all declarations are global (known in the complete Petri net), even though it is attached to a specific page.

Type A *type* is a label that is associated with a place. Every place must have exactly one type label which denotes the sort (which can be built from built-in sorts and sort constructs or from user defined sorts) of the tokens on that place.

HLMarking A *marking* is a label that is attached to a place and indicates the places initial marking. The marking is represented by a ground-term⁸, which must be a multiset over the place's type. Note that this label may be omitted, in which case the initial marking is considered to be empty. There can be at most one label of this kind.

Condition A *condition* is a label that can be attached to a transition. The condition is a term of type boolean and can contain variables. There can be at most one condition, and if it is missing, it is assumed to be true.

HLAnnotation An *arc annotation* is also a term that may contain variables. The term must be a multiset term over the type of the place

⁷In the tree editor, illegal arcs could be created, but this would not pass validation.

⁸A ground-term is a term that does not contain variables.

to which the arc is attached. Every arc should have exactly one arc annotation⁹.

Note: ISO/IEC 15909-2 also allows declarations directly attached to the net. This is not supported by the ePNK, because this breaks one of PNML’s own principles and could not be shown graphically. If there was strong need for this feature, it could be implemented though. Please, let us know.

All labels of HLPNGs are structural labels (see Sect. 3.4.2), which means that the user can edit it and leave it syntactically incorrect. Of course, this will not pass validation; but, it is possible to save nets with incorrect labels and load them again, so that the labels can be corrected another time. Since labels can refer to symbols that are defined other places, it is also important to link these labels before validation (see Sect. 3.3.4), which can be done by the pop-up menu on the net in the tree-editor: “ePNK” → “Link Labels (globally)”.

PNML does actually not define a concrete syntax for declarations and terms. This is up to the tool. Therefore, the ePNK comes with its own concrete syntax, which resembles the one of CPNTools, but is not identical! Before going into the details of the syntax, we briefly discuss some examples.

The following shows several declarations of variables, sorts and operators. Each of them could be in a separate declaration label, but they could also be contained in a single declaration:

```
vars
  x:NAT;

sorts
  A = MS(BOOL);

ops
  f(x:INT, y:INT) = x * y,
  g() = 1;

sorts B = (A*INT), C = (B*B);
```

⁹Actually, ISO/IEC 15909-2 would allow that this label is missing. This does not make much sense though, since in most cases there is no reasonable standard interpretation if the label is missing.

First, a variable x of built-in sort NAT is defined. Then a user-defined sort A is defined, which is a multiset over the built-in sort BOOL. Then, two named operations are defined, f and g . The operation f takes two parameters of type INT and the operation g does not have parameters. Note that named operations, basically, are abbreviations and, therefore, do not allow any recursion (see [5]). In the end, two other user-defined sorts are defined: B is a product of A and the built-in sort INT, and C is a pair over sort B. Note that also for sort declarations, recursion is not allowed.

The right-hand side of the sort declarations above give you an idea of the syntax for sorts already. There are some built-in sort like BOOL, INT, NAT, POS and DOT. From these, we can built products or multiset sorts.

Here are some examples of terms (using the above declarations):

```
x'f(x,x) ++ 1'x ++ x'g() ++ 1'5
```

```
1'(dot,1) ++ 1'(dot,1*1)
```

```
x > 1 and x < 5
```

The first is a multiset term over the sort INT, which could be used in arc inscriptions (if the attached place is of type INT). The second is a ground term over the product of built-in sort DOT with INT, where DOT is a sort that represents a type with a single element dot. The last term is a term of sort BOOL, which could be a condition.

The precise syntax is defined by the following grammar (that actually is a simplified version of the grammar that was used for generating the parser). The terminals ID, INT, NAT, STRING in this grammar represent legal identifiers and legal representations of integer numbers, non-negative integer numbers and string constants.

Listing 3.1 shows the part of the grammar for declarations. Listing 3.2 shows the part of the grammar for terms. Note that this part of the grammar is ambiguous for making it a bit more readable. The ambiguities are resolved by assigning a binding priority to the different operators – moreover all operators are left-associative. Every line in the declaration of BinOp represents operators on the same level of priority, where the first line has the least binding-power and the last the highest. The unary operators (actually there is only one) have the highest binding power of all. Note that there are also some operators like the cardinality, which use circumfix notation: if m is some multiset $|m|$ denotes the cardinality of that multiset. This operator has the same binding power as parentheses.

Listing 3.1: Grammar for declarations

```

Declarations :
  ( 'sorts' SortDecl ( ',' SortDecl )* ';' |
    'vars' VariableDecl ( ',' VariableDecl )* ';' |
4   'ops' OperatorDecl ( ',' OperatorDecl )* ';' |
    'sortsymbols' ArbitrarySort ( ',' declaration )* ';' |
    'opsymbols' ArbitraryOperator ( ',' ArbitraryOperator )*
    );

9 SortDecl :
  NamedSort;

NamedSort :
  ID '=' Sort;
14

VariableDecl :
  ID ':' Sort;

OperatorDecl :
19 NamedOperator;

NamedOperator :
  ID '(' ( VariableDecl ( ',' VariableDecl )* )? ')' '=' Term;

24 Sort :
  BuiltInSort | MultiSetSort | ProductSort | UserSort;

MultiSetSort :
  'MS' '(' Sort ')';
29

ProductSort :
  '(' ( Sort ( '*' Sort )* )? ')';

UserSort :
34 ID;

ArbitrarySort :
  ID;

39 ArbitraryOperator :
  ID ":" ( Sort ( "," Sort )* )? "->" Sort;

```

Listing 3.2: Grammar for terms

```

Term :
    Term BinOp Term |
    UnOp Term      |
    BasicTerm;

5
BinOp :
    // all binary operators are left-associative
    'or' | 'implies' | // lowest priority
    'and' |
10    '>' | '>=' | '<' | '<=' | 'contains' | // all comparison ops
        '<r' | '<=r' | '>r' | '>=r' | // on same level
        '<p' | '>p' | //
        '<s' | '<=s' | '>s' | '>=s' | //
    '==' | '!=' |
15    '++' | '--' |
    '<' |
    '+' | '-' |
    '*' | '**' | '/' | '%' ; // highest priority

20 UnOp :
    'not' ; // higher priority than all binary operators

BasicTerm :
    Variable |
25    UserOperator |
    OtherBuiltInOperator |
    BuiltInConst |
    '(' Term ')' | // a sub-term in parentheses
    '(' Term ( ',' Term )+ ')' ; // a tuple

30 Variable :
    ID;

UserOperator :
35    ID '(' ( Term ( ',' Term )* )? ')' ;

OtherBuiltInOperator :
    '|' BasicTerm '|' | '#(' Term ',' Term ')' |
    CyclicEnumsBuiltInOperator | PartitionsBuiltInOperator |
40    StringsBuiltInOperator | ListsBuiltInOperator;

```

Listings 3.3 and 3.4 show the part of the grammar for built-in sorts and constants. Note that every number constant will implicitly be assigned the

Listing 3.3: Grammar for sorts and constants (1)

```

BuiltInSort :
  Dot | Boolean | Number FiniteEnumeration | CyclicEnumeration |
  FiniteIntRange | StringSort | ListSort ;

5 BuiltInConst :
  DotConstant | BooleanConstant | MultisetConstant |
  NumberConstant | FiniteIntRangeConstant |
  StringConstant | ListConstant ;

10 MultisetConstant :
  'all' ':' Sort |
  'empty' ':' Sort;

Dot :
15   'DOT';

DotConstant :
  'dot';

20 Boolean :
  'BOOL';

BooleanConstant :
  'true' | 'false';

25 Number :
  'INT' | 'NAT' | 'POS' ;

NumberConstant :
30   INT (':' Number)?;

```

best fitting sort: INT, NAT, or POS. If a positive integer, say 5 should have the type INT instead, this can be represented by 5:INT, which is like a type cast in object-oriented programming languages.

In addition to these syntactical constraints, the terms must also be correctly typed, which we do not discuss here in detail.

For HLPNGs, there are many constraints. Like for PTNets, arcs may only run from places to transitions or from transitions to places. All of the

Listing 3.4: Grammar for sorts and constants (2)

```

FiniteEnumeration : 'enum' '{ ID ( ',' ID)* }' ;

CyclicEnumeration : 'cyclic' '{ ID ( ',' ID)* }' ;

5 CyclicEnumsBuiltInOperator :
    'succ' '(' Term ')' | 'pred' '(' Term ')' ;

FiniteIntRange : '[' INT '..' INT ']' ;

10 FiniteIntRangeConstant : INT FiniteIntRange ;

Partition :
    'partition' Sort 'in' ID
    '{ PartitionElement ( ',' PartitionElement)* }';

15 PartitionElement : ID ':' Term ( ',' Term)* ;

PartitionsBuiltInOperator : 'partition' ':' ID '(' Term ')';

20 StringSort : "STRING" ;

StringsBuiltInOperator :
    "concatstring" "(" Term "," Term ")" |
    // note that we do not have append (does not make sense)
25 "stringlength" "(" Term ")" |
    "substring" ":" NAT "," NAT "(" Term ")" ;

StringConstant : STRING ;

30 ListSort : "LIST" ":" Sort;

ListsBuiltInOperator :
    "concatlists" "(" Term "," Term ")" |
    "appendtolist" "(" Term "," Term ")" |
35 "listlength" "(" Term ")" |
    "sublist" ":" NAT "," NAT "(" Term ")" |
    "memberat" ":" NAT "(" Term ")" |
    "makelist" ":" Sort "(" (Term ( "," Term)* )? ")" ;

40 ListConstant : "emptylist" ":" Sort ;

```

other additional constraints concern the correctness of the labels of HLP-NGs. The following list gives an overview:

1. Every place must have a (correct) type.
2. Every declaration must be syntactically correct and correctly typed.
3. All declarations must properly resolve (must not be recursive).
4. Every term (in markings, arc annotations, and conditions) must be syntactically correct and correctly typed.
5. The marking of a place must be a ground term and must be a multiset over the sort of the place.
6. The arc annotation must be a term that is a multiset over the place's sort.
7. Every condition must be a term of sort `BOOL`.
8. All declarations must have a distinct name (actually, this causes a warning only since this is a condition on concrete syntax, which is not part of PNML).
9. Parameters of an operation declaration should have distinct names (actually, this causes a warning only since this is a condition on concrete syntax, which is not part of PNML).

3.6 Limitations and pitfalls

The current version of the ePNK (0.9.1) has some limitations and some problems, which will be discussed in this section in order to avoid some unpleasant surprises. If you identify other issues or problems, please, let us know.

3.6.1 Dirty-flag and saving

Technically, the tree-editor and the graphical editor for pages are different editors. Up to now, they are working together, but they are not tightly integrated yet. This has the effect that the “dirty-bit” of the editor (indicated by a star in the editor tab) is not always up to date and is not reliable. It is recommended that all save operations are done from the tree-editor and that the tree-editor is always open and the last to be closed. Before finally closing

the tree-editor, make some minor change (e.g. in a net's name) and then save it. Do not close the tree-editor on a net when there are still graphical editors open¹⁰.

In particular when a PNX-file is saved and closed with a graphical editor open, you will not be able to load it again. If you experience that problem, you can work around this by opening the corrupted file in a text-editor and delete all the diagram information of a graphical editor¹¹.

3.6.2 Graphical features

The graphical editor supports many features, such as different fonts, and colours for labels, colours and linewidth for nodes, and different versions of curved arcs. Up to now, this information is not stored in the PNML files (and some of it is not even supported by PNML at all).

Right now, the only graphical information that is stored in the saved PNML or PNX file is the following:

- Position and size of nodes
- Position (resp. relative position) of labels (the size of labels is not supported by PNML).
- Intermediate points of arcs (but only as polyline; the different kinds of curved lines of GMF are not supported by PNML anyway; the bezier curves that are supported by PNML are not supported by GMF).

You should not put too much effort in unsupported features, since they will be gone when you open the page next time. It is planned, to implement a toolspecific feature that would keep the graphics exactly as it was edited in the ePNK, though. But this would not be usable by other tools, since it is tool specific.

The ePNK will read all graphical attributes that are supported by PNML; but, except for the ones discussed above, they are not yet shown graphically. And some of the graphically attributes, like CSS attributes for fonts, colours, etc. are not checked for validity; they will be written exactly the same way they were, when loading the PNML file (see also Sect. 3.6.5).

¹⁰I will, eventually, implement a control that prevents the tree-editor from being closed as long as there are graphical editors and also make sure that the dirty-bit is properly updated. But, this does not have the highest priority on my worklist right now.

¹¹I plan to implement a toolspecific extension for the ePNK, which can save the diagram information of GMF in the PNML model; this will fix this problem.

3.6.3 Petri net types

As mentioned in Sect. 3.3.2, the type of a Petri net should be added right after the creation of a Petri net, and the type of the Petri net should never be changed later – except if you know exactly what you are doing. Otherwise, it could happen that the produced PNML is invalid.

For HLPNGs, it is no problem to change its kind any time, since this kind has an effect on the validation only, but no effect on the serialisation of the net to a PNML file.

3.6.4 Wrapping labels

All labels in ePNK can have line-breaks. In the graphical editor, the line-breaks can be added by pressing the CNTRL and ENTER at the same time.

3.6.5 Graceful PNML interpretation

The PNML files that are produced by the ePNK are PNML conformant as defined in ISO/IEC 15909-2. The only exception is, when some illegal graphical attributes are read from an existing PNML file; these attributes will not be touched by the ePNK, and therefore written again. But, if a PNML file is created by using the ePNK only and if it validates correctly, the saved file is PNML conformant.

The ePNK, however, is not a PNML validator. It reads PNML (and PNML-like) documents and writes them again in a graceful manner. This way, it is possible to save documents that do not properly validate and load them again. For example, when some elements do not have ids, the references to these elements are stored as XPath references, which is not conformant to PNML. If the ids are added later, and validation is successful, this will produce conformant PNML documents again.

3.6.6 Deviation from PNML

There is, however, one feature of HLPNGs that the ePNK is not able to read. These are labels that are directly attached to the net and not to one of its pages, which we call *net labels*. Since net labels are against PNML's own principles and they do not appear to be important anyway, net labels are not supported by the ePNK for now¹².

¹²If there is some evidence that this feature is important, the ePNK could be extended to support it – with some reluctance though.

The only way to read a PNML document with net labels with the ePNK right now is to open it in some text editor and add a pair of `<page>` and `<\page>` tags around all net labels. Then, the ePNK would be able to open it (with the original net labels on a separate page – without id and page name, which should be added then).

Another deviation of the ePNK from PNML is the following: In the ePNK, all declarations of HLPNGs can have a name attribute, which comes from the fact that `Declaration` implements the interface for symbol definitions (`SymbolDef`). As a consequence, also the `Unparsed` declaration can have a name attribute. In ISO/IEC 15909-2, `Unparsed` does not have a name attribute. Since the ePNK can read and write PNML files without these attributes and the ePNK currently does not create instances of `Unparsed`, this does not result in problems in practice. If the ePNK would read a unparsed element without a name attribute, a validation would show this as an error.

Chapter 4

Developers' guide

In this chapter, we discuss how to extend the ePNK, by defining new functionality, by defining new Petri net types, or by defining new tool-specific extensions. For all these extensions, the ePNK provides *extension-points* so that the extensions can be made without changing the actual code of the ePNK¹. Actually, the ePNK does not even provide own extension-points for adding functionality: The existing eclipse extension-points are good enough for that for now².

We start with Sect. 4.2, which shows how to add some functionality to the ePNK, which could be implementing a model-checker, or some other analysis or verification function, it could be a function that reads a net in PNML and produces some net in some other format, or a function that generates a net that is stored in the PNML format.

In Section 4.3, we discuss how to add a new Petri net type to the ePNK. Simple net types can almost completely be generated from a model, for more complex ones, such as high-level Petri nets, a mapping to XML must be defined.

At last, we discuss how to add tool-specific information to the ePNK in Sect. 4.4.

All of these extensions will be discussed by the help of some examples, which are deployed together with the ePNK. In these examples, we assume that the reader is familiar to the main principles and ideas of eclipse, its plug-in architecture, and eclipse plug-in development (see for example, the

¹Technically, you would not even need to see the code of the ePNK, but looking at it might help understand the ideas and principles behind the ePNK.

²Eventually, there might be some ePNK-specific extension-point for adding functionality for getting a more uniform “user-impression” of the added functionality. Since this is mostly a GUI-issue, this does not have the highest priority right now.

“Platform Plug-in Developer Guide” which is part of eclipse: “Help” → “Help contents” or one of the many eclipse resources [14, 4]); Sect. 4.1 will give a brief overview though. We go a bit more into the details of EMF and explain some of the steps that need to be done in EMF more explicitly, but it is also recommended to read up on some details on EMF [2] before starting with own development projects.

4.1 Eclipse: a development platform for the ePNK

As briefly discussed in Sect. 3.1 already, eclipse is an Integrated Development Environment (IDE). Here, we briefly explain how to set up the eclipse environment so that you can work on your own extensions and do the tutorials of this chapter – assuming that you have installed eclipse and the ePNK as explained in Chapter 1 already.

4.1.1 Installing the ePNK source projects

As a developer you would probably want the source code for the projects, though you would probably not need to change it, and it is recommended not to do so. You can download a zip-file “ePNK-developers-helios-version-0.9.1.zip” containing all projects with the source code from the ePNK homepage: <http://www2.imm.dtu.dk/~eki/projects/ePNK/>

In your eclipse development workspace, select “File” → “Import...” and, in the dialog, select “Existing Projects into Workspace” from category “General”. In the opened “Import” dialog select “Select archive file” and press “Browse” to select the “ePNK-developers-helios-version-0.9.1.zip” you obtained from the ePNK homepage. Then press “Finish”.

In your workspace you have now many plug-in projects, which are the basis for developing new plug-ins and, in particular, extensions of the ePNK. Therefore, this workbench is also called the *development workbench*. The reason for introducing an additional attribute to the term *workbench* here is that we will have another workbench once we start our tool with the new extensions, which is another instance of eclipse again: This workbench is called the *runtime workbench* since this is when the ePNK with your new extensions is running (and it will very much look like the original ePNK as discussed in Chapter 3). This runtime workbench can be started from the development workbench by “Run” → “Run Configurations...” and then selecting “Eclipse Application” and pressing the “New” icon and then “Run”. Later, it would be enough to press the “Run” button in the tool bar. For

now, however, we do not start the runtime workbench since we need to implement some new functionality first.

4.1.2 Installing Ecore Tools

As mentioned already, a major part of defining new Petri net types is creating ecore models with the new concepts of the Petri net type. There is an extension of eclipse that is called “Ecore Tools SDK”, which helps with creating the models. Installing this extension will also install all other tools necessary for generating code from the models.

To install this extension, start the Install dialog³ select the standard eclipse update site and select “Ecore Tools SDK (...)” and follow through the installation process.

After restarting eclipse, you can check whether the installation was successful: Open the project `org.pnml.tools.epnk` and, within that project, open the folder `model`. The files with extensions “.ecore”, “.ecorediag”, and “.genmodel” should have special icons now. If you open click on “PNML-CoreModel.ecorediag”, you should see a ecore model in a class diagram like graphical notation (resembling the PNML core model from Fig. 2.1).

4.2 Adding functions

In this section, we discuss how to add new functionality to the ePNK. To this end, the API that is generated from the PNML core model and the Petri net type definitions are discussed – giving you an idea of the general principles, which can be found in [2]. Via this API, you can access and also modify nets in a uniform way. Moreover, this section shows how to open, create and write PNML files from another program.

To this end, this section will discuss how to plug in functionality into the ePNK (or to eclipse in general) in three different ways.

- Section 4.2.1 shows how to implement a new eclipse view (cf. 3.1), which gives an overview of a PNML file that is selected in one of the eclipse resource explorer views.
- Section 4.2.2 shows how to implement a *wizard* for creating a PNML file (actually, we changed a wizard that was automatically created by the eclipse “new plug-in project wizard”). This wizard will create a simple PNML file with a P/T-System that contains a mutual exclusion

³Reminder: Select “Help” → “Install New Software...”.

algorithm for a user selected number of agents. Each of these agents will be shown on a different page.

- Section 4.2.3 shows how to implement a simple pop-up menu on a selected Petri net (in the tree editor), which starts a model checker, asking the user for some formulas to be checked, and then checking the formulas on the net. Since model checking can take quite some time, the model checker will run in the background and can be aborted by the user. This uses eclipse's concept of jobs. On the side, this shows how to use some of eclipse's user dialog functions.

4.2.1 Accessing a PNML file and its contents: A file overview

In this section, we discuss how to implement a new (and very simple) view, that will give an overview of the contents of a file that is selected in the explorer. Figure 4.1 shows a screenshot of the result. For the selected file “hlpng-gmf.pnml” in the “Project Explorer”, the “ePNK File Overview” in the bottom left, shows that the selected file is a Petri net document, which contains 3 Petri nets, a high-level net, a P/T-net, and an empty net – where the overview shows the respective “type tag” of PNML. The name of the first net is “A high-level next example”; the other two nets do not have a name.




Figure 4.1: The ePNK with the “File Overview” view

This view and its functionality is implemented in the plug-in project `org.pnml.tools.epnk.functions.tutorial` and we will go through this project now⁴. The implementation of the view is contained in a single class: `PNMLFileOverviewView` in the package

```
org.pnml.tools.epnk.functions.tutorial.overviewview.
```

We briefly explain the general structure of this view class, which is extracted in Listing 4.1, where we deleted imports and comment, which can be looked up in the source code. The class extends the eclipse `ViewPart`, which actually makes it a view and it implements the `ISelectionListener`, which allows our view to obtain the information on the current selection of the user in the workbench. Note that there is no explicit constructor. The reason is that the view will be set up via the `createPartControl` method: In the first three lines of that method, a viewer (which represents the content of that view) is initialized, and so-called providers will enable the view to properly show the contents. But, we do not discuss the details here. In the last two lines of the `createPartControl` method, our viewer registers itself with the eclipse selection mechanism as a *selection listener* and then creates the information that should be shown for the current user selection. We will discuss the respective method `selectionChanged` below. Note that there are two other methods. The `setFocus` just forwards the focus properly to the content of the view, once the view is focused. More important is the `dispose` method: the implementation of this method makes sure that our view removes itself as a selection listener once it is disposed (which typically would happen if the user decides to close the view).

Once the view has registered itself as a selection listener with the eclipse workbench, its `selectionChanged` method will be called whenever there is a change in the user's selection. In the implementation, of this method, the kind of the current selection is analysed and it is checked whether the first selected element is a file (i. e. whether it implements the interface `IFile`). If so, the method `getOverviewInfo` for computing the actual content of the file overview is called; this produces an array of `Strings`, which then will be set as the new content of that view – and, this way, shown to the user.

This `getOverviewInfo` is probably the most interesting part here, since it shows how to open and access a PNML or a PNX file (we do not even need to make a difference). The implementation of this method is shown in Listing 4.2 Up to line 7, it is checked whether the file extension is either “pnml” or “pnx” (the two file extensions, the ePNK has registered for the ePNK) and the path to that file is extracted and a URI is created. Actually,

⁴Note that this project also contains the implementation of the wizard in Sect. 4.2.2.

Listing 4.1: Class PNMLFileOverviewView: Infrastructure

```
package org.pnml.tools.epnk.functions.tutorials.overviewview;

import ...

5 public class PNMLFileOverviewView extends ViewPart
  implements ISelectionListener {

  private TableViewer viewer;

10 public void createPartControl(Composite parent) {
    viewer = new TableViewer(parent);
    viewer.setContentProvider(new ArrayContentProvider());
    viewer.setLabelProvider(new LabelProvider());
    getSite().getPage().addSelectionListener(this);
15 selectionChanged(null, getSite().getPage().getSelection());
  }

  public void setFocus() {
    viewer.getControl().setFocus();
20 }

  public void dispose() {
    super.dispose();
    getSite().getPage().removeSelectionListener(this);
25 }

  public void selectionChanged(IWorkbenchPart part,
    ISelection selection) {
    if (selection instanceof IStructuredSelection) {
30 IStructuredSelection structured =
      (IStructuredSelection) selection;
      Object first = structured.getFirstElement();
      if (first instanceof IFile) {
        viewer.setInput(getOverviewInfo((IFile) first));
35 } } }

  public String[] getOverviewInfo(IFile file) {
    ...
  }
40 }
```

Listing 4.2: Class PNMLFileOverviewView: Accessing the file

```

public String[] getOverviewInfo(IFile file) {
    String[] result = {"No ePNK file selected"};
    String extension = file.getFileExtension();
    if (extension != null &&
5      (extension.equals("pnml" ) || extension.equals("pnx"))) {
        String path = file.getLocationURI().toString();
        URI uri = URI.createURI(path);

        ResourceSet resourceSet = new ResourceSetImpl();
10      Resource resource = null;
        try {
            resource = resourceSet.getResource(uri, true);
        } catch (Exception e) {
            result[0] = "File could not be read.";
15      return result;
        }

        List<EObject> contents = resource.getContents();
        if (contents != null && contents.size() > 0) {
20      EObject object = contents.get(0);
            if (object instanceof PetriNetDoc) {
                PetriNetDoc document = (PetriNetDoc) object;
                List<PetriNet> nets = document.getNet();
                int no = nets.size();
25      result = new String[no + 1];
                result[0] = "The Petri Net Document contains "
                    + no + (no == 1 ? " net" : " nets:");
                no = 1;
                for (PetriNet net : nets) {
30      String name = net.getName() != null ?
                    net.getName().getText() : "unknown";
                    String type = net.getType() != null ?
                        net.getType().toString() : "unknown";
                    result[no++] = " " + name + ": " + type;
35      }
                } else
                    result[0] = "The file does not contain a PetriNetDoc.";
            } else
                result[0] = "The file does not contain any element.";
40  }
    return result;
}

```

eclipse provides also user dialogs and file dialogs that would allow us to ask the user for a file name that would be returned as a URI; here we abused the selection mechanism and the file to get hold of some legal URI of a pnml or pnx file. Therefore, the code that comes now, could be used at any other point when a program wants to read and access some file, once we have a String with the path of the file: This starts with creating a *resource set* and, within that resource set creating a *resource* with the given URI, which is the first parameter of the `getResource` method; the second parameter indicates whether cross-references to other resources should be resolved lazily or not (which is not relevant here). Note that in EMF, a resource or file should always be accessed (and created, see Sect. 4.2.2 for more information) in this way via a resource set. After we successfully got the resource, we can obtain its content by the `getContents` method which returns a list of its top-level objects – in case of PNML, this list should have length 1. Being defensive, we check whether the contents exists and whether its first element is an instance of `PetriNetDoc`. If so, we go systematically through all the contained nets, get their names and their PNML types and add a String with that information to the String array with the result. In the other cases, we return some error messages. Note that we do not need to close the file, or do anything else after we have obtained the information we need.

Let us have a closer look at how the contents of the Petri net document is accessed, once we have obtained a `PetriNetDoc` object. For any reference and attribute of the `PNMLCoreModel` there are respective getter and setter methods. For example, if we have a name label, the `getName` method will return the String with that name, and with `setName` we could set it – but we do not do that here. For attributes and features with a multiplicity greater than one, this is slightly different. For example a `PetriNetDocument` can contain many nets; therefore, `getNet` will return a list of nets, which then is iterated over to get the individual nets. And by adding a net to this list, this Petri net will be added to the Petri net document (see 4.2.2 for examples).

As stated above, the class `PNMLFileOverviewView` implements the “ePNK File Overview” as we have seen it in Fig. 4.1. But, if we just implement this class, it would not show up in eclipse, because eclipse would not know that it exists. In order to make the view known to eclipse we need to define it as an extension: This is done in the project’s “plugin.xml”. Double clicking on the plugin.xml file, will give you a convenient editor for defining and editing the extensions you want to define. Explaining the actual extensions are a bit easier with the XML fragment this editor produced. The fragment relevant for our overview view is shown in Listing 4.3. It says that an extension for

Listing 4.3: Defining the extension in “plugin.xml”

```

<extension
2   point="org.eclipse.ui.views">
  <category
    name="ePNK"
    id="org.pnml.tools.epnk.views.category">
  </category>
7  <view
    allowMultiple="false"
    category="org.pnml.tools.epnk.views.category"
    class="org.pnml. . . .overviewview.PNMLFileOverviewView"
    icon="icons/PetriNetDoc.gif"
12   id="org.pnml.tools.epnk.extensions.tutorials.pnmloverview"
    name="ePNK File Overview">
  </view>
</extension>

```

extension point `org.eclipse.ui.views` is defined, which is a new eclipse view. The category defines, where and under which category the new view can be found. We define a category specific for the ePNK and then define the actual view to use it. The attributes of the views say, that a view of this kind can at most be open once, that it uses the above category, refers to the class which actually implements it, `PNMLFileOverviewView`, and defines an icon (used in the tab of that view) and a name for that view.

Note that in order to access some of the classes like `Resource`, `ResourceSet`, and some of the ePNK classes like `PetriNetDoc`, `PetriNet`, etc. in the implementation of the view, we would also need to define by which plug-in projects they are provided: If you open the file `plugin.xml`, you will find these projects in the tab “Dependencies”. But this is a more technical issue, which we do not go into the details.

Now, you could start the runtime-workbench and then the “ePNK File Overview” view could be opened by the user by “Window” → “Show View” → “Other...” and then selecting “ePNK File Overview” in the category “ePNK”. This will show the view in the workspace as shown in Fig. 4.1.

4.2.2 Writing PNML files: Generating multi-agent mutex

Next, we will discuss how to create new files and fill them with some content. In typical applications, the contents might come from a file in a format of

another specific tool, which should be converted to PNML. In our example, however, we programmatically generate a Petri net: my favourite semaphore mutex example. To make it slightly more interesting, the number of agents competing for the semaphore is a parameter. This function is implemented as an eclipse wizard and it was implemented by creating a new file wizard for pnml files automatically by the eclipse “New Plug Project” wizard choosing the “custom plug-in wizard” with the choice of the “New File Wizard” in the “Template Selection” dialog. But, this does not need to bother you too much. If you are interested in the manual changes made to the automatically generated code, you will find all the manual changes in the two classes in the package `org.pnml.tools.epnk.functions.tutorials.wizards` enclosed by comments like `// eki: ...`.

In the rest of this section, we focus on the explanation of the parts of the implementation that are concerned with the file creation. This functionality is implemented in the method `createPNMLFile(String path, int number)` of the class `MultiAgentMutexNetWizard` in package

```
org.pnml.tools.epnk.functions.tutorials.wizards
```

in the plug-in project `org.pnml.tools.epnk.functions.tutorials`. The parameter `path` is a String representation of a path to the file that should be created. The parameter `number` is the number of agents that should be created in the mutex net that will be generated.

Creating a Petri net programmatically is quite simple, but code intensive. Therefore, we have split up the creation process into several parts for the different elements, which will be discussed top-down from creating the document, the net, its pages, and the places, transitions, reference places, and arcs on them. We discuss these methods one after the other – and omit some boring ones in the end (you will find all details in the source code). Listing 4.4 shows the method that creates the file. First, it calls the method `createPetriNetDoc` that creates the Petri net document, which is discussed later. This is the content of the file that we want to write. Then, we convert the path into an URI. Then, we use the resource set again – this time to actually create the file. Surprisingly, enough this is already all we need to do. At this point, we can add the content to the resource. Note that it does not even matter whether the resource is a PNML file or a PNX file – eclipse will, dependent on the file extension, chose the right implementation of the resource so that either a PNML file or a PNX file is written once we save the resource in the end. But, we configured the wizard in such a way that the user can chose only the “pnml” extension.

Adding the contents follows the same principle that we have discussed already. With `getContents` we get a list of EMF objects (which would be

Listing 4.4: Method `createPNMLFile(String path, int number)`'

```

public void createPNMLFile(String path, int number) {
    PetriNetDoc doc = createPetriNetDoc(number);

    final URI uri = URI.createURI(path);
5   ResourceSet resourceSet = new ResourceSetImpl();
    final Resource resource = resourceSet.createResource(uri);
    EList<EObject> contents = resource.getContents();
    contents.add(doc);
    try {
10   resource.save(null);
    } catch (IOException e) {
        // Do nothing for now if file could not be saved.
    }
}

```

empty, since the resource was newly created right now); then we add the Petri net document to this list. The only thing left is to save the resource, which is done by calling the `save` method. Note that the `save` method has a parameter, that could be used to configure the way a file is saved. But, `null` is fine here – and you should only change this, if you know exactly what you are doing.

So, let us dive a bit deeper: into the method `createPetriNetDoc`, which takes one parameter only – the number of agents. This method is shown in Listing 4.5. In the second line, a new Petri net document is created. Note that this is not done with the usual `new` construct. Rather, the factory for the PNML core model which is obtained by `PnmlcoremodelFactory.eINSTANCE` is used for this purpose. It is part of the EMF philosophy that we should not know anything about the actual implementation of classes. And EMF strongly recommends to create new objects only via these factories. Note that the new net and its type are also created by factories – since the type is plugged in, a the factory of that plug-in is used for this purpose `PtnetFactory` resp. its instance. After creating the net, its id is set, by the `setId` method. Note that this could be any string, but it is our responsibility to make sure that all id's are different (if we create them programmatically). Then the net is added to the list of nets of that document: to this end, we get the list of all nets of the document via `getNet` on which we call the `add` method. There is no way to directly add a net to a document. The type of the net can, again, be set with the `setType` method, since the type does not

Listing 4.5: Method `createPetriNetDoc(int number)`'

```
1 public PetriNetDoc createPetriNetDoc(int number) {
    PetriNetDoc doc = PnmlcoremodelFactory.
        eINSTANCE.createPetriNetDoc();

    PetriNet net = PnmlcoremodelFactory.eINSTANCE.createPetriNet();
6 net.setId("n1");
    doc.getNet().add(net);
    PetriNetType type = PtnetFactory.eINSTANCE.createPTNet();
    net.setType(type);

11 Name nameLabel = PnmlcoremodelFactory.eINSTANCE.createName();
    nameLabel.setText("Mutual exclusion");
    net.setName(nameLabel);

    Page page = createPage(type, "pg0", "semaphor page");
16 EList<Page> pages = net.getPage();
    pages.add(page);

    Place semaphor = this.createPlace(
        type, "semaphor", "semaphor", 1, 380, 140);
21 page.getObject().add(semaphor);

    for (int i=1; i<= number; i++) {
        page = createAgentPage(type, semaphor, i);
        pages.add(page);
26 }

    return doc;
}
```

allow for multiple values.

After that, a name label is created, its text value is set, and the name label is added to the net.

Next a new page is created by calling a separate method, which is added to the list of pages of the net, and the place semaphore is created as the single object on this page. To this end, we use the `createPlace` method.

In the for loop at the end of the method for each agent, there will be created one page with the net for each agent.

All the other methods follow the same principles, and there is not too much interesting to see in them. Therefore, we finish with discussing the method `createAgentPage`, which is shown in Listing 4.7. This method creates 3 places, one reference place (referring to the semaphore that was created on the first page above), 3 transitions, and 8 arcs. What makes this method a bit more interesting is the graphical information that is added to the arcs: some intermediate point, which makes the net look a bit nicer. If you have a closer look at the `createTransition`, `createPlace`, and `createRefPlace` you will find similar constructs for defining the position and size of the nodes, and the position of the labels associated with them. But this is straight-forward and follow the exact principles of ISO/IEC 15909-2 (see [5]).

In the runtime workbench (or a version of the ePNK in which this plugin is installed), you could invoke this function as follows: Go to the resource explorer – or any other explorer – of the workbench, press the right mouse button and select “New” → “Other...”, select “Multi-agent Mutex Net Wizard” in the category “ePNK”. Then, a dialog opens in which you can chose a folder⁵ (“container”) in which this file should be created, a “file name” (which must have extension “pnml”), and the number of agents. Note that, normally, the file creation wizard would overwrite existing files. This “multi-agent” wizard, however, was changed in such a way that existing files won’t be overwritten accidentally.

4.2.3 Long-running functions: A model checker

In this section, we discuss the implementation of a model checker for P/T-Nets, which are interpreted as Condition/Event-Systems here. To this end we use a simple library for symbolic model checking that was developed for teaching purposes: *Model Checking in Education (MCiE)*⁶. This library is

⁵If you have selected exactly one folder when you invoke the wizard, the fields of this dialog will be pre-set.

⁶see <http://www2.cs.uni-paderborn.de/cs/kindler/Lehre/MCiE/>

Listing 4.6: Method `createAgentPage()`'

```

1 public Page createAgentPage(PetriNetType type, Place sem, int i) {
    Page page = createPage(type, "pg"+i, "agent"+i);

    Place idle = createPlace(type, "idl"+i, "idl"+i, 1, 100, 220);
    Place pending = createPlace(type, "pen"+i, "pen"+i, 0, 100, 60);
6   Place critical = createPlace(type, "cri"+i, "cri"+i, 0, 300, 140);
    RefPlace semRef = createRefPlace("sem"+i, "sem", sem, 380, 140);

    Transition t1 = createTransition(type, "t1."+i, "t1."+i, 40, 140);
    Transition t2 = createTransition(type, "t2."+i, "t2."+i, 220, 60);
11   Transition t3 = createTransition(type, "t3."+i, "t3."+i, 220,220);

    Arc a1 = createArc(type, "a1."+i, idle, t1);
    Arc a2 = createArc(type, "a2."+i, t1, pending);
    ...
16   Arc a6 = createArc(type, "a6."+i, t3, idle);

    Arc a7 = createArc(type, "a7."+i, semRef, t2);
    Coordinate coordinate =
        PnmlcoremodelFactory.eINSTANCE.createCoordinate();
21   coordinate.setX(300);
    coordinate.setY(60);
    ArcGraphics arcGraphics =
        PnmlcoremodelFactory.eINSTANCE.createArcGraphics();
    arcGraphics.getPosition().add(coordinate);
26   a7.setGraphics(arcGraphics);

    Arc a8 = createArc(type, "a8."+i, t3, semRef);
    ...
    a8.setGraphics(arcGraphics);
31

    EList<Object> contents = page.getObject();
    contents.add(idle);
    contents.add(pending);
    ...
36   contents.add(t3);
    contents.add(a1);
    ...
    contents.add(a8);

41   return page;
}

```

deployed as part of the ePNK tutorials.

Within this developers' guide, we will not go into the details of model checking and its theoretical foundation, since this is not the point of this tutorial at all. For more information on model checking, we refer to a standard text book on model checking [3]. The point of this tutorial is to show how some function can be installed as a *pop-up* menu and *action* on a Petri net that is open in the tree editor. The actual action can be a function that might take some time and, therefore, should not block the graphical user interface of eclipse while it is running. To this end, the ePNK provides a way that makes it easy to use the eclipse *jobs*, which are running in the background – but of course provide the possibility to show a result.

The model checking functionality is implemented in the plug-in project `org.pnml.epnk.functions.modelchecking`. The main function (the actual model checking job) is implemented in the class `ModelcheckingJob` in package `org.pnml.epnk.functions.modelchecking.action`. The action initiating the model checking job is `ModelcheckingAction` in the same package.

Since the class `ModelcheckingAction` and the way it is integrated to the ePNK is quite simple, we start with explaining that first. It is shown in Listing 4.7. This class extends the `AbstractEPNKAction`, which is an ePNK convenience class implemented to make it easy to add a new action. We overwrite two methods: `isEnabled` and `createJob`. The method `isEnabled` checks whether the action is applicable for the selected Petri net. In our example, it checks whether the Petri net has type and whether this type is `PTNet`. The other method `createJob` creates the actual job, which is an instance of `ModelcheckingJob` with a Petri net and a defaultInput (a default formula in the user dialog in our case). This class extends the ePNK's convenience class `AbstractEPNKJob` and will be discussed later.

In order to make the action `ModelcheckingAction` know to eclipse and appear in the popup menu in the “ePNK” category, we need to define an extension. Listing 4.8 shows the part of the “plugin.xml” file that defines this extension.

At last, we have a look at the class `ModelcheckingJob`, which is implementing the user dialogs (asking the user for temporal formulas), converting the Petri net into ROBDDs, doing the actual model checking, and showing the result to the user again. In addition to the constructor, we need to implement (overwrite methods of `AbstractEPNKJob`) the following methods: `prepare()`, `getInput()`, `run()`, `showResult()`, and `canceling()`. Here is what these methods do:

Listing 4.7: The action class ModelcheckingAction'

```
package org.pnml.tools.epnk.functions.modelchecking.action;

3 import
    org.pnml.tools.epnk.actions.framework.jobs.AbstractEPNKAction;
import org.pnml.tools.epnk.actions.framework.jobs.AbstractEPNKJob;

import org.pnml.tools.epnk.pnmlcoremodel.PetriNet;
8 import org.pnml.tools.epnk.pnmlcoremodel.PetriNetType;
import org.pnml.tools.epnk.pntypes.ptnet.PTNet;

public class ModelcheckingAction extends AbstractEPNKAction {

13  @Override
    public boolean isEnabled(PetriNet petrinet) {
        if (petrinet != null) {
            PetriNetType type = petrinet.getType();
            return type != null && type instanceof PTNet;
18     }
        return false;
    }

    @Override
23     protected AbstractEPNKJob createJob(PetriNet petrinet,
        String defaultInput) {
        return new ModelcheckingJob(petrinet, defaultInput);
    }

28 }
}
```

Listing 4.8: Defining the popup action for the model checking action

```
<extension
2   point="org.eclipse.ui.popupMenus">
  <objectContribution
    id="org.pnml.tools.epnk.functions.modelchecking.contribution1"
    objectClass="org.pnml.tools.epnk.pnmlcoremodel.PetriNet">
    <menu
7      id="org.pnml.tools.epnk.actions.standardmenu"
        label="ePNK"
        path="additions">
      <separator
12     name="group1">
    </separator>
  </menu>
  <action
17     class="org.pnml.tools.epnk. ... .action.ModelcheckingAction"
        enablesFor="1"
        id="org.pnml.tools.epnk.functions.modelchecking"
        label="Model checker"
        menubarPath="org.pnml.tools. ... .standardmenu/group1">
    </action>
  </objectContribution>
22 </extension>
```

Constructor: Sets up all the data structures needed during the job; typically, this will be storing the default input. In our model checker example, we also set up some mappings, for mapping places of the Petri net to variables of the MCiE library, and mappings from transitions to formulas defining their behaviour, and some other information. The code of the constructor is shown in Listing 4.9.

Listing 4.9: The constructor of `ModelcheckingJob`'

```

public ModelcheckingJob(PetriNet petrinet, String defaultInput) {
    super(petrinet, "ePNK: Model checking job");
3   if (defaultInput != null) {
        defaultformula = defaultInput;
    }

    place2variable = new HashMap<Place,Variable>();
8   place2primedvariable = new HashMap<Place,Variable>();
    transitions    = new Vector<Formula>();
    placeNames     = new HashSet<String>();
    duplicateNames = false;
}

```

prepare(): This method is handling the user dialogs before the actual job starts. In our case, it asks the user for some CTL-formulas and, also allows the user to correct the input, if the formulas are syntactically incorrect – or to abort the action. The code for this user dialog is shown in Listing 4.10. Since this is the standard way of doing this in eclipse, we do not go into the details of this part though. The only relevant part for the ePNK is that the job will not be continued, if the `prepare()` method returns false – in the implementation of the model checking job, this is done, when the user presses cancel in one of the dialogs (line 11/12 and line 33/34).

In our model checker job, the `prepare` method will try to convert the Petri net into formulas defining the behaviour of the transitions and the initial marking. And on the way, it will be checked whether there are duplicate names of places, so that a warning can be issued. Listing 4.11 shows the part of the `prepare` method converting the initial marking into a state formula. The basic idea is that, in this formula, a variable corresponding to the place occurs exactly once. It occurs negated, if the place is not marked and it occurs without

Listing 4.10: The user dialog of the `prepare` method

```
...  
3  InputDialog dlg = new InputDialog(  
    null,  
    "ePNK: Model checker",  
    "Enter a comma separated list of temporal formulas please:",  
    defaultformula,  
8    null);  
    dlg.open();  
  
    if(dlg.getReturnCode() != Window.OK)  
        return false;  
13    defaultformula = dlg.getValue();  
  
    do {  
        try {  
18            Parser parser = new Parser(new StringReader(defaultformula));  
            formulas = parser.parseFormulaList();  
            parser.parseEnd();  
        } catch (Exception e) {  
            formulas = null;  
23            dlg = new InputDialog(  
                null,  
                "ePNK: Model checker",  
                "Syntax error in formula: \n\r" +  
                e.toString() + "\n\r" +  
28                "Fix the error please or press cancel:",  
                defaultformula,  
                null);  
            dlg.open();  
  
33            if(dlg.getReturnCode() != Window.OK) // Didn't click on OK!  
                return false;  
            defaultformula = dlg.getValue();  
        }  
    } while (formulas == null);
```

Listing 4.11: Building the formula for the initial marking (in `prepare()`)

```
FlatAccess flat = new FlatAccess(getPetriNet());

3  init = new Constant(1);
  for (org.pnml.tools.epnk.pnmlcoremodel.Place p : flat.getPlaces()) {
    if (p instanceof Place) {
      Place place = (Place) p;
      registerPlace(place);
8
      PTMarking marking = place.getInitialMarking();
      if (marking != null && marking.getText().getValue() > 0) {
        init = new BinaryOp(BinaryOp.AND,
13         init,
          place2variable.get(place));
      } else {
        init = new BinaryOp(BinaryOp.AND,
18         init,
          new UnaryOp(UnaryOp.NOT,place2variable.get(place)));
      }
    }
  }
}
```

negation, if the place is marked (with at least one token⁷). All these negated and un-negated variables will be connected by and operations and the formulas are represented in MCiE's data structure. What is more interesting here is that the ePNK provides a way to access a net that consist of pages with reference nodes in a flattened way. This convenience class of the ePNK is called `FlatAccess`, which can be initialized with a Petri net of any type. Then the instance `flat` is used to get all places of the net, independently of the pages they occur on. Likewise, `flat` provides methods to access all the transitions and to get all the input and output arcs of a place or transition (including the ones of the reference nodes referring to them). This way, it is easy to obtain the pre- and post-sets, without being bothered with the page structure. For some more examples of the use of these methods, you can have a look into the code that converts transitions into formulas, which however is not discussed here.

The last part of the `prepare` method, is converting the formulas into OBDD-representation and creating a transition system out of these formulas. This is shown in Listing 4.12. Again, this is specific to MCiE. But, there are two parts that are important for the `prepare` method in general: With `this.setName()`, we can give the job a specific name, which is used in eclipse's jobs view. In our example, we say that it is a model checking job, add the name of the net and the formula which the user entered. The last important part is that the `prepare` method returns `true` in order to indicate that the preparation successfully terminated, and the actual job can be run (in the background) now.

`getInput()` This method is called by the action, to get and store as default for the next call, the user input. In our case, the formula that was entered by the user during the prepare phase (in its String representation as entered by the user) is returned.

`run()` This method implements the part of the job that will be run in the background. In our case, this is the model checking job. This methods is actually quite simple (most of the programming work lies in the preparation). It is shown in Listing 4.13. But, it is the most computation intensive part, which is why we are using the job to run it in the background. Note that, in this method, we also prepare the `result` already in a String that will be shown to the user. But, there should not be any user dialog in the `run` method, since this method is

⁷Remember that we abuse P/T-Nets for representing Condition/Event-Systems.

Listing 4.12: Finishing the prepare method

```

Name name = getPetriNet().getName();
String netref = "";
if ( name != null && name.getText() != null) {
    netref = " on net " + name.getText();
5 }

this.setName("Model checking job" + netref + ": " + defaultformula);

Context context = new Context();
10 ROBDD is = init.toROBDD(context);
    ROBDD ts[] = new ROBDD[transitions.size()];
    ChangeSet css[] = new ChangeSet[transitions.size()];

for (int i = 0; i < ts.length; i++) {
15     ts[i] = transitions.get(i).toROBDD(context);
        css[i] = new ChangeSet(context);
        transitions.get(i).addChangedVariables(css[i]);
    }
    transitionssystem = new Transitionssystem(context,is,ts,css);
20
return true;

```

Listing 4.13: The run method

```

protected void run() {
    result = "Model checking results:\n\r";
    for (int i = 0; i < formulas.length; i++) {
4     ROBDD obdd = formulas[i].toROBDD(transitionssystem);
        result = result + " " + formulas[i] + ": " +
            transitionssystem.isValid(obdd) + "\n\r";
    }
}

```

run in a separate thread in the background – and user dialogs would require to be called from a dedicated GUI thread.

`showResult()` This is the method that will be called for showing the result to the user. And, the infrastructure from `AbstractEPNKJob` will make sure that it will be called from the dedicated GUI thread again. Therefore, we can use all eclipse dialogs for showing the result. Listing 4.14 shows the implementation of this method. The result String, which was prepared during the `run` method is shown to the user by initiating an information dialog.

Listing 4.14: Code for showing the final result

```

protected void showResult() {
2   MessageDialog.openInformation(
        null,
        "ePNK: Model checker",
        result
    );
7 }

```

`canceling()` This method is a call-back mechanism that allows the user to abort jobs. In the case of computation intensive jobs such as model checking, to abort the computation and not to let that thread continue in the background is very important; otherwise this thread would consume all the computation power until it finishes on its own – which could take extremely long. Therefore, MCiE provides a mechanism for aborting model checking operations on some model, by invoking `abort()` – from a different thread of course. Our implementation of the `canceling()` method invokes this `abort()` method to actually finish the model checking – possibly with some delay. This is shown in Listing 4.15, where `transitionsystem` is the one that was constructed in the `prepare` method and on which the model checking is done.

Together, the `ModelcheckingAction` and the `ModelcheckingJob`, plugged in via the “plugin.xml” implement the complete model checker. It can be invoked in the run-time workbench, by right-clicking on a Petri net of type `PTNet`, and then selecting “ePNK” → “Model checker”. Then, a user dialog asks for the formulas that should be checked. After that, the model checking runs as a background job: eclipse indicates a running (non-system) job

Listing 4.15: Code for aborting the model checking job

```

protected void canceling() {
    if (transitionsystem != null) {
3      transitionsystem.abort();
    }
}

```

at the bottom of the workspace by a small icon with a running bar in the “progress area”. Once the job is finished, the icon will show up with an exclamation mark. When the user clicks on it, the result dialog is shown. If the eclipse progress area and icon are too small for you, you can also open the eclipse progress view (“Window” → “Show View” → “Other...” and then select “Progress” in category “General”). This will show all currently running jobs, and the ones that are finished. When clicking on a finished job, the result is shown. Jobs that are still running, can be aborted by pressing the red “abort” button.

For the syntax of the temporal formulas, we refer to the documentation of the MCiE library <http://www2.cs.uni-paderborn.de/cs/kindler/Lehre/MCiE/> and its example formulas or have a look into the documentation of the parser package. Places will be represented as MCiE variables in the formula. But you need to make sure that places in the Petri net have legal MCiE variable names (in particular, there should not be white spaces or special characters in them). One speciality is that the binary temporal operators such as EU, AR, are represented in infix notation like $p1EU p2$ and not in the more usual notation $E[p1Up2]$.

4.3 Adding Petri net types

As mentioned several times already, it is one of the main features of the ePNK that new Petri net types can be plugged in. In this section, we discuss how to plug in a new Petri net type. In Sect. 4.3.1, we start with a simple version, for which we, basically, need to provide a model with the extensions only; as an example, we use P/T-Systems (PTNet), which come as an integral part of the ePNK; but it is defined with ePNK’s type definition mechanism.

For more complex Petri net types, we can also define the mapping from the concepts of the Petri net type in the model to their XML representation; and, for Petri nets with structural labels, a *parser* and a *linker* for the

labels must be provided. The parser is needed to convert a textual label in its concrete syntax to the structure; the linker is needed for linking the use of symbols in some label to their definition in others. We use the example of high-level Petri nets (HLPNG) for discussing the relevant details in Sect. 4.3.2.

4.3.1 Simple Petri net type definitions: PTNet

The definition of P/T-Systems follows almost exactly the general idea outlined already in Sect. 2.2.2 and the ecore model that we use in the Petri net type definition, is almost a copy of the one that we have seen in Fig. 2.2 on page 8 already. It remains to discuss some of the differences in these models, and to discuss the steps to make the type known to the ePNK (in short to plug it into the ePNK).

4.3.1.1 The model

The Petri net type `PTNet` is defined in the plugin project

`org.pnml.tools.epnk.pntypes.`

The main part is the ecore model in “PTnet.ecore”, where the diagram information is contained in “PTNet.ecorediag”. You can open this diagram by double-clicking on it in the resource explorer⁸ (see Sect. 4.1.2). The diagram is shown in Fig. 4.2.




Figure 4.2: The ecore model for the Petri net type `PTNet`

⁸If for some reason, you have opened this diagram with another editor, you can open it with the diagram editor again by selecting the file, right-clicking on it and then selecting “Open With” → “Ecore Diagram Editing”.

There are only some minor, but important differences, to the model from Fig. 2.2 on page 8. We discuss these differences below:

1. There is a class `PTNet` in the ecore model, which did not exist in the conceptual model. The reason is that packages are not very tangible in programming and in the eclipse plugin mechanisms. Therefore, we define a Petri net type as an explicit class, which inherits from `PetriNetType` from the PNML core model (`pnmlcoremodel`). It is the class `PTNet`, which will be plugged in as the extension to the ePNK later. Moreover, this class implements some methods that help the ePNK to access the information about its labels; the details, however, do not need to concern us right now.
2. There are two new classes `Place` and `Arc`, which inherit from the classes `Place` and `Arc` from the `pnmlcoremodel`. And it is these new classes to which the additional labels are attached. The reason for this is, that ecore does not have the concept of merging packages. Instead, the extended information for the specific Petri net type is attached to the derived classes in this new package. There could be also a class for `Page` and `Transition`, but we do not need them here, since in P/T-Systems only places and arcs need to have additional labels.

Note that the name of these two classes are the same as in the PNML core package, which is not ambiguous since the new classes are in a new package. For now, we assume that the names of these classes are the same as in the PNML core model⁹.

3. The additional classes for labels, `PTMarking` and `PTAnnotation`, are attached to the new class `Place` and `Arc` as in the conceptual model via a composition – only the directive “refines” is missing, due to the missing concept of merging packages in ecore. The features `text` are directly represented as an attribute of type `NonNegativeInteger` and `PositiveInteger`, which are predefined data types of the ePNK that represent the respective XML Schema data types, which are used in ISO/IEC 15909-2. The cardinality for the `text` attributes is 1 in both cases – the same as in the conceptual model.
4. The new labels `PTMarking` and `PTAnnotation` are derived from the PNML core model class `Label` and not, as in the conceptual model, from `Annotation`. The reason is that up to now, the ePNK does not

⁹In principle, the names could be different; but this would require some extra programming, which we do not discuss in this manual for now.

distinguish between attribute and annotation labels (and in ISO/IEC 15909-2 all concrete labels are annotations).

5. A last difference is that there is no OCL constraint in this model. In the ePNK, constraints will be plugged in in a different way: as EMF constraints (see Sect. 4.3.1.4).

Such a diagram and model can be created by the “Ecore Tools”, which will not be discussed here (see the “EMF Ecore Tools Developer Guide” in the “Eclipse Help” and the web pages for some information).

4.3.1.2 Generating the code

Also the code generation from that ecore model follows the EMF standard procedure. But, we will briefly go through the process of generating all the relevant code below.

Before we can generate the code from the model, we need to create the so-called *generator model* (“genmodel”). This generator model will contain some information on how the code should be generated. For example, the “genmodel” contains the information to which project and which packages, the java code for the model should be generated. The “genmodel” also allows us to configure the generation of the EMF tree editor; for example, we can state whether a features can be changed, whether it should be shown as a child element or as a property, etc. See [2] for more details. A new “genmodel” can be created from an ecore model by selecting the ecore model (“PTnet.ecore” in our case), clicking the right mouse button, and selecting “New” → “Other...” in the pop-up menu and then, in the “New” dialog choosing “EMF Generator Model” in the category “Eclipse Modeling Framework”. In the case of a new Petri net type, we will have references to other models like the PNML core model and their “genmodel”; in the wizard for creating the “genmodel”, make sure that you do not chose these other models as root models, but that you add (and select) the respective generator models in the lower part as “Referenced Generator models” instead. You do not need to make any changes in the “genmodel”, but we recommend that you change the “Base package” property to some reasonable path.

From the “genmodel”, we can create the code for the model (*model code*), and the code with the infrastructure for all editors, which is called *edit code*. But, we do not need to generate the editor code. This is done by opening the “genmodel”, and then selecting (after clicking the right mouse button) “Generate Model Code” and “Generate Edit Code”. After that, you will

find¹⁰ the code for the model in the “src” folder of the project with the model and “genmodel”. Moreover, the “plugin.xml” will make the model and its code known to eclipse by an extension:

```
org.eclipse.emf.ecore.generated_package
```

The edit code will be generated in a project with the same name extended by a suffix “.edit”. We do not need to change anything in the edit code¹¹.

We need to make a minor change in the model code, though, which is discussed in the next section.

4.3.1.3 Adding the Petri net type to the ePNK

After the above steps, the code for the new model is known to eclipse. But, the ePNK will not know that it is a Petri net type. To this end, we need to define another extension. Before, we can do that, we need to make a minor change in the automatically generated code. We need to make the constructor of the class that represents the new Petri net type public; in our example, this concerns the constructor of the class `PTNetImpl`, which can be found in the package `org.pnml.tools.epnk.pntypes.ptnet.impl`.

Listing 4.16 shows this class. You can see the constructor, which is public now. The manual change is indicated by the `@generated NOT` tag¹². The other manual change is the addition of the `toString()` method. This defines the value of the PNML type attribute for nets for that particular Petri net. Here, we use the one from ISO/IEC 15909-2 for P/T-Systems.

With this public constructor, we can plugin the `PTNetImpl` as a new Petri net type to the ePNK now. To this end, we use the extension point `org.pnml.tools.epnk.pntd` in the “plugin.xml”. Listing 4.17 shows the relevant part from the “plugin.xml” file in the project

```
org.pnml.tools.epnk.pntypes.
```

The attribute `point` refers to the ePNK type definition extension point, the `id` is a unique id for the new type, and `name` is some conclusive name (we use the one from ISO/IEC 15909-2). The type refers to the class that implements the new type; in our example, this is our `PTNetImpl` class. In general, the class that is chosen here must extend the class `PetriNetTypeImpl` from

¹⁰If you do not say otherwise in the “genmodel”.

¹¹In the `org.pnml.tools.epnk.pntypes.edit` project with the “edit code” for PTNets, some of the automatically generated icons in the folder `icons/obj16` have been replaced by some more appropriate images, but this is just a matter of usability.

¹²Actually, we could just delete the tag `@generated`, but it is easier to search for and keep track of manual changes, if they are tagged with `@generated NOT`.

Listing 4.16: The class PTNetImpl with manual changes

```
package org.pnml.tools.epnk.pntypes.ptnet.impl;

import org.eclipse.emf.ecore.EClass;
import org.pnml.tools.epnk.pnmlcoremodel.impl.PetriNetTypeImpl;
5 import org.pnml.tools.epnk.pntypes.ptnet.PTNet;
import org.pnml.tools.epnk.pntypes.ptnet.PtnetPackage;

// @generated
10 public class PTNetImpl extends PetriNetTypeImpl implements PTNet {

    /**
     * @generated NOT
     * @author eki
15 */
    public PTNetImpl() {
        super();
    }

20 /**
     * @generated
     */
    @Override
    protected EClass eStaticClass() {
25     return PtnetPackage.Literals.PT_NET;
    }

    // @generated NOT
    // @author eki
30 @Override
    public String toString() {
        return "http://www.pnml.org/version-2009/grammar/ptnet";
    }

35 }
```

Listing 4.17: The extension PTNetImpl

```

<extension
  id="org.pnml.tools.epnk.pntypes.ptnet"
  name="PTNets"
  point="org.pnml.tools.epnk.pntd">
5  <type
      class="org.pnml.tools.epnk.pntypes.ptnet.impl.PTNetImpl"
      description="Place/Transition Nets">
  </type>
</extension>

```

the PNML core model code, and have a public constructor (that is why we needed the manual change). The description can contain a longer description of the new net type – for P/T-Systems, we guessed that no further explanation would be needed.

4.3.1.4 Adding constraints

As mentioned earlier, it is not allowed in P/T-Systems to have arcs that run from places to places or from transitions to transitions. In the conceptual model of PTNets, this was included as an OCL constraint in the UML model already. In the ePNK, this constraint must be added separately, which is done by the standard mechanisms of EMF Validation in the “plugin.xml”.

Listing 4.18 shows the part of the “plugin.xml” that defines this constraint. The actual OCL constraint is defined in the bottom in the XML CDATA part. This OCL expression resembles the one from the conceptual UML model, but is slightly different syntactically – which is due to the specific technology. Moreover the “headline” that states the context `Arc` is missing, since in the EMF Validation technology, the context is explicitly set by the `target` element, which you can find immediately above, which is the `Arc` of the `ptnet` package (the URL is defined in the model). The declaration of the events is necessary here, since we made this constraint a *live constraint*, which means that the editors will make sure not to violate it during editing. To this end, the editor needs to know the change of which features might violate the constraint; in our example, this is setting the source or the target of an arc.

The rest of this constraint definition is a bit more technical. But, we go briefly through it. The extension that we actually define is a *constraint provider*, which consists of the package it refers to and the constraints. In

Listing 4.18: Adding a constraint for PTNets

```

1 <extension point="org.eclipse.emf.validation.constraintProviders">
  <constraintProvider cache="true">
    <package
      namespaceUri="http://org.pnml.tools/epnk/pnmlcoremodel">
    </package>
6
    <constraints categories="org.pnml.tools.epnk.validation">
      <constraint
        id=
11      "org.pnml.tools.epnk.pntypes.ptnet.validation.PT_TP_ArcsOnly"
        lang="OCL"
        mode="Live"
        name="PT or TP Arcs only"
        severity="ERROR"
        statusCode="301">
16      <message>
        The arc {0} must run from a place to a transition or vice versa.
        </message>
        <description>
        Arcs between two places or transitions are forbidden in
21 P/T-nets (see Clause 5.3.1 of ISO/IEC 15909-2).
        </description>
        <target
          class="Arc:http://org.pnml.tools/epnk/types/ptnet">
          <event name="Set">
26          <feature name="source"/>
          <feature name="target"/>
          </event>
          </target>
        <![CDATA[
31      ( self.source.oclIsKindOf(pnmlcoremodel::PlaceNode) and
        self.target.oclIsKindOf(pnmlcoremodel::TransitionNode) ) or
        ( self.source.oclIsKindOf(pnmlcoremodel::TransitionNode) and
        self.target.oclIsKindOf(pnmlcoremodel::PlaceNode) )
        ]]>
36      </constraint>
      </constraints>
    </constraintProvider>
  </extension>

```

our case, the package is the PNML core model – even though it is for a specific Petri net types. The reason is that the validation always starts from the PNML core model. The constraints are defined in a category; the ePNK defines its own category, which is used here:

```
org.pnml.tools.epnk.validation
```

Each constraint must have a unique `id`, must state the language it is defined in (OCL in our example), have a `name`, a `severity`, and a `statusCode`. The status code can be freely chosen; the ePNK uses codes starting with a 3 for constraints concerning Petri net types. The mode can be *live* or *batch*; in live mode, the graphical editor will watch them and not allow edit operations that would violate them – this is what we chose in our example. Other constraints, like correctness of structured labels, might be defined to be in batch mode; the graphical editor will allow for syntactically incorrect labels, but this will be reported when validating the net. The last information in the constraint is a `message`, which is shown to the end user when the constraint is violated. The parameter `{0}` refers to the object that violates the constraint (in its String representation) – for constraints other than OCL, there could be more parameters. Moreover, there is a longer `description` of the constraint.

As mentioned above, the constraint can be formulated in different languages. It could, for example, be in Java, which would require to implement a java class. There are some examples of such java constraints in the HLPNG definition. Java is more convenient for implementing more complex constraints.

Note that we did not define any mapping from the concepts defined in the core model of P/T-Systems to their representation in XML. The reason is that, the standard mapping is good enough: the name of the composition in which the label is contained is the XML element, and the text feature of the label is mapped to the XML element `<text>` (see Fig. 2.1 on page 10 for an example). A mapping to XML needs to be defined only when the standard mapping is not enough, or when we have structured labels, which will be discussed in the next section.

4.3.2 Petri net type definitions in general: HLPNG

In this section, we discuss some more advanced mechanisms that can be used for defining new Petri net types. These mechanism will be discussed by the help of the Petri net type definition of high-level Petri nets (HLPNGs). Therefore, we start with an overview of the concepts of HLPNGs, from the

implementation point of view (for the conceptual part we refer to Sect. 3.5.2 and for a detailed discussion of all models and concepts, we refer to [5]).

4.3.2.1 Overview of HLPNGs

As discussed in Sect. 3.5.2, HLPNGs have different kinds of complex (or structured) labels: *declarations* of variables, sorts, and operators; *types* defining the sort of the tokens of a place, *marking* which are multiset terms defining the initial marking, *conditions* as transition guards, and *arc annotations* that define which tokens are consumed, resp. produced when a transition fires. What is more, the labels cannot be considered isolated from each other any more – some labels, like markings, arc annotations, or conditions may use *symbols* that are defined in other labels – in particular, in the declarations.

Figure 4.3 shows the ecore model defining these concepts of HLPNGs, which can be found in the folder `model` in project¹³

```
org.pnml.tools.epnk.pntypes.hlpngs.pntd
```

This model follows the same principles as the model for PTNets, which was discussed in Sect. 4.3.1.1. The main differences are, that the defined Petri net type `HLPNG` extends a more advanced class `StructuredPetriNetType`, and all labels extend `StructuredLabel`, which are part of the PNML core model. These two classes provide the infrastructure needed for parsing the textual labels and for establishing the links between these labels. This structure will be discussed in Sect. 4.3.2.2.

The actual content of all these labels is defined in their containment `structure`; note that we use `Term` as the contents for `HLMarking`, `Condition`, and `HLAnnotation`, since all of them are terms – just with different additional constraints imposed on them (see Sect. 3.5.2 and Sect. 4.3.2.3).

The detailed structure and concepts of terms, sorts, and declarations, are defined in several other models. Since these details are not too relevant for understanding the definition of structured Petri net types, we discuss only the main part of that model. This part of the model is shown in Fig. 4.4 – this as well as the diagrams of all the other models can be found in the plugin `org.pnml.tools.epnk.pntypes.hlpngs.datatypes`. For a detailed discussion of these models and their concepts, we refer to [5]. There is only

¹³This is the plugin in which HLPNGs are plugged into the ePNK; since HLPNGs are quite complex, and require many models, and also the implementation of a parser, the underlying definitions are defined in different other projects; all of these projects have a name with prefix `org.pnml.tools.epnk.pntypes.hlpngs` – some of them are generated automatically from models or from a grammar.




Figure 4.3: The ecore model for HLPNGs

one important difference, which are the classes `SymbolDef` and `SymbolUse`, which do not occur in the models of ISO/IEC 15909-2. These two classes are the ePNKs infrastructure for dealing with the definition of symbols and their use in a uniform and generic way. They are part of the PNML core model concerning structured Petri net types, which will be discussed in the next section.

4.3.2.2 Structured Petri net types and structured labels

As mentioned above, the ePNK provides some general interfaces and infrastructure for defining structured Petri net types that extracts the general concepts of more complex types. This is, again captured in models (and the code generated from them).

The model for structured Petri net types can be found in the `model` folder of the ePNK core project `org.pnml.tools.epnk`:

```
PNMLStructuredPNTypemodel
```

The diagram is shown in Fig. 4.5. We know the classes `PetriNetType` and `Label` as well as the interface `ID`, which is used for all ePNK elements that have an id, already from the PNML core model. The abstract class `StructuredLabel` extends the class `Label`, it has an attribute `text`, which stores the content of this label as a text String. The actual structural contents will be defined by classes that extend it (we have seen some examples




Figure 4.4: The ecore model for the main concepts of HLPNGs

in Fig. 4.3 already). Since, the ePNK cannot not know these concrete implementations, classes extending the structural label must make the reference to this structural contents known to the ePNK. This is achieved by the method `getStructuralFeature()`; as long as the feature for the structure is called 'structure' in the model, we do not need to do anything in the implementation (the ePNK will access this feature in a reflective way); only if for some reason, the model chooses a different name, this method must be implemented manually. Moreover, every class for a structural label must provide a method for parsing a String – a representation of this label in concrete syntax; an implementation of this method may return null, if the text cannot be parsed, or it must return some object (to be precise an `EObject` which is the EMF version of objects) with all the substructure of that label. In particular, that object must have a type that is compatible with the




Figure 4.5: The model for structured Petri net types

labels structural feature. This method must be implemented manually for every new extension since the ePNK cannot guess the concrete syntax.

The abstract class `StructuredPetriNetType` has one additional method, which must provide a `Linker` for linking the uses of some symbols to their definitions, which are captured by classes `SymbolDef` and `SymbolUse`. A `SymbolDef` has an `ID` and has a name, which will be used to refer to it (the `id` is internal to PNML and the ePNK). This name will be used in `SymbolUse`, again as attribute name, to refer to the definition. The feature that actually refers to the definition, can be accessed via the method `getRefFeature()`. Since the ePNK, does not know anything about how to make these connections, the Petri net type needs to provide the linker (via the method `getLinker()`). `Linker` is an interface: a single method `getGlobalLinks()`, which takes a Petri net and returns a `SymbolUseMapping`, which is also an interface. Conceptually, the `SymbolUseMapping` maps a `SymbolUse` to its definition `SymbolDef`. All the symbol uses for which there exists a mapping, can be obtained (as a list) via the method `getSymbolUse()`; and for each symbol use, the method `getSymbolDef()` will return the definition of that symbol.

With this infrastructure, the ePNK can deal with all kinds of structural labels, provided they implement the required methods. We will have a look at an implementation examples next: We consider the label `Condition` in the Petri net type definition for HLPNGs again (see Fig. 4.3) – the other labels are similar. Its structural feature is the containment `structure` to class `Term`. Since this is the standard name for structured labels, we do not need to override the method `getStructuralFeature`. But, we need to implement the parse method. The parsers for all labels of HLPNGs were automatically generated by Xtext, and made available in a singleton class `HLPNGParser` in package

```
org.pnml.tools.epnk.pntypes.hlpngs.datatypes.concretesyntax
```

in a project with the same name. For parsing a term, `HLPNGParser` provides a method `parseTerm(String)`. This singleton and its method `parseTerm` is used in the implementation of `ConditionImpl` (you will find it in the package

```
org.pnml.tools.epnk.pntypes.hlpng.pntd.hlpngdefinition.impl
```

in project `org.pnml.tools.epnk.pntypes.hlpng.pntd`).

Since linking is across all the different labels of a net, there is only a single linker for every net. For HLPNGs, this is implemented in package `org.pnml.tools.epnk.pntypes.hlpngs.datatypes.concretesyntax`. Linking in project

```
org.pnml.tools.epnk.pntypes.hlpngs.datatypes.concretesyntax
```

This class is `HLPNGLinker`; basically it goes through the complete Petri net twice; in the first round, it creates a symbol table of all symbol definitions; in the second round, this symbol table will be used to look up the definitions for every symbol use, which is stored in the `SymbolMapping`, which implements the `SymbolUseMapping` that we discussed above.

To make this linker known to the ePNK, the class `HLPNGImpl` in package `org.pnml.tools.epnk.pntypes.hlpng.pntd.hlpngdefinition.impl` implements the method `getLinker()`: it returns an instance of `HLPNGLinker`. Note that, in the class `HLPNGImpl`, we also need to make the constructor public and implement the `toString()` method (as discussed in Sect. 4.3.1.3).

4.3.2.3 Constraints

For HLPNGs, we needed to implement quite many constraints. As an example for a java constraint, we discuss just one example here. The rest of them would not provide much insight into the mechanisms of the ePNK – though they might provide some insights to the inner workings of HLPNGs themselves. There is also one OCL constraint, which forbids connecting

places with places and transitions with transitions. But, this exactly as for PTNets, which is why we do not discuss it here again.

All constraints for HLPNGs are defined in the project

```
org.pnml.tools.epnk.pntypes.hlpng.pntd,
```

the java constraint implementations can be found in the package

```
org.pnml.tools.epnk.pntypes.hlpng.pntd.validation
```

We discuss the constraint that transition conditions must have type boolean, which is implemented in class `ConditionIsBoolType`. Listing 4.19 shows this class. This constraint extends the class `AbstractModelConstraint` from EMF Validation and implements the method `validate`. From the validation context, it obtains the target object, which should be a transition (see later). But, we are defensive and check that explicitly. Then, we obtain the condition label of that transition, and if it is not null, get the term (its structure). Then, we check whether the sort of the term is boolean¹⁴. If it is not, we return a failure status via the validation context, and add the transition and the textual label to an array of objects (which is used in the error message to be defined later). Otherwise, we return a success status. Note that the EMF Validation Framework makes sure that this `validate` method is called for all transitions of a selected Petri net, Petri net document or page, once it is properly plugged in, which is discuss below.

Plugging in a java constraint is very similar to plugging in OCL constraints. The relevant fragment of the “plugin.xml” is shown in Listing 4.20. The main differences are that the attribute `lang` is “Java” now and the attribute `class` refers to the class `ConditionIsBoolType`, which was discussed above. As target class, the transitions of HLPNGs are defined (that is why we could assume that the target object is a transition). Another difference is that this is no live constraint, but a batch constrain. This means, that it might be violated during editing, and a violation will only be detected on explicit validation. Since this is a batch constraint, we do not need to declare any events in the target.

Another difference to the OCL constraint is, that we can refer to several parameters in the message now. What the different parameters are, depends on the return value of the validation method. In our case, this was the transition (or its String representation) and the text of the label.

The ellipses (“...”) indicate that the constraint that we have discussed here, is just one of many other constraint, which are not discussed here.

¹⁴The implementation of `getSort()` for terms is actually quite complex, but we do not discuss that here.

Listing 4.19: The constraint that conditions have type boolean

```

1 package org.pnml.tools.epnk.pntypes.hlpng.pntd.validation;

import org.eclipse.core.runtime.IStatus;
import org.eclipse.emf.ecore.EObject;
import org.eclipse.emf.validation.AbstractModelConstraint;
6 import org.eclipse.emf.validation.IValidationContext;
import
    org.pnml.tools.epnk.pntypes.hlpng.pntd.hlpngdefinition.Condition;
import
    org.pnml.tools.epnk.pntypes.hlpng.pntd.hlpngdefinition.Transition;
11 import
    org.pnml.tools.epnk.pntypes.hlpngs.datatypes.booleans.Bool;
import org.pnml.tools.epnk.pntypes.hlpngs.datatypes.terms.Sort;
import org.pnml.tools.epnk.pntypes.hlpngs.datatypes.terms.Term;

16 public class ConditionIsBoolType extends AbstractModelConstraint {

    public IStatus validate(IValidationContext ctx) {
        EObject object = ctx.getTarget();

21     if (object instanceof Transition) {
        Transition transition = (Transition) object;
        Condition condition = transition.getCondition();
        if (condition != null) {
            Term term = condition.getStructure();
26         if (term != null) {
            Sort sort = term.getSort();
            if (sort != null) {
                if (!(sort instanceof Bool)) {
                    return ctx.createFailureStatus(
31                        new Object[] {transition,
                                    condition.getText()});
                }
            }
        }
36     }
    }
    return ctx.createSuccessStatus();
}
}

```

Listing 4.20: Adding the constraint for conditions

```
<extension point="org.eclipse.emf.validation.constraintProviders">
  <constraintProvider cache="true">
    <package
      namespaceUri="http://org.pnml.tools/epnk/pnmlcoremodel">
5    </package>

    <constraints categories="org.pnml.tools.epnk.validation">
      ...
      <constraint
10      lang="Java"
      class="org.pnml. ... .validation.ConditionIsBoolType"
      severity="ERROR"
      mode="Batch"
      name="Condition is of type boolean"
15      id="org.pnml. ... .validation.ConditionIsBoolType"
      statusCode="314">
      <target class=
"Transition:http://org.pnml.tools/epnk/pntypes/hlpng/pntd/hlpng"/>
      <description>
20      The condition must be of type BOOL.
      </description>
      <message>
      The condition {1} of transition {0} is not of type BOOL.
      </message>
25      </constraint>
      ...
    </constraints>
  </constraintProvider>
</extension>
```

4.3.2.4 XML Mappings

In the sections above, we have discussed how to define a Petri net type and all its concepts and constraints. For saving it in PNML, it is also necessary to define how these concepts are represented in XML – at least if the “standard mappings” do not work.

In this section, we will discuss how these mappings are defined. Conceptually, these mappings are tables (in ISO/IEC 15909-2, these tables are given in Clause 7.3.1). In the ePNK, these “tables are programmed” as part of the new Petri net type¹⁵.

We explain the concepts of these “programmed tables” by discussing some of the mappings for HLPNGs. The tables for a new Petri net type are programmed, by overwriting the method

`registerExtendedPNMLMetaData(ExtendedPNMLMetaData metadata)` of `PetriNetType`; the parameter `metadata` represents the table, to which we can add entries when this method is called.

Let us have a look at some examples. Listing 4.21 shows an excerpt of the `registerExtendedPNMLMetaData()` method of the class `HLPNGImpl`. Each of the `metadata.add` statements defines one table entry, which defines the mapping of one specific feature of the ecore model to an XML element (we will see later how to map an ecore attribute to an XML attribute). The three statements shown in Listing 4.21 define how the structure feature of the labels `Type`, the `HLMarking`, and the `Condition` are mapped to the XML element `<structure>`. We discuss the first one, the `Type`, in more detail:

- The first parameter, denotes the feature that is mapped, which is the composition from the class `Type` to the class `Sort` (see Fig. 4.3). The source and target classes are mentioned explicitly as second and third parameter again. We refer to the feature and the two classes via singleton classes that describes the elements of the packages (`HLPNGdefinition` and `Terms`). These *package classes*, provide access to all the classes and features within a package (see [2] for more details). Note that `HlpngdefinitionPackage.eINSTANCE` refers to the package `hlpngdefinition` and `TermsPackage.eINSTANCE` to the package `terms`.
- As mentioned above, the second parameter denotes the class to which the feature belongs (it could be a sub-class of `Type` in principle); this

¹⁵It might be, that a future version will provide a means to plugin these tables directly in some form; but since “programming the tables” is not too difficult, this does not have a high priority.

Listing 4.21: Mappings for type, marking, and condition extensions

```
1 public void registerExtendedPNMLMetaData(  
    ExtendedPNMLMetaData metadata) {  
    ...  
  
    metadata.add(  
6      HlpngdefinitionPackage.eINSTANCE.getType_Structure(),  
      HlpngdefinitionPackage.eINSTANCE.getType(),  
      TermsPackage.eINSTANCE.getSort(),  
      "structure",  
      null,  
11     HLPNGFactory.getHLPNGFactory());  
  
    metadata.add(  
      HlpngdefinitionPackage.eINSTANCE.getHLMarking_Structure(),  
      HlpngdefinitionPackage.eINSTANCE.getHLMarking(),  
16     TermsPackage.eINSTANCE.getTerm(),  
      "structure",  
      null,  
      HLPNGFactory.getHLPNGFactory());  
  
21    metadata.add(  
      HlpngdefinitionPackage.eINSTANCE.getCondition_Structure(),  
      HlpngdefinitionPackage.eINSTANCE.getCondition(),  
      TermsPackage.eINSTANCE.getTerm(),  
      "structure",  
26     null,  
      HLPNGFactory.getHLPNGFactory());  
  
    ...  
}
```

is often called the *container* class.

- The third parameter denotes the class that the feature refers to (this could also be a sub class of `Sort`); this is often called the *object* class.
- The fourth parameter, defines the XML representation, the string that will be used as XML element in the serialisation of this feature (in our example “structure”).
- The fifth parameter could refer to a XML attribute, that might be necessary for creating an ecore object from the XML element (we will discuss an example later). In most cases, these XML attributes are not needed, since the XML element (and the context in which it occurs) provide enough information for creating the ecore element from it.
- The last parameter refers to a factory that is capable of creating an ecore instance of the respective class from the XML element and – if provided – the XML attribute.

You can see that these table entries can be used in two directions. In the one direction, it tells how to serialise a Petri net to its XML syntax; in the other direction, it tells how to create the model elements from the XML syntax. In the latter case, the factories play an important role. Listing 4.22 shows the interface that all these factories must implement. The methods `canCreateObject` and `createObject` have the same parameters, which basically reflect the entries of the table that we discussed above. Only the third and sixth one (the factory itself) are missing. And, there is an additional parameter (`provider`), which will provide access to the values of all attributes of the currently read XML element (in case the factory needs to read them for creating an the object). The method `canCreateObject` is used to find out whether the factory is able to create an object from the provided information, the `createObject` method is used to actually create it. The `createAttributeObject` is used to create an object for some XML attribute. The implementation of these factories is straight-forward and a bit boring – we do not discuss the details here. You can have a look into the class `HLPNGFactory` in package `org.pnml.tools.epnk.pntypes.hlpng.pntd.hlpngserialisation.factory` in the project `org.pnml.tools.epnk.pntypes.hlpng.pntd` to get some inspiration. What is more, with an extension that came into version 0.9.0 of the ePNK, the factory can be set to `null`. In which case the standard mechanism for creating an object of the target class will be used; therefore, we

Listing 4.22: Interface Factory

```
package org.pnml.tools.epnk.pnmlcoremodel.serialisation;

import org.eclipse.emf.ecore.EObject;
import org.eclipse.emf.ecore.EStructuralFeature;
5 public interface IPNMLFactory {

    public boolean canCreateObject(
        EStructuralFeature feature,
10     Object container,
        String element,
        String attribute,
        IAttributeProvider provider);

15 public EObject createObject(
        EStructuralFeature feature,
        Object container,
        String element,
        String attribute,
20     IAttributeProvider provider);

    public Object createAttributeObject(
        Object object,
        String attribute,
25     IAttributeProvider provider);

}
```

will need factories in very rare and special cases only. Typically, the factory can be set to `null`¹⁶

Listing 4.23: Mappings of an attribute

```

metadata.addAttributeMapping(
    BooleansPackage.eINSTANCE.getBooleanConstant_Value(),
3    BooleansPackage.eINSTANCE.getBooleanConstant(),
    "value",
    HLPNGFactory.getHLPNGFactory());

```

Listing 4.23 shows an example¹⁷ of how a feature of the model can be mapped to an XML attribute. In this example, the value of the boolean constant is mapped to the XML attribute `value`. This is where the method `createAttributeObject` of the factory comes into play.

The discussion above, gives a general idea of how these tables and mappings work. All this, however, could have been achieved with the existing mechanisms of EMF: Extended Metadata. But, the mappings of PNML have some specialities that could not be expressed in the Extended Metadata concepts. Therefore, we introduced yet another extension in the ePNK. In the rest of this section, we will discuss some of these special situations.

Let us consider the serialisation of the simple term $x \text{ ' } f(x, x)$, where x is a variable and f a user defined operator. The PNML representation is shown in Listing 4.24. In addition to being a bit verbose, there is one thing that is special about this mapping. There is a XML element `<subterm>` for the association from the top-level term (number of) to its subterm, which are represented as two other XML elements, `<variable>` and `<useroperator>`. The XML element `<subterm>` defines to which feature of the term the XML element that is contained in it should go. The XML element inside (e.g. `<variable>`) defines the type that object should have.

The problem here, is that there is an intermediate XML element that has no object as counter part in the model – it represents an association. We call them association elements. The mapping for these association elements is shown in Listing 4.25. The first entry is actually as we have seen it before. The only difference is that the factory produces an instance of a new class

¹⁶Note that except for two features, which were used to test this new mechanism, the mappings for HLPNGs have not been updated yet; therefore, you will find factories all over these mappings. But, this has historic reasons only and will eventually be changed (making the mappings more maintainable).

¹⁷Actually, this is the only one of HLPNGs.

Listing 4.24: PNML representation of $x'f(x,x)$

```

<numberof>
  <subterm>
3    <variable refvariable="5"/>
  </subterm>
  <subterm>
    <useroperator declaration="1">
      <subterm>
8        <variable refvariable="5"/>
      </subterm>
      <subterm>
        <variable refvariable="5"/>
      </subterm>
13    </useroperator>
  </subterm>
</numberof>

```

Listing 4.25: Mappings of associations to XML elements

```

metadata.add(TermsPackage.eINSTANCE.getOperator_Subterm(),
  TermsPackage.eINSTANCE.getOperator(),
  TermsPackage.eINSTANCE.getTerm(),
  "subterm",
5  null,
  HLPNGFactory.getHLPNGFactory());

metadata.add(null,
  HlpngserialisationPackage.eINSTANCE.getTermAssoc(),
10  TermsPackage.eINSTANCE.getVariable(),
  "variable",
  null,
  HLPNGFactory.getHLPNGFactory());

15 metadata.add(null,
  HlpngserialisationPackage.eINSTANCE.getTermAssoc(),
  TermsPackage.eINSTANCE.getUserOperator(),
  "useroperator",
  null,
20  HLPNGFactory.getHLPNGFactory());

```

TermAssoc, which has the nature of a term but, actually, represents an association to a term. We will discuss that class in more detail later. The two other mappings, define the mapping of variables and user operators to XML, and these are different, since they do not refer to any feature at all. They just refer to a container class and a contained class. The container class is the class **TermAssoc**, which will make sure that the variable resp. user operator will be added to the subterm feature of the operator on the level above¹⁸.

The class **TermAssoc** does not need to be programmed. This class, as well as the other classes for representing associations, could completely be generated from a model. This model is shown in Fig. 4.6. These classes extend a specific class of our model (the one to which the respective association should go), and the general class for **AssocClass**, which is defined by the ePNK, and implements all the necessary functionality. Note that these




Figure 4.6: The package `hlpngserialisation`

classes will not occur in the model anymore, once it is completely loaded – they are only used while a PNML file is loaded.

In the case of subterms, every subterm occurs in a separate `<subterm>` element – even if a term has several subterms, there is one subterm element for each of them (see Listing 4.24). In the case of parameters of an operation

¹⁸There would actually be another way of doing this, in a slightly more elegant way when using “standard features”, which will be discussed later in this section.

declaration, this is different: Listing 4.26 shows the PNML representation of the declaration of a named operator $f(x:\text{INT}, y:\text{INT}) = x * y$. Here,

Listing 4.26: PNML structure for declaration $f(x:\text{INT}, y:\text{INT}) = x * y$

```

<namedoperator id="1" name="f">
  <parameter>
    <variabledecl id="2" name="x">
      <integer/>
5    </variabledecl>
    <variabledecl id="3" name="y">
      <integer/>
      </variabledecl>
    </parameter>
10  <def>
    <mult>
      <subterm>
        <variable refvariable="2"/>
        </subterm>
15    <subterm>
        <variable refvariable="3"/>
        </subterm>
      </mult>
    </def>
20 </namedoperator>

```

all variable declarations occur in the same `<parameter>` element. We called these *bundled association elements*. The table entries for this mapping are shown in Listing 4.27. The first one, is almost the same as for association elements, and the Factory `HLPNGFactory` would create an instance of `VariableDeclAssoc` for an XML element `<parameter>`. The new last parameter `true` says, that this is a bundled association. The second table entry defines the mappings for variable entries, which is independent of the context, which is why the first to parameters are `null`. We call this a *context independent element mapping*.

This context independent element mapping can be applied in any other context. In combination with another special case of mappings which we call *standard feature*, this is a very powerful mechanism. For example, for `Declarations` and sub-element for which context element mapping exists (in the example, there would be variable declarations, sort declarations, and operator declarations), all these elements should be added to this standard feature. The table entry shown in Listing 4.28 defines the composition

Listing 4.27: Mapping bundled association elements

```

metadata.add(
    TermsPackage.eINSTANCE.getNamedOperator_Parameters(),
    TermsPackage.eINSTANCE.getNamedOperator(),
    TermsPackage.eINSTANCE.getVariableDecl(),
5   "parameter",
    null,
    HLPNGFactory.getHLPNGFactory(),
    true);

10 metadata.add(
    null,
    null,
    TermsPackage.eINSTANCE.getVariableDecl(),
    "variabledecl",
15   null,
    HLPNGFactory.getHLPNGFactory());

```

declaration as the standard feature of Declarations. Note that there

Listing 4.28: Defining a standard feature

```

metadata.add(TermsPackage.eINSTANCE.getDeclarations_Declaration(),
    TermsPackage.eINSTANCE.getDeclarations(),
    TermsPackage.eINSTANCE.getDeclaration(),
4   null,
    null,
    null);

```

is no mapping to XML here. A standard feature of an element just says that, whenever there comes some context independent element that is not mapped explicitly to a feature, this element should be added to the standard feature of the model. Of course, there should only be one standard feature – otherwise there would be some ambiguities.

4.4 Adding tool specific information

As discussed in Sect. 2.2.1, the PNML allows tool specific information to be added to all elements of Petri nets – indicated by a special XML element

`<toolspecific>`. The ePNK will read and write any tool specific information, and, in principle, its context could be accessed and modified via the class `AnyType`, which is defined in the plugin `org.eclipse.emf.ecore`. But this is tedious and, basically, means navigating in the elements XML structure.

Therefore, the ePNK provides an extension point to plugin tool specific extensions, so that they can be accessed and modified via an API specific to the extension which can be defined in terms of a model. We will discuss how to use this extension point by the help of an example: the token positions, which are part of the standard. We have seen an example already in Fig. 2.3 and Listing 2.1 on page 10.

This tool specific extension is defined in the project

`org.pnml.tools.epnk.toolspecific.tokenpositions`

Most of this code and the “plugin.xml” is automatically generated from the model “Tokenpositions.ecore” in the folder “model”. This model is shown in Fig. 4.7. The new classes are `PNMLToolInfo` and `Tokengraphics`. The




Figure 4.7: The model for tool specific extension tokenpositions

class `PNMLToolInfo` represents the actual tool specific information: it must implement the PNML core model interface `ToolInfo`. The actual contents of this tool specific information is `Tokengraphics`, which consists of one or many coordinates; the class `Coordinate` comes from the PNML core model.

From this model, the code can be generated in the same way as describes in Sect. 4.3.1.2. First, the “genmodel” must be created, and from the “genmodel”, the model code and the edit code must be generated.

After the code generation, the only thing left to do is to manually create a factory for this tool specific extension, and use this factory for plugging it

into the ePNK. The factory for our extension is shown in Listing 4.29. The

Listing 4.29: Factory for the tool specific extension

```
package org.pnml.tools.epnk.toolspecific.tokenpositions.factory;

import org.pnml.tools.epnk.pnmlcoremodel.ToolInfo;
4 import org.pnml.tools.epnk.toolspecific.extension.
    ToolspecificExtensionFactory;
import org.pnml.tools.epnk.toolspecific.tokenpositions.
    TokenpositionsFactory;

9
public class TokenpositionsExtensionFactory
    implements ToolspecificExtensionFactory {

    private final static String toolname = "org.pnml.tool";
14 private final static String toolversion = "1.0";

    public ToolInfo createToolInfo(String tool, String version) {
        // ToolInfo object does not depend on these values:
        return createToolInfo();
19    }

    public ToolInfo createToolInfo() {
        return TokenpositionsFactory.eINSTANCE.createPNMLToolInfo();
    }

24 public String getToolName() {
    return toolname;
    }

29 public String getToolVersion() {
    return toolversion;
    }

}
```

factory implements the ePNK interface `ToolspecificExtensionFactory`, which consists of four methods. The two methods `createToolInfo` create an instance of this tool specific extension; the method with the two `String` parameters, `tool` and `version` is used, when there is the tool name and a version given, which might return instances of different classes – in our

example, however, the version number is irrelevant. The two other methods, must return the tool name for that extension and its version, which, in our example, are encoded in constants.

Listing 4.30 shows the fragment of the “plugin.xml” that is needed to plug in the token position extensions to the ePNK. In addition to the name and the id, the attribute `class` defines to the factory for the tool specific extensions; this class must implement the interface `ToolspecificExtensionFactory`. Moreover, there is a brief description

Listing 4.30: Plugging in the token position extension

```

<extension
2  id="org.pnml.tools.epnk.toolspecific.tokenpositions"
  name="Token Positions"
  point="org.pnml.tools.epnk.toolspecific">
  <type
    class="org.pnml. ... .factory.TokenpositionsExtensionFactory"
7  description="The tool specific extension for token positions">
  </type>
</extension>

```

of this extension.

Note that there is no way of explicitly defining the XML syntax of these extensions yet. The standard XMI serialisation will be used. Eventually, there will be a mapping mechanism similar to the one for Petri net types.

4.5 Overview of the ePNK and its API

In this section, we give a brief overview of the different parts of the ePNK, its project structure and the API. As mentioned earlier, developers would probably not need to change anything in these projects, but the overview helps to better understand the ideas behind the ePNK, the necessary dependencies (that need to be included in new projects via the “plugin.xml”) and the functions that are available in the ePNK API, that could be used by developers in the extensions. Note that we do not really discuss the API for the code that is generated for the models (model code) since it is mostly straight-forward (and more an EMF issue).

Like all extension of eclipse, the ePNK is organized in many eclipse projects, which together make the ePNK. Most of these projects are so-called plugin projects, which are most relevant for developers, since these

are the projects to look up the API and to which extensions need to refer (in form of dependencies). In addition, there are some projects, which are only containing documentation (like this manual), there are some projects that do not contain any code, but from which other projects are generated, and there are so-called features, which define collections of plugin projects in order to deploy them. And there is a project for generating the ePNK update site from the features.

In this manual, we focus on the plugin projects of the ePNK, the most important of which are listed below:

`org.pnml.tools.epnk` This is the core project of the ePNK. In this project, you will find the PNML core model, some additional models, and the *model code*, that was generated from them. All the models, can be found in the folder `model`. The PNML core model is contained in `PNMLCoreModel.ecore`; in order to avoid clutter in the graphical diagrams, the core model is actually split up into three parts: `PNMLCoreModel.ecorediag` contains the most important concepts; `PNMLCoreModelGraphics.ecorediag` contains the graphical features of PNML; and `PNMLCoreModelProxies.ecorediag` contains some extensions to the PNML core model that are necessary to maintain labels in the graphical editor of the ePNK by so-called *label proxies* and *page label proxies*. These elements, however, are not of any concern for normal developers.

There are three other models in this project: `PNMLDataTypes.ecore` defines the data types for non-negative and positive numbers, which are used instead of the respective `XMLSchemaDataTypes` here. `PNMLStructuredPNTTypeModel.ecore` defines the concepts of structured Petri net types (see Sect. 4.3.2.2). `Serialisation.ecore` provides some general structure that will be needed for the XML serialisation of so-called association elements (see Sect. 4.3.2.4).

In this project, there is also one convenience class `FlatAccess` in package `org.pmm.tools.epnk.helpers`, which might turn out to be helpful in practice. This allows to handle a Petri net that is distributed over several pages as if it was flat (see Sect. 4.2.3 for an example).

In this plugin, also the two extension points of the ePNK are defined: one for defining new Petri net types (PNTD), another for defining new toolspecific extensions.

In addition to the model code, also the so-called edit code and editor code are generated from these models, which together define the tree

editor for PNML (see below).

Note that, in this project, also the constraint context and the constraint category `org.pnml.tools.epnk.validation`, to which all other constraints for new Petri net types should be added, are defined here (see Sect. 4.3.1.4 and 4.3.2.3).

`org.pnml.tools.epnk.edit` This project contains the edit code that was generated from the models in project `org.pnml.tools.epnk`. Though most of this code was automatically generated from the models, there are several manual changes, that enable generically dealing with plugged in Petri net type definitions.

Moreover, the generated standard EMF images in the folder `icons` were replaced by nicer ones.

`org.pnml.tools.epnk.editor` This project contains the editor code for the EMF tree editor for PNML that was generated from the models in project `org.pnml.tools.epnk`. In this project, there are only a few, but crucial extensions, that made it possible to integrate the EMF tree editor with the graphical editor for pages.

`org.pnml.tools.epnk.pntypes` This project contains the model and the generated model code for P/T-nets (PTNet), as well as the extension that plugs in this type to the ePNK. The model code is completely generated from the models, except for two changes in class `PTNetImpl` as discussed in Sect. 4.3.1.

`org.pnml.tools.epnk.pntypes.edit` This is the edit project that was generated from the model `PTnet.ecore` of project `org.pnml.tools.epnk.pntypes`. There are no manual changes in the generated code; – only the icons in the folder `icons` were replaced by nicer ones.

`org.pnml.tools.epnk.toolspecific.tokenpositions` In this project, the tool specific extension for token positions (as defined in ISO/IEC 15909-2) is defined (see Sect. 4.4). From the model `Tokenposition.ecore`, the model code and the edit project `org.pnml.tools.epnk.toolspecific.tokenpositions.edit` was generated (which does not contain any manual changes – not even nicer icons).

`org.pnml.tools.epnk.actions` This project defines the standard actions of the ePNK, which are the popup menus for adding missing ids and for linking the labels of structured Petri net types (see Sect. 4.3.2.2).

Moreover, the classes `AbstractEPNKAction` and `AbstractEPNKJob` are defined in this project, which are convenience classes to make it easier to define actions for the ePNK that run in the background (see Sect. 4.2.3).

`org.pnml.tools.epnk.functions.tutorials` This project contain the two functions that were discussed in Sect. 4.2.1 and Sect. 4.2.2, which can serve as a guideline for defining own extension projects.

`org.pnml.tools.epnk.functions.modelchecking` This project contains the model checker extension for the ePNK that is discussed in Sect. 4.2.3. Note that the project `MCiE` is used in this model checker extension only (`MCiE` is not relevant for anything else in the ePNK – except if you want to implement you own model checker based on `MCiE`).

`org.pnml.tools.epnk.diagram` Contains the code for the GMF-generated graphical editor for pages of Petri nets. This code was generated from the GMF models in project `org.pnml.tools.epnk.gmf`. But, there are major manual changes for making this graphical editor generic and for integrating it with the tree editor.

`org.pnml.tools.epnk.gmf.integration` This project defines the popup menus for starting the graphical editor on a page that is selected in a tree editor or in a graphical editor.

The plugin projects with the prefix `org.pnml.tools.epnk.pntypes.hlpng` resp. `org.pnml.tools.epnk.pntypes.hlpngs` together are used for defining high-level Petri nets. The following list gives an overview in a bottom up way – ending with the actual Petri net type definition for HLPNGs.

`org.pnml.tools.epnk.pntypes.hlpngs.datatypes` In this project, all the models that define the concepts of sorts, operators, variables and terms for HLPNGs are contained. In particular, there is a model `HLPNGDataTypes.ecore` for the general structure of terms, declarations and the built-in sorts and operators that occur in all versions of HLPNGs. And there are many more models and diagrams for specific versions of HLPNGs (Clauses 5.3.2–5.3.12 of ISO/IEC 15909).

`org.pnml.tools.epnk.pntypes.hlpngs.datatypes.concretesyntax`

This project is an Xtext project that defines the grammar for the concrete syntax for the different labels of HLPNGs, from which a parser is generated. Here, we do not discuss the details of generating the parser. The manually written class `HLPNGParser` accesses the automatically generated parsing operations and provides methods for parsing every kind of label of HLPNGs. The other important manually written class is `HLPNGLinker`, which provides the global Linker for labels (as discussed in Sect. 4.3.2.2).

Note that the packages `... .concretesyntax.generator` and `... .concretesyntax.ui` were automatically created when Xtext projects are generated. They are among other things relevant during the code generation for the parser. They are not relevant for the use of the ePNK.

`org.pnml.tools.epnk.pntypes.hlpng.pntd` This project actually combines all the parts discussed above into a Petri net type definition for HLPNGs. The main model is `HLPNGDefinition.ecore`, which was shown in Fig. 4.4 already. The other model `HLPNGSerialisation.ecore` defines the auxiliary classes that temporarily store XML elements that represent associations (see Sect. 4.3.2.4) and are used and referred to in the XML mappings.

The global linker from project `HLPNGLinker` is made available, by manually implementing the method `getLinker` in class `HLPNGImpl`. The `parse` method for the different structured labels are also manually implemented – they refer to the different `parserMethods` of class `HLPNGParser`.

4.6 Deploying extensions

In this section, we will briefly discuss how own extensions of the ePNK could be deployed, so that others can use it. Typically, an extension will comprise several plug in, which would be combined into an eclipse *feature*, which are a special kind of eclipse project. This way, the extension comes deployable.

A feature in turn, can be used in an eclipse *update project* which can be used to create your own *update site*, so that your plugin could be installed from this site, similar to the way you installed the ePNK.

Since features and update site are standard eclipse concepts, however, we do not explain the details here. For now, looking up the keywords “fea-

ture” and “update site” in the eclipse help (or googling for them) should be enough.

If you have a feature that might be interesting for a wider audience, you can also contact us, so that we can make it available via the ePNK update site.

4.7 Future plans

In this section, we give an overview of some extensions that are planned for the ePNK, which will increase its flexibility, make it easier to use, and provide additional extension mechanisms. The order indicates the priorities, and might be roughly the order in which the features are implemented:

- Full support all graphical features of PNML in the graphical editor (maybe except bezier curves for arcs, which needed to be implemented manually, since GMF does not support them).
- Add a toolspecific feature that exactly saves the graphics of a Petri net as it appears in the graphical ePNK editor (even if no PNML representation exists).
- Smoother integration of EMF-tree editor with the GMF-editor.
- Add the attributes mechanism of PNML together with a mechanism to define the graphical appearance of net objects.
- Implement the serialisation of the structural labels of HLPNGs to concrete syntax.
- Define an extension point for defining the concrete syntax of HLPNGs (and possibly other more complex Petri net types).
- Portation of the code-generation for extended P/T-Nets to the ePNK (and define the extended Petri net type for that).
- Implement an explicit ePNK extension point for functionality, in order to unify the access to ePNK functions.
- Define an extension point for the XML mapping that directly takes a table instead of “programming the table”.
- Long-term: Support modular PNML.

Chapter 5

PNML: Observations and suggestions

During the work on the ePNK, several issues came up concerning PNML. In this chapter, these issues will be summarized and some suggestions on how to improve future versions of PNML will be made.

5.1 Toolspecific extensions: Type attribute

The toolspecific extensions should probably be equipped with an optional attribute type: The rationale is that the same tool could have different tool-specific extensions. With the additional attribute, we would not force the tool to mix all its extensions in a single class (or define the type in an extension of the tool name, which would be also a mess).

5.2 Relaxing requirement for ids everywhere

PNML requires an id for every element in the document. In the ePNK, I have made the ids optional, when they are not needed for referring to an element. An id is only required, if there is a reference to such an element (and actually the document can still be serialized using XPath expressions without ids at all, which however is not compatible with PNML);

For future versions of PNML, it should be considered to make ids optional in case elements cannot be referenced at all or when elements are not referenced from other elements in the net.

5.3 Naming in PNML standard

The feature `variableDecl` in the UML model of `Variable` in the `terms` package is later mapped to `refvariable` in XML. This could actually be the same name in both cases (the ePNK model used “`refvariable`” in the model). For the definition of PNML as an exchange format, this does not make any difference – but, it is smoother conceptually, technically, and from a readability point of view, since we do not need an extra mapping.

In the `terms` package, there is a class `BuiltInConst`. Later (e.g. in packages `dots`, `booleans`, `integers`) it is referred to as `BuiltInConstants`. This needs to be aligned. In ePNK, `BuiltInConst` is used in all cases.

In package `booleans`, `Equality` and `Inequality` are directly derived from `Operators`. Though, it does not make a difference for PNML as a transfer format, this is probably an error. In the ePNK model, `Equality` and `Inequality` are derived from `BuiltInOperators`, which makes explicit that these are built-in operators in the model.

In ISO/IEC 15909-2, Table 8 (p. 35) the declaration feature of a variable is mapped to the attribute `variabledecl`; in the RELAX NG grammar, however, this is represented as “`refvariable`” (Annex B.2.1, p. 54) and also in the example of `Symmetric Nets` (Annex C, e.g. p. 93). This is probably a type and should be aligned! In the ePNK “`refvariable`” was used.

Also in Table 8, the `Integer::Addition` class is mapped to “`add`”. But, in the RELAX/NG grammars this is called “`addition`”. The ePNK uses “`addition`”, and “`subtraction`” to be consistent with the RELAX/NG grammar. But “`addition`” and “`subtraction`” is not consistent with the naming of “`div`”, “`mod`”, and “`mult`”. This could be aligned.

5.4 Mapping of models to XML

The sorts `ProductSort` and `MultisetSort` (and their children) are not mapped to XML in the same style as `Terms` (there is no element for the feature (`elementSort` or `basis`), but just an element for the sort itself. That is a bit awkward and we might want to align this.

The same applies for `NamedSorts`.

In order to map these features of PNML (resp. HLPNG) to XML, the ePNK introduced the concept of *standard features*.

5.5 Explicit result type for NamedOperators

A NamedOperator should have an explicit sort for its return value. Right now, we rely on that the definition of NamedOperators is not cyclic – which is explicitly required. But, if this is not checked, other calculations might run into infinite loops, if not done properly. This could be done more easily if the return type of an operator was made explicit.

5.6 Opposite references for source and target

Maybe, the source and target references of arcs in the PNML Core Model should have opposites, so that the arcs can be accessed from the nodes more easily (they should be transient though, since they do not need to be serialized). I will make this change in the ePNK anyway (technically, this has nothing to do with ISO/IEC 15909-2, since it does not mandate the PNML Core Model, but only the XML format – but, it might be relevant for Part 3).

5.7 Name labels for arcs

I still do not see any reason, why arcs of a Petri net should have a name.

5.8 Remove the possibility of net labels

We should consider to remove the labels that are directly contained in a net (and not on page) from the PNML Core Model; this does not make much sense conceptually and makes things more difficult to implement (see Sect. 3.6.6).

5.9 Attributes with default values

The RELAX/NG grammars of ISO/IEC 15909-2 require that all attributes of all elements must of must always be there. Many other formats allow omitting attributes with default values. We should discuss, whether ISO/IEC 15909 should also allow to omit attributes that have a default value (and define what the default values are).

In the current implementation, the ePNK serializes all attributes as defined in ISO/IEC 15909-2, but it is no problem changing that (and for all

attributes, the ePNK defines some default values in the models – since, strangely enough, this was the only way to enforce their serialisation in the EMF technology).

5.10 Recursive sort and operator definitions

Up to now, PNML does not allow defining recursive sorts and has limited expressibility for defining recursive functions. There could be more constructs for that; but this would require to have a more careful definition of this (right now the semantics of sorts is to flatten them, which would not work for recursive types anymore).

We also might want unions for defining recursive sorts in an easy way. But, this would impose a great burden on all implementations of PNML even if they do not use it – therefore, we should be very careful with that (it could be part of a even more general kind of HLPNGs).

5.11 Names for `ArbitrayDeclarations`

As mentioned in Sect. 3.6.6 already, `Unparsed` declarations of ISO/IEC do not have a name attribute. Since, any declaration is a symbol definition, of some kind, it would be reasonable, if this would be aligned. Therefore, `Unparsed` should have a name attribute of type `String`.

5.12 Incorrect types in the HLPNG models

In ISO/IEC 15909-2, some classes in the packages `FiniteIntRanges`, `Strings`, and `Lists`, have attributes that are technically incorrect. They are on the wrong meta-level: they refer to a PNML/HLPNG type instead of a data type in UML. This should be changed

Here is a list of the required changes:

- `FiniteIntRange`: `start` and `end` should refer to `int` (`EInt`)
- `FiniteIntRangeConstant`: `value` should refer to `int` (`EInt`)
- `Substring`: `start` and `length` should refer to `NonNegativeInteger` (is defined as a separate basic datatype in the ePNK)
- `Sublist`: `start` and `length` should refer to `NonNegativeInteger`

5.13. REFERENCE TO SORT IN FINITEINTRANGECONSTANTS 101

- **MemberAtIndex**: `index` should refer to `NonNegativeInteger` (this attribute is capitalized in the wrong way in Fig. 19, which should also be changed)

Moreover, the type `PrimitiveType::Integer` that is used in the `value` feature of `NumberConstant` in the `Integers` package is not defined within ISO/IEC 15909-2. This should be changed (in the ePNK the type `EInt` is used).

5.13 Reference to Sort in FiniteIntRangeConstants

The `range` feature of `FiniteIntRangeConstant` refers to `FiniteIntRange`. This makes the concrete syntax for these constants (and the implementation of a parser for this concrete syntax) very inconvenient. It would be much easier, if this `range` feature referred to `Sort` in general, and there was an additional constraint that this sort is (or refers to) a `FiniteIntRange`.

5.14 Omitting redundant opposit references

In the package for `Partitions`, there are the features `refpartition` and `refpartitionelement`, which are redandant and are not serialised. It would be better to delete them form these models (in the API generatded by EMF, the respective elements can be accessed via the `eContainer`, if necessary).

5.15 Useless operator: Append

The `Append` operator in package `Strings`, does not make any sense, if there is no data type for characters. So, either a data type for characters should be introduced, or `Append` should be removed.

Note: In the ePNK, the `Append` it is included for compatibility reasons; but terms using it, will never be typed correctly and cannot be resolved. If this operator is used, there will be validation errors, but the respective nets can be serialised.

5.16 Incorrect constraints for List Append

The constraint for the `Append` operator in `Lists` is incorrect (and not complete): It requires both arguments to be lists; but if this was meant seriously, the `Append` operator would not be different from the list concatenation. One

of the arguments of the **Append** should be of the basis sort of the list (representing a single element that is appended to the list). In the ePNK, the first argument is the of the basis sort for now.

5.17 Implicit and constant parameteres

The operators **Substring** of package **Strings** and **MemberAtIndex** and **Sublist** of package **Lists** have constant implicit parameters for the start, length, resp. index. These should rather be subterms of the respective types (expressed as constraints); maybe this was the intention with the wrong types (see 5.12), but the RELAX/NG grammar clearly states that these are not subterms.

In the ePNK, this is implemented according to ISO/IEC 15909-2, but in a future version, we will (optionally) allow resp. subterms; having only constants here, is not too useful.

5.18 Mismatch between tables and grammars

In the mapping of the PNML core model and type specific models to XML in Clause 7 (Table 8) of ISO/IEC 15909-2, the **FiniteIntRangeConstant** has a PNML attribute **range** which is an IDREF. This does not correspond to the RELAX/NG grammar and the UML model. This should rather be a (default) sub-element of type **Sort** (this is what was implemented in the ePNK and is suggested by the RELAX/NG gammar). In ISO/IEC 15909-2, the entry "range:IDREF" in Table 8 for **FiniteIntRangeConstant** should be deleted.

Bibliography

- [1] J. Billington, S. Christensen, K. van Hee, E. Kindler, O. Kummer, L. Petrucci, R. Post, C. Stehno, and M. Weber. The Petri Net Markup Language: Concepts, technology, and tools. In W. van der Aalst and E. Best, editors, *Application and Theory of Petri Nets 2003, 24th International Conference*, volume 2679 of *LNCS*, pages 483–505. Springer, June 2003.
- [2] F. Budinsky, D. Steinberg, E. Merks, R. Ellersick, and T. J. Grose. *Eclipse Modeling Framework*. The Eclipse Series. Addison-Wesley, 2nd edition edition, Apr. 2006.
- [3] E. Clarke, O. Grumberg, and D. Peled. *Model checking*. MIT Press, 1999.
- [4] E. Clayberg and D. Rubel. *Eclipse Plug-ins*. The Eclipse Series. Addison-Wesley, 3rd edition edition, 2008.
- [5] L. Hillah, E. Kindler, F. Kordon, L. Petrucci, and N. Treves. A primer on the Petri Net Markup Language and ISO/IEC 15909-2. In K. Jensen, editor, *10th Workshop on Coloured Petri Nets (CPN 09)*, pages 101–120, Oct. 2009.
- [6] L.-M. Hillah, F. Kordon, L. Petrucci, and N. Trèves. PNML Framework: An extendable reference implementation of the petri net markup language. In J. Lilius and W. Penczek, editors, *Petri Nets*, volume 6128 of *Lecture Notes in Computer Science*, pages 318–327. Springer, 2010.
- [7] ISO/IEC. Systems and software engineering – High-level Petri nets, Part 2: Transfer Format. International Standard ISO/IEC 15909-2:2011, Feb. 2011.
- [8] M. Jünger, E. Kindler, and M. Weber. The Petri Net Markup Language. *Petri Net Newsletter*, 59:24–29, Oct. 2000.

- [9] M. Jünger, E. Kindler, and M. Weber. Towards a generic interchange format for Petri nets – position paper. In R. Bastide, J. Billington, E. Kindler, F. Kordon, and K. H. Mortensen, editors, *Meeting on XML/SGML based Interchange Formats for Petri Nets*, pages 1–5, June 2000.
- [10] E. Kindler. Der Petrinetz-Kern: Ein Traum wird wahr. In H. Ehrig, W. Reisig, and H. Weber, editors, *Move-On-Workshop der DFG-Forschergruppe Petrinetz-Technologie*, pages 121–124. Technische Universität Berlin, 1997.
- [11] E. Kindler. The Petri Net Markup Language and ISO/IEC 15909-2: Concepts, status, and future directions. In E. Schnieder, editor, *Entwurf komplexer Automatisierungssysteme, 9. Fachtagung*, pages 35–55, May 2006. invited paper.
- [12] E. Kindler and J. Desel. Der Traum von einem universellen Petrinetz-Werkzeug — Der Petrinetz-Kern. In J. Desel, A. Oberweis, and E. Kindler, editors, *3. Workshop Algorithmen und Werkzeuge für Petrinetze*, number 341 in Forschungsberichte. Institut AIFB, Universität Karlsruhe, Oct. 1996.
- [13] E. Kindler and M. Weber. The Petri Net Kernel – an infrastructure for building Petri net tools. *Software Tools for Technology Transfer*, 3(4):486–497, July 2001.
- [14] The Eclipse Foundation. The eclipse platform. <http://www.eclipse.org>.
- [15] M. Weber and E. Kindler. The Petri Net Markup Language. In H. Ehrig, W. Reisig, G. Rozenberg, and H. Weber, editors, *Petri Net Technologies for Modeling Communication Based Systems*, volume 2472 of *LNCS*, pages 124–144. Springer, 2003.