



Chapitre d'actes

2010

Open Access

This version of the publication is provided by the author(s) and made available in accordance with the copyright holder(s).

AIPiNA: an Algebraic Petri Net Analyzer

Buchs, Didier; Hostettler, Steve Patrick; Marechal Marin, Alexis Ayar; Risoldi, Matteo

How to cite

BUCHS, Didier et al. AIPiNA: an Algebraic Petri Net Analyzer. In: Tools and Algorithms for the Construction and Analysis of Systems - TACAS 2010, 16th International Conference. Javier Esparza and Rupak Majumdar (Ed.). Paphos (Cyprus). [s.l.] : Springer, 2010. p. 349–352. (Lecture Notes in Computer Science (LNCS)) doi: 10.1007/978-3-642-12002-2_29

This publication URL: <https://archive-ouverte.unige.ch/unige:5622>

Publication DOI: [10.1007/978-3-642-12002-2_29](https://doi.org/10.1007/978-3-642-12002-2_29)

AlPiNA: An Algebraic Petri Net Analyzer[★]

Didier Buchs, Steve Hostettler, Alexis Marechal and Matteo Risoldi

Software Modeling and Verification laboratory
University of Geneva, Route de Drize 7, CH-1227 Carouge, Switzerland,
<http://smv.unige.ch>

Abstract. *AlPiNA* is a graphical editor and model checker for a class of high-level Petri nets called *Algebraic Petri Nets*. Its main purpose is to perform reachability checks on complex models. It performs symbolic model checking based on ΣDD , an efficient evolution in the Decision Diagrams field, using novel techniques such as *algebraic clustering* and *algebraic unfolding*. AlPiNA offers a user-friendly interface, and is easily extensible.

1 Introduction

This article introduces the AlPiNA model checking tool. AlPiNA allows checking reachability properties on Algebraic Petri Nets (APN) models, a class of High Level Petri Nets. It encodes state spaces symbolically as Decision Diagrams [6], which reduces memory consumption and computation time that are major obstacles to the practical use of model checking. Users can specify properties to verify using a dedicated language, and they can provide additional information on the model to improve model checking performance. In the current iteration of AlPiNA, we focus on reachability properties for several reasons – among others, the fact that many interesting properties can be expressed as reachability properties as proven in CPN Tools [7].

AlPiNA has two main goals. The first goal is improving model checking performance by leveraging the Decision Diagrams framework and the innovative concepts of *algebraic clustering* and *algebraic net unfolding*. Algebraic clustering reduces the memory footprint of state space calculation by semi-automatically decomposing the system in independent processes. Partial algebraic net unfolding allows reducing the complexity of the data type unfolding. The second goal of AlPiNA is coupling this high performance with a user friendly interface. The user can specify models and properties with a graphical and textual editor. We propose to separate the model and performance-related information. This gives the users a high-level view of the model, freeing them from the need to use low-level formalisms in a complex way.

The article is structured as follows. Section 2 quickly illustrates the theoretical foundations of AlPiNA. Section 3 describes the tool’s architecture and shows some benchmarks. Finally, the tool’s current status and perspectives are discussed.

[★] This project was partially funded by the COMEDIA project of the Hasler foundation, ManCom initiative project number 2107.

2 Theoretical Foundations of AIPiNA

Concurrency and non-determinism are the major causes of exponential state space explosion [12]. This happens when model components have few causal dependencies with each other and therefore evolve almost independently. Because of the exponential nature of the model checking problem, the state space rapidly becomes intractable as the number of components increases. To overcome this, the state space encoding must have a lower complexity than the explicit enumeration of states. We extend the approach initiated by McMillan [4] called Symbolic Model Checking, which exploits maximal sharing of state elements. In APNs, values are instances of algebraic abstract data types (ADT), therefore they require a more powerful encoding of the state space than Binary Decision Diagrams [1]. Because of this, we defined an evolution of Decision Diagrams (DD) [6] called Σ DD [3].

Clusters (i.e. sets of states) maximize the sharing induced by encoding with DDs [6]. For example, all the places of a Petri net that represent a process and its resources are grouped together. In this case, the cluster is called a *topological cluster* [9] since it is solely based on the Petri net topology. In high-level Petri nets, because of the level of abstraction, places can represent *classes* of similar processes and resources. *Algebraic clusters* [2] allow the user to group process instances with their resources. AIPiNA automatically derives clustering from this grouping. The more independent the resulting clusters are, the more efficient the symbolic representation will be. In the best case, the memory consumption is logarithmic to the number of states.

Since AIPiNA uses APNs, it has to manipulate universally quantified variables. An interesting way of improving performance is to perform an *algebraic net unfolding* [2]. It instantiates the variables of the system in a pre-processing phase, before state space exploration. By doing this, it becomes possible to compile the model with bindings that satisfy the transition guards. Unfolding may significantly increase the speed of the state space construction when the data domains are finite or bounded. Still, it is not always possible or even desirable to perform unfolding for two reasons. The first reason is that a bound may be difficult to figure out: if the bound is too small, the validation becomes incorrect; if it is too large, unfolding may become very expensive and model checking itself intractable. The second reason is that sometimes it is useless to unfold a data domain if only a few of its values are effectively used.

To tackle this problem, we propose to perform *partial unfolding*, i.e. choosing only a subset of the domains. The choice whether a domain should be part of the unfolding is a trade-off between the possible speed gain and the cost of the unfolding itself. Its computational complexity is $O(n^c)$ where n is the size of the largest data domain and c the largest number of input arcs.

In AIPiNA we generate the state space using an algorithm called saturation [5]. The algorithm benefits from the clustering of the state space to *fire* all transitions local to a component before firing inter-component transitions. All the transitions local to a given cluster are only applied to the subset of the state space relevant to the cluster, avoiding superfluous computations. A detailed technical description of the encoding as well as the notion of *algebraic cluster* and *algebraic unfolding* has been given in [2].

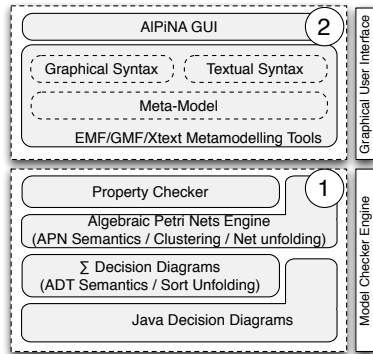


Fig. 1. Architecture Overview

		AIPiNA				Maria		Helena	
		Partial Unfold.		Total Unfold.					
Model Size	States #	Mem (MB)	Time (s)	Mem (MB)	Time (s)	Mem (MB)	Time (s)	Mem (MB)	Time (s)
Distributed Database									
10	197E3	10	0.8	12.4	1.3	47	44.3	24	9
15	7.2E7	33	2.6	41	5.8	-	-	1.4E3	7.5E3
35	5.8E17	544	69.4	789	278	-	-	-	-
Dining Philosophers									
10	186E4			1.9	0.15	375	141	11	5
15	2.5E9			2.6	0.18	-	-	409	822
300	1.2E188			162	48.5	-	-	-	-
Slotted Ring									
5	53856			4.9	0.2	23	4.3	10	5
10	8.3E9			55.6	1.7	-	-	-	-
15	1.5E15			330	9.8	-	-	-	-
Leader Election									
15	399E4			27.7	1.4	795	361	107	142
50	1.7E21			702	76	-	-	-	-

Fig. 2. State space generation

3 Tool description

AIPiNA’s architecture can be seen in Fig. 1. AIPiNA is composed of a *Model Checker Engine* (1) and a *Graphical User Interface* (2) built on top of it.

AIPiNA’s architecture can be seen in Fig. 1. AIPiNA is composed of the *Model Checker Engine* (1) and a *Graphical User Interface* (2) built on top of it.

The foundation of the *Model Checker Engine* is the symbolic representation offered by DD structures, as presented in the previous section. The first two layers of our engine refer to libraries that handle DD structures. The third layer is a bridge between the APN semantics and the underlying layers. It performs optimisations such as algebraic clustering and net unfolding. On top of the engine block, we find the property checker layer, that uses the state space generated by the previous layer to compute the properties satisfaction. These two layers communicate with the GUI block, they receive the models and return the generated state space and properties satisfaction results.

The second block of the AIPiNA architecture is the GUI. We used the Eclipse Tools from the Eclipse Modeling Project (EMP) [8] to create a user friendly interface, following the MDA directives. The first layer is the metamodels specifications, created with EMF. With these metamodels, we created a graphical concrete syntax using GMF for the Petri Nets editor, and a textual concrete syntax using XText for the textual editors. This schema allows us to create an extensible and modular tool.

AIPiNA has good memory consumption and processing time as shown in Fig. 2. It outperforms by an order of magnitude two widely used high level Petri nets model checkers, Maria [10] and Helena [11]. This figure shows the results obtained for some well known examples in the model checking field. The “-” symbol indicates that a result could not be computed¹. Every example shows that the techniques we present in this tool can produce excellent results when applicable. The distributed database example shows also that *partial net unfolding* can give better results than total unfolding. The blank cells indicate that the test has not been run. Indeed, *partial net unfolding* is not

¹ These benchmarks were computed using a 4 GB ram, 2.5 GHz Core 2 Duo Macbook Pro. The source code can be downloaded at <http://alpina.unige.ch>

useful when the algebra are to small, which is the case for the *Dining Philosophers*, the *Slotted Ring* and the *Leader election*.

4 Current Status and Perspectives

Compared to other high-level model checkers, ALPiNA has the advantage of treating state spaces larger by orders of magnitude while being user friendly. Users benefit from the efficiency based on the Decision Diagrams technology in a transparent manner. They can also easily specify *algebraic clustering* and *algebraic net unfolding* to improve model checking performance. Thanks to this, ALPiNA outperforms Maria and Helena when the model has strong concurrency.

All the features mentioned in this paper have been implemented in ALPiNA. A public release can be found at <http://alpina.unige.ch>. The tool has a user-friendly interface, taking full advantage of the EMF tools features. We are currently working on the next version which should bring modularity to the formalism and CTL support. Moreover, we will improve user guidance while defining the algebraic clustering.

References

1. R. E. Bryant. Graph-based algorithms for boolean function manipulation. In *Transactions on Computers, C-35*, pages 677–691. IEEE, 1986.
2. D. Buchs and S. Hostettler. Managing complexity in model checking with decision diagrams for algebraic petri net. In D. Moldt, editor, *Pre-proceedings of the International Workshop on Petri Nets and Software Engineering*, pages 255–271, 2009. Available at <http://smv.unige.ch/publications/pdfs/pnse09.pdf>.
3. D. Buchs and S. Hostettler. Sigma Decision Diagrams: Toward efficient rewriting of sets of terms. In A. Corradini, editor, *TERMGRAPH 2009: Preliminary proceedings of the 5th International Workshop on Computing with Terms and Graphs*, number TR-09-05 in TERMGRAPH workshops, pages 18–32. Università di Pisa, 2009. Available at <http://smv.unige.ch/publications/pdfs/termgraph09.pdf>.
4. J. R. Burch, E. M. Clarke, K. L. McMillan, D. L. Dill, and L. J. Hwang. Symbolic model checking: 10^{20} states and beyond. *Inf. Comput.*, 98(2):142–170, 1992.
5. G. Ciardo, G. Lüttgen, and R. Siminiceanu. Efficient symbolic state-space construction for asynchronous systems. In *Application and Theory of Petri Nets 2000 (Proc. 21th Int. Conf. on Applications and Theory of Petri Nets, Aarhus, Denmark)*, *Lecture Notes in Computer Science 1825*, pages 103–122. Springer-Verlag, 2000.
6. J.-M. Couvreur and Y. Thierry-Mieg. Hierarchical decision diagrams to exploit model structure. In *FORTE*, pages 443–457, 2005.
7. CPN Group. CPN tools. <http://wiki.daimi.au.dk/cpntools/cpntools.wiki>.
8. Eclipse. Eclipse modeling project. <http://www.eclipse.org/modeling/>.
9. A. Hamez, Y. Thierry-Mieg, and F. Kordon. Hierarchical set decision diagrams and automatic saturation. In *Petri Nets*, pages 211–230, 2008.
10. M. Mäkelä. Modular reachability analyzer. <http://www.tcs.hut.fi/Software/maria/>.
11. C. Pajault and S. Evangelista. High level net analyzer. <http://helena.cnam.fr/>.
12. A. Valmari. The state explosion problem. In *Lectures on Petri Nets I: Basic Models, Advances in Petri Nets, the volumes are based on the Advanced Course on Petri Nets*, pages 429–528, London, UK, 1998. Springer-Verlag.

A Appendix: oral presentation

A.1 Organisation of the presentation

The presentation will be organised as a live demo mixed with the occasional slide explaining some of the underlying concepts. Following is the list of steps we will go through, illustrated by screenshots of the actual demo we will perform.

- A quick introduction explaining the subject and the goals of the presentation. We are going to introduce a tool for model checking algebraic Petri nets. The goal of the presentation will be showing the audience the fundamental features of the tool through examples. We will focus on the usability features rather than the theoretical aspects.
- We will start by creating a new empty project using the ALPiNA wizard (Fig. 3). The project creation screen also shows that there are some examples bundled with ALPiNA's distribution.

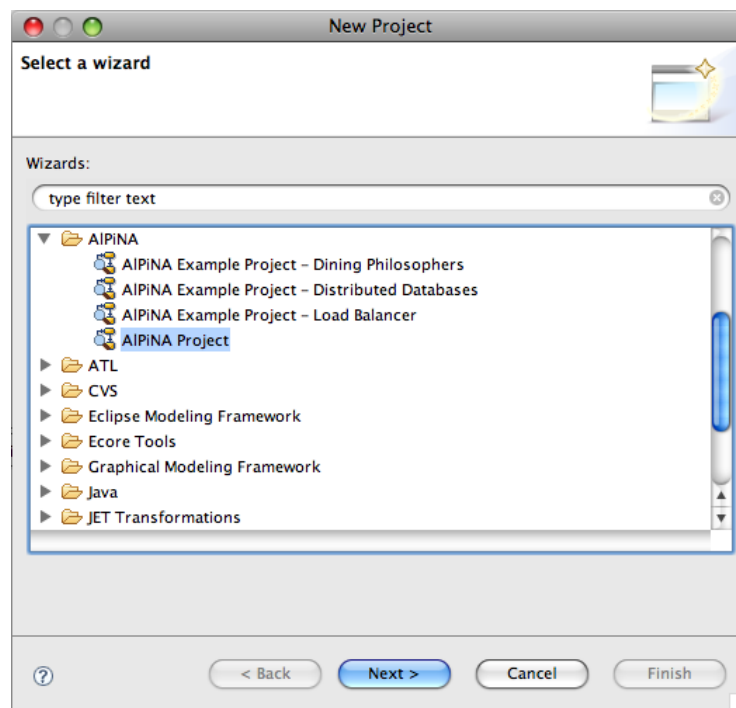


Fig. 3. Creation of a new ALPiNA project

- We will show how to use AIPiNA's graphical editor by creating a very simple model of a counter. Using the tool palette, we will create a small net with a place, a transition and two arcs (Fig. 4). This will illustrate how easy and intuitive it is to create a net using the graphical syntax provided.

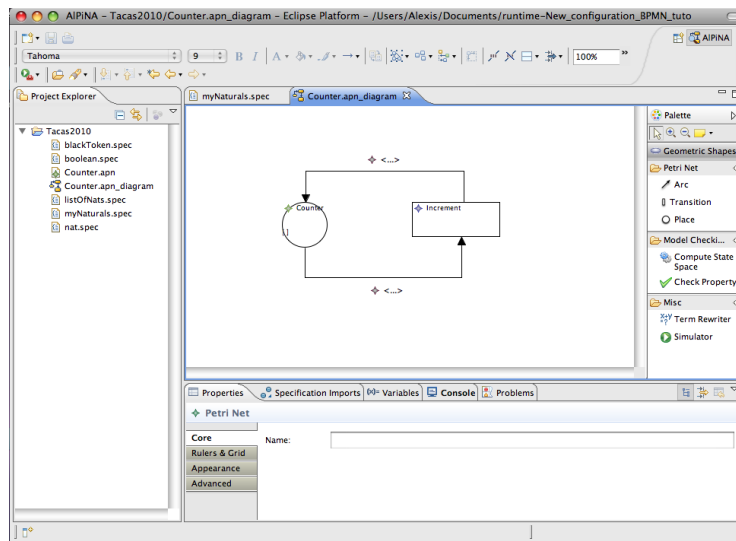


Fig. 4. Creation of a simple counter network

- Then we will show how to use the textual editor to create a small data type, representing natural numbers (only the generator values plus the “greater than” operation will be implemented). This will also show the syntax highlighting and autocompletion features of the editor, which are also features that improve the usability of AIPiNA (Fig. 5).

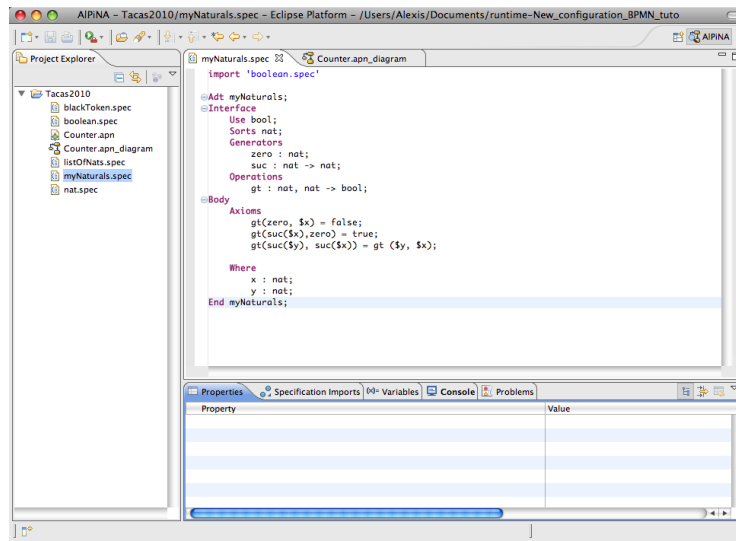


Fig. 5. Creation of an Algebraic Data Type

- We will complete the network to use the newly defined data type and to introduce a guard on the transition, which limits the counter to the value 5 (it should stop after the token in the *Counter* place reaches the value of five). Fig. 6

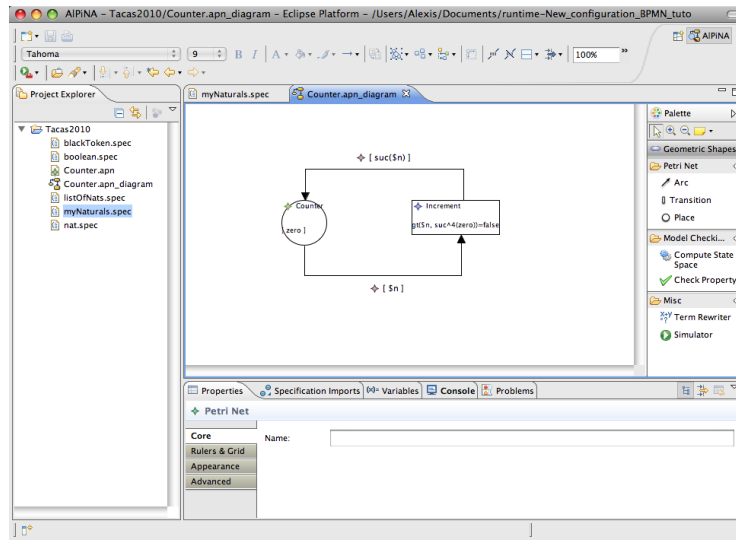


Fig. 6. Complete counter network

- We will show how the “Compute State Space” tool allows a one-click calculation of the state space, giving statistics on the number of states and on the performance of the computation (Fig. 7).

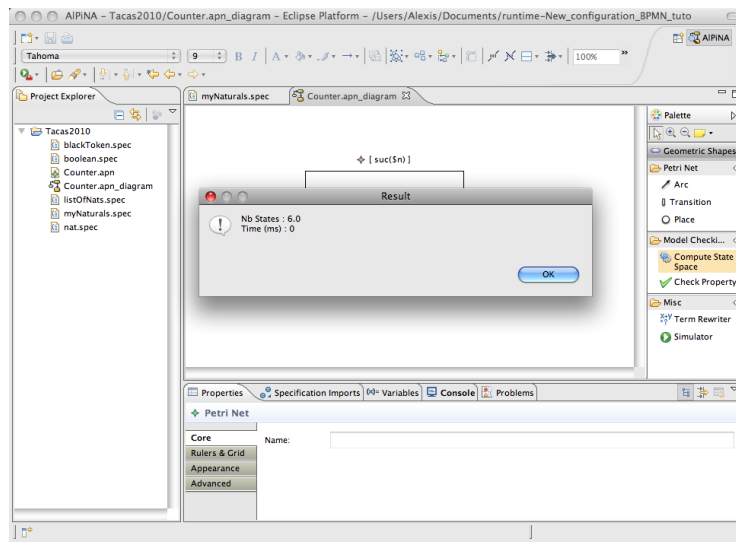


Fig. 7. State space calculation

- To wrap up the example, we will show how to use the editor to define properties of the model, and how the “Check Property” tool provides a one-click check of the properties, with counterexamples (Fig. 8 and 9). This first part of the presentation should take about one third of the allocated presentation time.




Fig. 8. Property definition

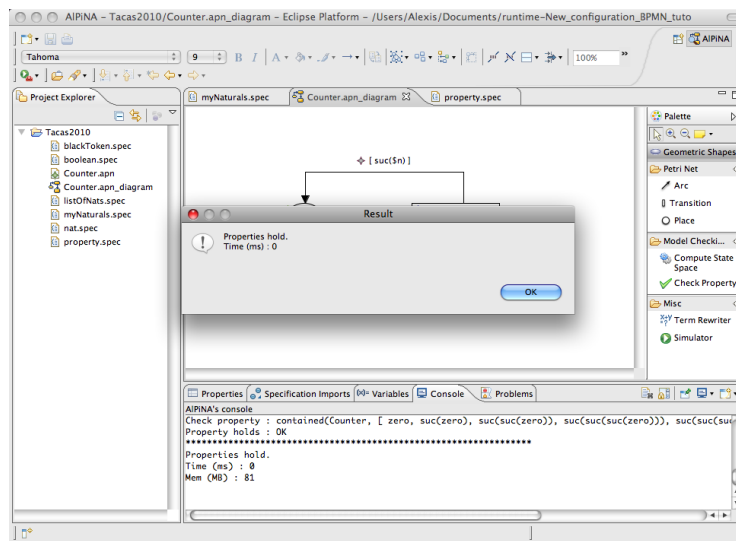


Fig. 9. Property check

- We will show properties for this second example. This will show more complex properties, such as the intersection of place contents (Fig. 11).

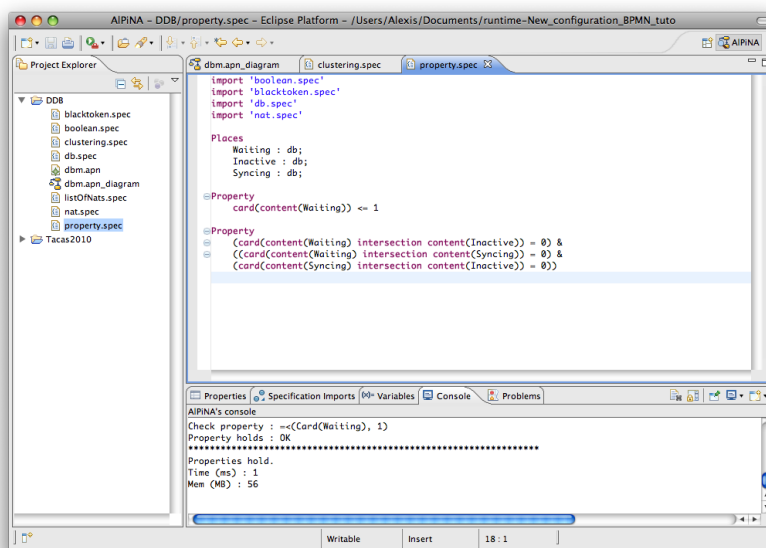


Fig. 11. Distributed database properties

- We will show the definition of clustering in the editor for this model, as it had been explained in the slides. This will illustrate the clustering language (Fig. 12).

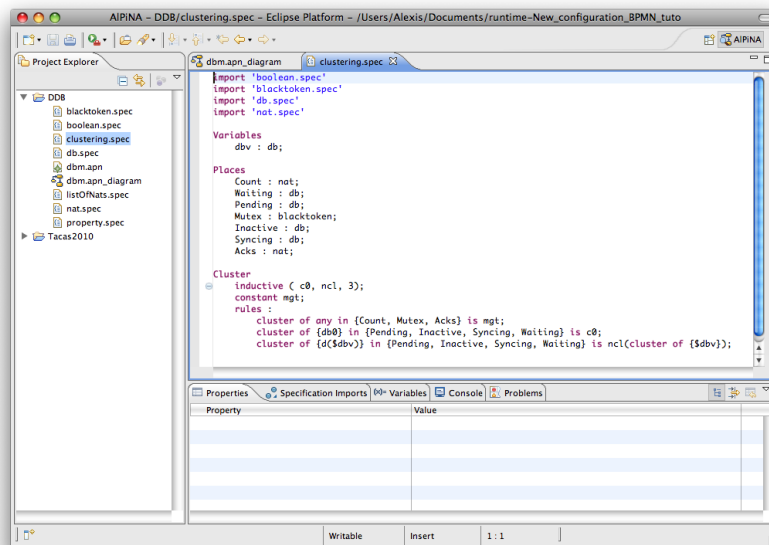


Fig. 12. Distributed database clustering

- Finally, we will show how the state space computation for this model allows unfolding, as explained in the slides. We will show the dialog allowing a quick selection of which type of unfolding should be used for each data type (Fig. 13).

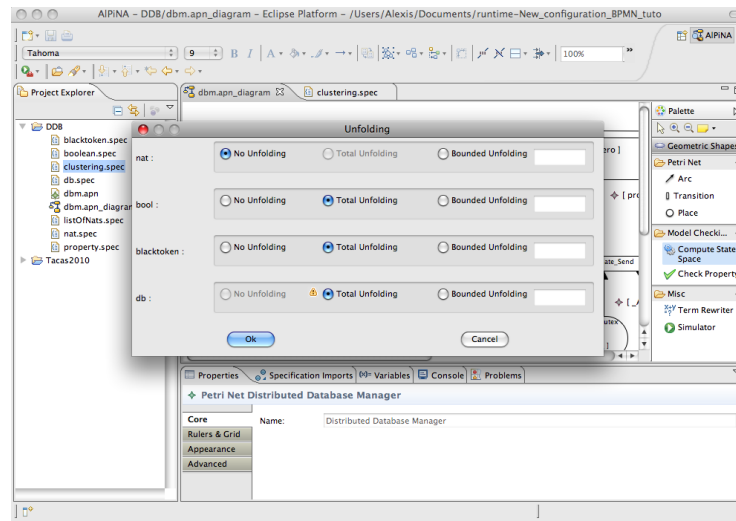


Fig. 13. Distributed database unfolding

- We will briefly show benchmark results (like those seen in the article, Fig. 2) to discuss how clustering and unfolding impact the model checking performance on various complex models.
- We will wrap up the presentation by resuming what was shown and listing future perspectives of AIPiNA. Pointers to the download page and the online tutorial will be given.

A.2 Status of the tool

AIPiNA's current release is 0.9. The tool is in active development and new releases appear regularly. The current version already implements all the features that are discussed in the article and in the presentation. For the moment AIPiNA is an academic research project and it is being used by a small community of researchers and students. To ensure the quality and stability of the tool, the development process is taking advantage of batteries of unit tests, integration tests and a bug tracking system.

A.3 Links

Webpage <http://alpina.unige.ch>

Current stable release <http://alpina.unige.ch/downloads>

Nightly builds <http://alpina.unige.ch:9000/hudson>

Online tutorial <http://alpina.unige.ch/documentation>
Bug tracking system <http://alpina.unige.ch:8000>