

David E. Lightfoot
Clemens A. Szyperski (Eds.)

LNCS 4228

Modular Programming Languages

7th Joint Modular Languages Conference, JMLC 2006
Oxford, UK, September 2006
Proceedings



Springer

David E. Lightfoot Clemens A. Szyperski (Eds.)

Modular Programming Languages

7th Joint Modular Languages Conference, JMLC 2006
Oxford, UK, September 13-15, 2006
Proceedings

Volume Editors

David E. Lightfoot
Oxford Brookes University
School of Technology
Department of Computing
Oxford, OX33 1HX, UK
E-mail: dlightfoot@brookes.ac.uk

Clemens A. Szyperski
One Microsoft Way
Redmond WA 98052, USA
E-mail: cszypers@microsoft.com

Library of Congress Control Number: 2006932028

CR Subject Classification (1998): D.3, D.2, D.1, D.4, F.3

LNCS Sublibrary: SL 1 – Theoretical Computer Science and General Issues

ISSN	0302-9743
ISBN-10	3-540-40927-0 Springer Berlin Heidelberg New York
ISBN-13	978-3-540-40927-4 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media

springer.com

© Springer-Verlag Berlin Heidelberg 2006
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 11860990 06/3142 5 4 3 2 1 0

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, UK

Takeo Kanade

Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler

University of Surrey, Guildford, UK

Jon M. Kleinberg

Cornell University, Ithaca, NY, USA

Friedemann Mattern

ETH Zurich, Switzerland

John C. Mitchell

Stanford University, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

Oscar Nierstrasz

University of Bern, Switzerland

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

University of Dortmund, Germany

Madhu Sudan

Massachusetts Institute of Technology, MA, USA

Demetri Terzopoulos

University of California, Los Angeles, CA, USA

Doug Tygar

University of California, Berkeley, CA, USA

Moshe Y. Vardi

Rice University, Houston, TX, USA

Gerhard Weikum

Max-Planck Institute of Computer Science, Saarbruecken, Germany

Preface

On behalf of the Steering Committee we are pleased to present the proceedings of the 2006 Joint Modular Languages Conference (JMLC), organized by Oxford Brookes University, Oxford, UK and held at Jesus College, Oxford. The mission of JMLC is to explore the concepts of well-structured programming languages and software and those of teaching good design and programming style. JMLC 2006 was the seventh in a series of successful conferences with themes including the construction of large and distributed software systems, and software engineering aspects in new and dynamic application areas.

We were fortunate to have a dedicated Program Committee comprising 41 internationally recognized researchers and industrial practitioners. We received 36 submissions and each paper was reviewed by at least three Program Committee members (four for papers with an author on the Program Committee). The entire reviewing process was supported by the OpenConf system. In total, 23 submissions were accepted along with two invited papers and are included in this proceedings volume.

For the successful local organization of JMLC we thank Muneera Masterson, Ali McNiffe and Fiona Parker and other staff and student helpers of Oxford Brookes University as well as Rosemary Frame and Jo Knighton and other staff of Jesus College, Oxford. The proceedings you now hold were published by Springer and we are grateful for their support. Finally, we must thank the many authors who contributed the high-quality papers contained within these proceedings.

September 2006

David Lightfoot
Clemens Szyperski

Organization

JMLC 2006 was the seventh conference in a successful series, with past events held in:

1987 in Bled, Slovenia;
1990 in Loughborough, UK;
1994 in Ulm, Germany;
1997 in Linz, Austria;
2000 in Zürich, Switzerland;
2003 in Klagenfurt, Austria.

Steering Committee

László Böszörményi, University of Klagenfurt, Austria
Michael Franz, UC Irvine, USA
Jürg Gutknecht, ETH Zürich, Switzerland
David Lightfoot, Oxford Brookes University, UK (Program Co-chair
and Local Organizer)
Hanspeter Mössenböck, University of Linz, Austria
Clemens Szyperski, Microsoft, USA (Program Co-chair)
Niklaus Wirth, ETH Zürich emeritus, Switzerland

Program Committee

Jonathan Aldrich, CMU, USA
Pierre America, Philips Research, Netherlands
Uwe Assmann, TU Dresden, Germany
Nick Benton, Microsoft Research Cambridge, UK
László Böszörményi, University of Klagenfurt, Austria
Gilad Bracha, Sun Java Software, USA
Michael E. Caspersen, Aarhus University, Denmark
Craig Chambers, University of Washington, USA
Michael Franz, UC Irvine, USA
K. John Gough, Queensland UT, Australia
Dominik Gruntz, Fachhochschule Aargau, Switzerland
Jürg Gutknecht, ETH Zürich, Switzerland
Thomas Henzinger, EPF Lausanne, Switzerland
Nigel Horspool, University of Victoria, Canada
Zoltán Horváth, Budapest University (ELTE), Hungary
Mehdi Jazayeri, TU Vienna, Austria
Helmut Ketz, Fachhochschule Reutlingen, Germany

VIII Organization

Brian Kirk, Robinson Associates, UK
Christoph Kirsch, University of Salzburg, Austria
Jens Knoop, TU Vienna, Austria
Kai Koskimies, TU Tampere, Finland
Liu Ling, University of Shanghai, China
Jochen Ludewig, University of Stuttgart, Germany
Jan Madey, University of Warsaw, Poland
Ole Lehrmann Madsen, Aarhus University, Denmark
Roland Mittermeir, University of Klagenfurt, Austria
Hanspeter Mössenböck, University of Linz, Austria
Pieter Muller, Esmertec, Switzerland
Judit Nyeky, Budapest University (ELTE), Hungary
Martin Odersky, EPF Lausanne, Switzerland
Jens Palsberg, Purdue University, USA
Frank Peschel-Gallee, Microsoft, USA
Gustav Pomberger, University of Linz, Austria
Wolfgang Pree, University of Salzburg, Austria
Paul Reed, Padded Cell Software, UK
Paul Roe, Queensland UT, Australia
Markus Schordan, TU Vienna, Austria
Brian Shearing, The Software Factory, UK
Pat Terry, Rhodes University, South Africa
Fyodor Tkachov, Institute for Nuclear Research (RAS), Russia
Wolfgang Weck, Independent Software Architect, Switzerland
Mark Woodman, Middlesex University, UK

Additional Reviewer

Christian Wimmer, University of Linz, Austria

Sponsoring Institutions

Microsoft Research
Sun Microsystems
Robinson Associates
dpunkt Verlag

Lecture Notes in Computer Science

For information about Vols. 1–4077

please contact your bookseller or Springer

Vol. 4228: D.E. Lightfoot, C.A. Szyperski (Eds.), *Modular Programming Languages*. X, 415 pages. 2006.

Vol. 4206: P. Dourish, A. Friday (Eds.), *UbiComp 2006: Ubiquitous Computing*. XIX, 526 pages. 2006.

Vol. 4193: T.P. Runarsson, H.-G. Beyer, E. Burke, J.J. Merelo-Guervós, L. D. Whitley, X. Yao (Eds.), *Parallel Problem Solving from Nature – PPSN IX*. XIX, 1061 pages. 2006.

Vol. 4192: B. Mohr, J.L. Traeff, J. Worringer, J. Dongarra (Eds.), *Recent Advances in Parallel Virtual Machine and Message Passing Interface*. XVI, 414 pages. 2006.

Vol. 4188: P. Sojka, I. Kopeček, K. Pala (Eds.), *Text, Speech and Dialogue*. XIV, 721 pages. 2006. (Sublibrary LNAI).

Vol. 4186: C. Jesshope, C. Egan (Eds.), *Advances in Computer Systems Architecture*. XIV, 605 pages. 2006.

Vol. 4185: R. Mizoguchi, Z. Shi, F. Giunchiglia (Eds.), *The Semantic Web – ASWC 2006*. XX, 778 pages. 2006.

Vol. 4184: M. Bravetti, M. Nuñez, G. Zavattaro (Eds.), *Web Services and Formal Methods*. X, 289 pages. 2006.

Vol. 4180: M. Kohlhase, *OMDoc – An Open Markup Format for Mathematical Documents [version 1.2]*. XIX, 428 pages. 2006. (Sublibrary LNAI).

Vol. 4178: A. Corradini, H. Ehrig, U. Montanary, L. Ribeiro, G. Rozenberg (Eds.), *Graph Transformations*. XII, 473 pages. 2006.

Vol. 4176: S.K. Katsikas, J. Lopez, M. Backes, S. Gritzalis, B. Preneel (Eds.), *Information Security*. XIV, 548 pages. 2006.

Vol. 4175: P. Bücher, B.M.E. Moret (Eds.), *Algorithms in Bioinformatics*. XII, 402 pages. 2006. (Sublibrary LNBI).

Vol. 4169: H.L. Bodlaender, M.A. Langston (Eds.), *Parameterized and Exact Computation*. XI, 279 pages. 2006.

Vol. 4168: Y. Azar, T. Erlebach (Eds.), *Algorithms – ESA 2006*. XVIII, 843 pages. 2006.

Vol. 4165: W. Jonker, M. Petkovic (Eds.), *Secure, Data Management*. X, 185 pages. 2006.

Vol. 4163: H. Bersini, J. Carneiro (Eds.), *Artificial Immune Systems*. XII, 460 pages. 2006.

Vol. 4162: R. Kráľovič, P. Urzyczyn (Eds.), *Mathematical Foundations of Computer Science 2006*. XV, 814 pages. 2006.

Vol. 4159: J. Ma, H. Jin, L.T. Yang, J.J.-P. Tsai (Eds.), *Ubiquitous Intelligence and Computing*. XXII, 1190 pages. 2006.

Vol. 4158: L.T. Yang, H. Jin, J. Ma, T. Ungerer (Eds.), *Autonomic and Trusted Computing*. XIV, 613 pages. 2006.

Vol. 4156: S. Amer-Yahia, Z. Bellahsene, E. Hunt, R. Unland, J.X. Yu (Eds.), *Database and XML Technologies*. IX, 123 pages. 2006.

Vol. 4155: O. Stock, M. Schaerf (Eds.), *Reasoning, Action and Interaction in AI Theories and Systems*. XVIII, 343 pages. 2006. (Sublibrary LNAI).

Vol. 4153: N. Zheng, X. Jiang, X. Lan (Eds.), *Advances in Machine Vision, Image Processing, and Pattern Analysis*. XIII, 506 pages. 2006.

Vol. 4152: Y. Manolopoulos, J. Pokorný, T. Sellis (Eds.), *Advances in Databases and Information Systems*. XV, 448 pages. 2006.

Vol. 4151: A. Iglesias, N. Takayama (Eds.), *Mathematical Software – ICMS 2006*. XVII, 452 pages. 2006.

Vol. 4150: M. Dorigo, L.M. Gambardella, M. Birattari, A. Martinoli, R. Poli, T. Stützle (Eds.), *Ant Colony Optimization and Swarm Intelligence*. XVI, 526 pages. 2006.

Vol. 4149: M. Klusch, M. Rovatsos, T.R. Payne (Eds.), *Cooperative Information Agents X*. XII, 477 pages. 2006. (Sublibrary LNAI).

Vol. 4148: J. Vounckx, N. Azemard, P. Maurine (Eds.), *Integrated Circuit and System Design*. XVI, 677 pages. 2006.

Vol. 4146: J.C. Rajapakse, L. Wong, R. Acharya (Eds.), *Pattern Recognition in Bioinformatics*. XIV, 186 pages. 2006. (Sublibrary LNBI).

Vol. 4144: T. Ball, R.B. Jones (Eds.), *Computer Aided Verification*. XV, 564 pages. 2006.

Vol. 4139: T. Salakoski, F. Ginter, S. Pyysalo, T. Pahikkala, *Advances in Natural Language Processing*. XVI, 771 pages. 2006. (Sublibrary LNAI).

Vol. 4138: X. Cheng, W. Li, T. Znati (Eds.), *Wireless Algorithms, Systems, and Applications*. XVI, 709 pages. 2006.

Vol. 4137: C. Baier, H. Hermanns (Eds.), *CONCUR 2006 – Concurrency Theory*. XIII, 525 pages. 2006.

Vol. 4136: R.A. Schmidt (Ed.), *Relations and Kleene Algebra in Computer Science*. XI, 433 pages. 2006.

Vol. 4135: C.S. Calude, M.J. Dinneen, G. Păun, G. Rozenberg, S. Stepney (Eds.), *Unconventional Computation*. X, 267 pages. 2006.

Vol. 4134: K. Yi (Ed.), *Static Analysis*. XIII, 443 pages. 2006.

Vol. 4133: J. Gratch, M. Young, R. Aylett, D. Ballin, P. Olivier (Eds.), *Intelligent Virtual Agents*. XIV, 472 pages. 2006. (Sublibrary LNAI).

Vol. 4132: S. Kollias, A. Stafylopatis, W. Duch, E. Oja (Eds.), *Artificial Neural Networks – ICANN 2006, Part II*. XXXIV, 1028 pages. 2006.

- Vol. 4131: S. Kollias, A. Stafylopatis, W. Duch, E. Oja (Eds.), *Artificial Neural Networks – ICANN 2006, Part I*. XXXIV, 1008 pages. 2006.
- Vol. 4130: U. Furbach, N. Shankar (Eds.), *Automated Reasoning*. XV, 680 pages. 2006. (Sublibrary LNAI).
- Vol. 4129: D. McGookin, S. Brewster (Eds.), *Haptic and Audio Interaction Design*. XII, 167 pages. 2006.
- Vol. 4128: W.E. Nagel, W.V. Walter, W. Lehner (Eds.), *Euro-Par 2006 Parallel Processing*. XXXIII, 1221 pages. 2006.
- Vol. 4127: E. Damiani, P. Liu (Eds.), *Data and Applications Security XX*. X, 319 pages. 2006.
- Vol. 4126: P. Barahona, F. Bry, E. Franconi, N. Henze, U. Sattler, *Reasoning Web*. X, 269 pages. 2006.
- Vol. 4124: H. de Meer, J.P. G. Sterbenz (Eds.), *Self-Organizing Systems*. XIV, 261 pages. 2006.
- Vol. 4121: A. Biere, C.P. Gomes (Eds.), *Theory and Applications of Satisfiability Testing - SAT 2006*. XII, 438 pages. 2006.
- Vol. 4119: C. Dony, J.L. Knudsen, A. Romanovsky, A. Tripathi (Eds.), *Advanced Topics in Exception Handling Components*. X, 302 pages. 2006.
- Vol. 4117: C. Dwork (Ed.), *Advances in Cryptology - CRYPTO 2006*. XIII, 621 pages. 2006.
- Vol. 4116: R. De Prisco, M. Yung (Eds.), *Security and Cryptography for Networks*. XI, 366 pages. 2006.
- Vol. 4115: D.-S. Huang, K. Li, G.W. Irwin (Eds.), *Computational Intelligence and Bioinformatics, Part III*. XXI, 803 pages. 2006. (Sublibrary LNBFI).
- Vol. 4114: D.-S. Huang, K. Li, G.W. Irwin (Eds.), *Computational Intelligence, Part II*. XXVII, 1337 pages. 2006. (Sublibrary LNAI).
- Vol. 4113: D.-S. Huang, K. Li, G.W. Irwin (Eds.), *Intelligent Computing, Part I*. XXVII, 1331 pages. 2006.
- Vol. 4112: D.Z. Chen, D. T. Lee (Eds.), *Computing and Combinatorics*. XIV, 528 pages. 2006.
- Vol. 4111: F.S. de Boer, M.M. Bonsangue, S. Graf, W.-P. de Roever (Eds.), *Formal Methods for Components and Objects*. VIII, 447 pages. 2006.
- Vol. 4110: J. Díaz, K. Jansen, J.D.P. Rolim, U. Zwick (Eds.), *Approximation, Randomization, and Combinatorial Optimization*. XII, 522 pages. 2006.
- Vol. 4109: D.-Y. Yeung, J.T. Kwok, A. Fred, F. Roli, D. de Ridder (Eds.), *Structural, Syntactic, and Statistical Pattern Recognition*. XXI, 939 pages. 2006.
- Vol. 4108: J.M. Borwein, W.M. Farmer (Eds.), *Mathematical Knowledge Management*. VIII, 295 pages. 2006. (Sublibrary LNAI).
- Vol. 4106: T.R. Roth-Berghofer, M.H. Göker, H. A. Güvenir (Eds.), *Advances in Case-Based Reasoning*. XIV, 566 pages. 2006. (Sublibrary LNAI).
- Vol. 4105: B. Günsel, A.K. Jain, A. M. Tekalp, B. Sankur (Eds.), *Multimedia, Content Representation, Classification and Security*. XIX, 804 pages. 2006.
- Vol. 4104: T. Kunz, S.S. Ravi (Eds.), *Ad-Hoc, Mobile, and Wireless Networks*. XII, 474 pages. 2006.
- Vol. 4103: J. Eder, S. Dustdar (Eds.), *Business Process Management Workshops*. XI, 508 pages. 2006.
- Vol. 4102: S. Dustdar, J.L. Fiadeiro, A. Sheth (Eds.), *Business Process Management*. XV, 486 pages. 2006.
- Vol. 4099: Q. Yang, G. Webb (Eds.), *PRICAI 2006: Trends in Artificial Intelligence*. XXVIII, 1263 pages. 2006. (Sublibrary LNAI).
- Vol. 4098: F. Pfenning (Ed.), *Term Rewriting and Applications*. XIII, 415 pages. 2006.
- Vol. 4097: X. Zhou, O. Sokolsky, L. Yan, E.-S. Jung, Z. Shao, Y. Mu, D.C. Lee, D. Kim, Y.-S. Jeong, C.-Z. Xu (Eds.), *Emerging Directions in Embedded and Ubiquitous Computing*. XXVII, 1034 pages. 2006.
- Vol. 4096: E. Sha, S.-K. Han, C.-Z. Xu, M.H. Kim, L.T. Yang, B. Xiao (Eds.), *Embedded and Ubiquitous Computing*. XXIV, 1170 pages. 2006.
- Vol. 4095: S. Nolfi, G. Baldassare, R. Calabretta, D. Marocco, D. Parisi, J.C. T. Hallam, O. Miglino, J.-A. Meyer (Eds.), *From Animals to Animats 9*. XV, 869 pages. 2006. (Sublibrary LNAI).
- Vol. 4094: O. H. Ibarra, H.-C. Yen (Eds.), *Implementation and Application of Automata*. XIII, 291 pages. 2006.
- Vol. 4093: X. Li, O.R. Zaïane, Z. Li (Eds.), *Advanced Data Mining and Applications*. XXI, 1110 pages. 2006. (Sublibrary LNAI).
- Vol. 4092: J. Lang, F. Lin, J. Wang (Eds.), *Knowledge Science, Engineering and Management*. XV, 664 pages. 2006. (Sublibrary LNAI).
- Vol. 4091: G.-Z. Yang, T. Jiang, D. Shen, L. Gu, J. Yang (Eds.), *Medical Imaging and Augmented Reality*. XIII, 399 pages. 2006.
- Vol. 4090: S. Spaccapietra, K. Aberer, P. Cudré-Mauroux (Eds.), *Journal on Data Semantics VI*. XI, 211 pages. 2006.
- Vol. 4089: W. Löwe, M. Südholt (Eds.), *Software Composition*. X, 339 pages. 2006.
- Vol. 4088: Z.-Z. Shi, R. Sadananda (Eds.), *Agent Computing and Multi-Agent Systems*. XVII, 827 pages. 2006. (Sublibrary LNAI).
- Vol. 4087: F. Schwenker, S. Marinai (Eds.), *Artificial Neural Networks in Pattern Recognition*. IX, 299 pages. 2006. (Sublibrary LNAI).
- Vol. 4085: J. Misra, T. Nipkow, E. Sekerinski (Eds.), *FM 2006: Formal Methods*. XV, 620 pages. 2006.
- Vol. 4084: M.A. Wimmer, H.J. Scholl, Å. Grönlund, K.V. Andersen (Eds.), *Electronic Government*. XV, 353 pages. 2006.
- Vol. 4083: S. Fischer-Hübner, S. Furnell, C. Lambri-noudakis (Eds.), *Trust and Privacy in Digital Business*. XIII, 243 pages. 2006.
- Vol. 4082: K. Bauknecht, B. Pröll, H. Werthner (Eds.), *E-Commerce and Web Technologies*. XIII, 243 pages. 2006.
- Vol. 4081: A. M. Tjoa, J. Trujillo (Eds.), *Data Warehousing and Knowledge Discovery*. XVII, 578 pages. 2006.
- Vol. 4080: S. Bressan, J. Küng, R. Wagner (Eds.), *Database and Expert Systems Applications*. XXI, 959 pages. 2006.
- Vol. 4079: S. Etalle, M. Truszczyński (Eds.), *Logic Programming*. XIV, 474 pages. 2006.

Table of Contents

Separating Concerns with Domain Specific Languages	1
<i>Steve Cook</i>	
Event-Based Programming Without Inversion of Control	4
<i>Philipp Haller, Martin Odersky</i>	
Programming Language Concepts for Multimedia Application Development	23
<i>Oliver Lampl, Elmar Stellnberger, László Böszörményi</i>	
Implicit and Dynamic Parameters in C++	37
<i>Christian Heinlein</i>	
Reconciling Virtual Classes with Genericity	57
<i>Erik Ernst</i>	
Oberon Script: A Lightweight Compiler and Runtime System for the Web	73
<i>Ralph Sommerer</i>	
Efficient Layer Activation for Switching Context-Dependent Behavior	84
<i>Pascal Costanza, Robert Hirschfeld, Wolfgang De Meuter</i>	
Object-Oriented Language Processing	104
<i>Pietu Pohjalainen</i>	
A Framework for Modular Linking in OO Languages	116
<i>Sean McDirmid, Wilson C. Hsieh, Matthew Flatt</i>	
Flexible Type-Safe Linking of Components for Java-Like Languages	136
<i>Davide Ancona, Giovanni Lagorio, Elena Zucca</i>	
Towards a Formal Semantics for AspectJ Weaving	155
<i>Nadia Belblidia, Mourad Debbabi</i>	
Symbolic Analysis of Imperative Programming Languages	172
<i>Bernd Burgstaller, Bernhard Scholz, Johann Blieberger</i>	
Array-Structured Object Types for Mathematical Programming	195
<i>Felix Friedrich, Jürg Gutknecht</i>	

MetaModelica: A Unified Equation-Based Semantical and Mathematical
Modeling Language 211
Adrian Pop, Peter Fritzson

A Component Language for Structured Parallel Programming 230
Luc Bläser

Internal and External Token-Based Synchronization in
Object-Oriented Languages 251
Franz Puntigam

A New Component-Oriented Programming Language with the
First-Class Connector 271
Bo Chen, ZhouJun Li, HuoWang Chen

A Component Plug-In Architecture for the .NET Platform 287
*Reinhard Wolfinger, Deepak Dhungana, Herbert Prähofer,
Hanspeter Mössenböck*

Improve Component-Based Programs with Connectors 306
Joachim H. Fröhlich, Manuel Schwarzhinger

Automatic Object Colocation Based on Read Barriers 326
Christian Wimmer, Hanspeter Mössenböck

Nearly Optimal Register Allocation with PBQP 346
Lang Hames, Bernhard Scholz

Fast Profile-Based Partial Redundancy Elimination 362
R. Nigel Horspool, David J. Pereira, Bernhard Scholz

The Dining Philosophers Problem Revisited 377
Jürg Gutknecht

A Mobile Agent Service-Oriented Scripting Language Encoded
on a Process Calculus 383
Hervé Paulino, Luís Lopes

A Case Study in Concurrent Programming with Active Objects 403
Ulrike Glavitsch, Thomas M. Frey

Author Index 415

Separating Concerns with Domain Specific Languages

Steve Cook

Microsoft UK Ltd, Cambridge
steve.cook@microsoft.com

Abstract. I'll talk about the separation of concerns in the development of large distributed enterprise systems, how to manage it using domain specific languages, and how to build these languages. This brief note outlines some of the topics I'll cover.

1 Separation of Concerns

Most developments in programming language design are intended to improve the ability of the programmer to separate the expression of different concerns. This has progressively led to the development of language features such as procedures, abstract data types, objects, polymorphic type systems, aspects, and so on.

We're now moving into an era when the normal case of software development is distributed and heterogeneous, with the internet playing a pivotal role. It's simply not practical today to use a single programming language to create all aspects of a large and complex computing system. Different technologies are used to implement user-interfaces, business subsystems, middleware, databases, workflow systems, sensors, etc. Enterprise programming stacks include as first-class participants a variety of inter-related programming and scripting languages, databases, metadata and configuration files. Most programming projects involve interoperating with what is already there, which requires interfacing to existing technology stacks.

In such a world, concerns such as the structure and organization of business data and processes inherently span multiple technologies. A given business concept will show up in the user interface, in the formats used to communicate between components, in the interfaces offered from one component to another, in the schemas for databases where business data is stored, and in the programming logic for manipulating all of the above. Even the simplest change, such as changing the name of the business concept, impacts all of these pieces. Such concerns cannot possibly be effectively separated by improving programming language design. How then can we approach this problem?

2 Development Using Domain Specific Languages

A promising approach has been described variously as "Language-Oriented Programming" [1], "Language Workbenches" [2], "Generative Programming" [3] and "Model Driven Engineering" [4]. All of these phrases essentially describe the same

pattern: given a class of problems, design a special-purpose language – a Domain Specific Language or DSL - to solve it.

A simple (and old) example of this pattern is the language of regular expressions. For example, using the .Net class `System.Text.RegularExpressions.Regex`, the regular expression “`(?<user>[^\@]+)@(?<host>.+)`” applied to a string of characters will find email addresses in it, and for each address found, appropriately extract the user and host variables. Programming such a procedure directly in a general-purpose language is a significantly larger and more error-prone task.

In developing complex enterprise systems, it is increasingly the case that graphical languages can be used to express certain concerns most effectively. Business workflows, business data, and system, application and data centre configuration are obvious candidates for graphical representation. Also textual languages, while effective for inputting large programs, may not be the most effective medium for displaying, analyzing and interpreting these programs.

Putting these ingredients together provides the motivation for an emerging class of graphical language-processing tools, which includes the DSL Tools from Microsoft [5], the Generic Modeling Environment (GME) from Vanderbilt University [6], and commercial examples from MetaCase, Xactium and others. These tools enable the language author to design and implement the abstract and concrete syntax for a DSL, together with the ancillaries needed to integrate the language into a development process.

Of course it is not sufficient simply to design what a DSL looks like; it is also necessary to give its expressions meaning, which in practical terms means to generate executable artifacts from it: these will most likely be programs in more general-purpose languages, together with configuration files, scripts and metadata, that can be deployed together to implement the intention of the developer.

As soon as generation is introduced into the development process, there is the possibility of developers changing the generated artifacts. Uncontrolled, this will break the process: the source form of the DSL will become out of date and useless. Alleviating this issue is one of the main challenges of making DSLs successful. Not all artifacts can be generated from DSLs, so it is essential to be able to interface generated artifacts with hand-coded ones: various language techniques such as partial classes [7] can enable this.

3 Software Factories

A DSL can provide a means to simplify the development of one area of concern. But the development of large distributed applications involves the integration of multiple areas of concern, with multiple stakeholders manipulating the system via multiple different viewpoints.

Managing the complexity of such a development involves delivering appropriate languages, tools and guidance to individual participants in the process at the right place and time. Enabling this is the province of Software Factories [8], an approach

to software development that focuses on the explicit identification of viewpoints in the development process, the definition of DSLs, tools and guidance to support these viewpoints, and the delivery of these capabilities to individuals during the enactment of the process.

References

1. Dimitriev, S. Language-Oriented Programming: The Next Programming Paradigm, <http://www.onboard.jetbrains.com/isl/articles/04/10/lop/>
2. Fowler, M. Language Workbenches: The Killer App for Domain Specific Languages? <http://martinfowler.com/articles/languageWorkbench.html>
3. Czarnecki, K. and Eisenecker, U.W. Generative Programming – Methods, Tools and Applications. Addison-Wesley (2000).
4. Bézivin J., Jouault F, and Valduriez P. On the Need for Megamodels. Proceedings of the OOPSLA/GPCE: Best Practices for Model-Driven Software Development workshop, 19th Annual ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications (2004).
5. DSL Tools Workshop. <http://msdn.microsoft.com/vstudio/DSLTools/>
6. Akos Ledeczi, Miklos Maroti, Arpad Bakay, Gabor Karsai, Garrett, J., Thomason, C., Nordstrom, G., Sprinkle, J. and Volgyesi, P. The Generic Modeling Environment. Proceedings of WISP'2001, May, 2001. <http://www.isis.vanderbilt.edu/Projects/gme/GME2000Overview.pdf>
7. C# programming guide, <http://msdn2.microsoft.com/en-us/library/wa80x488.aspx>
8. Greenfield, J., Short, K., Cook, S., Kent, S. Software Factories: Assembling Applications with Patterns, Models, Frameworks and Tools. Wiley (2004).

Event-Based Programming Without Inversion of Control

Philipp Haller and Martin Odersky

École Polytechnique Fédérale de Lausanne (EPFL)
1015 Lausanne, Switzerland

1 Introduction

Concurrent programming is indispensable. On the one hand, distributed and mobile environments naturally involve concurrency. On the other hand, there is a general trend towards multi-core processors that are capable of running multiple threads in parallel.

With *actors* there exists a computation model which is especially suited for concurrent and distributed computations [16,1]. Actors are basically concurrent processes which communicate through *asynchronous message passing*. When combined with *pattern matching* for messages, actor-based process models have been proven to be very effective, as the success of Erlang documents [3,25].

Erlang [4] is a dynamically typed functional programming language designed for programming real-time control systems. Examples of such systems are telephone exchanges, network simulators and distributed resource controllers. In these systems, large numbers of concurrent processes can be active simultaneously. Moreover, it is generally difficult to predict the number of processes and their memory requirements as they vary with time.

For the implementation of these processes, operating system threads and threads of virtual machines, such as the Java Virtual Machine [22], are usually too heavyweight. The main reasons are: (1) Over-provisioning of stacks leads to quick exhaustion of virtual address space and (2) locking mechanisms often lack suitable contention managers [12]. Therefore, Erlang implements concurrent processes by its own runtime system and not by the underlying operating system [5].

Actor abstractions as lightweight as Erlang's processes have been unavailable on popular virtual machines so far. At the same time, standard virtual machines are becoming an increasingly important platform for exactly the same domain of applications in which Erlang—because of its process model—has been so successful: Real-time control systems [23,27].

Another domain where virtual machines are expected to become ubiquitous are applications running on mobile devices, such as cellular phones or personal digital assistants [20]. Usually, these devices are exposed to severe resource constraints. On such devices, only a few hundred kilobytes of memory is available to a virtual machine and applications.

This has important consequences: (1) A virtual machine for mobile devices usually offers only a restricted subset of the services of a common virtual machine

for desktop or server computers. For example, the KVM¹ has no support for reflection (introspection) and serialization. (2) Programming abstractions used by applications have to be very lightweight to be useful. Again, thread-based concurrency abstractions are too heavyweight. Furthermore, programming models have to cope with the restricted set of services a mobile virtual machine provides.

A common alternative to programming with threads are event-driven programming models. Programming in explicitly event-driven models is very difficult [21].

Most programming models support event-driven programming only through *inversion of control*. Instead of calling blocking operations (e.g. for obtaining user input), a program merely registers its interest to be resumed on certain *events* (e.g. an event signaling a pressed button, or changed contents of a text field). In the process, *event handlers* are installed in the execution environment which are called when certain events occur. The program never calls these event handlers itself. Instead, the execution environment dispatches events to the installed handlers. Thus, control over the execution of program logic is “inverted”.

Virtually all approaches based on inversion of control suffer from the following two problems: First, the interactive logic of a program is fragmented across multiple event handlers (or classes, as in the state design pattern [13]). Second, control flow among handlers is expressed implicitly through manipulation of shared state [10].

To obtain very lightweight abstractions without inversion of control, we make actors *thread-less*. We introduce *event-based actors* as an implementation technique for lightweight actor abstractions on *non-cooperative* virtual machines such as the JVM. Non-cooperative means that the virtual machine provides no means to explicitly manage the execution state of a program.

The central idea is as follows: An actor that waits in a receive statement is not represented by a blocked thread but by a closure that captures the rest of the actor’s computation. The closure is executed once a message is sent to the actor that matches one of the message patterns specified in the receive. The execution of the closure is “piggy-backed” on the thread of the sender. If the receiving closure terminates, control is returned to the sender as if a procedure returns. If the receiving closure blocks in a second receive, control is returned to the sender by throwing a special exception that unwinds the receiver’s call stack.

A necessary condition for the scheme to work is that receivers never return normally to their enclosing actor. In other words, no code in an actor can depend on the termination or the result of a receive block. We can express this non-returning property at compile time through Scala’s type system. This is not a severe restriction in practice, as programs can always be organized in a way so that the “rest of the computation” of an actor is executed from within a receive. To the best of our knowledge, event-based actors are the first to (1) allow reactive behavior to be expressed without *inversion of control*, and (2) support arbitrary blocking operations in reactions, at the same time. Our actor library outperforms other state-of-the-art actor languages with respect to message passing

¹ See <http://java.sun.com/products/cldc/>.

speed and memory consumption by several orders of magnitude. Our implementation is able to make use of multi-processors and multi-core processors because reactions can be executed simultaneously on multiple processors. By extending our event-based actors with a portable runtime system, we show how the essence of distributed Erlang [31] can be implemented in Scala. Our library supports virtually all primitives and built-in-functions which are introduced in the Erlang book [4]. The portability of our runtime system is established by two working prototypes based on TCP and the JXTA² peer-to-peer framework, respectively.

All this has been achieved without extending or changing the programming language. The event-based actor library is thus a good demonstrator of Scala's abstraction capabilities. Beginning with the upcoming release 2.1.7, it is part of the Scala standard distribution³.

Other Related Work. Actalk [8] implements actors as a library for Smalltalk-80 by extending a minimal kernel of pure Smalltalk objects. Their implementation is not event-based and Smalltalk-80 does not support parallel execution of concurrent actors on multi-processors (or multi-core processors).

Actra [29] extends the Smalltalk/V virtual machine with an object-based real-time kernel which provides lightweight processes. In contrast, we implement lightweight actors on unmodified virtual machines.

Chrysanthakopoulos and Singh [11] discuss the design and implementation of a channel-based asynchronous messaging library. Channels can be viewed as special state-less actors which have to be instantiated to indicate the types of messages they can receive. Instead of using heavyweight operating system threads they develop their own scheduler to support continuation passing style (CPS) code. Using CLU-style iterators blocking-style code is CPS-transformed by the C# compiler.

SALSA (Simple Actor Language, System and Architecture) [30] extends Java with concurrency constructs that directly support the notion of actors. A pre-processor translates SALSA programs into Java source code which in turn is linked to a custom-built actor library. As SALSA implements actors on the JVM, it is somewhat closer related to our work than Smalltalk-based actors or channels. Moreover, performance results have been published which enables us to compare our system with SALSA, using ports of existing benchmarks.

Timber is an object-oriented and functional programming language designed for real-time embedded systems [6]. It offers message passing primitives for both synchronous and asynchronous communication between concurrent *reactive objects*. In contrast to event-based actors, reactive objects cannot call operations that might block indefinitely. Instead, they install call-back methods in the computing environment which executes these operations on behalf of them.

Frugal objects [14] (FROBs) are distributed reactive objects that communicate through typed events. FROBs are basically actors with an event-based computation model, just as our event-based actors. The goals of FROBs and

² See <http://www.jxta.org/>.

³ Available from <http://scala.epfl.ch/>.