Michael Hobbs
Andrzej M. Goscinski
Wanlei Zhou (Eds.)

LNCS 3719

Distributed and Parallel Computing

6th International Conference on Algorithms and Architectures for Parallel Processing, ICA3PP Melbourne, Australia, October 2005, Proceedings

Springer

Michael Hobbs Andrzej M. Goscinski
Wanlei Zhou (Eds.)

Distributed and Parallel Computing

6th International Conference on Algorithms and Architectures for Parallel Processing, ICA3PP
Melbourne, Australia, October 2-3, 2005
Proceedings

Springer

Volume Editors

Michael Hobbs
Andrzej M. Goscinski
Wanlei Zhou
Deakin University
School of Information Technology
Geelong, Victoria 3217
Australia
E-mail: {mick,ang,wanlei}@deakin.edu.au

Library of Congress Control Number: 2005933043

CR Subject Classification (1998): D, F.1-3. C, I.6

ISSN 0302-9743
ISBN-10 3-540-29235-7 Springer Berlin Heidelberg New York
ISBN-13 978-3-540-29235-7 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

springeronline.com

© Springer-Verlag Berlin Heidelberg 2005
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 11564621 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
New York University, NY, 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

There are many applications that require parallel and distributed processing to allow complicated engineering, business and research problems to be solved in a reasonable time. Parallel and distributed processing is able to improve company profit, lower costs of design, production, and deployment of new technologies, and create better business environments. The major lesson learned by car and aircraft engineers, drug manufacturers, genome researchers and other specialist is that a computer system is a very powerful tool that is able to help them solving even more complicated problems. That has led computing specialists to new computer system architecture and exploiting parallel computers, clusters of clusters, and distributed systems in the form of grids. There are also institutions that do not have so complicated problems but would like to improve profit, lower costs of design and production by using parallel and distributed processing on clusters.

In general to achieve these goals, parallel and distributed processing must become the computing mainstream. This implies a need for new architectures of parallel and distributed systems, new system management facilities, and new application algorithms. This also implies a need for better understanding of grids and clusters, and in particular their operating systems, scheduling algorithms, load balancing, heterogeneity, transparency, application deployment, which is of the most critical importance for their development and taking them by industry and business.

ICA3PP has been a premier conference that has brought together researchers and practitioners from academia, industry and governments around the world to advance the theories and technologies of parallel and distributed computing. Previously, ICA3PP conferences have been held successfully in Brisbane, Singapore, Melbourne, Hong Kong and Beijing.

ICA3PP 2005 returned to Melbourne with the main focus on the most critical areas of parallel and distributed computing: operating systems and middleware, fault-tolerant systems, scheduling and load balancing, algorithms, tools and environments, and communication and networks.

In total, the conference received 98 papers from researchers and practitioners from 15 countries. Each paper was reviewed by at least three internationally renowned referees and selected based on their originality, significance, correctness, relevance, and clarity of presentation. Among the high quality submissions, 27 long papers and 25 short papers were accepted. All of the selected papers are included in the proceedings. After the conference, the proceedings editors will recommend some high quality papers from the conference to be published in a special issue of an international journal.

We are delighted to be able to host two well-known international scholars, Professor Ian Foster and Professor Zhiwei Xu, who delivered the keynote speeches.

试读结束，需要全本PDF请购买 www.ertongbook.com

We would like to take this opportunity to thank all the authors for their submissions to the conference. Many of them have traveled some distance to participate in the conference. We also thank the Program Committee members and additional reviewers for their efforts in reviewing the large number of papers. Thanks also go to the local conference organizers for their great support.

Last but not least, we would like to express our gratitude to all of the organizations who have supported our efforts to bring the conference to fruition. We are grateful to the IEEE Technical Committee on Scalable Computing for the cooperation; and to Deakin University, NICTA and Alexander Technology for their sponsorships and assistance.

October 2005 Michael Hobbs, Andrzej Goscinski and Wanlei Zhou
Melbourne,

Organization

This conference was organized by the School of IT, Deakin University, Australia; and Martin Lack & Associates, Australia. Sponsorship was provided by the School of IT, Deakin University, Australia; the Faculty of Science and Technology, Deakin University, Australia; National ICT Australia (NICTA); and Alexander Technology, Australia.

Conference Chairs

Andrzej M. Goscinski (Deakin University, Australia)
Wanlei Zhou (Deakin University, Australia)

Program Chair

Michael J. Hobbs (Deakin University, Australia)

Program Committee

Jemal Abawajy (Deakin University, Australia)
David Abromson (Monash University, Australia)
Akkihebbal Ananda (National University of Singapore, Singapore)
Bill Appelbe (RMIT University, Australia)
Mark Baker (Portsmouth University, UK)
Amnon Barak (Hebrew University of Jerusalem, Israel)
Arndt Bode (Technical University of Munich, Germany)
Peter Brezany (University of Vienna, Austria)
Marian Bubak (AGH University of Science and Technology., Cracow, Poland)
Raj Buyya (University of Melbourne, Australia)
Jianning Cao (Hong Kong Polytechnic University, Hong Kong)
Samuel Chanson (Hong Kong University of Science and Technology, Hong Kong)
Jianer Chen (Texas A&M University, USA)
Francis Chin (University of Hong Kong, Hong Kong)
Toni Cortes (University of Politecnica de Catalunya, Spain)
Brian d'Auriol (University of Texas at El Paso, USA)
Xiaote Deng (City University of Hong Kong, Hong Kong)
Robert Dew (Deakin University, Australia)
Jack Dongarra (University of Tennessee, USA)
Ding-Zhu Du (University of Minnesota, USA)
Wen Gao (Inst. of Computing Technology, China)
Al Giest (Oak Ridge Nation Labs, USA)
Minyi Guo (University of Aizu, Japan)

Salim Hariri (Syracuse University, USA)
Louis Hertzberger (University of Amsterdam, The Netherland)
Shi-Jinn Horng (National Taiwan University of Science and Technology, Taiwan)
Ali Hurson (Pennsylvania State University, USA)
Weijia Jia (City University of Hong Kong, Hong Kong)
Xiaohua Jia (City University of Hong Kong, Hong Kong)
Hai Jin (Huazhong University of Science and Technology, China)
Peter Kacsuk (MTA SZTAKI Research Inst., Hungry)
Krishna Kavi (The University of North Texas, USA)
Zvi Kedem (New York University, USA)
Wayne Kelly (Queensland University of Tech., Australia)
Tohru Kikuno (Osaka University, Japan)
Jacek Kitowski (University of Mining and Metallurgy, Poland)
Domenico Laforenza (ISTI-CNR, Italy)
Laurent Leferve (INRIA, Lyon, France)
Keqin Li (State University of New York at New Paltz, USA)
Zhi Yong Liu (National Natural Science Foundation, China)
Thomas Ludwig (University of Heidelberg, Germany)
George Mohay (Queensland University of Tech., Australia)
Christine Morin (IRISA/INRIA, France)
Edgar Nett (Otto-von-Guericke University, Germany)
Yi Pan (Georgia State University USA)
Marcin Paprzycki (Oklahoma State University)
Sushil K. Prasad (Georgia State University USA)
Rajeev Raje (Purdue University, USA)
Michel Raynal (IRISA, France)
Justin Rough (Deakin University, Australia)
Srinivas Sampalli (Dalhousie University, Canada)
Edwin Sha (University of Texas at Dallas, USA)
Behrooz Shirazi (University of Texas Arlington, USA)
Jackie Silcock (Deakin University, Australia)
Beth Simon (University of San Diego, USA)
Chengzheng Sun (Griffith University, Australia)
Jiachang Sun (Institute of Software, China)
Vaidy Sunderam (Emory University, Atlanta, USA)
Yong-Meng Teo (National University of Singapore, Singapore)
Alistair Veitch (Hewlett Packard Labs, California, USA)
Greg Wickham (GrangeNet, Australia)
Jie Wu (Florida Atlantic University, USA)
Yue Wu (University of Electronic Science and Technology, China)
Roman Wyrzykowski (Czestochowa University of Technology, Poland)
Jingling Xue (University of New South Wales, Australia)
Zhiwei Xu (Inst. of Computing Technology, China)
Laurence T. Yang (St. Francis Xavier University, Canada)
Chung-Kwong Yuen (National University of Singapore, Singapore)
Si Q Zheng (University of Texas Dallas, USA)
Weimin Zheng (Tsinghua University, China)
Wei Zhao (Texas A&M University, USA)
Jun Zou (Chinese University of Hong Kong, China)
Albert Zoymaya (University of Sydney, Australia)

Lecture Notes in Computer Science

For information about Vols. 1–3638

please contact your bookseller or Springer

Vol. 3739: W. Fan, Z. Wu, J. Yang (Eds.), Advances in Web-Age Information Management. XXII, 930 pages. 2005.

Vol. 3738: V.R. Syrotiuk, E. Chávez (Eds.), Ad-Hoc, Mobile, and Wireless Networks. XI, 360 pages. 2005.

Vol. 3734: S. Jain, H.U. Simon, E. Tomita (Eds.), Algorithmic Learning Theory. XII, 490 pages. 2005.

Vol. 3731: F. Wang (Ed.), Formal Techniques for Networked and Distributed Systems - FORTE 2005. XII, 558 pages. 2005.

Vol. 3728: V. Paliouras, J. Vounckx, D. Verkest (Eds.), Integrated Circuit and System Design. XV, 753 pages. 2005.

Vol. 3726: L.T. Yang, O.F. Rana, B. Di Martino, J. Dongarra (Eds.), High Performance Computing and Communcations. XXVI, 1116 pages. 2005.

Vol. 3725: D. Borrione, W. Paul (Eds.), Correct Hardware Design and Verification Methods. XII, 412 pages. 2005.

Vol. 3724: P. Fraigniaud (Ed.), Distributed Computing. XIV, 520 pages. 2005.

Vol. 3722: D. Van Hung, M. Wirsing (Eds.), Theoretical Aspects of Computing – ICTAC 2005. XIV, 614 pages. 2005.

Vol. 3719: M. Hobbs, A.M. Goscinski, W. Zhou (Eds.), Distributed and Parallel Computing. XI, 448 pages. 2005.

Vol. 3718: V.G. Ganzha, E.W. Mayr, E.V. Vorozhtsov (Eds.), Computer Algebra in Scientific Computing. XII, 502 pages. 2005.

Vol. 3717: B. Gramlich (Ed.), Frontiers of Combining Systems. X, 321 pages. 2005. (Subseries LNAI).

Vol. 3715: E. Dawson, S. Vaudenay (Eds.), Progress in Cryptology – Mycrypt 2005. XI, 329 pages. 2005.

Vol. 3714: H. Obbink, K. Pohl (Eds.), Software Product Lines. XIII, 235 pages. 2005.

Vol. 3713: L. Briand, C. Williams (Eds.), Model Driven Engineering Languages and Systems. XV, 722 pages. 2005.

Vol. 3712: R. Reussner, J. Mayer, J.A. Stafford, S. Overhage, S. Becker, P.J. Schroeder (Eds.), Quality of Software Architectures and Software Quality. XIII, 289 pages. 2005.

Vol. 3711: F. Kishino, Y. Kitamura, H. Kato, N. Nagata (Eds.), Entertainment Computing - ICEC 2005. XXIV, 540 pages. 2005.

Vol. 3710: M. Barni, I. Cox, T. Kalker, H.J. Kim (Eds.), Digital Watermarking. XII, 485 pages. 2005.

Vol. 3708: J. Blanc-Talon, W. Philips, D. Popescu, P. Scheunders (Eds.), Advanced Concepts for Intelligent Vision Systems. XXII, 725 pages. 2005.

Vol. 3707: D.A. Peled, Y.-K. Tsay (Eds.), Automated Technology for Verification and Analysis. XII, 506 pages. 2005.

Vol. 3706: H. Fuks, S. Lukosch, A.C. Salgado (Eds.), Groupware: Design, Implementation, and Use. XII, 378 pages. 2005.

Vol. 3703: F. Fages, S. Soliman (Eds.), Principles and Practice of Semantic Web Reasoning. VIII, 163 pages. 2005.

Vol. 3702: B. Beckert (Ed.), Automated Reasoning with Analytic Tableaux and Related Methods. XIII, 343 pages. 2005. (Subseries LNAI).

Vol. 3701: M. Coppo, E. Lodi, G. M. Pinna (Eds.), Theoretical Computer Science. XI, 411 pages. 2005.

Vol. 3699: C.S. Calude, M.J. Dinneen, G. Păun, M. J. Pérez-Jiménez, G. Rozenberg (Eds.), Unconventional Computation. XI, 267 pages. 2005.

Vol. 3698: U. Furbach (Ed.), KI 2005: Advances in Artificial Intelligence. XIII, 409 pages. 2005. (Subseries LNAI).

Vol. 3697: W. Duch, J. Kacprzyk, E. Oja, S. Zadrożny (Eds.), Artificial Neural Networks: Formal Models and Their Applications – ICANN 2005, Part II. XXXII, 1045 pages. 2005.

Vol. 3696: W. Duch, J. Kacprzyk, E. Oja, S. Zadrożny (Eds.), Artificial Neural Networks: Biological Inspirations – ICANN 2005, Part I. XXXI, 703 pages. 2005.

Vol. 3695: M.R. Berthold, R. Glen, K. Diederichs, O. Kohlbacher, I. Fischer (Eds.), Computational Life Sciences. XI, 277 pages. 2005. (Subseries LNBI).

Vol. 3694: M. Malek, E. Nett, N. Suri (Eds.), Service Availability. VIII, 213 pages. 2005.

Vol. 3693: A.G. Cohn, D.M. Mark (Eds.), Spatial Information Theory. XII, 493 pages. 2005.

Vol. 3692: R. Casadio, G. Myers (Eds.), Algorithms in Bioinformatics. X, 436 pages. 2005. (Subseries LNBI).

Vol. 3691: A. Gagalowicz, W. Philips (Eds.), Computer Analysis of Images and Patterns. XIX, 865 pages. 2005.

Vol. 3690: M. Pěchouček, P. Petta, L.Z. Varga (Eds.), Multi-Agent Systems and Applications IV. XVII, 667 pages. 2005. (Subseries LNAI).

Vol. 3688: R. Winther, B.A. Gan, G. Dahll (Eds.), Computer Safety, Reliability, and Security. XI, 405 pages. 2005.

Vol. 3687: S. Singh, M. Singh, C. Apte, P. Perner (Eds.), Pattern Recognition and Image Analysis, Part II. XXV, 809 pages. 2005.

Vol. 3686: S. Singh, M. Singh, C. Apte, P. Perner (Eds.), Pattern Recognition and Data Mining, Part I. XXVI, 689 pages. 2005.

Vol. 3685: V. Gorodetsky, I. Kotenko, V. Skormin (Eds.), Computer Network Security. XIV, 480 pages. 2005.

Vol. 3684: R. Khosla, R.J. Howlett, L.C. Jain (Eds.), Knowledge-Based Intelligent Information and Engineering Systems, Part IV. LXXIX, 933 pages. 2005. (Subseries LNAI).

Vol. 3683: R. Khosla, R.J. Howlett, L.C. Jain (Eds.), Knowledge-Based Intelligent Information and Engineering Systems, Part III. LXXX, 1397 pages. 2005. (Subseries LNAI).

Vol. 3682: R. Khosla, R.J. Howlett, L.C. Jain (Eds.), Knowledge-Based Intelligent Information and Engineering Systems, Part II. LXXIX, 1371 pages. 2005. (Subseries LNAI).

Vol. 3681: R. Khosla, R.J. Howlett, L.C. Jain (Eds.), Knowledge-Based Intelligent Information and Engineering Systems, Part I. LXXX, 1319 pages. 2005. (Subseries LNAI).

Vol. 3679: S.d.C. di Vimercati, P. Syverson, D. Gollmann (Eds.), Computer Security – ESORICS 2005. XI, 509 pages. 2005.

Vol. 3678: A. McLysaght, D.H. Huson (Eds.), Comparative Genomics. VIII, 167 pages. 2005. (Subseries LNBI).

Vol. 3677: J. Dittmann, S. Katzenbeisser, A. Uhl (Eds.), Communications and Multimedia Security. XIII, 360 pages. 2005.

Vol. 3676: R. Glück, M. Lowry (Eds.), Generative Programming and Component Engineering. XI, 448 pages. 2005.

Vol. 3675: Y. Luo (Ed.), Cooperative Design, Visualization, and Engineering. XI, 264 pages. 2005.

Vol. 3674: W. Jonker, M. Petković (Eds.), Secure Data Management. X, 241 pages. 2005.

Vol. 3673: S. Bandini, S. Manzoni (Eds.), AI*IA 2005: Advances in Artificial Intelligence. XIV, 614 pages. 2005. (Subseries LNAI).

Vol. 3672: C. Hankin, I. Siveroni (Eds.), Static Analysis. X, 369 pages. 2005.

Vol. 3671: S. Bressan, S. Ceri, E. Hunt, Z.G. Ives, Z. Bellahsène, M. Rys, R. Unland (Eds.), Database and XML Technologies. X, 239 pages. 2005.

Vol. 3670: M. Bravetti, L. Kloul, G. Zavattaro (Eds.), Formal Techniques for Computer Systems and Business Processes. XIII, 349 pages. 2005.

Vol. 3669: G.S. Brodal, S. Leonardi (Eds.), Algorithms – ESA 2005. XVIII, 901 pages. 2005.

Vol. 3668: M. Gabbrielli, G. Gupta (Eds.), Logic Programming. XIV, 454 pages. 2005.

Vol. 3666: B.D. Martino, D. Kranzlmüller, J. Dongarra (Eds.), Recent Advances in Parallel Virtual Machine and Message Passing Interface. XVII, 546 pages. 2005.

Vol. 3665: K. S. Candan, A. Celentano (Eds.), Advances in Multimedia Information Systems. X, 221 pages. 2005.

Vol. 3664: C. Türker, M. Agosti, H.-J. Schek (Eds.), Peer-to-Peer, Grid, and Service-Orientation in Digital Library Architectures. X, 261 pages. 2005.

Vol. 3663: W.G. Kropatsch, R. Sablatnig, A. Hanbury (Eds.), Pattern Recognition. XIV, 512 pages. 2005.

Vol. 3662: C. Baral, G. Greco, N. Leone, G. Terracina (Eds.), Logic Programming and Nonmonotonic Reasoning. XIII, 454 pages. 2005. (Subseries LNAI).

Vol. 3661: T. Panayiotopoulos, J. Gratch, R. Aylett, D. Ballin, P. Olivier, T. Rist (Eds.), Intelligent Virtual Agents. XIII, 506 pages. 2005. (Subseries LNAI).

Vol. 3660: M. Beigl, S. Intille, J. Rekimoto, H. Tokuda (Eds.), UbiComp 2005: Ubiquitous Computing. XVII, 394 pages. 2005.

Vol. 3659: J.R. Rao, B. Sunar (Eds.), Cryptographic Hardware and Embedded Systems – CHES 2005. XIV, 458 pages. 2005.

Vol. 3658: V. Matoušek, P. Mautner, T. Pavelka (Eds.), Text, Speech and Dialogue. XV, 460 pages. 2005. (Subseries LNAI).

Vol. 3657: F.S. de Boer, M.M. Bonsangue, S. Graf, W.-P. de Roever (Eds.), Formal Methods for Components and Objects. VIII, 325 pages. 2005.

Vol. 3656: M. Kamel, A. Campilho (Eds.), Image Analysis and Recognition. XXIV, 1279 pages. 2005.

Vol. 3655: A. Aldini, R. Gorrieri, F. Martinelli (Eds.), Foundations of Security Analysis and Design III. VII, 273 pages. 2005.

Vol. 3654: S. Jajodia, D. Wijesekera (Eds.), Data and Applications Security XIX. X, 353 pages. 2005.

Vol. 3653: M. Abadi, L. de Alfaro (Eds.), CONCUR 2005 – Concurrency Theory. XIV, 578 pages. 2005.

Vol. 3652: A. Rauber, S. Christodoulakis, A M. Tjoa (Eds.), Research and Advanced Technology for Digital Libraries. XVIII, 545 pages. 2005.

Vol. 3651: R. Dale, K.-F. Wong, J. Su, O.Y. Kwong (Eds.), Natural Language Processing – IJCNLP 2005. XXI, 1031 pages. 2005. (Subseries LNAI).

Vol. 3650: J. Zhou, J. Lopez, R.H. Deng, F. Bao (Eds.), Information Security. XII, 516 pages. 2005.

Vol. 3649: W.M. P. van der Aalst, B. Benatallah, F. Casati, F. Curbera (Eds.), Business Process Management. XII, 472 pages. 2005.

Vol. 3648: J.C. Cunha, P.D. Medeiros (Eds.), Euro-Par 2005 Parallel Processing. XXXVI, 1299 pages. 2005.

Vol. 3646: A. F. Famili, J.N. Kok, J.M. Peña, A. Siebes, A. Feelders (Eds.), Advances in Intelligent Data Analysis VI. XIV, 522 pages. 2005.

Vol. 3645: D.-S. Huang, X.-P. Zhang, G.-B. Huang (Eds.), Advances in Intelligent Computing, Part II. XIII, 1010 pages. 2005.

Vol. 3644: D.-S. Huang, X.-P. Zhang, G.-B. Huang (Eds.), Advances in Intelligent Computing, Part I. XXVII, 1101 pages. 2005.

Vol. 3643: R. Moreno Díaz, F. Pichler, A. Quesada Arencibia (Eds.), Computer Aided Systems Theory – EUROCAST 2005. XIV, 629 pages. 2005.

Vol. 3642: D. Ślezak, J. Yao, J.F. Peters, W. Ziarko, X. Hu (Eds.), Rough Sets, Fuzzy Sets, Data Mining, and Granular Computing, Part II. XXIII, 738 pages. 2005. (Subseries LNAI).

Vol. 3641: D. Ślezak, G. Wang, M. Szczuka, I. Düntsch, Y. Yao (Eds.), Rough Sets, Fuzzy Sets, Data Mining, and Granular Computing, Part I. XXIV, 742 pages. 2005. (Subseries LNAI).

Vol. 3639: P. Godefroid (Ed.), Model Checking Software. XI, 289 pages. 2005.

Table of Contents

Improving Concurrent Write Scheme in File Server Group
Fengjung Liu, Chu-sing Yang 1

A Comparative Performance Study of Distributed Mutual Exclusion Algorithms with a Class of Extended Petri Nets
Alexander Kostin, Ljudmila Ilushechkina, Erhan Basri 11

A Practical Comparison of Cluster Operating Systems Implementing Sequential and Transactional Consistency
Stefan Frenz, Renaud Lottiaux, Michael Schoettner, Christine Morin, Ralph Goeckelmann, Peter Schulthess 23

Clock Synchronization State Graphs Based on Clock Precision Difference
Ying Zhao, Wanlei Zhou, Yingying Zhang, E.J. Lanham, Jiumei Huang 34

A Recursive-Adjustment Co-allocation Scheme in Data Grid Environments
Chao-Tung Yang, I-Hsien Yang, Kuan-Ching Li, Ching-Hsien Hsu 40

Reducing the Bandwidth Requirements of P2P Keyword Indexing
John Casey, Wanlei Zhou 50

A Deadline and Budget Constrained Scheduling Algorithm for eScience Applications on Data Grids
Srikumar Venugopal, Rajkumar Buyya 60

A Survivability Model for Cluster System
Khin Mi Mi Aung, Kiejin Park, Jong Sou Park 73

Localization Techniques for Cluster-Based Data Grid
Ching-Hsien Hsu, Guan-Hao Lin, Kuan-Ching Li, Chao-Tung Yang 83

GridFTP and Parallel TCP Support in NaradaBrokering
Sang Boem Lim, Geoffrey Fox, Ali Kaplan, Shrideep Pallickara, Marlon Pierce 93

2-Layered Metadata Service Model in Grid Environment
Muzhou Xiong, Hai Jin, Song Wu 103

试读结束，需要全本PDF请购买 www.ertongbook.com

pKSS: An Efficient Keyword Search System in DHT Peer-to-Peer Network
Yin Li, Fanyuan Ma, Liang Zhang 112

A Comparative Study at the Logical Level of Centralised and Distributed Recovery in Clusters
Andrew Maloney, Andrzej Goscinski 118

Toward Self Discovery for an Autonomic Cluster
Eric Dines, Andrzej Goscinski 125

Mining Traces of Large Scale Systems
Christophe Cérin, Michel Koskas 132

Setup Algorithm of Web Service Composition
YanPing Yang, QingPing Tan, Yong Xiao 139

Self Healing and Self Configuration in a WSRF Grid Environment
Michael Messig, Andrzej Goscinski 149

Study on Life Cycle Model of Dynamic Composed Web Services
Chen Yanping, Li Zengzhi, Jin Qinxue, Wang Chuang 159

Fault-Tolerant Dynamic Job Scheduling Policy
J.H. Abawajy 165

An Efficient Dynamic Load-Balancing Algorithm in a Large-Scale Cluster
Bao-Yin Zhang, Ze-Yao Mo, Guang-Wen Yang, Wei-Min Zheng 174

Job Scheduling Policy for High Throughput Grid Computing
J.H. Abawajy 184

High Performance Task Scheduling Algorithm for Heterogeneous Computing System
E. Ilavarasan, P. Thambidurai, R. Mahilmannan 193

Execution Environments and Benchmarks for the Study of Applications' Scheduling on Clusters
Adam K.L. Wong, Andrzej M. Goscinski 204

Data Distribution Strategies for Domain Decomposition Applications in Grid Environments
Beatriz Otero, José M. Cela, Rosa M. Badia, Jesús Labarta 214

Inter-round Scheduling for Divisible Workload Applications
DongWoo Lee, R.S. Ramakrishna 225

Scheduling Divisible Workloads Using the Adaptive Time Factoring Algorithm
Tiago Ferreto, César De Rose 232

Adaptive Policy Triggering for Load Balancing
Dan Feng, Lingfang Zeng 240

Parallel Algorithms for Fault-Tolerant Mobile Agent Execution
Jin Yang, Jiannong Cao, Weigang Wu, Cheng-Zhong Xu 246

Design and Multithreading Implementation of the Wave-Front Algorithm for Constructing Voronoi Diagrams
Grace J. Hwang, Joseph M. Arul, Eric Lin, Chung-Yun Hung 257

A Proposal of Parallel Strategy for Global Wavelet-Based Registration of Remote-Sensing Images
Haifang Zhou, Yu Tang, Xuejun Yang, Hengzhu Liu 267

Performance Analysis of a Parallel Sort Merge Join on Cluster Architectures
Erich Schikuta 277

Parallel Clustering on the Star Graph
M. Fazeli, H. Sarbazi-Azad, R. Farivar 287

Hierarchical Parallel Simulated Annealing and Its Applications
Shiming Xu, Wenguang Chen, Weimin Zheng, Tao Wang, Yimin Zhang 293

Multi-color Difference Schemes of Helmholtz Equation and Its Parallel Fast Solver over 3-D Dodecahedron Partitions
Jiachang Sun 301

GridMD: Program Architecture for Distributed Molecular Simulation
Ilya Valuev 309

Visuel: A Novel Performance Monitoring and Analysis Toolkit for Cluster and Grid Environments
Kuan-Ching Li, Hsiang-Yao Cheng, Chao-Tung Yang, Ching-Hsien Hsu, Hsiao-Hsi Wang, Chia-Wen Hsu, Sheng-Shiang Hung, Chia-Fu Chang, Chun-Chieh Liu, Yu-Hwa Pan 315

Introduction to a New Tariff Mechanism for Charging for Computer Power in the Grid
Sena Seneviratne, David Levy 326

Host Load Prediction for Grid Computing Using Free Load Profiles
Sena Seneviratne, David Levy 336

Active Link: Status Detection Mechanism for Distributed Service Based on Active Networks
Zhan Tao, Zhou Xingshe, Liao Zhigang, Chen Yan 345

Performance Monitoring for Distributed Service Oriented Grid Architecture
Liang Peng, Melvin Koh, Jie Song, Simon See 351

Distributed Defense Against Distributed Denial-of-Service Attacks
Wei Shi, Yang Xiang, Wanlei Zhou 357

Security and Safety Assurance Architecture: Model and Implementation (Supporting Multiple Levels of Criticality)
Li Zhongwen 363

Modeling and Analysis of Worm and Killer-Worm Propagation Using the Divide-and-Conquer Strategy
Dan Wu, Dongyang Long, Changji Wang, Zhanpeng Guan 370

An Efficient Reliable Architecture for Application Layer Anycast Service
Shui Yu, Wanlei Zhou 376

A Distributed Approach to Estimate Link-Level Loss Rates
Weiping Zhu 386

Evaluation of Interconnection Network Performance Under Heavy Non-uniform Loads
C. Izu, J. Miguel-Alonso, J.A. Gregorio 396

Analytical Models of Probability Distributions for MPI Point-to-Point Communication Times on Distributed Memory Parallel Computers
D.A. Grove, P.D. Coddington 406

Communication Data Multiplexing in Distributed Simulation
Jong Sik Lee 416

Novel Adaptive Subcarrier Power and Bit Allocation Using Wavelet Packet Parallel Architecture
Ren Ren, Shihua Zhu 422

A Low–Level Communication Library for Java HPC
Sang Boem Lim, Bryan Carpenter, Geoffrey Fox, Han-Ku Lee 429

Object-Oriented Design and Implementations of 3G-324M Protocol Stack
Weijia Jia, Haohuan Fu, Ji Shen 435

Efficient Techniques and Hardware Analysis for Mesh-Connected Processors
Wu Jigang, Thambipillai Srikanthan, Schröder Heiko 442

Author Index 447

Improving Concurrent Write Scheme in File Server Group

Fengjung Liu[1] and Chu-sing Yang[2]

[1] Department of Management Information Systems, Tajen University, Pingtung, 907, Taiwan
fjliu@mail.tajen.edu.tw
[2] Department of Computer Science and Engineering, National Sun Yat-sen University, Kaohsiung, 80424, Taiwan
csyang@cse.nsysu.edu.tw

Abstract. The explosive growth of the Web contents has led to increasing attention on scalability and availability of file system. Hence, the ways to improve the reliability and availability of system, to achieve the expected reduction in operational expenses and to reduce the operations of management of system have become essential issues. A basic technique for improving reliability of a file system is to mask the effects of failures through replication. Consistency control protocols are implemented to ensure the consistency among replicas. In this paper, we leveraged the concept of intermediate file handle to cover the heterogeneity of file system and proposed an efficient data consistency control scheme supporting dependence checking among writes and management of out-of-ordered requests for file server group. Finally, the results of experiments proved the efficiency of the proposed consistency control mechanism. Above all, easy to implement is our main design consideration.

1 Introduction

The explosive growth of the Web contents has led to increasing attention on scalability and availability of file system. A basic technique for improving reliability of file system is to mask the effects of failures by replication. There are two major approaches of building a highly reliable file system: hardware replication approach and software replication approach. In a distributed environment, it is not always necessary to use special hardware for improved reliability. The computers connected by the high-speed network are a natural resource of duplicates. The software replication approach replicates file systems on workstations in the network. Consistency control protocols are designed to ensure the consistency among replicas.

2 Related Works

Many Distributed File systems, such as intermezzo[1], Coda[2], Deceit[3], Ficus [4], RNFS[5] and Pangaea[6], implemented reliable file system services through software

M. Hobbs, A. Goscinski, and W. Zhou. (Eds.): ICA3PP 2005, LNCS 3719, pp. 1 – 10, 2005.
© Springer-Verlag Berlin Heidelberg 2005

replication approach. In particular, FSG[7,9,10], RNFS and Deceit are NFS-based systems. JetFile[8] and Coda are the instances of multicast-based file systems.

2.1 Overviews of Network File System

The Network File System, NFS, is the most popular distributed file system developed by SUN Microsystems. Each server computer can serve an arbitrary number of the sub-tree in its local file system. Clients are able to mount the exported sub-trees, linking then to its own file system, using the same semantics valid while mounting physical local devices. Its main features are

I. It is a "stateless" protocol. A server does not need to maintain any protocol state information about any of its clients to function correctly.
II. The NFS protocol is idempotent. Because of this and the statelessness property of NFS, what a client has to do for recovery from the crashed server is simply trying the failed RPC until the server reappears.
III. A client accesses a file using a handle, called *fHandle,* obtained from the server as a result of a LOOKUP operation.
IV. Updates are synchronous with respect to failures. If a write RPC completes, the client is assumed that data has been written. Again, this is a property not met by all NFS implementations.

Because of the stateless property of NFS, it could reduce the overhead of recovery after system crashing and make our implementation easier.

2.2 Multicast

In IP multicast [11] there are 2^{28} (2^{112} in IPv6) distinct multicast channels. Channels are named with IP addresses from a subnet of the IP address space. IP packets are only delivered with best effort. To multicast a packet, the sender uses the name of the multicast channel as the IP destination address.

Scalable Reliable Multicast, SRM, [12] is designed to meet only the minimal definition of reliable multicast, i.e., eventual delivery of all data to all group members. As opposed to ISIS [13], SRM does not enforce any particular delivery order. Delivery order is to some extent orthogonal to reliable delivery and can instead be enforced on top of SRM. SRM is logically layered above IP multicast and also relies on the same lightweight delivery model. To be scalable, it does not make use of any negative or positive packet acknowledgements, nor does it keep any knowledge of receiver group membership. Applications will often be able to recover after a period pf packet loss by only requesting to current data. Thus, it is not always necessary to catch up on every missed application data packets.

3 System Design

In system design, we assumed the fail-stop property be approximated by the kernel and the hardware. Additionally, we also assumed that the servers are connected by a Local Area Network and the network is not subject to partitions.

试读结束，需要全本PDF请购买 www.ertongbook.com