



Preliminary Schedule
TUESDAY, JUNE 24

08:30-09:00

Tutorial Registration

09:00-12:30

(coffee etc.
provided at 10:30-
11:00)

Parallel Tutorials

Tutorial 1: Spam and Denial of Information Attacks and Defenses
Calton Pu, Georgia Institute of Technology, USA

Tutorial 2: Software Patterns for Fault Tolerance
Robert Hanmer, Alcatel-Lucent, USA

Tutorial 3: Database Replication and Clustering
Rui Oliveira, Universidade do Minho, Portugal, Luis Rodrigues, IST/INESC-ID, Portugal and Jose Pereira, Universidade do Minho, Portugal

12:30-14:00

Lunch (not provided)

14:00-17:30

(coffee etc.
provided at 15:30-
16:00)

Parallel Tutorials

Tutorial 4: Secure Event Dissemination in Publish-Subscribe Networks
Ling Liu, Georgia Institute of Technology, USA

Tutorial 5: Architecting Dependable Systems: Preventing and Tolerating Faults
Rogério de Lemos, University of Kent, UK and Cristina Gacek, Newcastle University, UK

Tutorial 6: Soft Errors: technology Trends, System Effects, and Protection Techniques
Subhasish Mitra, Stanford University, USA and Pia Sanda, IBM, USA

Tutorial 1

Tutorial: Spam and Denial of Information Attacks and Defenses

Calton Pu

School of Computer Science
Georgia Institute of Technology

Spam is a significant information security threat that has affected many email and web users. We will outline the impact of spam in several communications media, including email, web, phishing, and social networks. Further, we use spam as an illustrative instance of the general class of denial of information (DOI) attacks, which can be seen as the information analog of denial of service (DoS) attacks. Examples of significant research challenges in DOI attacks include: camouflage and randomization attacks, adversarial (adaptive) classification, evolution of spam production and delivery. We will discuss research results in DOI defense techniques that successfully address some of these challenges, e.g., in the resistance to adversarial adaptation such as camouflage and randomization attacks. The tutorial concludes with a discussion on current opportunities and challenges in spam and DOI research.

About the speaker

Calton Pu received his PhD from University of Washington in 1986. He has taught at Columbia University and Oregon Graduate Institute, currently holding the position of Professor and John P. Imlay, Jr. Chair in Software in the School of Computer Science, Georgia Institute of Technology. He has published more than 250 papers in journals, book chapters, conference proceedings, and refereed workshops in several system-related areas, including operating systems, transaction processing, systems reliability, security, and Internet data management. He is working on spam and denial of information (with several academic and industry partners), service computing (with IBM Research), and automated system management (with HP Labs). He has served on more than 100 program committees for more than 50 international conferences and workshops, including PC co-chair of SRDS, ICDE, CoopIs, DOA, and general co-chair for CIKM, ICDE, CEAS, and SCC.

Research Collaborators and Tutorial Contributors: **Mr. Steve Webb**, **Mr. Byungki Byun**, and **Prof. Chin-Hui Lee** of Georgia Tech, **Prof. Kang Li** and **Prof. Lakshminish Ramaswamy** of University of Georgia, **Prof. James Caverlee** of Texas A&M University.

Tutorial 2

Software Patterns for Fault Tolerance

Robert S. Hanmer
Alcatel-Lucent

This is an overview tutorial that introduces software patterns and how they can be used to communicate the principles of reliability. Software patterns have been discussed in the software design and development community for more than a decade. Practitioners find them useful to learn new ways to design their software and also references to remember things that they already know. Patterns describe solutions that have withstood the test of time in production systems.

Three main topics will be covered in this tutorial. The first is a basic introduction to patterns, the worldwide pattern community and pattern resources. Patterns work together to generate solutions to problems that are non-obvious through pattern languages; pattern languages and other collections of patterns are the second topic to be covered. The third is the use of pattern languages to document software and system architectures. Illustrative patterns and pattern languages will come from software fault tolerance.

At the end of this tutorial the participants will be able to describe patterns, know where to find pattern resources, understand the basics of pattern languages and how patterns work together, and also be introduced to a variety of patterns for software fault tolerance.

About the speaker

Robert S. Hanmer is a consulting member of technical staff in the Technical Component Management area in Alcatel-Lucent's Procurement and Sourcing area. He is based in Naperville, Illinois. Current responsibilities include developing software sourcing strategies for a number of technology cells including availability enhancing middleware. Previous positions within Lucent Technologies and Bell Laboratories have included development, architecture and evaluation of highly reliable systems focusing especially on the areas of reliability and performance. He is active in the software patterns community, including serving as program chair at several pattern conferences and teaching pattern writing classes. He has authored or co-authored two book chapters of fault tolerance patterns and a book on patterns for software fault tolerance. He is a member of the IEEE Computer Society, the Hillside Group, and is a Senior Member of the ACM. He received his B.S. and M.S. degrees in Computer Science from Northwestern University in Evanston, Illinois, USA.

Contact info: Robert S. Hanmer, Consulting Member of Technical Staff, Alcatel-Lucent, 2000 Lucent Lane 2H-207, PO Box 3033, Naperville, IL 60566-7033 USA, hanmer@alcatel-lucent.com, Voice: +1 630 979 4786

Tutorial 3

Database Replication and Clustering

Rui Oliveira

Dep. Informática/CCTC, Universidade do Minho
rco@di.uminho.pt

Luís Rodrigues

IST/INESC-ID
ler@ist.utl.pt

José Pereira

Dep. Informática/CCTC, Universidade do Minho
jop@di.uminho.pt

Database replication and clustering are increasingly widespread solutions to scalability and high availability challenges in current information systems. The tutorial offers a complete overview of these techniques, from theory to practice. It starts by making an overview of the main challenges involved and the algorithms and engineering techniques to address them. Then, it provides an in depth description of the architecture of a modern open-source database replication system and a hands-on guide to its deployment. Therefore, the tutorial aims both at database management system and middleware developers, as well as those making use of database technology in planning and deploying large scale applications.

About the speakers

Rui Oliveira graduated (1991) in Electrotechnic and Computers Engineering by Universidade do Porto, has a Master (1994) in Computer Science from the Universidade do Minho and a PhD (2000) from the Swiss Federal Institute of Technology in Lausanne. He is associate Professor in Computer Science at Universidade do Minho, teaching dependable distributed systems to Masters and Doctoral Programmes. He currently leads the Computer Science and Technology Center of Universidade do Minho. Rui Oliveira is the Project Manager of the GORDA project an european research project devoted to consistent database replication systems. His current research interests include fault-tolerant and large scale distributed systems, distributed data management and peer-to-peer computing. He is a member of ACM and IEEE.

Luís Rodrigues graduated (1986), has a Master (1991) and a PhD (1996) in Electrotechnic and Computers Engineering, by the Instituto Superior Técnico da Universidade Técnica Lisboa (IST). He obtained the "Agregação" in Informatics (2003)

by the Universidade de Lisboa. He is full Professor at Departamento de Engenharia Informática, of Instituto Superior Técnico da Universidade Técnica Lisboa (IST) and a senior researcher and Distributed Systems Group of INESC-ID. Previously he was at the Faculdade de Ciências da Universidade de Lisboa where he was a founding member of the LASIGE laboratory and leader of the Distributed Algorithms and Network Protocols group. His current research interests include fault-tolerant and real-time distributed systems, group membership and communication, replicated data management, publish-subscribe systems, peer-to-peer computing and mobile computing. He has more than 100 publications in these areas. He is co-author of two textbooks on distributed computing. He is a member of the Ordem dos Engenheiros, ACM, and IEEE.

José Pereira graduated (1995) and has a Master (1998), and a PhD (2002) from the Universidade do Minho. He is an assistant Professor in Computer Science at Universidade do Minho, teaching distributed systems foundations, transactional systems, and data replication topics to Masters and Doctoral Programmes. He is currently the Technical Manager of the GORDA – Open Replication of Database european research project. His current research interests include dependable distributed systems, namely, on database replication and large scale group communication.

Tutorial 4

Secure Event Dissemination in Publish-Subscribe Networks

Ling Liu
College of Computing
Georgia Institute of Technology
lingliu@cc.gatech.edu

With the increasing popularity of publish-subscribe networks and event-driven computing architecture, event dissemination in publish-subscribe networks is a target of adversaries. There are many security threats to event-dissemination networks. Adversary can prevent users from getting correct data from pub-sub nodes by modifying the contents of the messages, or spoofing the identity of the publishers, subscribers, or pub-sub routing nodes. An adversary can block communication between a publisher and subscribers by creating false routing information, or simply generating jamming signals. An adversary can gain control of an entire pub-sub tree by spoofing the identity of a publisher. An adversary can compromise a routing node, get all information from that node, and can even re-program it to behave like a malicious node. The design and implementation of a secure and dependable publish-subscribe network must simultaneously address three research challenges: (1) Vulnerability of event dissemination in pub-sub networks to eavesdropping, unauthorized access, spoofing, replay, and denial of service attacks; (2) Challenges of providing secure content-based event dissemination while providing complex publication-subscription matching in the pub-sub networks, and supporting anonymity of publishers and subscribers; and (3) Added security risk of individual nodes being compromised to behave like a malicious node.

This tutorial will discuss some of the latest techniques that have been proposed to address these research challenges. The tutorial is designed to be self-contained, and gives the essential background for anyone interested in learning about the concept, the alternative models and techniques for secure event dissemination, and the general principles and techniques for design and development of a secure and efficient publish-subscribe architecture for scalable and dependable event dissemination. The main objective of this tutorial is to provide an in-depth coverage of the design and implementation issues in building a dependable and secure publish-subscribe systems and applications, the key trade-offs in secure event dissemination, as well as the limitations of current approaches.

About the speaker

Dr. Ling Liu is an Associate Professor in the College of Computing at Georgia Institute of Technology. There she directs the research programs in Distributed Data Intensive Systems Lab (DiSL), examining performance, security, privacy, and data management issues in building large scale distributed computing systems. Dr. Liu and the DiSL research group have been working on various aspects of distributed data intensive systems, ranging from event driven architecture, secure and scalable processing of complex events. Decentralized overlay networks, including publish-subscribe networks, to mobile computing and location based services, sensor network and event stream processing. She has published over 200 international journal and conference articles in the areas of Internet Computing systems, Internet data management, distributed systems, and information security. Dr. Liu is currently on the editorial board of several international journals, including IEEE Transactions on Knowledge and Data Engineering, IEEE Transactions on Service Computing, International Journal of Peer-to-Peer Networking and Applications (Springer), International Journal of Web Services Research, Wireless Network Journal (WINET). Dr. Liu is the recipient of the best paper award of ICDCS 2003, the best paper award of WWW 2004, a recipient of 2005 Pat Goldberg Memorial Best Paper Award. Dr. Liu's research is primarily sponsored by NSF, DARPA, DoE, and IBM. For more information see <http://www.cc.gatech.edu/~lingliu>.

Tutorial 5

Architecting Dependable Systems: Preventing and Tolerating Faults

Rogério de Lemos
University of Kent, UK

Cristina Gacek
Newcastle University, UK

The aim of this tutorial is to provide an insight on how the structuring of software systems at the architectural level is fundamental for the development of dependable systems. Taking as a basis the different dependability means, the intention is to show how dependability should be considered at the architectural level, and the impact this should have when developing dependable systems. The main objectives of this tutorial are the following:

- to establish the major principles associated with software architectures and dependability that are relevant when reasoning about faults at the architectural level;
- to introduce and discuss existing approaches for architecting dependable systems, particularly, in the context of fault prevention and tolerance;
- to identify the main challenges that lie ahead when considering the structuring of dependable systems at the architectural level.

At the end of the tutorial, the participants should have a better appreciation of the challenges, problems and solutions that are currently associated with the structuring of dependable systems at the architectural level. These should include methods, techniques, and tools that are relevant in the context of dependability means, mainly, rigorous design, and fault tolerance. The material of the tutorial will be presented in the context of several case studies, including, embedded and service oriented systems, and we will be using UML2.0, ACME, and AADL description languages to support our examples.

About the speakers

Rogério de Lemos is a Lecturer in Computing Science at the University of Kent (UK). Before joining the Computing Laboratory at the University of Kent he was a Senior Research Associate at the Centre for Software Reliability (CSR) at the University of Newcastle upon Tyne (UK). He has participated in several conference PCs, including HASE 2001 and HASE 2002; FMRTFT 2002; EDCC-4, EDCC-6 and EDCC-7; ISADS 2003; WICSA 2005, WICSA 2007, and WICSA 2008; LADC 2005; ICSE 2006 Emerging Results Track; SRDS 2006; and ICDCS 2006; and the PC co-chair of LADC

2003. Until recently, he was a member of the Steering Committee of LADC, and he is on the editorial board of the Journal of Hybrid Systems. He gave a tutorial at SCTF 1997 on safety analysis of critical software, and tutorials at LADC 2005 and ICSE 2006, DSN 2007, LADC 2007 on software architectures for dependable systems. He has co-edited four books and a journal special issue on Architecting Dependable Systems.

Cristina Gacek is a Lecturer at the School of Computing Science of the Newcastle University (UK). She has extensive work experience both as a researcher and as a practitioner. She received a Ph.D. in Computer Science from the University of Southern California (USC - USA), where she worked as a research assistant. Dr. Cristina Gacek was the leader of the Software Architectures group at the Fraunhofer Institute for Experimental Software Engineering (IESE), and worked both for TRW (USA) and IBM Brasil. She has participated in several conference program committees, was the program committee chair for ICSR-7, and has co-organized many workshops, including several on Architecting Dependable Systems. She has given a tutorial on software architectures for product lines at ICSR-6 and on software architectures for dependable systems at DSN 2007. Cristina is an IEEE and ACM member and has several scientific publications. She is a co-editor of four books and a journal special issue on Architecting Dependable Systems.

Tutorial 6

Soft Errors: Technology Trends, System Effects, and Protection Techniques

Subhasish Mitra
Stanford, US
subh@stanford.edu

Pia Sanda
IBM, US

Radiation-induced soft errors are getting worse in digital systems manufactured in advanced technologies. Stringent data integrity and availability requirements of enterprise computing and networking applications demand special attention to soft errors in sequential elements and combinational logic. This tutorial will discuss the impact of technology scaling on soft error rates, circuit-level modeling of soft errors, architectural impact of soft errors, challenges associated with evaluation of run-time behaviors of systems in the presence of soft errors, actual data on system behaviors in the presence of soft errors, metrics for quantifying soft error vulnerabilities, design of architectures with Built-in-Soft-Error-Resilience techniques, and actual case studies. Two of the presenters co-founded a new workshop on soft errors (SELSE 2005-2007). Lessons learnt from these workshops will also be included in the tutorial.

Targeted audience: Researchers and practitioners interested in reliable systems modeling, architecture, design, CAD.

About the speakers

Subhasish Mitra is an Assistant Professor in the Departments of Electrical Engineering and Computer Science of Stanford University. His research interests include robust system design, VLSI design and test, computer architecture and design for emerging nanotechnologies. Prior to joining Stanford, he was a Principal Engineer at Intel Corporation where he developed enabling technologies for robust system design – Design for Excellence (Reliability, Testability, and Debug) – that have been deployed in several products. He received Ph.D. in Electrical Engineering from Stanford University.

Prof. Mitra has co-authored more than 90 technical papers and several patents, and invented design and test techniques that have seen wide-spread proliferation in the chip design industry. His X-Compact technique for test compression is used by major (more than 40) Intel products including microprocessors, chipsets, and communications chips, and is supported by major CAD tools. His recent work on imperfection-immune circuit

design using carbon nanotubes, jointly with his students and collaborators, has been highlighted in the MIT Technology Review, EE Times, and several other publications. Prof. Mitra's major honors include the National Science Foundation CAREER Award, Terman Fellow at Stanford University, IEEE Circuits and Systems Society Donald O. Pederson Award for the best paper published in the IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, a Divisional Recognition Award from Intel "for a Breakthrough Soft Error Protection Technology," a Best Paper Award at the Intel Design and Test Technology Conference for his work on Built-In Soft Error Resilience, and the Intel Achievement Award, Intel's highest corporate honor, "for the development and deployment of a breakthrough test compression technology that improved scan test cost by an order of magnitude." Prof. Mitra has held several consulting positions, and served on the committees of several IEEE and ACM sponsored conferences and workshops as co-founder, general and program chair, and organizing and program committee member.

Pia Sanda received the Ph.D. degree in physics from Cornell University, Ithaca, NY, in surface Raman scattering. She was a Manager in the VLSI Design Area, IBM T. J. Watson Research Center, Yorktown Heights, NY. She began her career at IBM in imaging science. In silicon technology, she designed and built 0.1- μ m channel length CMOS FET's using phase shift lithography and contributed to the device and cell design for the 256-Mb DRAM. She has been engaged in designing high-performance circuits for microprocessors and has recently explored new avenues for test and improved semiconductor manufacturability, such as the new PICA measurement technique. She is currently the Program Director of Soft Error Management at IBM.