

A design environment for mobile applications

Stephen Gilmore Valentin Haenel Jane Hillston
Jennifer Tenzer

Laboratory for Foundations of Computer Science,
The University of Edinburgh, Scotland.

January 5, 2006

Abstract

In this paper we show how high-level UML models of mobile computing applications can be analysed for classical performance measures such as throughput. The approach proceeds by compiling the UML model into a representation in the formally-defined modelling language of PEPA nets. The compilation process and subsequent performance analysis based on numerical solution of a Continuous-Time Markov Chain is supported by a software tool, the Choreographer design platform. Choreographer interoperates with popular UML tools by reading and writing UML models in the XML Metadata Interchange format (XMI).

1 Introduction

Mobile code and mobile computing applications provide some of the most difficult implementation challenges for application developers today. Mobile code applications must be robust against changing execution environments and computing platforms [1] and may be expected to work with low bandwidth, intermittently unavailable network connections. Mobile computing offers further challenges because mobile computing devices are typically weak computing devices such as PDAs and telephones. It has additional concerns unknown in the immobile computing application area, such as the need to conserve battery energy [2].

The problems of development of mobile applications have been addressed within theoretical computer science by the creation of formal languages which are used to study essential properties

of mobility such as the management of names, the acquisition of new capabilities (or loss of existing ones) and movement of the locus of computation. Sometimes mobility is expressed via other ideas, believed to be more primitive and hence more fundamental, such as the name passing and channel passing used in the π -calculus [3]. Another approach is to represent code movement with a metaphor such as the nested named containers used to represent authorisation domains in the ambient calculus [4]. In the PEPA nets approach [5] mobility is represented more literally by the passing of objects with behaviour as the tokens of a coloured Petri net. The PEPA nets language and others such as the Stochastic π -calculus [6] have the additional benefit that they are quantified modelling languages which can be used for performance analysis of systems in addition to determining behavioural properties such as freedom from deadlock. For mobile computing the quantitative information is especially relevant: small devices do not have plentiful computing capacity so describing resource consumption and predicting resource-related problems is particularly helpful.

One idea to connect the high-level models used by modellers with the low-level implementations produced by developers might be to try to interface with software development directly by converting stochastic process models into programs [7], or programs into stochastic process models. However, an alternative approach, and the one which we follow here, is to connect with another language used in the software development process, namely the high-level modelling language, UML. Specifically, we extract a PEPA net model from a UML activity diagram, analyse the PEPA net and report the results back as a modified activity diagram.

Structure of this paper: Section 2 describes how UML constructs are mapped into the constructs of the PEPA nets modelling language. Section 3 discusses the Choreographer design platform. Section 4 presents an example. Section 5 discusses related work. Future work and conclusions follow in sections 6 and 7.

2 Mapping

In this section we give describe the mapping from UML activity diagrams to PEPA nets. We use activity diagrams following the approach to the representation of mobility developed by Baumeister *et al.* within the AGILE project [8, 9]. In this approach activities which result in an object

changing location are decorated with the stereotype $\langle\langle\text{move}\rangle\rangle$. Furthermore, each object box is decorated with the tag $\langle\langle\text{atLoc} = \cdot\rangle\rangle$, which indicates the location of the object at that point within the sequence of activities. An example of such an activity diagram is shown in Figure 3.

A PEPA net is a structured performance modelling formalism. It is a coloured Petri net in which the net level represents the infrastructure of the system. Thus the places of the Petri net correspond to the different locations or computational contexts present in the system, while the transitions correspond to state changes which result in a different distribution of computational elements across contexts. The tokens of the PEPA net are PEPA components — process algebra terms which capture the logical behaviour of the computational elements, including those actions which cause them to change context. PEPA components are also used to define the contexts of the system. Each place may have associated with it storage locations for a number of computational elements and static components. The static components are additional PEPA components which capture the facilities of that context. Unlike tokens, static components may not perform actions which trigger movement around the net: they are restricted to remain within one place. The initial marking of the PEPA net captures the local state of each PEPA token and static component, and their initial location within the system.

The UML activity diagrams which we consider clearly exhibit both locations and movements — these correspond to places and transitions at the net level of the PEPA net. Therefore the first step of our mapping is to create a net in which there is one place for each location which appears on the right hand side of an $\langle\langle\text{atLoc} = \cdot\rangle\rangle$ tag in the activity diagram. Second, we create a net level transition corresponding to each $\langle\langle\text{move}\rangle\rangle$. Such a transition will have an input arc corresponding to each object flow into the activity and an output arc for each object flow out of the activity. Moreover the places connected to these arcs are determined by the locations of the corresponding objects. This completes the net level definition of the PEPA net, but it remains to instantiate the PEPA equations which define the tokens of the PEPA nets and the static components specifying the contexts associated with each place.

The tokens of the PEPA net correspond to the objects of the activity diagram: we define one token for each object. The activities which the token undertakes are defined by the structure of the activity diagram, disregarding activities which do not involve this object, i.e. which do not have an association with this object. Each activity in the activity diagram is mapped to an activity in

the PEPA definition. In the diagrams we consider there are only two possibilities for the associations between the activities of an object. They may be sequentially ordered which corresponds to the prefix operation in the PEPA definitions. Otherwise they may be alternative behaviours, represented either with an explicit decision diamond, or implicitly by having two outgoing associations from an activity. In either case this situation is captured by a choice in the PEPA component, with the two activities as the alternatives.

It remains to define any static components which may be required in the model. A static component is only required if there is an activity which does not have an associated object flow. In this case the activity will not be defined within any of the tokens of the PEPA net but must still nevertheless have a representation in the PEPA net. In this case the activity is mapped to the activity of a static component associated with one of the places of the PEPA net; which place is determined by considering the last location to which a move was made. The static component is then defined to have each of the non-move activities which are associated with that location ordered by prefix and choice as explained above.

The final step is to define the initial marking of the PEPA net and the context of each place. The capacity of a place to host a token is denoted by a cell or slot of appropriate type. Thus each place is given a cell corresponding to each object which exhibits that location in the activity diagram. There is a cooperation between the cells if the corresponding objects share an activity. Additionally, if there are activities at a location which do not have an associated object/token the corresponding static component must be included in the definition of the place. Such a component will cooperate with the cells at that place on any shared activities.

The complete translation is summarised in the table below:

Activity diagram	PEPA net
location	net-level place
⟨⟨move⟩⟩ activity	net-level transition
object	PEPA token
activity with associated object	activity of the corresponding token
activity without associated object	activity of appropriate static component
first recorded location of object	place of the token in initial marking
location of object-less activity	place of the static component

3 Implementation

Within the Design Environments for Global ApplicationS project (DEGAS) we implemented software tools to realise the above-described mapping of UML activity diagrams to PEPA nets, in addition to providing an implementation of the PEPA nets modelling language in the PEPA Workbench for PEPA nets [5], an extension of our existing PEPA Workbench [10]. The solvers and analysers are connected to the UML drawing tools by software connectors known as *extractors* (for input) and *reflectors* (for output). The solvers, analysers, extractors and reflectors are federated into an integrated design environment called Choreographer.

The Choreographer design platform contains analysers for security and performance analysis. The security analysis is routed via the LySa [11] process calculus, implemented on top of the Succinct Solver suite [12], a state-of-the-art static analysis tool. This aspect of the use of Choreographer is not discussed further here. The interested reader is referred to [13, 14] for further details.

Choreographer contains dynamic analysers for PEPA [15] models and PEPA nets [5]. It uses the Java edition of the PEPA Workbench [16] to solve the CTMC representations which are generated by the PEPA Workbench for PEPA nets. The overall design of the Choreographer platform is shown in Figure 1.

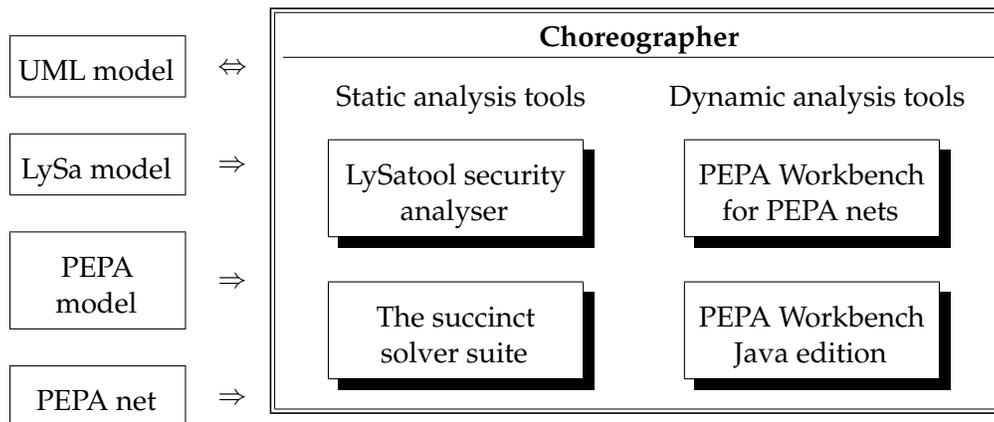


Figure 1: The Choreographer design platform

The other extractors for PEPA and Lysa in Choreographer [17, 18] both use a DOM [19] tree as internal representation of the UML model. For the PEPA Net Extractor/Reflector a different approach has been chosen and the UML model is stored in the Metadata Repository (MDR) by

NetBeans [20].

The MDR project is aimed at persistent storage and manipulation of metadata. It is based on the Meta Object Facility (MOF) standard [21]. A MOF metamodel can be imported into MDR in form of an XMI document. On the basis of the metamodel MDR generates Java interfaces which are compliant with the Java Metadata Interface (JMI) standard [22].

Once a metamodel is present in the repository it can be instantiated. A metamodel instance can be “filled” by a model in XMI format which conforms to the metamodel specification. The model can be manipulated via the Java interfaces that have been generated for the metamodel or by MDR’s reflective API which is independent of the metamodel. MDR also supports the export of models in the repository back to XMI.

In order to process a UML model with MDR, an XMI file containing the UML metamodel must be imported into MDR. For the Extractor/Reflector we have chosen the UML metamodel version 1.4 [23], because it is the basis of the Poseidon UML tool which is used in the DEGAS project. After the import a new instance of the metaclass “UmlPackage” is created. Any UML model that conforms to the imported UML metamodel can now be read into this new instance.

The solution based on MDR has two main advantages in comparison to the DOM approach. First, different UML tools can be supported because MDR is not bound to a particular XMI version or tool-specific saving format. The XMI output of any UML tool that is compliant with the UML metamodel can be read into MDR. Second, MDR’s interfaces for accessing and manipulating the UML model reduce the amount of code that has to be written.

The drawback of the MDR approach is that graphic and structural aspects of the UML model have to be treated separately. The UML metamodel only specifies the structure of UML models and does not contain any elements that allow storage of the diagram layout. So far UML tool vendors have to find their own solutions for saving this information.

The Poseidon tool stores layout data in additional elements of the XMI file. MDR does not recognise these additional elements because they do not appear in the UML metamodel. In order to extract from a Poseidon project, the part of the UML model which conforms to the UML metamodel first has to be separated from the rest. In the Extractor/Reflector module the critical elements are removed by a Poseidon preprocessor. After a successful reflection of the results from the PEPA Workbench for PEPA Nets the layout information is added again to the modified model

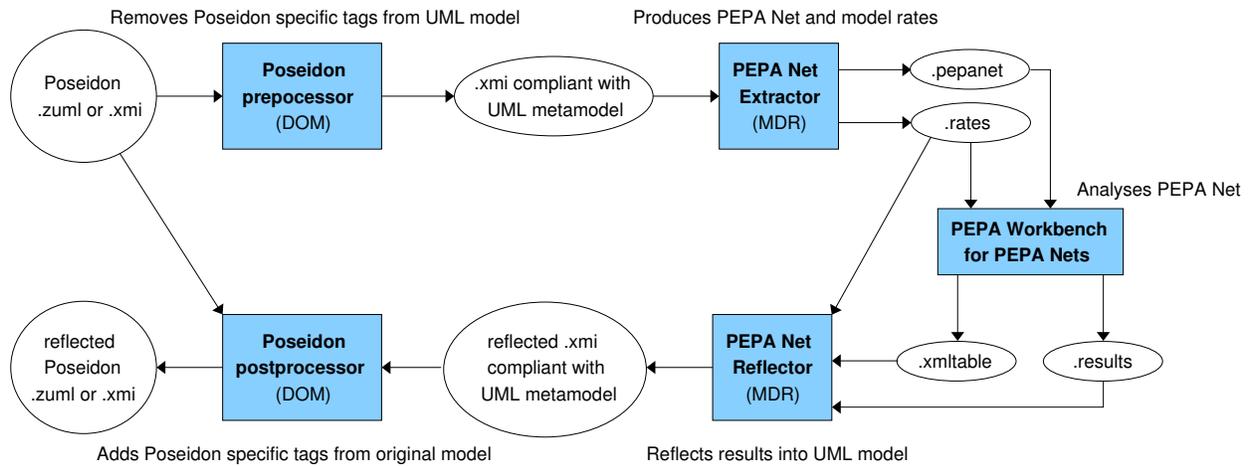


Figure 2: The extraction and reflection process

by a Poseidon postprocessor. An overview of this process and the intermediate outputs are shown in Figure 2. Other UML tools which save the layout data in a different way require their own pre- and postprocessors.

Poseidon provides support for saving the UML model without diagram data which is similar to our preprocessor. We have decided not to use this functionality because we want to reuse the layout data of the original model for the reflected UML model where possible. If the layout data from Poseidon is not saved at all, the layout of the UML diagrams is lost. As shown in Figure 2 the Poseidon postprocessor requires the reflected *.xmi* file and the original Poseidon project as input. It merges the new structural information of the *.xmi* with the old layout data from Poseidon.

4 Example

In this section we present an example use of the Choreographer design platform to investigate the throughput of activities in a UML activity diagram. The example which we model represents both physical and logical mobility. The scenario is of a PDA user on board a moving train connecting to a remote Web site and loading pages of dynamically-generated HTML content. As the train moves the connection to a (stationary) transmitter must be handed over to the next transmitter which is in range. The activities of the model are depicted in Figure 3.

The activities which lead up to the handover event are downloading a file, detecting a weak

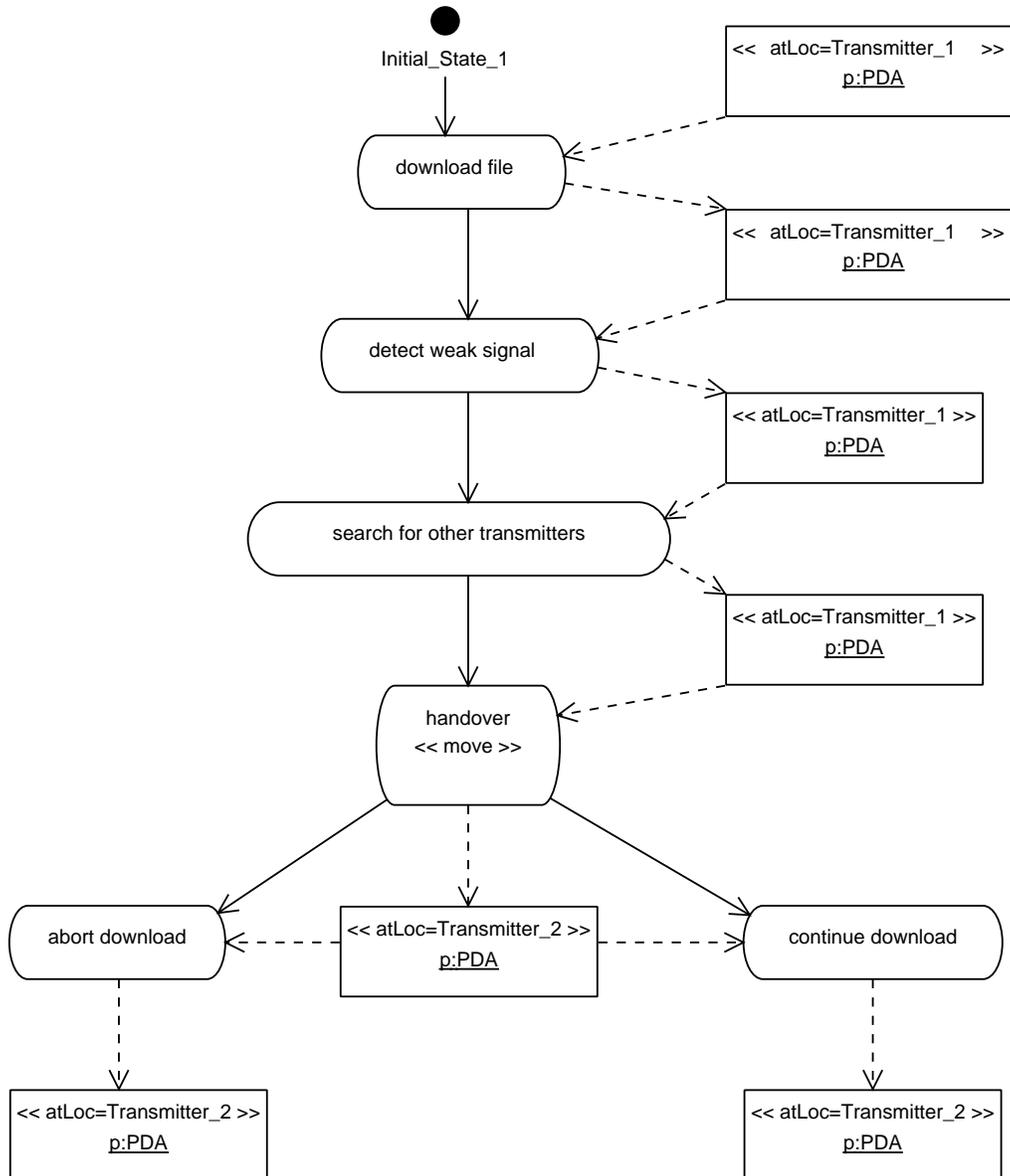


Figure 3: Activity diagram showing the movement and state changes in the model of a PDA user

signal and searching for other transmitters, as described in [24]. The handover activity is a movement activity, as marked by the stereotype $\langle\langle\text{move}\rangle\rangle$.

The handover must happen (because the train is moving) but it is not certain to succeed. We model two possible outcomes of the handover: either the handover succeeds (and the file download continues) or it fails (and the download is aborted). We set the relative probabilities of these two events to be the same: it is as likely that the connection will be dropped as it is that it will survive. In either case, the movement of the train has brought the device into range of the second transmitter.

The analysis process implemented by Choreographer extracts a PEPA net model from the input UML activity diagram and analyses this with the PEPA Workbench for PEPA nets [10, 5]. The results are then reflected back to the UML level, as shown in Figure 4.

When the results are viewed in the Poseidon for UML tool the activities described by the modeller are now annotated with the throughput information computed by the PEPA Workbench, as shown in Figure 5.

With an activity diagram the modelling focus is on activities, and so the performance results which are written back to the diagram also centre on activities, recording throughput. However, a UML project will typically contain diagrams of several different types, each modelling an aspect of the whole system under consideration. State diagrams, a variant on Harel's statecharts [25], would be used to provide a separate record of the behaviour of objects of each of the classes which are being modelled. The purpose of a state diagram is to expose the states of interest in the description of the behaviour of a class and here a different performance measure is more appropriate, namely the steady-state probabilities of the states.

We discuss this aspect of the functionality of Choreographer by considering another part of the mobile PDA user model, simply the request/response view of a client connecting with a Web server. The client and server models are shown in Figures 6 and 7. States in the state diagram notation are represented by boxes with rounded edges. Transitions between states are labelled by the name of the activity which causes the transition. A rate (not shown) is associated with every activity.

The client is modelled at a very high level of abstraction, simply stating that it generates HTTP requests, waits for a response from the server, and then engages in local processing before gener-

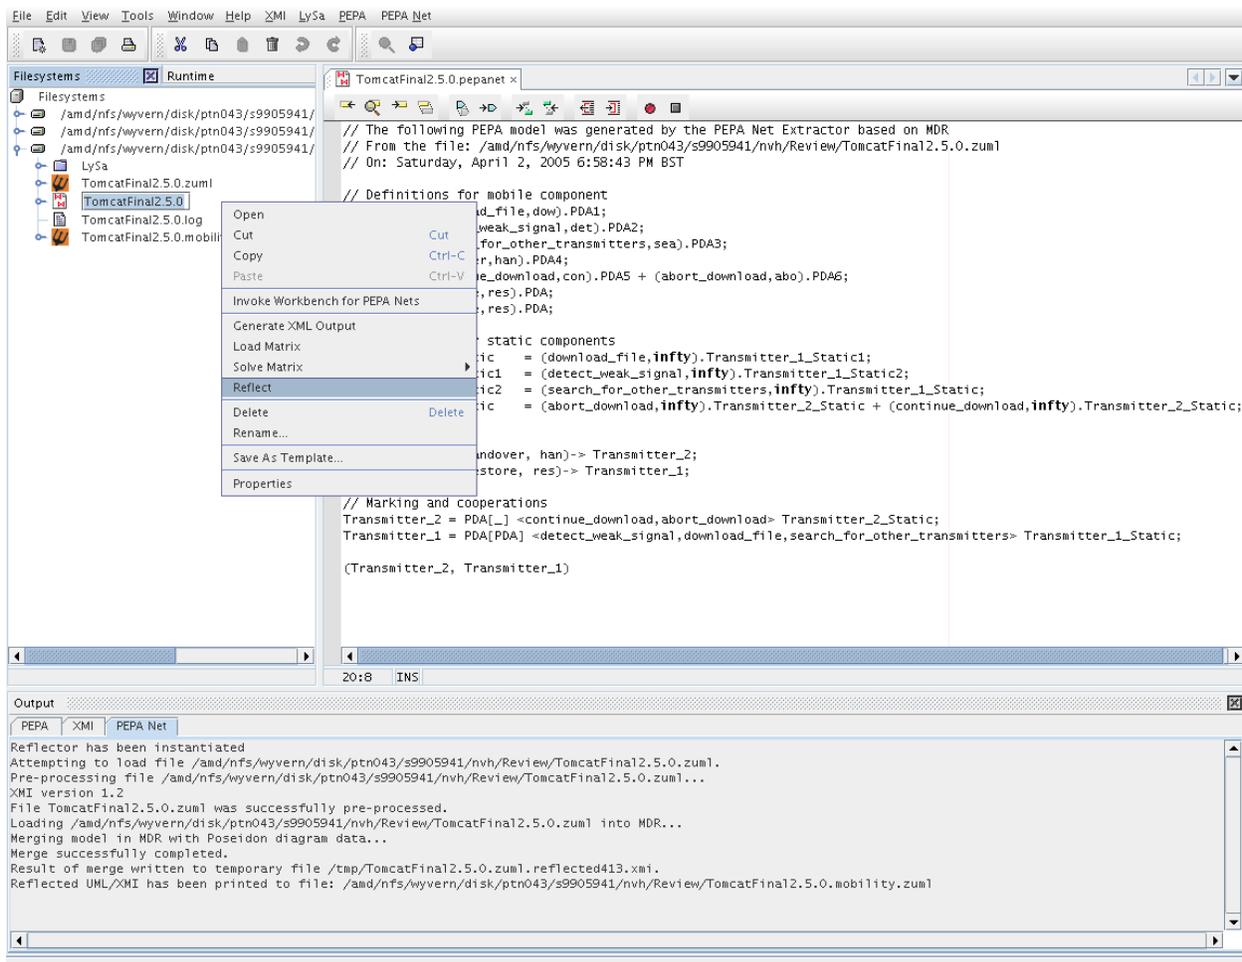


Figure 4: Reflecting the analysed PEPA net model back to an activity diagram representation using the Choreographer design platform

ating the next request.

The model of the server commits more detail about how the server is implemented. The model represents a Tomcat web server [26] which accepts requests for JSP pages, locates the JSP source, translates the page into Java source code, compiles the Java source to Java servlets, and executes the servlet which dynamically generates an HTML page which is sent back to the client as the response to their request. The Tomcat web server includes a simple, but very profitable, optimisation. After the initial locate-translate-compile-execute cycle has completed, the servlet remains resident in memory in the Web container, ready for subsequent invocations. The second and sub-

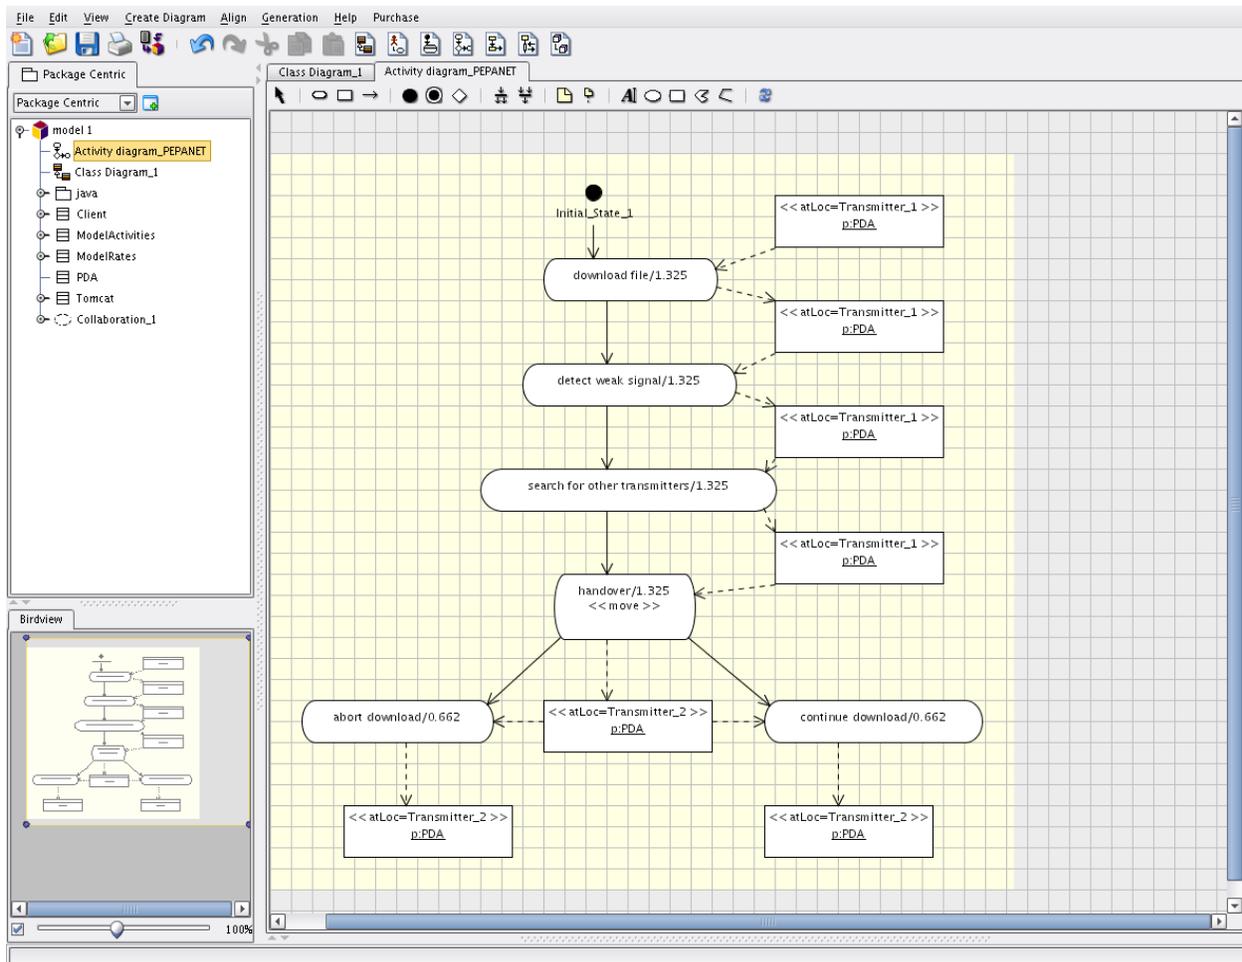


Figure 5: The activity diagram annotated with throughput results viewed in the Poseidon for UML modelling tool

sequent requests for this JSP page are directed to the pre-loaded servlet, bypassing the lengthy translate and compile activities.

We estimated the parameters to the Tomcat model by timing a range of JSP pages running on our Tomcat server, and solved the model with and without the locate servlet optimisation. This quantified the optimisation from the user's point of view in terms of the reduction in the delay spent waiting for the response from the server.

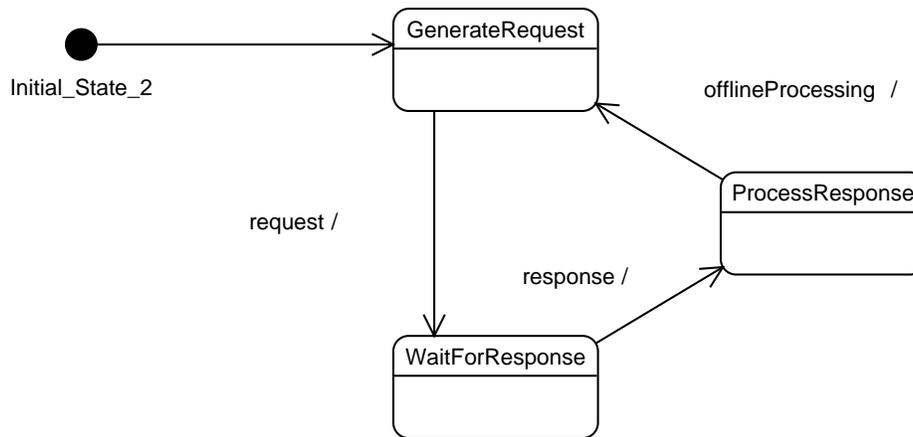


Figure 6: Client state diagram showing the generation of HTTP requests and the local processing of responses when they are received

5 Related work

Space limitations preclude a comprehensive treatment of related work here so we highlight only a few works in the related literature. Our present work builds on our earlier work on mapping UML activity diagrams to process algebras [27]. Early works on the derivation of stochastic process algebra models from UML diagrams include [28, 29]. Other works have targeted different performance analysis methods [30]. A closely related approach to ours is taken in [31], in the UML-based performance analysis of a Web-based micro-business service.

Other authors take the approach of modifying the UML notation in order to express the properties of greatest interest, by their mobility [8] or performance-related [32]. While the reason for these extensions is evident and appealing, the significant loss is that the modified UML can no longer be processed with the UML tools. Thus, we have been able to interoperate with standard, unaltered UML tools whereas others cannot.

Other performance modelling projects based on XMI processing include the UML- Ψ project, a simulator for UML models [33]. Our analysis tools operate via the numerical solution of Markov chains, with all of the advantages and disadvantages which this approach entails (exact solution is an advantage, susceptibility to state-space explosion a disadvantage). Simulation has quite differ-

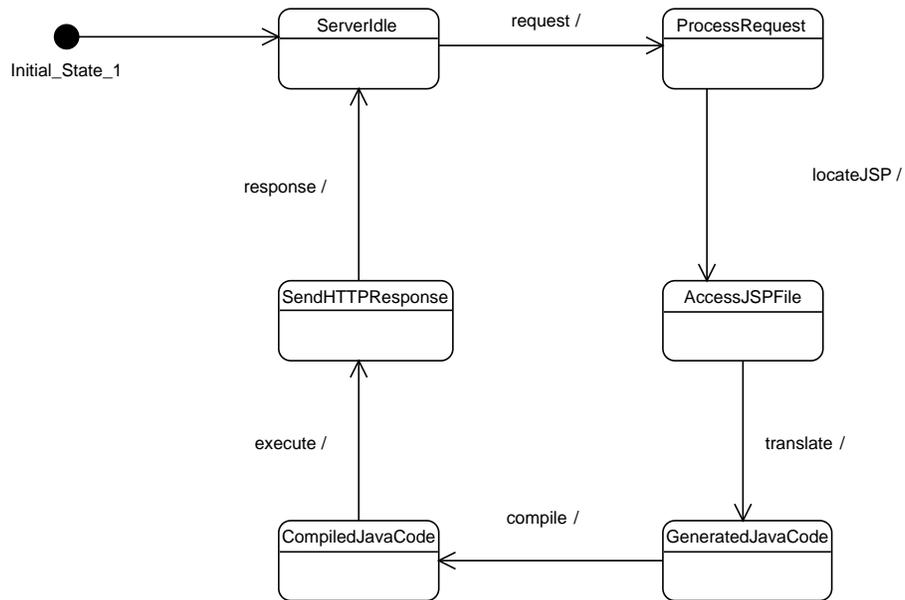


Figure 7: Server state diagram modelling the JSP translate-compile-execute lifecycle before the addition of direct servlet lookup

ent analysis characteristics (approximate solutions require the calculation of confidence intervals, but large state-space size is tolerated) meaning that the approach to simulation of UML models taken in UML- Ψ could complement ours well.

6 Future work

The activity diagrams which are covered by the current version of the PEPA net Extractor/Reflector module have to follow some restrictions. Future versions of the module should try to weaken these restrictions and extend the set of activity diagrams that are extractable without manual modification by the user. For example, advanced UML activity diagram features such as merge, decision, join and fork nodes could be considered.

Another possibility for extending the PEPA Net Extractor is to use more information of the UML model in order to specify the resulting PEPA net more precisely. In activity diagrams tags that define which action is performed by which static component could be introduced to the UML

model. It may also be useful to base the extraction on more than one UML diagram type. Interaction diagrams, for instance, would permit explicit definition of which components cooperate with each other. This becomes particularly important if several mobile and static components are considered at one place, which is not possible in the current version of the module.

The value of using the PEPA modelling language as an intermediate language on the route to automated performance analysis of UML models by numerical solution of continuous-time Markov chains has not yet been fully exploited. PEPA is supported by a suite of performance analysis tools including the PRISM probabilistic model-checker [34], the Imperial PEPA Compiler (ipc) [35] and the Möbius multi-paradigm, multi-formalism performance modelling platform [36, 37]. The availability of well-engineered, sophisticated analysis tools such as these is the primary reason why we were not tempted to implement UML performance analysis tools directly, but rather go through the intermediate PEPA representation.

We have previously connected our extractors and reflectors for the PEPA stochastic process algebra [17] to the PRISM model-checker and applied this to a real-world analysis problem from the domain of mobile telephony [38]. We have previously analysed PEPA net models using PRISM [39]. However, much more could be done to have tighter integration with tools such as Möbius, PRISM and ipc accessible to the Choreographer design platform.

7 Conclusions

The most direct way in which design environments can assist developers of complex applications is by providing a representation language with linguistic constructs at the right level of expressivity, together with analysis procedures to check properties of the design. Used together with a UML modelling tool such as Poseidon for UML, the Choreographer design platform meets these needs for developers of mobile code and mobile computing applications. With little overhead the modelling language allows the modeller to precisely record the mobile and immobile components of the system and to distinguish location-changing events from changes of computational state. The extractor-workbench-reflector tool chain powers the performance analysis of high-level model descriptions, returning results in the language in which they were submitted.

Even with such features, a design environment would be unsuitable to use if the encoding of a diagram into a net was sufficiently complex that its implementation would be likely to have

significant flaws. Fortunately the mapping of our mobility-aware diagrams into PEPA nets is sufficiently clean that this problem does not arise. In the implementation of the extractor software also we followed a principled implementation plan, utilising fully the metadata repository which was available to us, with the added benefits from this increasing our confidence in the realisation of the mapping in software.

Acknowledgements: The authors are supported by the DEGAS (Design Environments for Global ApplicationS) project IST-2001-32072 funded by the FET Proactive Initiative on Global Computing. Perdita Stevens provided helpful comments on earlier drafts of this paper. The UML diagrams in the paper were created with Poseidon for UML Community Edition. The Choreographer design platform is available for download from <http://www.lfcs.ed.ac.uk/choreographer> as a Java application which has been successfully tested on Windows and Red Hat Linux systems.

References

- [1] Alfonso Fuggetta, Gian Pietro Picco, and Giovanni Vigna. Understanding code mobility. *IEEE Trans. Softw. Eng.*, 24(5):342–361, 1998.
- [2] G. Chen, R. Shetty, M. Kandemir, N. Vijaykrishnan, M. J. Irwin, and M. Wolczko. Tuning garbage collection for reducing memory system energy in an embedded Java environment. *ACM Transactions on Embedded Computing Systems*, 1(1):27–55, 2002.
- [3] Robin Milner. *Communicating and Mobile Systems: the π -Calculus*. Cambridge University Press, 1999.
- [4] L. Cardelli and A. Gordon. Mobile ambients. *Theoretical Computer Science*, 240:177–213, 2000.
- [5] S. Gilmore, J. Hillston, M. Ribaud, and L. Kloul. PEPA nets: A structured performance modelling formalism. *Performance Evaluation*, 54(2):79–104, October 2003.
- [6] Corrado Priami. Stochastic π -calculus. *Computer Journal*, 38(7):578–589, 1995.
- [7] S. Gilmore, J. Hillston, and D.R.W. Holton. From SPA models to programs. In M. Ribaud, editor, *Proceedings of the Fourth Annual Workshop on Process Algebra and Performance Modelling*, pages 179–198. Dipartimento di Informatica, Università di Torino, CLUT, July 1996.
- [8] Hubert Baumeister, Nora Koch, Piotr Kosiuczenko, Perdita Stevens, and Martin Wirsing. UML for global computing. In *Global Computing: Programming Environments, Languages, Security, and Analysis of Systems, IST/FET International Workshop, GC 2003*, pages 1–24, Rovereto, Italy, February 2003.

- [9] L. Andrade, P. Baldan, H. Baumeister, R. Bruni, A. Corradini, R. De Nicola, J. L. Fiadeiro, F. Gaducci, S. Gnesi, P. Hoffman, N. Koch, P. Kosiuczenko, A. Lapadula, D. Latella, A. Lopes, M. Loreti, M. Massink, F. Mazzanti, U. Montanari, C. Oliveira, R. Pugliese, A. Tarlecki, M. Wermelinger, M. Wirsing, and A. Zawlocki. AGILE: Software architecture for mobility. In Dirk Pattinson Martin Wirsing and Rolf Hennicker, editors, *Recent Trends in Algebraic Development Techniques—16th International Workshop, WADT 2002*, volume 2755 of LNCS, Frauenchiemsee, Germany, September 2002. Springer.
- [10] S. Gilmore and J. Hillston. The PEPA Workbench: A Tool to Support a Process Algebra-based Approach to Performance Modelling. In *Proceedings of the Seventh International Conference on Modelling Techniques and Tools for Computer Performance Evaluation*, number 794 in Lecture Notes in Computer Science, pages 353–368, Vienna, May 1994. Springer-Verlag.
- [11] M. Buchholtz, H. Riis Nielson, and F. Nielson. A calculus for control flow analysis of security protocols. *International Journal of Information Security*, 2(3–4):145–167, 2004.
- [12] F. Nielson, H. Riis Nielson, H. Sun, M. Buchholtz, R. R. Hansen, H. Pilegaard, and H. Seidl. The Succinct Solver Suite. In Kurt Jensen and Andreas Podelski, editors, *Proc. TACAS’04*, volume 2988 of *Lecture Notes in Computer Science*, pages 251–265. Springer-Verlag, 2004.
- [13] S. Gilmore, V. Haenel, L. Kloul, and M. Maidl. Choreographing security and performance analysis. DEGAS project paper. Submitted for publication, December 2004.
- [14] Mikael Buchholtz, Stephen Gilmore, Valentin Haenel, and Carlo Montangero. End-to-end integrated security and performance analysis on the DEGAS Choreographer platform. To appear in the proceedings of Formal Methods 2005. Editors: John Fitzgerald, Ian Hayes and Andrzej Tarlecki, January 2005.
- [15] J. Hillston. *A Compositional Approach to Performance Modelling*. Cambridge University Press, 1996.
- [16] N.V. Haenel. *User Guide for the Java Edition of the PEPA Workbench—Tabasco release*. LFCS, Edinburgh, October 2003.
- [17] C. Canevet, S. Gilmore, J. Hillston, M. Prowse, and P. Stevens. Performance modelling with UML and stochastic process algebras. *IEE Proceedings: Computers and Digital Techniques*, 150(2):107–120, March 2003.
- [18] M. Buchholtz, C. Montangero, L. Perrone, and S. Semprini. For-LySa: UML for authentication analysis. In C. Priami and P. Quaglia, editors, *Proceedings of the second workshop on Global Computing*, volume 3267 of *Lecture Notes in Computer Science*, pages 92–105. Springer Verlag, 2004.
- [19] Document Object Model (DOM) Specification. Available from the W3C at <http://www.w3.org/DOM/>.

- [20] NetBeans Metadata Repository (MDR). Website at <http://mdr.netbeans.org/>.
- [21] Meta-Object Facility (MOF). Available from the OMG at <http://www.omg.org>.
- [22] Java Metadata Interface (JMI) Specification, JSR 40. Available from the Java Community Process at <http://www.jcp.org/>.
- [23] OMG Unified Modeling Specification, version 1.4, september 2001. Available from the OMG at <http://www.omg.org>.
- [24] Héctor Velayos and Gunnar Karlsson. Techniques to reduce IEEE 802.11b MAC layer handover time. Technical Report TRITA-IMIT-LCN R 03:02, KTH, Royal Institute of Technology, Stockholm, April 2003.
- [25] D. Harel. Statecharts: a visual approach to complex systems. *Science of Computer Programming*, 8:231–274, 1987.
- [26] Apache Jakarta Tomcat. Web site available at <http://jakarta.apache.org/tomcat/>.
- [27] C. Canevet, S. Gilmore, J. Hillston, L. Kloul, and P. Stevens. Analysing UML 2.0 activity diagrams in the software performance engineering process. In *Proceedings of the Fourth International Workshop on Software and Performance*, pages 74–78, Redwood Shores, California, USA, January 2004. ACM Press.
- [28] R. Pooley. Using UML to derive stochastic process algebra models. In N. Davies and J. Bradley, editors, *Proceedings of the Fifteenth UK Performance Engineering Workshop*, pages 23–33, Department of Computer Science, The University of Bristol, July 1999.
- [29] P. Mitton and R. Holton. PEPA performability modelling using UML statecharts. In N. Thomas and J. Bradley, editors, *Proceedings of the Sixteenth UK Performance Engineering Workshop*, pages 19–33, Department of Computer Science, The University of Durham, July 2000.
- [30] Vittorio Cortellessa and Raffaella Mirandola. Deriving a queueing network based performance model from UML diagrams. In *WOSP '00: Proceedings of the second international workshop on Software and performance*, pages 58–70, New York, NY, USA, 2000. ACM Press.
- [31] Katerina Pokozy-Korenblat, Corrado Priami, and Paola Quaglia. Performance analysis of a UML micro-business case study. In Corrado Priami and Paola Quaglia, editors, *Global Computing*, volume 3267 of *Lecture Notes in Computer Science*, pages 107–126. Springer, 2004.
- [32] David N. Jansen and Holger Hermanns. QoS modelling and analysis with UML-statecharts: the StoCharts approach. *SIGMETRICS Perform. Eval. Rev.*, 32(4):28–33, 2005.

- [33] Moreno Marzolla and Simonetta Balsamo. UML-PSI: The UML Performance Simulator. In *First International Conference on The Quantitative Evaluation of Systems (QEST'04)*, pages 340–341, Enschede, the Netherlands, 2004. IEEE Computer Society Press.
- [34] M. Kwiatkowska, G. Norman, and D. Parker. Probabilistic symbolic model checking with PRISM: A hybrid approach. In J.-P. Katoen and P. Stevens, editors, *Proc. 8th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'02)*, volume 2280 of *LNCS*, pages 52–66. Springer, April 2002.
- [35] J.T. Bradley, N.J. Dingle, S.T. Gilmore, and W.J. Knottenbelt. Derivation of passage-time densities in PEPA models using IPC: The Imperial PEPA Compiler. In G Kotsis, editor, *Proceedings of the 11th IEEE/ACM International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunications Systems*, pages 344–351, University of Central Florida, October 2003. IEEE Computer Society Press.
- [36] G. Clark, T. Courtney, D. Daly, D. Deavours, S. Derisavi, J. M. Doyle, W. H. Sanders, and P. Webster. The Möbius modeling tool. In *Proceedings of the 9th International Workshop on Petri Nets and Performance Models*, pages 241–250, Aachen, Germany, September 2001.
- [37] G. Clark and W.H. Sanders. Implementing a stochastic process algebra within the Möbius modeling framework. In L. de Alfaro and S. Gilmore, editors, *Proceedings of the first joint PAPM-PROBMIV Workshop*, volume 2165 of *Lecture Notes in Computer Science*, pages 200–215, Aachen, Germany, September 2001. Springer-Verlag.
- [38] Stephen Gilmore and Leïla Kloul. A unified tool for performance modelling and prediction. *Reliability Engineering and System Safety*, 89(1):17–32, July 2005.
- [39] S. Gilmore, J. Hillston, L. Kloul, and M. Ribaud. Software performance modelling using PEPA nets. In *Proceedings of the Fourth International Workshop on Software and Performance*, pages 13–24, Redwood Shores, California, USA, January 2004. ACM Press.