Edmund Clarke Orna Grumberg Somesh Jha Yuan Lu Helmut Veith

There is currently an increasing demand for concurrent programs. Checking the correctness of concurrent programs is a complex task due to the interleavings of processes. Sometimes, violation of the correctness properties in such systems causes human or resource losses; therefore, it is crucial to check the correctness of such systems. Two main approaches to software analysis are testing and formal verification. Testing can help discover many bugs at a low cost. However, it cannot prove the correctness of a program. Formal verification, on the other hand, is the approach for proving program correctness. Model checking is a formal verification technique that is suitable for concurrent programs. It aims to automatically establish the correctness (expressed in terms of temporal properties) of a program through an exhaustive search of the behavior of the system. Model checking was initially introduced for the purpose of verifying finite?state concurrent programs, and extending it to infinite?state systems is an active research area. In this thesis, we focus on the formal verification of parameterized systems. That is, systems in which the number of executing processes is not bounded a priori. We provide fully-automatic and parameterized model checking techniques for establishing the correctness of safety properties for certain classes of concurrent programs. We provide an open?source prototype for every technique and present our experimental results on several benchmarks. First, we address the problem of automatically checking safety properties for bounded as well as parameterized phaser programs. Phaser programs are concurrent programs that make use of the complex synchronization construct of Habanero Java phasers. For the bounded case, we establish the decidability of checking the violation of program assertions and the undecidability of checking deadlock?freedom. For the parameterized case, we study different formulations of the verification problem and propose an exact procedure that is guaranteed to terminate for some reachability problems even in the presence of unbounded phases and arbitrarily many spawned processes. Second, we propose an approach for automatic verification of parameterized concurrent programs in which shared variables are manipulated by atomic transitions to count and synchronize the spawned processes. For this purpose, we introduce counting predicates that related counters that refer to the number of processes satisfying some given properties to the variables that are directly manipulated by the concurrent processes. We then combine existing works on the counter, predicate, and constrained monotonic abstraction and build a nested counterexample?based refinement scheme to establish correctness. Third, we introduce Lazy Constrained Monotonic Abstraction for more efficient exploration of well?structured abstractions of infinite?state non?monotonic systems. We propose several heuristics and assess the efficiency of the proposed technique by extensive experiments using our open?source prototype. Lastly, we propose a sound but (in general) incomplete procedure for automatic verification of safety properties for a class of fault?tolerant distributed protocols described in the Heard?Of (HO for short) model. The HO model is a popular model for describing distributed protocols. We propose a verification procedure that is guaranteed to terminate even for unbounded number of the processes that execute the distributed protocol.

Informatics - 10 Years Back, 10 Years Ahead presents a unique collection of expository papers on major current issues in the field of computer science and information technology. The 26 contributions written by leading researchers on personal invitation assess the state of the art of the field by looking back over the past decade, presenting important results, identifying relevant open problems, and developing visions for the decade to come. This book marks two remarkable and festive moments: the 10th anniversary of the International Research and Conference Center for Computer Science in Dagstuhl, Germany and the 2000th volume published in the Lecture Notes in Computer Science series.

This book constitutes the refereed proceedings of the 10th International Conference on Hybrid Systems: Computation and Control, HSCC 2007, held in Pisa, Italy in April 2007. The 44 revised full papers and 39 revised short papers presented together with the abstracts of 3 keynote talks were carefully reviewed and selected from 167 submissions. Among the topics addressed are models of heterogeneous systems, computability and complexity issues, real-time computing and control, embedded and resource-aware control, control and estimation over wireless networks, tools for analysis, verification, control, and design, programming languages support and implementation, applications, including automotive, communication networks, avionics, energy systems, transportation networks, biology and other sciences, manufacturing, and robotics.

The First CADE in the Third Millennium This volume contains the papers presented at the Eighteenth International C- ference on Automated Deduction (CADE-18) held on July 27–30th, 2002, at the University of Copenhagen as part of the Federated Logic Conference (FLoC 2002). Despite a large number of deduction-related conferences springing into existence at the end of the last millennium, the CADE conferences continue to be the major forum for the presentation of new research in all aspects of automated deduction. CADE-18 was sponsored by the Association for Auto- ted Reasoning, CADE Inc., the Department of Computer Science at Chalmers University, the Gesellschaft fur ? Informatik, Safelogic AB, and the University of Koblenz-Landau. There were 70 submissions, including 60 regular papers and 10 system - scriptions. Each submission was reviewed by at least ?ve program committee members and an electronic program committee meeting was held via the Int- net. The committee decided to accept 27 regular papers and 9 system descr- tions. One paper switched its category after refereeing, thus the total number of system descriptions in this volume is 10. In addition to the refereed papers, this volume contains an extended abstract of the CADE invited talk by Ian Horrocks, the joint CADE/CAV invited talk by Sharad Malik, and the joint CADE-TABLEAUX invited talk by Matthias Baaz. One more invited lecture was given by Daniel Jackson.

This volume contains the proceedings of the Fourth Biennial Conference on F- mal Methods in Computer-Aided Design (FMCAD). The conference is devoted to the use of mathematical methods for the analysis of digital hardware c- cuits and systems. The workreported in this bookdescribes the use of formal mathematics and associated tools to design and verify digital hardware systems. Functional veri?cation has become one of the principal costs in a modern computer design e?ort. FMCAD provides a venue for academic and industrial researchers and practitioners to share their ideas and experiences of using - screte mathematical modeling and veri?cation. Over the past 20 years, this area has grown from just a few academic researchers to a vibrant worldwide com- nity of people from both academia and industry. This volume includes 23 papers selected from the 47 submitted papers, each of which was reviewed by at least three program committee members. The history of FMCAD dates backto 1984, when the earliest meetings on this topic occurred as part of IFIP WG10.2.

Foundations of Information Technology in the Era of Network and Mobile Computing is presented in two distinct but interrelated tracks: -Algorithms, Complexity and Models of Computation; -Logic, Semantics, Specification and Verification. This volume contains 45 original and significant contributions addressing these foundational questions, as well as 4 papers by outstanding invited speakers. These papers were presented at the 2nd IFIP International Conference on Theoretical Computer Science (TCS 2002), which was held in conjunction with the 17th World Computer Congress, sponsored by the International Federation for Information Processing (IFIP), and which convened in Montréal, Québec, Canada in August 2002.

This book constitutes the thoroughly refereed post-proceedings of the Third International Workshop on Formal Approaches to Testing of Software, FATES 2003, held in Montreal, Quebec, Canada, on October 6th, 2003. The 18 revised full papers presented were carefully selected from 43 submissions during two rounds of reviewing and improvement. The papers are organized in topical sections on program testing and analysis, test theory and test derivation algorithms, and test methods and test tools.

Assertion-based design is a powerful new paradigm that is facilitating quality improvement in electronic design. Assertions are statements used to describe properties of the design (I.e., design intent), that can be included to actively check correctness throughout the design cycle and even the lifecycle of the product. With the appearance of two new languages, PSL and SVA, assertions have already started to improve verification quality and productivity. This is the first book that presents an “under-the-hood” view of generating assertion checkers, and as such provides a unique and consistent perspective on employing assertions in major areas, such as: specification, verification, debugging, on-line monitoring and design quality improvement.

The refereed proceedings of the 15th International Conference on Computer Aided Verification, CAV 2003, held in Boulder, CO, USA in July 2003. The 32 revised full papers and 9 tool papers presented were carefully reviewed and selected from a total of 102 submissions. The papers are organized in topical sections on bounded model checking; symbolic model checking; games, trees, and counters; tools; abstraction; dense time; infinite state systems; applications; theorem proving; automata-based verification; invariants; and explicit model checking.

Artificial Intelligence continues to be one of the most exciting and fast-developing fields of computer science. This book presents the 177 long papers and 123 short papers accepted for ECAI 2016, the latest edition of the biennial European Conference on Artificial Intelligence, Europe’s premier venue for presenting scientific results in AI. The conference was held in The Hague, the Netherlands, from August 29 to September 2, 2016. ECAI 2016 also incorporated the conference on Prestigious Applications of Intelligent Systems (PAIS) 2016, and the Starting AI Researcher Symposium (STAIRS). The papers from PAIS are included in this volume; the papers from STAIRS are published in a separate volume in the Frontiers in Artificial Intelligence and Applications (FAIA) series. Organized by the European Association for Artificial Intelligence (EurAI) and the Benelux Association for Artificial Intelligence (BNVKI), the ECAI conference provides an opportunity for researchers to present and hear about the very best research in contemporary AI. This proceedings will be of interest to all those seeking an overview of the very latest innovations and developments in this field.

Annotation This book documents the scientific outcome and constitutes the final report of the Japanese research project on discovery science. During three years more than 60 scientists participated in the project and developed a wealth of new methods for knowledge discovery and data mining. The 52 revised full papers presented were carefully reviewed and span the whole range of knowledge discovery from logical foundations and inductive reasoning to statistical inference and computational learning. A broad variety of advanced applications are presented including knowledge discovery and data mining in very large databases, knowledge discovery in network environments, text mining, information extraction, rule mining, Web mining, image processing, and pattern recognition.

This book constitutes the refereed proceedings of the 13th International Conference on Computer Aided Verification, CAV 2001, held in Paris, France in July 2001. The 33 revised full papers presented were carefully reviewed and selected from 106 regular paper submissions; also included are 13 reviewed tool presentations selected from 27 submissions. The book offers topical sections on model checking and theorem proving, automata techniques, verification core technology, BDD and decision trees, abstraction and refinement, combinations, infinite state systems, temporal logics and verification, microprocessor verification and cache coherence, SAT and applications, and timed automata.

This book presents 19 revised invited keynote lectures and revised tutorial lectures given at the 4th International Symposium on Formal Methods for Components and Objects, FMCO 2005, Amsterdam, November 2005. The book provides a unique combination of ideas on software engineering and formal methods that reflect the current interest in the application or development of formal methods for large scale software systems such as component-based systems and object systems.

The new edition of an introduction to multiagent systems that captures the state of the art in both theory and practice, suitable as textbook or reference. Multiagent systems are made up of multiple interacting intelligent agents—computational entities to some degree autonomous and able to cooperate, compete, communicate, act flexibly, and exercise control over their behavior within the frame of their objectives. They are the enabling technology for a wide range of advanced applications relying on distributed and parallel processing of data, information, and knowledge relevant in domains ranging from industrial manufacturing to e-commerce to health care. This book offers a state-of-the-art introduction to multiagent systems, covering the field in both breadth and depth, and treating both theory and practice. It is suitable for classroom use or independent study. This second edition has been completely revised, capturing the tremendous developments in multiagent systems since the first edition appeared in 1999. Sixteen of the book's seventeen chapters were written for this edition; all chapters are by leaders in the field, with each author contributing to the broad base of knowledge and experience on which the book rests. The book covers basic concepts of computational agency from the perspective of both individual agents and agent organizations; communication among agents; coordination among agents; distributed cognition; development and engineering of multiagent systems; and background knowledge in logics and game theory. Each chapter includes references, many illustrations and examples, and exercises of varying degrees of difficulty. The chapters and the overall book are designed to be self-contained and understandable without additional material. Supplemental resources are available on the book's Web site. Contributors Rafael Bordini, Felix Brandt, Amit Chopra, Vincent Conitzer, Virginia Dignum, Jürgen Dix, Ed Durfee, Edith Elkind, Ulle Endriss, Alessandro Farinelli, Shaheen Fatima, Michael Fisher, Nicholas R. Jennings, Kevin Leyton-Brown, Evangelos Markakis, Lin Padgham, Julian Padget, Iyad Rahwan, Talal Rahwan, Alex Rogers, Jordi Sabater-Mir, Yoav Shoham, Munindar P. Singh, Kagan Tumer, Karl Tuyls, Wiebe van der Hoek, Laurent Vercouter, Meritxell Vinyals, Michael Winikoff, Michael Wooldridge, Shlomo Zilberstein

The authors have here put together the first reference on all aspects of testing and validating service-oriented architectures. With contributions by leading academic and industrial research groups it offers detailed guidelines for the actual validation process. Readers will find a comprehensive survey of state-of-the-art approaches as well as techniques and tools to improve the quality of service-oriented applications. It also includes references and scenarios for future research and development.

This book constitutes the thoroughly refereed post-proceedings of the Third International Conference on Formal Modeling and Analysis of Timed Systems, FORMATS 2005, held in Uppsala, Sweden in September 2005 in conjunction with ARTIST2 summer school on Component Modelling, Testing and Verification, and Static analysis of embedded systems. The 19 revised full papers presented together with the abstracts of 3 invited talks were carefully selected from 43 submissions. The papers cover work on semantics and modeling of timed systems, formalisms for modeling and verification including timed automata, hybrid automata, and timed petri nets, games for verification and synthesis, model-checking, case studies and issues related to implementation, security and performance analysis.

This open access two-volume set constitutes the proceedings of the 26th International Conference on Tools and Algorithms for the Construction and Analysis of Systems, TACAS 2020, which took place in Dublin, Ireland, in April 2020, and was held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2020. The total of 60 regular papers presented in these volumes was carefully reviewed and selected from 155 submissions. The papers are organized in topical sections as follows: Part I: Program verification; SAT and SMT;

Timed and Dynamical Systems; Verifying Concurrent Systems; Probabilistic Systems; Model Checking and Reachability; and Timed and Probabilistic Systems. Part II: Bisimulation; Verification and Efficiency; Logic and Proof; Tools and Case Studies; Games and Automata; and SV-COMP 2020.

Perhaps nothing characterizes the inherent heterogeneity in embedded sys tems than the ability to choose between hardware and software implementations of a given system function. Indeed, most embedded systems at their core repre sent a careful division and design of hardware and software parts of the system To do this task effectively, models and methods are necessary functionality. to capture application behavior, needs and system implementation constraints. Formal modeling can be valuable in addressing these tasks. As with most engineering domains, co-design practice defines the state of the it seeks to add new capabilities in system conceptualization, mod art, though eling, optimization and implementation. These advances -particularly those related to synthesis and verification tasks -direct1y depend upon formal under standing of system behavior and performance measures. Current practice in system modeling relies upon exploiting high-level programming frameworks, such as SystemC, Esterel, to capture design at increasingly higher levels of ab straction and attempts to reduce the system implementation task. While raising the abstraction levels for design and verification tasks, to be really useful, these approaches must also provide for reuse, adaptation of the existing intellectual property (IP) blocks.

This book constitutes the refereed proceedings of the 12th IFIP WG 10.5 Advanced Research Working Conference on Correct Hardware Design and Verification Methods, CHARME 2003, held in L'Aquila, Italy in October 2003. The 24 revised full papers and 8 short papers presented were carefully reviewed and selected from 65 submissions. The papers are organized in topical sections on software verification, automata based methods, processor verification, specification methods, theorem proving, bounded model checking, and model checking and applications.

The SPIN workshop series brings together researchers and practitioners int- ested in explicit state model checking technology as it is applied to the veri?- tion of software systems. Since 1995, when the SPIN workshop series was instigated, SPIN workshops have been held on an annual basis at Montr ? eal (1995), New Brunswick (1996), Enschede (1997), Paris (1998), Trento (1999), Toulouse (1999), Stanford (2000), andToronto(2001). Whilethe?rstSPINworkshopwasastand-aloneevent,later workshopshavebeenorganizedasmoreorlesscloselya?liatedeventswithlarger conferences, in particular with CAV (1996), TACAS (1997), FORTE/PSTV (1998), FLOC (1999), World Congress on Formal Methods (1999), FMOODS (2000), and ICSE (2001). This year, SPIN 2002 was held as a satellite event of ETAPS 2002, the European Joint Conferences on Theory and Practice of Software. The co-location of SPIN workshops with conferences has proven to be very successful and has helped to disseminate SPIN model checking technology to wider audiences. Since 1999, the proceedings of the SPIN workshops have appeared in Springer-Verlag's "Lecture Notes in Computer Science" series. The history of successful SPIN workshops is evidence for the maturing of model checking technology, not only in the hardware domain, but increasingly also in the software area. While in earlier years algorithms and tool development 1 around the SPIN model checker were the focus of this workshop series, the scopehasrecentlywidenedtoincludemoregeneralapproachestosoftwaremodel checking. Current research in this area concentrates not so much on completely verifyingsystemmodels,butratheronanalyzingsourcecodeinordertodiscover software faults.

A graduate-level textbook that presents a unified mathematical framework for modeling and analyzing cyber-physical systems, with a strong focus on verification. Verification aims to establish whether a system meets a set of requirements. For such cyber-physical systems as driverless cars, autonomous spacecraft, and air-traffic management systems, verification is key to building safe systems with high levels of assurance. This graduate-level textbook presents a unified mathematical framework for modeling and analyzing cyber-physical systems, with a strong focus on verification. It distills the ideas and algorithms that have emerged from more than three decades of research and have led to the creation of industrial-scale modeling and verification techniques for cyber-physical systems.

"The recent years have brought a number of advances in the development of infinite state verification, using techniques such as symbolic or parameterized representations, symmetry reductions, abstractions, constraint-based approaches, combinations of model checking and theorem proving. The active state of research on this topic provides a good time-point to increase impact by bringing together leading scientists and practitioners from these individual approaches. This volume gives an overview of the current research directions and provides information for researchers interested in the development of mathematical techniques for the analysis of infinite state systems."

This is the first book presenting a broad overview of parallelism in constraint-based reasoning formalisms. In recent years, an increasing number of contributions have been made on scaling constraint reasoning thanks to parallel architectures. The goal in this book is to overview these achievements in a concise way, assuming the reader is familiar with the classical, sequential background. It presents work demonstrating the use of multiple resources from single machine multi-core and GPU-based computations to very large scale distributed execution platforms up to 80,000 processing units. The contributions in the book cover the most important and recent contributions in parallel propositional satisfiability (SAT), maximum satisfiability (MaxSAT), quantified Boolean formulas (QBF), satisfiability modulo theory (SMT), theorem proving (TP), answer set programming (ASP), mixed integer linear programming (MILP), constraint programming (CP), stochastic local search (SLS), optimal path finding with A*, model checking for linear-time temporal logic (MC/LTL), binary decision diagrams (BDD), and model-based diagnosis (MBD). The book is suitable for researchers, graduate students, advanced undergraduates, and practitioners who wish to learn about the state of the art in parallel constraint reasoning.

This book constitutes the refereed proceedings of the 8th International Conference on Formal Engineering Methods, ICFEM 2006, held in Macao, China, in November 2006. The 38 revised full papers presented together with three keynote talks were carefully reviewed and selected from 108 submissions. The papers address all current issues in formal methods and their applications in software engineering.

This book constitutes the thoroughly refereed post-proceedings of the 6th International Conference on Theory and Applications of Satisfiability Testing, SAT 2003, held in Santa Margherita Ligure, Italy, in May 2003. The 33 revised full papers presented together with 5 articles reporting results of the related SAT competition and QBF evaluation were carefully selected during two rounds of reviewing and improvement from 67 submissions. The whole spectrum of research in propositional and quantified Boolean formula satisfiability testing is covered including proof systems, search techniques, probabilistic analysis of algorithms and their properties, problem encodings, industrial applications, specific tools, case studies, and empirical results.

Today's embedded devices and sensor networks are becoming more and more sophisticated, requiring more efficient and highly flexible compilers. Engineers are discovering that many of the compilers in use today are ill-suited to meet the demands of more advanced computer architectures. Updated to include the latest techniques, The Compiler Design Handbook, Second Edition offers a unique opportunity for designers and researchers to update their knowledge, refine their skills, and prepare for emerging innovations. The completely revised handbook includes 14 new chapters addressing topics such as worst case execution time estimation, garbage collection, and energy aware compilation. The editors take special care to consider the growing proliferation of embedded devices, as well as the need for efficient techniques to debug faulty code. New contributors provide additional insight to chapters on register allocation, software pipelining, instruction scheduling, and type systems. Written by top researchers and

designers from around the world, The Compiler Design Handbook, Second Edition gives designers the opportunity to incorporate and develop innovative techniques for optimization and code generation.
This book constitutes the refereed proceedings of the 16th International Conference on Foundations of Software Technology and Theoretical Computer Science, FST&TCS '96, held in Hyderabad, India, in December 1996. The volume presents 28 revised full papers selected from a total of 98 submissions; also included are four invited contributions. The papers are organized in topical sections on computational geometry, process algebras, program semantics, algorithms, rewriting and equational-temporal logics, complexity theory, and type theory.
The 31st International Colloquium on Automata, Languages, and Programming (ICALP 2004) was held from July 12 to July 16 in Turku, Finland. This volume contains all contributed papers presented at ICALP 2004, together with the invitedlecturesbyPhilippeFlajolet(INRIA), RobertHarper(CarnegieMellon), Monika Henzinger (Google), Martin Hofmann (Munich), Alexander Razborov (Princeton and Moscow), Wojciech Rytter (Warsaw and NJIT), and Mihalis Yannakakis (Stanford). ICALP is a series of annual conferences of the European Association for Theoretical Computer Science (EATCS). The?rst ICALP took place in 1972 and the ICALP program currently consists of track A (focusing on algorithms, automata, complexity, and cryptography) and track B (focusing on databases, logics, semantics, and principles of programming). Inresponsetothecallforpapersthereceivedtheprogramcommitteeof379papers, 272 for track A and 107 for track B. This is the highest number of submitted papersinthehistoryofICALPconferences. Theprogramcommitteeselected97 papersforinclusionintothescienti?cprogram. Theprogramcommitteefortrack A met on March 27 and 28 in Barcelona and selected 69 papers from track A. TheprogramcommitteefortrackBselected28papersfromtrackBinthecourse of an electronic discussion lasting for two weeks in the second half of March. The selections were based on originality, quality, and relevance to theor- ical computer science. We wish to thank all authors who submitted extended abstracts for consideration, the program committee for its hard work, and all referees who assisted the program committee in the evaluation process.