Languages and Architectures for Pervasive Business Processes

Thomas Hildebrandt (hilde@itu.dk)
Programming, Logic and Semantics Group, IT University of Copenhagen
www.itu.dk/people/hilde/

November 18, 2005

1 Pervasive Business Processes

The vision for the future, pervasive it-support at our work places is that we can carry a mobile device anywhere allowing us seamlessly to cooperate, share information and coordinate tasks with our colleagues. A topical example is the electronic medical journals currently being developed. *Workflow management systems* (WFMS) incorporated in most ERP-systems such as Microsoft Business Solutions Axapta are likely to play a central role for this vision to become true. Already early 2000 it was estimated that there were more than 200 suppliers worldwide offering standalone workflow management software [44] and in Denmark the largest 290 companies are expected to invest more than 4 billion Dkr during the next few years in ERP systems [18].

A workflow management system is a generic it-system supporting coordination, resource allocation and analysis of the execution of *business processes* in organisations. In the current reference model for Workflow management systems (Figure 1), the core service is carried out by a so-called *workflow engine* on a central server. The workflow engine keeps track of the status and progress of active processes, referred to as *active cases*. Active cases are typically long-lived, as e.g. the process initiated by the first consultation of a pregnant woman to the birth of the child at the hospital, and is therefore persisted in a database. Workflow processes are described in a so-called *business or workflow process language*, either by a business analyst or a manager in the organisation and often using a graphical *process definition tool*. Through *client applications* workers can check the status and next possible actions on an active case and report if one of these actions is carried out. Besides the client applications, the workflow engine may interact directly with different active cases running at other workflow engines or other it-services, as for instance web services. Finally, the workflow

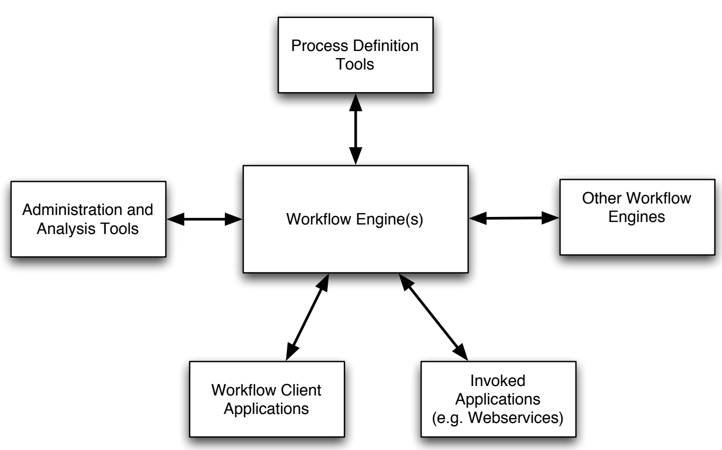

Figure 1: WFMS Reference Model [16].

engine may be connected to *analysis tools* allowing it to, e.g., analyse best-case and worst-case scenarios of business processes.

The current workflow management systems and business process languages are far from the vision of pervasive and seamless it-support. Researchers from the Programming, Logic and Semantics group and the Design of Organisational IT group at ITU have taken up the challenge to provide a foundation for pervasive business process languages. We focus on the following key issues.

Flexibility and Extensibility Only within the last few years are the first standards being proposed for business process languages. Early languages such as XLANG [41] from Microsoft and WSFL [27] from IBM have merged into the industry open standard BPEL4WS [2], only to be developed further into the Business Process Modelling Notation (BPMN) [5] proposed by the Business Process Management Initiative. Clearly, business process languages are only in their early days and will continue to evolve according to the development of technology and new requirements of users. Most of the proposed languages are based on XML or at least have an XML representation, in principle supporting extensibility and exchange of process descriptions between different systems. However, in reality the compatibility between systems is limited. Moreover, changes of the business process *language* presently require expensive changes in the core implementation of the workflow engine. The next generation architecture must support *safe and easy dynamic changes and extensions of the business process language itself.*

Mobility In the current architecture of business process management systems, process descriptions and the state of running cases are stored centrally in a relational database. The client applications thus must be connected to the server to stay synchronised with the status of active cases and be able to report execution of actions. The next generation of process languages and architectures should support *movement of process descriptions and running cases* to *mobile and possibly disconnected clients* or to *outsource process descriptions* to other companies. Movement of running cases may for instance be used when a doctor is visiting a patient at home, bring the case along. An example of outsourcing could be the movement of a medical case from one hospital to another.

Higher-order: Managing the Business Process Management One of the key ideas of employing a business process management system is to be able to manipulate and redesign processes. As an example, a medical work process may need to be adapted due to the introduction of new medical equipment. Ironically, none of the present business process languages are able to describe these so called "higher-order" work processes in which business processes and active processes are also the subjects of the process. The future business process languages should support the *higher-order work processes of the business analysts and managers who create and manipulate business processes and running cases.*

Context-dependency and Disconnected operation Finally, an emerging requirement of mobile it-systems is that of *location and context awareness* [26]. A location aware business process will be able to depend on its current location and context, e.g. adapting to different users or to different environments. Besides entirely new services such as "locating the nearest doctor" context-dependency is envisioned to be able to improve user-interfaces of small handheld devices by customising the user-interface to the present context. Also, *disconnected operation*, i.e. the ability to operate when disconnected from a central server can be regarded as a special case of context awareness, making the applications on a hand-held device aware of the present access to shared resources.

Targeting the issues identified above requires a *formalisation* of business process languages and a fresh look at the *architectures* and *uses* of workflow management systems.

2 Formalisations and architectures for pervasive business processes

The research project on formalisations and architectures for pervasive business processes at ITU aims to combine the research of the principal investigator, associate professor Thomas Hildebrandt, in *formal semantics of interacting processes* with focus on *extensible semantical frameworks* [6, 14, 13] and *higher-order mobile resources* [11, 7, 10] and the research of the co-investigators, assistant professor, Henning Niss and associate professors Jens Chr. Godskesen and Kjeld Schmidt, and external research collaborators, professor Robin Milner and professor Marta Kwiatkowska. The co-investigators and external collaborators add the important aspects of *implementation of distributed systems* (Niss), *tools for formal verification and analysis of processes* (Godskesen, Kwiatkowska), and *computer supported cooperative work* (Schmidt). The research will be interleaved with the development of prototype tools by research programmers in cooperation with developers at Microsoft Business Solutions (MBS) to demonstrate the innovations in practice and continuously measure the results against industry languages and case studies in CSCW. Thereby it is ensured that the proposed formalisations will be aligned with the technological possibilities in design of distributed systems, applicable to formal analysis and meet the requirements of users in practice.

The proposed research project is expected to contribute with important innovations for developing extensible mobile WFMSs and contribute to an international research effort on foundations for global ubiquitous computing in which professor Milner and professor Kwiatkowska are key researchers [15]. Below we shortly present our initial work and origin of the research project.

2.1 Formalisations of pervasive business processes

None of the present business process languages have a formal definition. Consequently, it is not clear how all processes described in the languages should be interpreted or how automatic it-support or code generated from the processes will behave, and it is virtually impossible to guarantee the stability of the business process management systems when process languages and the technology evolve. It is therefore a highly topical open challenge to provide a formalisation of the semantics of business process that can provide a sound and solid foundation for the next generation of business process languages. Such a formalisation should take advantage of the advanced features of the current theory of process models and applicable to formal analysis and verification tools. The *Process Modelling Group* [45] recently instigated by prof. Robin Milner aims to enhance the quality of commercial process software by putting existing theoretical work to use and to further the scientific understanding of process modelling. The group members includes key researchers, people working in software companies and standard bodies. From ITU the group includes Thomas Hildebrandt and Carsten Butz and from DIKU the group includes FIRST PhD student Christian Stefansen.

Pioneering work has provided a formal foundation describing general workflow processes in the Petri Net model [44] . Due to its graphical presentation the Petri Net model is a natural choice for formalising workflows and it has very rich support for automated tools (e.g. see [19]). However, it is currently up to debate if Petri Nets will be as suitable a foundation for the next generation of workflow process languages [46, 39]. Although it is possible, even the current proposed languages such as BPEL are not easily described as Petri Nets [33]. A number of researchers and practitioners have recently been advocating the use of process calculi, in particular the pi-calculus [38, 39], as a formal foundation for business processes. This includes a research project within the NEXT research project (jointly between DIKU, ITU and Microsoft Business Solutions) using first-order process calculi to formalize contracts [40].

Higher-order mobile processes Support for *mobile* and *higher-order* processes is in particular a reason for considering other formal models than Petri Nets. Only recently have extensions of the Petri Net model been proposed which are able to express mobile and higher-order process descriptions (see e.g. [17]) and the theory is far from well understood yet. Compared to this, calculi for mobility and higher-order processes have been around for 25 years since the seminal work on the pi-calculus by Milner [29] and the calculus of higher-order communicating systems (CHOCS) by Thomsen [42]. It is still a very active research area, not least because of the ubiquitous use of mobile computing. In the recent years Hildebrandt, Godskesen and Bundgaard have successfully extended the theory of higher-order mobile processes to cover mobility of both passive and active mobile processes and mobile connections as in the pi-calculus [11, 10, 7].

Extensibility The mathematical notation and semantics of process calculi limits their direct use by practitioners as opposed to the graphical notation and intuitive flow-semantics of e.g. Petri Nets. In the last decade a number of general and *extensible* semantical frameworks have been developed which may combine the benefits from several models, e.g. Petri Nets and process calculi. *Bigraphical reactive systems* developed by Milner [22, 30, 21, 32] is a recent promising framework that can be viewed as a specialization of the well-established extensible semantical framework of graph-rewriting [35] aimed at mobile communicating processes. Indeed, Milner and his PhD student Høgh-Jensen have described Petri Nets and the pi-calculus as bigraphical reactive systems [31, 20]. Bundgaard and Hildebrandt have shown that also higher-order process calculi for mobility can be described as bigraphical reactive systems [6]. This provides a good starting point for an extensible formalisation of exchange and manipulation of process descriptions and active cases.

Another promising extensible framework is the very rich theory of *presheaf models for concurrency* developed by Winskel and his PhD students as part of a larger framework of categorical models for concurrency [48]. Again, semantics of Petri Nets and the pi-calculus have been provided within this framework [8], as well as of more recent higher-order calculi for mobility [49]. Presheaf models suffer even more than process calculi from the present lack of an operational and graphical presentation. Research in the combination of presheaf models and bigraphical reactive systems could provide an enrichment of both frameworks.

Formal analysis and verification Besides being the basis for an extensible architecture, a formalisation of business processes opens up for the use of automatic and formal reasoning and analysis techniques, e.g. for predicting the behaviour of business processes. Advanced formal reasoning techniques and verification tools do exist for models such as Coloured Petri Nets [19], the pi-calculus [47] as well as quantitative models such as timed [43] and probabilistic automata [24]. However, due to the lacking formalisation of business processes the practical use of these tools within this area is still very limited [45, 33].

We aim to bridge the gap between current tools for formal analysis of concurrent interacting systems and real uses of business processes. In particular we aim to combine existing tools for verification of mobile, probabilistic and timed processes, notably the mobiliy workbench [47], PRISM [24, 28] and UPPAAL tools [43]. The *quantitative analysis* afforded by probabilistic model checking techniques are lacking for business processes and are key to provide for instance best-case and worst-case scenarios for processes where the outcomes of intermediate actions are only known with a certain probability.

2.2 Architectures for pervasive business processes

The centralised client-server architecture for business process management systems illustrated in Fig. 1 above is still dominating the landscape. The current business process languages all have an XML representation, which is the de-facto standard for exchange and sharing of semi-structured data between it-systems. However, the actual use of the XML representation is limited. In particular, the state of the running cases is persisted in proprietary formats in a (central) relational database. This limits the mobility of running cases between systems. Moreover, it enlarges the distance between the process description and the state of the running process which becomes complex to read and understand by anyone else than programmers.

Reactive XML as an extensible middleware for pervasive business processes Hildebrandt, Niss and their MSc student Martin Olsen have developed a rudimentary implementation of a coordination middleware called *Reactive XML* [12] that exploits XML as format for exchange and sharing of not only data and process descriptions but also *the state of active processes*. This is much more flexible than storing or exchanging the state of active processes by relational databases or e.g. as Java and C# bytecode. Moreover, Reactive XML is directly based on the extensible formal model of bigraphical reactive systems. Currently Martin Olsen is working on an implementation of a subset of BPEL4WS in Reactive XML as part of his MSc thesis. The basic idea is to represent active BPEL processes in an XML format close to BPEL but also containing run-time information. The processes is then persisted in an XML database and executed by *rewriting the XML* using so-called *reaction rules* which are also written in XML. Reaction rules can be regarded as simple transformations akin to XSLT. As for XSLT, the implementation of the rewriting logic is generic and only needs to be implemented once. This means that the language can be extended or changed simply by providing a new or changed set of reaction rules.

Value-orientation and Peer-to-peer storage The prototype of Reactive XML exploits a novel *value-oriented architecture* for a distributed *peer-to-peer XML database* called *XMLstore* developed at ITU and DIKU [9]. The XMLstore provides *transparent sharing of XML data* between all participating peers. Value-oriented storage means that no stored values are ever updated or deleted. Instead, new values are created and identical sub-documents are aggressively shared between documents. The technique results in a semantically well-understood storage model and the *efficient storage of complete histories of process executions and data*. Besides the explicit uses of historical data in workflow management systems our work has shown that it may be used for optimistic concurrency control and conflict resolution between concurrent changes of data and processes in a peer-to-peer network [12]. This would allow active cases to migrate between clients and allow a higher degree of concurrency (and thus flexibility) than present workflow management systems. We believe that a peer-to-peer architecture may be well suited for disconnected operation in mobile ad-hoc systems, in particular where large immutable sets of data is shared, as for instance *historical medical data* or *sensor data*, and conflicting updates are rare. This may be achieved in systems with a clear separation between the roles of the participants, as for instance between doctors and nurses at a hospital.

Disconnected operation has been studied widely in the area of operating systems [36, 23]. Such work has mainly focused on making *data*, for example files, available to a disconnected device and synchronizing changes, potentially also by other disconnected devices, to such data items when the device becomes connected again. Some work has been carried out on supporting disconnected operation of WFMSs, see e.g. [1, 25]. It appears that most approaches has been based on a pessimistic locking strategy for concurrency control [1, 25] which limits the degree of concurrency of coworkers. No approaches have been based on formal semantics of mobile processes.

Context-dependency Benefitting from research within related projects at the Laboratory for Context-dependent Mobile communication at ITU in which Hildebrandt, Niss and Schmidt participates we aim to propose novel but well-founded extensions to business process language to express context-dependency. In particular, programming languages for context-aware computing is studied as part of a project on Bigraphical Programming Languages [4] described briefly below.

2.3 Hospital work as test case

Insight in the practical use and development of it-systems are necessary to steer the development of formal models. Research in computer supported cooperative work has for the last decade studied the use of shared information and coordination in cooperative work [37]. Hospital work has been identified as a challenging case for a safe, extensible, mobile and context-dependent WFMS [3, 34]. Hospital work is highly and increasingly cooperative, involving a wide and expanding array of clinical and technical specialities. At the same time, hospital work is generally characterized by being location bound; that is, the location of the activities of clinical staff is essentially determined by the location of individual patients and often also by immobile equipment at specialized locations. Thus, the spatio-temporal alignment of mobile and immobile resources (patients, doctors, nurses, technical specialists, beds, tools, equipment), i.e., their configuration at a certain place at a certain time, is a key issue in coordinating hospital work. Moreover, the configuration is highly dynamic: New tools and equipment are continuously developed, along with new constraints and procedures for their use. Finally, the application domain is highly safety-critcal.

3 Related activities

The project relates to several research projects at ITU. Firstly, the project provides a bridge between the Bigraphical Programming Languages (BPL) project and the NEXT research project. The BPL project at ITU aims in collaboration with Robin Milner at developing programming languages for global ubiquitous mobile systems founded on the theoretical model of bigraphical reactive systems. The NEXT project is a research project on the software architectures for the next generation of business software in collaboration with Department of Computer Science, Copenhagen University (DIKU) and Microsoft Business Solutions (MBS). Secondly, ongoing CSCW research projects undertaken by ITU faculty, MSc- and PhD-students with Schmidt as one of the central researchers will provide cases against which the formalisations and architectures developed will be measured. The projects involve, besides MBS, hospitals and banks in Copenhagen.

On a regional and national level, the project will collaborate with the Plan-X project which is a collaboration between DIKU and ITU and coordinates the research and development of the XML Store framework. It will also involve collaboration with research groups at the Department of Computer Science, University of Aarhus (DAIMI) working on XML-technologies and process models.

On an international level, the project relates directly to the agenda of the Process Modelling Group and the UK Grand Challenge on Global Ubiquitous Computing [15] led by prof. Robin Milner and prof. Marta Kwiatkowska.

References

[1] G. Alonso, M. Günthör, M. Kamath, D. Agrawa, A. El Abbadi, and C. Mohan. Exotica/FDMC: Handling disconnected clients in a workflow management system. In *Proceedings of the Third International Conference on Cooperative Information Systems*, pages 99–110, Vienna, May 1995.

[2] Tony Andrews, Francisco Curbera, Hitesh Dholakia, Yaron Goland, Johannes Klein, Frank Leymann, Kevin Liu, Dieter Roller, Doug Smith, Satish Thatte, Ivana Trickovic, and Sanjiva Weerawarana. Business process execution language for web services. Technical report, 2003.

[3] Jakob E. Bardram and Claus Bossen. Moving to get aHead: Local mobility and collaborative work. In Kari Kuutti et al., editor, *ECSCW 2003: Proceedings of the Eighth European Conference on Computer Supported Cooperative Work*, pages 355–374. Kluwer Academic Publishers, September 2003.

[4] Lars Birkedal, Søren Debois, Ebbe Elsborg, Thomas Hildebrandt, and Henning Niss. Bigraphical models of context-aware systems. Technical Report ITU TR-2005-74, IT University of Copenhagen, November 2005. ISSN 1600-6100.

[5] The Business Process Management Initiative (BMPI). Business process modeling notation (BPMN) version 1.0, may 2004.

[6] Mikkel Bundgaard and Thomas Hildebrandt. Bigraphical semantics of higher-order mobile embedded resources with local names. In *Proceedings of the 1st International Workshop on Graph Transformation for Verification and Concurrency (GT-VC '05)*, 2005. Accepted for publication.

[7] Mikkel Bundgaard, Thomas Hildebrandt, and Jens Chr. Godskesen. A CPS encoding of name-passing in higher-order mobile embedded resources. *Theoretical Computer Science*, 2005. Accepted for publication in a special issue of TCS.

[8] G.-L. Cattani, I. Stark, and G. Winskel. A presheaf model for the pi-calculus. In *In the proc. of Category Theory in Computer Science 1997*, Springer Lecture Notes in Computer Science, Genova, 1997.

[9] Fritz Henglein and Henning Niss. Plan-X webpage, 2005. (XML Store is part of the Plan-X Project).

[10] Thomas Hildebrandt and Jens Chr. Godskesen. Extending Howe's method to typed higher-order mobile embedded resources. In *Proceedings of 25th International conference on FST&TCS*, 2005.

[11] Thomas Hildebrandt, Jens Chr. Godskesen, and Mikkel Bundgaard. Bisimulation congruences for Homer — a calculus of higher order mobile embedded resources. Technical Report TR-2004-52, IT University of Copenhagen, 2004.

[12] Thomas Hildebrandt, Henning Niss, Martin Olsen, and Jacob W. Winther. Distributed reactive XML. In Lubos Brim and Isabelle Linden, editors, *Proceedings of the 1st International Workshop on Methods and Tools for Coordinating Concurrent, Distributed and Mobile Systems (MTCoord05)*, 2005.

[13] Thomas Hildebrandt, Prakash Panangaden, and Glynn Winskel. A relational model of non-deterministic dataflow. *Journal of Mathematical Structures in Computer Science*, 14:613–628, 2004.

[14] Thomas Troels Hildebrandt. *Categorical Models for Concurrency: Independence, Fairness and Dataflow.* BRICS dissertation series, Ph. D. Thesis, University of Aarhus, February 2000.

[15] Tony Hoare and Robin Milner. Grand challenges in computing. Technical report, The British Computer Society, 2004.

[16] David Hollingsworth. Workflow management coalition: The workflow reference model. Technical Report TC00-1003, Workflow Management Coalition, 1995.

[17] Jörn W. Janneck and Robert Esser. Higher-order Petri net modelling: techniques and applications. In *CRPITS '12: Proceedings of the conference on Application and theory of petri nets*, pages 17–25, Darlinghurst, Australia, Australia, 2002. Australian Computer Society, Inc.

[18] Dan Jensen. Danske milliard investeringer i ERP på vej. *Computerworld*, August 19th 2004.

[19] K. Jensen. *Coloured Petri Nets. Basic Concepts, Analysis Methods and Practical Use, Volume 1-3.* Monographs in Theoretical Computer Science. Springer-Verlag, 1997. 2nd corrected printing.

[20] Ole Høgh Jensen. *Mobile Processes in Bigraphs.* PhD thesis, University of Aalborg, 2005. Forthcoming.

[21] Ole Høgh Jensen and Robin Milner. Bigraphs and transitions. In *Proceedings of the 30th ACM SIGPLAN-SIGACT symposium on Principles of programming languages (POPL)*, pages 38–49. ACM Press, 2003.

[22] Ole Høgh Jensen and Robin Milner. Bigraphs and mobile processes (revised). Technical Report UCAM-CL-TR-580, University of Cambridge – Computer Laboratory, February 2004.

[23] J. J. Kistler and M. Satyanarayanan. Disconnected operation in the Coda File System. In *Thirteenth ACM Symposium on Operating Systems Principles*, volume 25(5), pages 213–225. ACM Press, 1991.

[24] M. Kwiatkowska, G. Norman, and D. Parker. Prism 2.0: A tool for probabilistic model checking. In *In Proc. 1st International Conference on Quantitative Evaluation of Systems (QEST'04)*, pages 322–323. IEEE Computer Society Press, September 2004.

[25] Seungil Lee, Dongsoo Han, and Dongman Lee. Supporting voluntary disconnection in WFMSs. In *Proceedings of the Third International Symposium on Cooperative Database Systems for Advanced Applications (CODAS 2001)*, pages 147– 154, Beijing, China, April 2001.

[26] Ulf Leonhardt. *Supporting Location-Awareness in Open Distributed Systems.* Ph. D. Thesis, Imperial College, University of London, May 1998.

[27] Frank Leymann. Web services flow language. Technical report, 2001.

[28] G. Norman M. Kwiatkowska and D. Parker. Probabilistic symbolic model checking with prism: A hybrid approach. *International Journal on Software Tools for Technology Transfer (STTT)*, 6(2):128–142, September 2004.

[29] Robin Milner. *Communicating and Mobile Systems: the π-Calculus*. Cambridge University Press, 1993.

[30] Robin Milner. Axioms for bigraphical structure. Technical Report UCAM-CL-TR-581, University of Cambridge – Computer Laboratory, February 2004.

[31] Robin Milner. Bigraphs for Petri Nets. In Jörg Desel, Wolfgang Reisig, and Grzegorz Rozenberg, editors, *Lectures on Concurrency and Petri Nets: Advances in Petri Nets*, volume 3098 of *Lecture Notes in Computer Science*, pages 686–701. Springer-Verlag, 2004.

[32] Robin Milner. Pure bigraphs: Structure and dynamics. *Information and Computation*, 2005.

[33] Chun Ouyang, Wil M.P. van der Aalst, Stephen Breutel, Marlon Dumas, Arthur H.M. ter Hofstede, and Eric Verbeek. Formal semantics and analysis of control flow in WS-BPEL. Technical Report BPM-05-13, 2005.

[34] Madhu C. Reddy, Paul Dourish, and Wanda Pratt. Coordinating heterogeneous work: Information and representation in medical care. In Wolfgang Prinz et al., editor, *ECSCW 2001: Proceedings of the Seventh European Conference on Computer Supported Cooperative Work*, pages 239–258. Kluwer Academic Publishers, September 2001.

[35] Grzegorz Rozenberg, editor. *Handbook of Graph Grammars and Computing by Graph Transformations, Volume 1: Foundations*. World Scientific, 1997.

[36] Yasushi Saito and Marc Shapiro. Optimistic replication. *ACM Computing Surveys*, 37(1):42–81, 2005.

[37] Kjeld Schmidt and Liam J. Bannon. Taking CSCW seriously: Supporting articulation work. *Computer Supported Cooperative Work (CSCW): An International Journal*, 1(1-2):7–40, 1992.

[38] Howard Smith and Peter Fingar. *Business Process Management: The Third Wave*. Meghan-Kiffer Press, 2003.

[39] Howard Smith and Peter Fingar. Workflow is just a Pi process. *BPTrends*, january 2004.

[40] Christrian Stefansen. A declarative framework for enterprise information systems, 2005. Qualification Report.

[41] Satish Thatte. XLANG: Web services for business process design. Technical report, 2001.

[42] B. Thomsen. Plain CHOCS. *Acta Informatica*, 30:1–59, 1993.

[43] UPPAAL home. www.uppaal.com.

[44] Wil van der Aalst and Kees van Hee. *Workflow Management*. Cooperative Information System Series. MIT Press, 2002.

[45] Wil van der Aalst and Kees van Hee. *Workshop Proceedings, The Process Modelling Group*. Technische Universiteit Eindhoven, 6.-8. June 2005. www.process-modelling-group.org.

[46] Will M.P. van der Aalst. Why workflow is not just a pi-process. *BPTrends*, february 2004.

[47] Björn Victor and Faron Moller. The mobility workbench: A tool for the pi-calculus). Technical Report LFCS report ECS-LFCS-94-285, Laboratory for Foundations of Computer Science, February 1994.

[48] Glynn Winskel and Mogens Nielsen. Models for concurrency. In S. Abramsky, D. Gabbay, and T. S. E. Maibaum, editors, *Handbook of Logic in Computer Science*, volume IV. Oxford University Press, 1995.

[49] Glynn Winskel and Francesco Zappa Nardelli. New-hopla: a higher order process language with name generation. In *TCS2004, Third IFIP International Conference on Theoretical Computer Science*, 2004.