


Proceedings



25th IEEE Symposium on Reliable Distributed Systems

SRDS 2006

2-4 October 2006 • Leeds, United Kingdom

Sponsored by

IEEE Computer Society Technical Committee on Distributed Processing

Supported by

School of Computing, Leeds University

In cooperation with

IFIP WG 10.4 on Dependable Computing and Fault Tolerance




UNIVERSITY OF LEEDS



Los Alamitos, California

Washington • Tokyo

Proceedings



SRDS 2006

Table of Contents

Message from the Symposium Chair	ix
Message from the Technical Program Co-Chairs	x
Conference Committees	xi
External Reviewers	xiii

Keynote Address

Session 1: FT Peer-to-Peer Infrastructures

Reducing the Availability Management Overheads of Federated Content Sharing Systems	5
<i>Christopher Peery, Thu D. Nguyen, and Francisco Matias Cuenca-Acuna</i>	
Topology Sensitive Replica Selection	18
<i>Dmitry Brodsky, Michael J. Feeley, and Norman C. Hutchinson</i>	
Deleting Files in the Celeste Peer-to-Peer Storage System	29
<i>Gal Badishi, Germano Caronni, Idit Keidar, Raphael Rom, and Glenn Scott</i>	
Reliably Executing Tasks in the Presence of Untrusted Entities	39
<i>Antonio Fernández, Luis López, Agustín Santos, and Chryssis Georgiou</i>	

Session 2: Security and High Assurance Systems I

WRAPS: Denial-of-Service Defense through Web Referrals	51
<i>XiaoFeng Wang and Michael K. Reiter</i>	
A Client-Transparent Approach to Defend against Denial of Service Attacks.....	61
<i>Mudhakar Srivatsa, Arun Iyengar, Jian Yin, and Ling Liu</i>	
Proactive Resilience Revisited: The Delicate Balance between Resisting Intrusions and Remaining Available	71
<i>Paulo Sousa, Nuno Ferreira Neves, Paulo Verissimo, and William H. Sanders</i>	

Session 3: Autonomic, Pervasive Computing

Call Availability Prediction in a Telecommunication System: A Data Driven Empirical Approach.....	83
<i>Guenther Hoffmann and Mirosław Malek</i>	
FT-PPTC: An Efficient and Fault-Tolerant Commit Protocol for Mobile Environments	96
<i>Brahim Ayari, Abdelmajid Khelil, and Neeraj Suri</i>	
Modeling Distributed Computing System Reliability with DRBD	106
<i>Salvatore Distefano, Marco Scarpa, and Antonio Puliafito</i>	

Session 4: Fault-Tolerant Middleware and Protocols

DRIFT: Efficient Message Ordering in Ad Hoc Networks Using Virtual Flooding	119
<i>Stefan Pleisch, Thomas Clouser, Mikhail Nesterenko, and André Schiper</i>	
Generalised Repair for Overlay Networks.....	132
<i>Barry Porter, François Taïani, and Geoff Coulson</i>	
Decentralized Local Failure Detection in Dynamic Distributed Systems.....	143
<i>Nigamanth Sridhar</i>	

Session 5: Distributed Algorithms

Improvements and Reconsideration of Distributed Snapshot Protocols.....	155
<i>Adnan Agbaria</i>	
Weakly-Persistent Causal Objects in Dynamic Distributed Systems	165
<i>R. Baldoni, M. Malek, A. Milani, and S. Tucci Piergiovanni</i>	
Non-blocking Synchronous Checkpointing Based on Rollback-Dependency Trackability.....	411
<i>Tiemi C. Sakata and Islene C. Garcia</i>	
PLATO: Predictive Latency-Aware Total Ordering.....	175
<i>Mahesh Balakrishnan, Ken Birman, and Amar Phanishayee</i>	

Session 6: Security and High Assurance Systems II

Cryptree: A Folder Tree Structure for Cryptographic File Systems	189
<i>Dominik Grolimund, Luzius Meisser, Stefan Schmid, and Roger Wattenhofer</i>	
How to Safeguard Your Sensitive Data	199
<i>Bob Mungamuru, Hector Garcia-Molina, and Subhasish Mitra</i>	
Solving Consensus Using Structural Failure Models	212
<i>Timo Warns, Felix C. Freiling, and Wilhelm Hasselbring</i>	

Session 7: Performance and Quality Assessment of Distributed Algorithms

Performance Evaluation of a Fair Fault-Tolerant Mutual Exclusion Algorithm	225
<i>Julien Sopena, Luciana Arantes, and Pierre Sens</i>	
Experimental Comparison of Local and Shared Coin Randomized Consensus Protocols	235
<i>Henrique Moniz, Nuno Ferreira Neves, Miguel Correia, and Paulo Verissimo</i>	
Hidden Markov Models as a Support for Diagnosis: Formalization of the Problem and Synthesis of the Solution	245
<i>Alessandro Daidone, Felicita Di Giandomenico, Andrea Bondavalli, and Silvano Chiaradonna</i>	

Session 8: Fault Tolerance in Distributed Operating Systems

Consistent Replication of Multithreaded Distributed Objects	257
<i>Hans P. Reiser, Jörg Domaschka, Franz J. Hauck, Rüdiger Kapitza, and Wolfgang Schröder-Preikschat</i>	
Recovering from Distributable Thread Failures with Assured Timeliness in Real-Time Distributed Systems	267
<i>Edward Curley, Jonathan Anderson, Binoy Ravindran, and E. D. Jensen</i>	
SegmentShield: Exploiting Segmentation Hardware for Protecting against Buffer Overflow Attacks	277
<i>Takahiro Shinagawa</i>	

Session 9: Fault-Tolerant Services

A Scalable Services Architecture	289
<i>Tudor Marian, Ken Birman, and Robbert van Renesse</i>	
Fault-Tolerant and Scalable TCP Splice and Web Server Architecture	301
<i>Manish Marwah, Shivakant Mishra, and Christof Fetzer</i>	

Adaptive Batching for Replicated Servers.....	311
<i>Roy Friedman and Erez Hadad</i>	
Satem: Trusted Service Code Execution across Transactions	321
<i>Gang Xu, Cristian Borcea, and Liviu Iftode</i>	

Session 10: Panel

Panelists: *Edgar Nett, University of Magdeburg, Germany*
Santosh Shrivastava, University of Newcastle, UK
Kane H. Kim, University of California, Irvine, USA
John K. Davies, BAE Systems, UK

Coordination in Loosely Coupled Systems.....	337
<i>Santosh Shrivastava</i>	
Systematic Composition and Analyzability of Dependable Networked Embedded Computing Systems	339
<i>K. H. Kim</i>	
Open and Challenging Research Issues in Dependable Distributed Computing: A Personal View from the Defence Industry	341
<i>John K. Davies</i>	

Session 11: Distributed Fault-Tolerant Multimedia Systems

AVCast: New Approaches for Implementing Availability-Dependent Reliability for Multicast Receivers	345
<i>Thadpong Pongthawornkamol and Indranil Gupta</i>	
MOve: Design of an Application-Malleable Overlay	355
<i>Sébastien Monnet, Ramsés Morales, Gabriel Antoniu, and Indranil Gupta</i>	
An SNMP Based Failure Detection Service	365
<i>Matthias Wiesmann, Péter Urbán, and Xavier Défago</i>	

Session 12: Fault-Tolerant Database Systems

Lightweight Reflection for Middleware-Based Database Replication	377
<i>J. Salas, R. Jiménez-Peris, M. Patiño-Martínez, and B. Kemme</i>	
Improving DBMS Performance through Diverse Redundancy	391
<i>Vladimir Stankovic and Peter Popov</i>	
Managing Transaction Conflicts in Middleware-Based Database Replication Architectures	401
<i>F. D. Muñoz-Escoí, J. Pla-Civera, M. I. Ruiz-Fuertes, L. Irún-Briz, H. Decker, J. E. Armendáriz-Iñigo, and J. R. González de Mendivil</i>	

Author Index	421
---------------------------	-----