QEST '04

Quantitative Evaluation of Sys-Tems '04

First International Conference on the
Quantitative Evaluation of SysTems

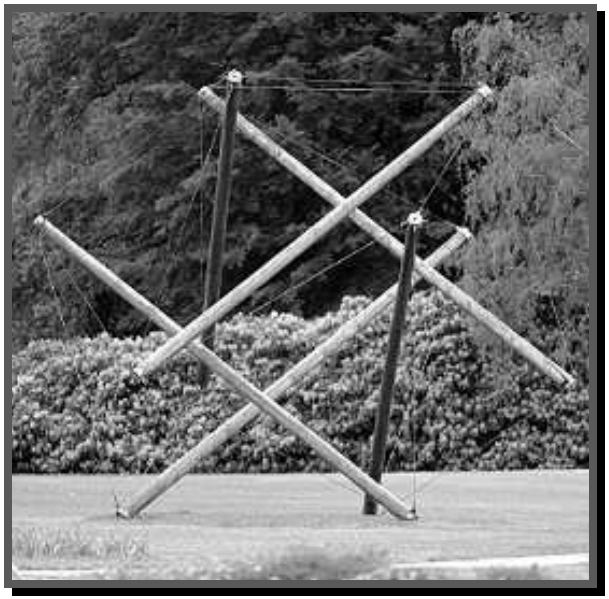

Final Programme

September, 27 – 30, 2004
University of Twente, The Netherlands

Welcome at the University of Twente!

It is my pleasure to welcome you at the University of Twente for the first International Conference on the Quantitative Evaluation of SysTems.

As you know, QEST is both an old and a new conference. Old as it is the successor of three successful previous conference series, new as it is the first time that it is held in its fully integrated form under its new name.

Looking at the programme, I think that QEST is already a conference we can be proud of. The programme committee has put together a very strong programme, out of the over 80 submitted papers. The selected papers, plus three excellent invited speakers, a variety of tool demonstrations and a strong full-day tutorial programme should make your visit to the "nicest campus university in the Netherlands" certainly worthwhile.

The organising committee has tried to make the programme also attractive from a social and informal point of view, with lunches at different locations, a visit to the Rijksmuseum Twenthe and a conference dinner in the park-restaurant "De Jaargetijden" in the center of Enschede.

I thank all authors of submitted papers, all reviewers for their hard work, the programme committee for their careful selection process and the local organisers for their enormous efforts. Of course, I thank you, for being here, for sharing your ideas about the quantitative evaluation of systems with all of us. Together, you make QEST 2004 happen!

Finally, I am glad to thank the sponsors of QEST 2004. The financial support of the Royal Netherlands Academy of Arts and Sciences (KNAW), the Netherlands, the Netherlands Organisation for Scientific Research (NWO), the Dutch research school Institute for Programming Research and Algorithmics (IPA), the Department of Electrical Engineering, Mathematics and Computer Science at the University of Twente, and the Centre for Telematics and Information Technology at the University of Twente helped us enormously in organising this event.

I wish you all a very successful and enjoyable conference!

Boudewijn R. Haverkort
General Chair QEST 2004

Location of the Conference

The tutorials will be held in room U1, U2 and U3 of building 11 “De Zilverling”. Please follow the signs after you have registered in “De Waaier” (see figure below).

The conference sessions are in building 12 “De Waaier”, in room 3 and 4.
Please note that it is prohibited to smoke in public areas of the university buildings.

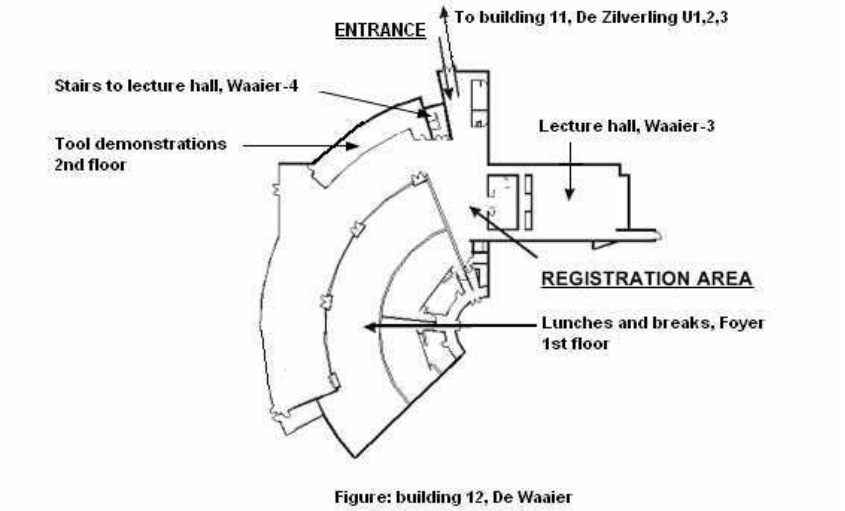

Figure: building 12, De Waaier

E-mail and internet facilities

It will be possible to access the internet during the conference. We will provide access to the campus-wide wireless network and to a number of computers with SSH and a web browser. To access the wireless network you need to bring your own laptop and wireless card. Before you can access the wireless network you will have to register the MAC address of the wireless network card. To register follow the instructions at http://qest2004.cs.utwente.nl/registration/regwireless.php. You can also mail us with the required information at qest2004@cs.utwente.nl.

Social event

On Wednesday September, 29 the social event will take place. We will go by bus to the Rijksmuseum Twenthe, where we will have an introduction into Dutch 18th century painting and industrial art. The 18th century has been a remarkable and a tremendously important period in Dutch history. All achievements of our society as it is today have their origin in this period.

We will leave by bus starting from 15:00. The visit at the museum will start at 15:30 with coffee and cake followed by a short introduction.
At 18:00, restaurant “De Jaargetijden” will gladly welcome us with a drink followed by Mediterranean buffet which will start at around 19:00.

Addresses

Twenthe Rijksmuseum
Lasondersingel 129-131
7514 BP Enschede
Phone: +31 (0) 53-435 8675

Restaurant De Jaargetijden
Parkweg 49
7513 CN Enschede
Phone: +31 (0) 53-431 3144

Time Table	Monday (Sep. 27)	Tuesday (Sep. 28)	Wednesday (Sep. 29)	Thursday (Sep. 30)
8:30	**REGISTRATION** Waaier, Entrance **08:30 - 17:30**	**REGISTRATION at 08:00** **Opening Session** Waaier-4 **08:30 - 09:00**	**Invited speaker** **M. Colajanni** Waaier-4 **08:30 - 09:30**	
9:00	**Tutorials T1.1. T2 T3.1** Zilverling-U1,2,3 **09:00 - 10:30**	**Invited speaker** **G. Koole** Waaier-4 **09:00 - 10:00**		**Invited Speaker** **T.A. Henzinger** Waaier-4 **09:00 - 10:00**
9:30			**Break** Waaier, Foyer **09:30 - 10:00**	
10:00		**Break** Waaier, Foyer **10:00 - 10:30**	**Session IV: Analysis, simulation and synthesis of MC** Waaier-4 **10:00 - 12:30**	**Break** Waaier, Foyer **10:00 - 10:30**
10:30	**Break** Waaier, Foyer **10:30 - 11:00**	**Session I: Case studies** Waaier-4 **10:30 - 12:30**		**Session VI: Model Checking** Waaier-4 **10:30 - 12:30**
	Tutorials T1.2 T2 T3.2 Zilverling-U1,2,3 **11:00 - 12:30**			**Parallel:** **Tool Demonstrations** Waaier, Foyer **10:30 - 12:30**
12:30	**Lunch** Waaier, Foyer **12:30 - 13:30**	**Lunch** Faculty Club **12:30 - 14:00**	**Lunch** (parallel: tools demo) Waaier, Foyer **12:30 - 13:30**	**Lunch** Faculty Club **12:30 - 14:00**

13:30	**Tutorials T4.1 T5 T6.1** Zilverling- U1,2,3 **13:30 - 15:00**	**Lunch till 14:00**	**Session V: Scheduling and optimization, 13:30 - 15:00** Waaier-4	**Lunch till 14:00**
14:00		**Session II: Time and continuous Petri nets** Waaier- 4 **14:00 - 15:30**		**Session VII: Symbolic techniques** Waaier-4 **14:00 - 15:30**
15:00	**Break** Waaier, Foyer **15:00 - 15:30**		**Start Social Event: assembly in front of "De Zilverling"** **15:00**	**Parallel:** **Tool Demonstrations** Waaier, Foyer **14:00 - 15:30**
15:30	**Tutorials T4.2 T5 T6.2** Zilverling-U1,2,3 **15:30 - 17:00**	**Break** Waaier, Foyer **15:30 - 16:00**		**Break** Waaier, Foyer **15:30 - 16:00**
16:00		**Session IIIa: Queuing theory** Waaier-4 **16:00 - 18:00**		**Session VIII: Semantics** Waaier-4 **16:00 - 17:00**
16:30		**Session IIIb: Tool presentations** Waaier-3 **16:00 - 18:00**		**Parallel:** **Tool Demonstrations** Waaier, Foyer **16:00 - 17:00**
17:00	**Reception** Faculty Club **17:00 - 18:30**			**Closing session** Waaier-4 **17:00 - 17:30**
18:00			**Dinner** **"De Jaargetijden"** **starting at 18:30**	
18:30	**Dinner at Faculty Club** (registration needed)			

Detailed conference programme

Tuesday 28 September, 2004

Opening
Building: De Waaier-4, 08:30 – 09:00

President of the University of Twente: prof.dr. F. van Vught

General chair QEST 2004: prof.dr.ir. B.R.H.M. Haverkort

PC chairs of QEST 2004:
prof. dr. G. Franceschinis, dr.ir. J.P. Katoen, prof. dr. M. Woodside

Invited speaker
Building De Waaier-4, 09:00 – 10:00
Session chair: Murray Woodside

G. Koole
Performance analysis and optimization in customer contact centers.

Break Building De Waaier, Foyer 10:00 – 10:30

Session I : Case studies
Building De Waaier-4, 10:30 – 12:30
Session chair: Markus Siegle

A performance evaluation tool for RAID disk arrays.
Alexander Thomasian, Chunqi Han, Gang Fu and Chang Liu

Comparing passive and active worm defenses.
Michael Liljenstam and David Nicol

Stochastic assessment of schedules for lacquer production.
Henrik Bohnenkamp, Holger Hermanns, Ric Klaren, Angelika Mader and Yaroslav Usenko

Performance Model Interchange Format (PMIF 2.0): XML definition and implementation.
Connie Smith and Catalina Llado

Lunch at the **Faculty Club** 12:30 – 14:00

Session II: Time and continuous Petri nets
Building De Waaier-4, 14:00 – 15:30
Session chair: Gianfranco Balbo

On performance bounds for interval time Petri nets.
Simona Bernardi and Javier Campos

On observability in timed continuous Petri net systems.
Jorge Julvez, Emilio Jimenez, Laura Recalde and Manuel Silva

ORIS: a tool for state space analysis of real-time preemptive systems.
G. Bucci, Luigi Sassoli and E. Vicario

Break Building De Waaier, Foyer 15:30 – 16:00

Session IIIa: Queuing theory
(parallel to session IIIb: Tool presentations)
Building De Waaier-4, 16:00 – 18:00
Session chair: Peter Buchholz

Bit-level and packet-level, or Pollaczec-Khintine formulae revisited.
Jose Incera and Gerardo Rubino

Sojourn time distributions in modulated G-queues with batch processing.
Peter Harrison and H. Zatschler

ETAQA truncation models for the MAP/MAP/1 departure process.
Armin Heindl, Qi Zhang and Evgenia Smirni

A new approach for the prediction of end-to-end performance of multimedia streams.
Gerardo Rubino and Martin Varela

Session IIIb: Tool presentations

(parallel to session IIIa: Queuing theory)
Building De Waaier, Foyer 16:00 – 18:00
Session chair: Andrew Miner

D-ITG: Distributed Internet Traffic Generator
S. Avallone, D. Emma, S. Guadagno, A. Pescapè, G. Ventre

Visualization of Analysis Results in the ProC/B Toolset
M. Arns, M. Fischer, C. Tepper, M. Völker

Performance Analysis with BenchIT: Portable, Flexible, Easy to Use
Guido Juckeland, Michael Kluge, Wolfgang E. Nagel, Stefan Pflüger

PRISM 2.0: A Tool for Probabilistic Model Checking
Marta Kwiatkowska, Gethin Norman, David Parker

Extended SWN solvers in GreatSPN
J.M. Ilié, S. Baarir, M. Beccuti, C. Delamare, S. Donatelli, C. Dutheillet, G. Franceschinis, R. Gaeta, P. Moreaux, S. Alciati

Modeling and Evaluation with Object Stochastic Activity Networks
Mohammad Abdollahi Azgomi and Ali Movaghar

The Möbius Modeling Environment: Recent developments
Tod Courtney, David Daly, Salem Derisavi, Shravan Gaonkar, Mark Griffith, Vinh Lam, William H. Sanders

TSMV: A Symbolic Model Checker for Quantitative Analysis of Systems
Nicolas Markey and Philippe Schnoebelen

WALTy: A User Behavior Tailored Tool for Evaluating Web Application Performance
Roberto Politi, Giancarlo Ruffo, Rossano Schifanella and Matteo Sereno

The ipc/HYDRA Tool Chain for the Analysis of PEPA Models
Jeremy T. Bradley and William J. Knottenbelt

TwoTowers 4.0: Towards the Integration of Security Analysis and Performance Evaluation
Alessandro Aldini and Marco Bernardo

SMART: the Stochastic Model checking Analyzer for Reliability and Timing
Gianfranco Ciardo and Andrew S. Miner

UML-PSI: The UML Performance Simulator
Moreno Marzolla and Simonetta Balsamo

XPRIT: an XML-based tool to translate UML diagrams into execution graphs and queueing networks
Vittorio Cortellessa, Michele Gentile and Marco Puzziti

Wednesday 29 September, 2004

Invited speaker

Building De Waaier-4, 08:30 – 09:30
Session chair: Giuliana Franceschinis

M. Colajanni
Emerging Internet-based services: New frontiers for performance models and applications

Break Building De Waaier, Foyer 09:30 – 10:00

Session IV: Analysis, simulation and synthesis of Markov Chains

Building De Waaier-4, 10:00 – 12:30
Session chair: Gianfranco Ciardo

Approximate computation of transient results for large Markov chains.
Peter Buchholz and William Sanders

A structured path-based approach for computing transient rewards of large CTMCs.
Vinh Lam, Peter Buchholz and William Sanders

Learning continuous time Markov chains from sample executions.
Koushik Sen, Mahesh Viswanathan and Gul Agha

On numerical problems in simulations of highly reliable Markovian systems.
Bruno Tuffin

Evaluation of reward analysis methods with MRMSolve 2.0.
Gabor Horvath, S. Racz, A. Tari and Miklos Telek

Lunch Building De Waaier, Foyer 12:30 – 13:30

Session V: Scheduling and optimization

Building De Waaier-4, 13:30 – 15:00
Session chair: William H. Sanders

Adaptive disk scheduling for overload management.
Alma Riska, Erik Reidel and Sami Iren

Improving the performance of online auction sites through closing time rescheduling.
Daniel Menasce and Vasudeva Akula

Analysis and algorithms for restart.
Aad van Moorsel and Katinka Wolter

Social Event

The social event starts at 15:00. Assemble in front of "De Zilverling" in time, from which busses leave. We will have dinner together in restaurant "De Jaargetijden", starting around 18:30.

Thursday 30 September, 2004

Invited speaker

Building De Waaier-4, 09:00 – 10:00
Session chair: Joost-Pieter Katoen

T.A. Henzinger
Trading memory for randomness

Break Building De Waaier, Foyer 10:00 – 10:30

Session VI: Model checking

Building De Waaier-4, 10:30 am – 12:30
Session chair: Ed Brinksma

Backward stochastic bisimulation in CSL model checking.
Jeremy Sproston and Susanna Donatelli

Partial order reduction for probabilistic systems.
Christel Baier, Marcus Groesser and Frank Ciesinski

Partial order reduction on concurrent probabilistic programs.
Pedro Ruben D'Argenio and Peter Niebert

Dependability checking with StoCharts: Is train radio reliable enough for trains?
David N. Jansen and Holger Hermanns

Lunch at the **Faculty Club** 12:30 – 14:00

Session VII: Symbolic techniques

Building De Waaier-4, 14:00 – 15:30
Session chair: Peter Kemper

Improving efficiency of implicit Markov chain state classification.
Andrew Miner and Shuxing Cheng

Saturation NOW.
Ming-Ying Chung and Gianfranco Ciardo

Saturation for a general class of models.
Andrew Miner

Break Building De Waaier-4 Foyer 15:30 – 16:00

Session VIII: Semantics

Building De Waaier-4, 16:00 – 17:00
Session chair: Christel Baier

Axiomatization of trace semantics for stochastic nondeterministic processes.
Augusto Parma and Roberto Segala

Approximate reasoning for real-time probabilistic processes.
Vineet Gupta, Radha Jagadeesan and Prakash Panangaden

Closing Session

Building De Waaier-4 17:00 – 17:30
Chair: William Sanders and Boudewijn Haverkort

Tutorial programme

T1: Software performance models

T1.1. : Software Model to Performance Model Transformations
Building De Zilverling, room U1, 09:00 – 10:30

Antinisca Di Marco, Vittorio Cortellessa, Paola Inverardi, University of L'Aquila, Italy

Abstract
It is widely recognized that in order to make performance validation an integrated activity along the software lifecycle it is crucial to be supported from automated approaches. Easiness to annotate software models with performance parameters and efficient translations of the annotated models into "ready-to-validate" models are the key challenges in this direction. Several methodologies have been introduced in the last few years to perform the annotation and translation tasks. This tutorial introduces the attendance to the main methodologies for annotating and transforming software models into performance models. We shortly introduce the topic, and then the methodology descriptions are driven from a common simple example, in order to show how they can be used in practice. Along the tutorial, the methodologies are classified following different dimensions and parameters. Quite recently the Object Management Group has published a Call for Proposal aimed at standardizing the model-to-model transformation domain. The standardization process seems to be in its final phase, and the different proposals are going to be merged into standard OMG specifications. We also intend to show how these approaches shall take advantage from the OMG standard to come.

Break Building De Waaier, Foyer 10:30 – 11:00

T1.2. : Performance Analysis Based on the UML SPT Profile
Building De Zilverling, room U1, 11:00 – 12:30

Dorina Petriu, Carleton University, Canada

Abstract
The "UML Profile for Schedulability, Performance and Time" (STP) standardized by OMG enables the use of UML models for quantitative predictions regarding schedulability and performance characteristics. The tutorial will discuss the capabilities and limitations of the present version of the Performance Profile, which is due for an upgrade to align it with UML 2.0. Moreover, the STP Profile needs to be harmonized with the emerging "UML Profile for Quality-of-Service and Fault Tolerance" which supports modelling a wide range of QoS concepts. The tutorial will continue by presenting a two-step approach developed at Carleton University for building automatically performance models from UML design specifications. In the first step, an annotated UML model describing a set of scenarios and their use of resources is transformed into an intermediate format, named Core Scenario Model (CSM). CSM is based on the STP profile, and contains only the relevant performance information extracted from different UML diagrams in a

form that can be checked for consistency. In the second step, the CSM model is transformed into a performance model. Different performance modelling formalisms are considered, such as layered queueing networks, timed Petri nets and simulation.

T2: Quantitative program analysis for security

Quantitative Program Analysis and Applications to Computer Security
Building De Zilverling, room U2, 09:00 – 12:30 (break: 10:30 – 11:00)

Alessandra di Pierro, University of Pisa, Italy and Herbert Wiklicky, Imperial College, UK

Abstract
Program analysis provides semantics based compile-time techniques for statically predicting safe and computable approximations to the set of behaviours arising dynamically when executing a program on a computer. Quantitative program analysis aims at developing techniques which provide approximate answers with some numerical estimate of the precision of the approximation. One useful source of numerical information for a quantitative program analysis is a probabilistic semantics and in particular the use of vector space or linear algebraic structures for modelling the computational domain. By exploiting the probabilistic information resulting from a probabilistic program analysis one can quantify the precision of the analysis and obtain as a result answers which are for example "approximate up to 35%". We will illustrate such a quantitative approach in the framework of Probabilistic Abstract Interpretation which provides a general methodology for constructing quantitative analyses. We will then demonstrate this technique by presenting some applications in computer security. We will concentrate on the problem of preventing unauthorised users from reading sensitive information, aka confidentiality, and show how a quantitative program analysis of confidentiality can be defined which gives a measure for the confinement of a system. We will also illustrate applications to the analysis of other security aspects which are crucial in a distributed setting, namely firewalls and network vulnerability.

T3 Structured analysis of Markov chains

T3.1. : Structured analysis of Markov models
Building De Zilverling, room U3, 09:00 – 10:30

Peter Buchholz, University of Dortmund, Germany

Abstract
Markov models are a commonly used model type to analyze all kinds of systems. Although the general approach of generating Markov chains from high level specifications and of analyzing the resulting Markov chain are well understood and easy to realize, practical problems occur due the so called state space explosion which means that the size of the state space of the Markov chain grows exponentially with the size of the model. Recently a remarkable step forward has been made in generating and storing huge state space and transition matrices with 10^{20} states and beyond using advanced data structures. Based on these representations functional properties of the system can be analyzed. Additionally, these compact storage schemes have been used for quantitative analysis via numerical computation of the stationary or transient distribution. However, although the new techniques enlarge the size of solvable state spaces significantly, the size is limited by the size of the probability vector which allows the analysis of systems with 10^7 states, but not with 10^{20} states. Thus, new approaches are required which avoid storing all states and still allow the computation of quantitative results. The tutorial gives an overview of different recently developed approaches which exploit the compact representation of the whole states space for quantitative analysis without storing the whole vector of state probabilities. In particular: approximation techniques based on fixed point computations, bounding techniques, and hybrid techniques combining simulation and transient numerical analysis are presented and compared. Furthermore, it is outlined how these techniques can be integrated in software tools for performance analysis and which research problems are currently open.

Break Building De Waaier, Foyer 10:30 – 11:00

T3.2. : Modeling and Analysis of Markov Chains Using Decision Diagrams
Building De Zilverling, room U3, 11:00 – 12:30

Gianfranco Ciardo, University of California at Riverside, USA

Abstract
Markov chains are a powerful and conceptually easy class of discrete-state stochastic processes that can be used to model a wide variety of systems. However, high modeling fidelity is often achieved only when the state space of the Markov chain is very large; making the numerical solution of these models a formidable task. In this tutorial, we review and compare several techniques to encode a very large matrix, such as the transition rate matrix of a continuous-time Markov chain. Then, we examine numerical solution algorithms that operate on these encodings.

T4: Timed models

T4.1. : Timed Models: From Theory to Implementation
Building De Zilverling, room U1, 13:30 – 15:00

Patricia Bouyer, LSV-CNRS and ENS of Cachan, France

Abstract
Many critical applications have explicit timing constraints. For example, behaviours of systems interacting with an environment (e.g. embedded systems) depend on quantitative timing constraints like response times, transmission delays, etc. For representing such timed systems, several timed models have been proposed. Since their introduction by Alur and Dill in the 90s, timed automata are one of the most-studied and most-established models for real-time systems. Numerous works have been devoted to the comprehension of timed automata and the major property of timed automata is probably that reachability is decidable, which implies in particular that it can be used for verification purposes. Based on this nice theoretical result, several model-checkers have been developed (for instance CMC, HyTech, Kronos and Uppaal) and a lot of case studies have been treated. Decidability properties for timed automata are proved based on the construction of the so-called region automaton: it abstracts finitely and in a correct way the behaviours of timed automata. However, in practice, such a construction is not implemented, because it suffers from an enormous combinatorics explosion, but algorithms searching on-the-fly through the automaton are preferred. These algorithms are based on the notion of zones and can be efficiently implemented using data structures like DBMs. Among these algorithms are the forward analysis algorithms that compute all the reachable configurations from the initial configurations. However, the exact forward computation does not always terminate. Abstractions have to be used to avoid this termination problem but one has to be careful when using such abstractions. In this tutorial, we will present basic results on timed automata as well as abstractions and algorithms that are used in practice.

Break Building De Waaier, Foyer 15:00 – 15:30

T4.2. : Time in Process Algebra: A Conceptual Overview
Building De Zilverling, room U1, 15:30 – 16:30

Mario Bravetti, University of Bologna, Italy

Abstract

It is widely recognized that dealing with time related aspects in process algebra is often crucial for the specification and analysis of complex real systems. Research work in this field has led to a rather huge literature, where several kinds of time have been taken into account: time may be either based on a discrete or continuous domain, time elapsing may be either probabilistically (so-called stochastic-time) or deterministically (so-called real-time) bounded, analysis may be based on model checking, model solving etc. In this tutorial we perform a conceptual dissertation about the treatment of the various kinds of time in transition systems where notions of composition are defined (as e.g. by defining a process algebra). Moreover we show the conceptual relationship between the notion of time considered and the kind of semantics (in the sense of classical process algebra literature) which must be adopted for representing such a notion of time in the composition operators. The presentation of the different kinds of time extensions is formally supported by making use of simple process algebras which extend a common core timed calculus. All together such extensions form a unifying theory for the representation of different kinds of time in process algebras.

T5: Stochastic Games

Stochastic games: a tutorial
Building De Zilverling, room U2 13:30 – 16:30 (break 15:00 – 15:30)

Marcin Jurdzinski, University of Warwick, UK

Abstract
Stochastic games are repeated games with probabilistic transitions and with payoffs that are functions of the whole play, i.e., the history of the players' interaction. The theory of stochastic games is surveyed as a basis for formal specification and algorithmic analysis of quantitative properties of systems. Potential applications include design, verification, and performance evaluation of computational systems, and formalization and analysis of models used in evolutionary and population biology, and in economics. A general framework of games with Borel measurable payoff functions is presented, as well as important special cases of discounted, limiting average, and omega-regular payoffs. Algorithmic methods for solving stochastic games are surveyed, including value iteration, strategy improvement, mathematical programming, and graph theoretic algorithms. Important subclasses of the general stochastic game model are discussed, including Markov decision processes, perfect-information games, and non-stochastic 2-player games. Finally, the quantitative mu-calculus is presented as a unifying formalism for the quantitative temporal specification and verification, and it is given a stochastic games semantics. This allows applying the rich theory of stochastic games to the semantic study and to the algorithmic applications of the quantitative mu-calculus in automated quantitative model checking.

T6: Probabilistic model checking and Random neural nets

T6.1. : Probabilistic Model Checking
Building De Zilverling, room U3, 13:30 – 15:00

Marta Kwiatkowska, University of Birmingham, UK

Abstract
Probability features increasingly often in software and hardware systems: it is used in distributed co-ordination and routing problems, to model fault-tolerance and performance, and to provide adaptive resource management strategies. Probabilistic model checking is an automatic procedure for establishing if a desired property holds in a probabilistic model, aimed at

verifying probabilistic specifications such as "leader election is eventually resolved with probability 1", "the chance of shutdown occurring is at most 0.01%"', and "the probability that a message will be delivered within 30 ms is at least 0.75". A probabilistic model checker calculates the probability of a given temporal logic property being satisfied, as opposed to validity. In contrast to conventional model checkers, which rely on reachability analysis of the underlying transition system graph, probabilistic model checking additionally involves numerical solutions of linear equations and linear programming problems. This tutorial will introduce the theory and the practical details of automatic verification of probabilistic systems against temporal logic specifications. It will cover discrete- and continuous-time Markov chains, Markov decision processes and probabilistic timed automata, as well as the temporal logics PCTL, CSL and PTCTL. The usefulness of the techniques will be demonstrated through a number of case studies analysing real-world probabilistic protocols performed with PRISM, a probabilistic model checker developed at the University of Birmingham. Examples will include the dynamic configuration protocol for IPv4 link-local addresses, self-stabilising algorithms, power management, and CSMA/CD protocol.

Break Building De Waaier, Foyer 15:00 – 15:30

T6.2. : Random Neural Networks and Applications
Building De Zilverling, room U3, 15:30 – 16:30

Gerardo Rubino, INRIA Rennes, France

Abstract
Random Neural Networks (RNN) are Neural Networks (NN) of a new kind, proposed by Gelenbe in the early 1990s. A specific property of RNN is that they are also Markovian open queuing networks with two different types of customers: positive and negative ones (negative ones actually behave as signals in the model; they are not observable, only their effects are). The previous sentence means that the same mathematical object can be seen as a particular NN or as an open network of queues. When we look at RNN as networks of queues, they are called G-networks. A critical feature of these models is that they belong to the product form family, that is, their steady-state distribution is a product of factors one for each node in the system. This is the key of their nice mathematical properties which allow for very efficient analysis.
As other NN, RNN have been used in many fields that can be roughly classified into two very different types: learning and optimization. In the first type of applications, the RNN is trained to perform some task. In the second type, the RNN is tuned to allow finding a pseudo-optimum of some complex function.

This tutorial will present the mathematical model and with some detail, one example of each of the two types of applications. Both examples come from the networking area and have been developed by the author and co-workers. Both examples address performance issues.

QEST 2004

General Chair

Boudewijn Haverkort
University of Twente
The Netherlands

Programme Chairs

Giuliana Franceschinis
Università del Piemonte Orientale
Italy

Joost-Pieter Katoen
University of Twente
The Netherlands

Murray Woodside
Carleton University
Canada

Tools Demo Chair

Andrew Miner
Iowa State University
USA

Tutorial Chair

Marco Bernardo
Università degli Studi di Urbino "Carlo Bo"
Italy

Local Organizers

Boudewijn Haverkort
Joost-Pieter Katoen
Joke Lammerink
Helen Kool
Jenny Applequist
Lucia Cloth
Pieter-Tjerk de Boer
Ric Klaren
Henrik Bohnenkamp
José Martinez

QEST 2004

QEST Ad Hoc Steering Committee

Gianfranco Balbo (Italy)
Andrea Bobbio (Italy)
Steve Bruell (USA)
Peter Buchholz (Germany)
Gianfranco Ciardo (USA)
Susanna Donatelli (Italy)
Giuliana Franceschinis (Italy)
Roberto Gorrieri (Italy)
Günter Haring (Austria)
Peter Harrison (UK)
Boudewijn Haverkort (The Netherlands)
Jane Hillston (UK)
Ravi K. Iyer (USA)
Joost-Pieter Katoen (The Netherlands)
Peter Kemper (Germany)
Marta Kwiatkowska (UK)
Raymond Marie (France)
Prakash Panangaden (Canada)
Ramon Puigjaner (Spain)
William H. Sanders (USA)
Roberto Segala (Italy)
William Stewart (USA)
Kishor Trivedi (USA)
C. Murray Woodside (Canada)

QEST 2004 Programme Committee

Christel Baier (Germany)
Gianfranco Balbo (Italy)
Bernard Berthomieu (France)
Jonathan Billington (Australia)
Peter Buchholz (Germany)
Mariacarla Calzarossa (Italy)
Javier Campos (Spain)
Lucy Cherkasova (USA)
Gianfranco Ciardo (USA)
Luca de Alfaro (USA)
Josée Desharnais (Canada)
Tony Field (UK)
Giuliana Franceschinis (Italy)
Reinhard German (Germany)
Robert L. Givan (USA)
Roberto Gorrieri (Italy)
Serge Haddad (France)
Günther Haring (Austria)
Peter Harrison (UK)
Jane Hillston (UK)
Joost-Pieter Katoen (The Netherlands)
Peter Kemper (Germany)
Pieter Kritzinger (South Africa)
Marta Kwiatkowska (UK)
Kim G. Larsen (Denmark)
Raymond Marie (France)
Annabelle McIver (Australia)
Daniel Menasce (USA)
Aad van Moorsel (UK)
Ramon Puigjaner (Spain)
Gerardo Rubino (France)
William H. Sanders (USA)
Philippe Schnoebelen (France)
Roberto Segala (Italy)
Evgenia Smirni (USA)
Miklos Telek (Hungary)
Kishor Trivedi (USA)
Moshe Vardi (USA)
Lisa Wells (Denmark)
Katinka Wolter (Germany)
C. Murray Woodside (Canada)

Sponsors

Centre for Telematics and Information Technology, University of Twente

Institute for Programming Research and Algorithmics

The Royal Netherlands Academy of Arts and Sciences

University of Twente, Department of Electrical Engineering, Mathematics and Computer Science

The Netherlands Organisation for Scientific Researc