10292 Abstracts Collection and Summary
Resilience Assessment and Evaluation
— Dagstuhl Seminar —

Alberto Avritzer[1], Katinka Wolter[2], Aad van Moorsel[3] and Marco Vieira[4]

[1] Siemens - Princeton, US
alberto.avritzer@siemens.com
[2] FU Berlin, DE
katinka.wolter@fu-berlin.de
[3] University of Newcastle, GB
aad.vanmoorsel@ncl.ac.uk
[4] University of Coimbra, PT
mvieira@dei.uc.pt

Abstract. From July 18 to July 23, 2010 the Dagstuhl Seminar 10292 "Resilience Assessment and Evaluation " was held in Schloss Dagstuhl – Leibniz Center for Informatics. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Resilience, dependability, distributed systems, complex systems, critical infrastructures, Cloud computing

10292 Summary – Resilience Assessment and Evaluation

Resilience of computing systems includes their dependability as well as their fault-tolerance and security. Resilience defines the ability of a computing system to perform proper service in the presence of all kinds of disturbances and to recover from any service degradation. These properties are immensely important in a world where many aspects of daily life depend on the correct, reliable and secure operation of computing systems. Considered systems include, but are not limited to, infrastructures, computer networks, including adhoc and mesh networks, web-based systems, or service-oriented architectures, embedded systems, manufacturing systems, control systems and more.

Much work has been done already on modelling, measuring and evaluating performance and dependability of systems but a lot remains to be done. The combination of performance and dependability, performability, has been widely studied. But aspect such as benchmarking performability, or the combination of security and performance are still ongoing work. Similarly, systems to be studied

Dagstuhl Seminar Proceedings 10292
Resilience Assessment and Evaluation
http://drops.dagstuhl.de/opus/volltexte/2010/2898

change over time, as do their characteristics. This creates the need to continue and expand work in this area.

The workshop has addressed methods and tools to describe, measure, evaluate, benchmark, guarantee and improve resilience as well as case-studies and experimental work concerning resilience of computing systems. We have aimed at collecting existing work into a monograph with profound investigation of the subject as well as discuss open problems and research directions for the future.

Joint work of: Avritzer, Alberto; van Moorsel, Aad; Wolter, Katinka

Critical Infrastructures

Felivita di Giandomenico (CNR - Pisa, IT)

Critical Infrastructures (CI) are complex and highly interdependent systems, networks and assets that provide essential services in our daily life. They span a number of key sectors, including energy, finance, authorities, telecommunications, information technology, supply services and many others.

Dependability Benchmark applied to Documents

Ana Maria Ambrosio (National Institute for Space Research - São José, BR)

This work is motivated by the evidences of low quality software requirements specifications and by the very high costs of the rework caused by poorly written requirements. It is proposed a Benchmark to help evaluating and comparing requirement specifications of on-board data handling software for satellite computers. As the target of the benchmark proposed is a document, there is no workload to run or faultload to inject. Instead, a checklist composed by questions replaces the workload and is used to obtain measures that portray specific characteristics of the Software Requirements Specification. The checklist is composed of questions which influences the representativeness of the benchmark. The questions are based on three different sources: the Packet Utilization Standard (PUS), the conformance and fault injection (COFI) testing methodology and results of field studies on errors found in requirement reviews. This research has been carried on in cooperation among the National Institute for Space Research (INPE), the Instituto Tecnológico de Aeronáutica (ITA) and the University of Coimbra.

Keywords: Dpendability benchmarking, satellite computer, software requirement

Software Aging and Rejuvenation for Increased Resilience: Modelling, Analsys and Applications

Alberto Avritzer (Siemens - Princeton, US)

We present an historical overview of the research on software aging and rejuvenation techniques. The chapter presents application to resilience improvement of software aging and rejuvenation on the dependability and security domains.

Keywords: Software aging, rejuvenation, software reliability, software

Joint work of: Avritzer, Alberto; Ricardo M. Czekster; Kishor Trivedi

Non-functional Modeling and Analysis of Context-Aware Mobile Software Systems

Vittorio Cortellessa (Univ. degli Studi di L'Aquila, IT)

Context-awareness is becoming a first class attribute of software systems. In fact, applications for mobile devices need to be aware of their context in order to adapt their structure and behavior and offer the best quality of service even in case the (software and hardware) resources are limited. Existing approaches for non-functional modeling and analysis fail to capture the characteristics related to the context, thus resulting not suited for this domain.

In this talk we introduce a framework for modeling and analyzing the non-functional properties of context-aware mobile software systems. The framework allows to model: the software architecture, the context management, the adaptable behaviors and the non-functional parameters. Such models can then be transformed intonon-functional models for analysis purposes. We tailor an integrated environment for modeling these elements in UML, and we show how to use it for non-functional analysis purposes (e.g. performance, reliability).

Keywords: Performance, reliability, context-aware software, UML

Joint work of: Luca Berardinelli, Vittorio Cortellessa, Antinisca Di Marco

A Framework for the Analysis of Interdependencies in Critical Infrastructures: The Case of Electric Power Systems

Felicita Di Giandomenico (CNR - Pisa, IT)

Electric Power Systems (EPS) are prominent representatives of critical infrastructures.

Existing EPS are composed by two cooperating infrastructures: the Electric Infrastructure (EI) for the electricity generation and transportation to the enduser customers, and its Information Technology based Control System (ITCS) in charge of monitoring and controlling the EI physical parameters and of triggering appropriate reconfigurations in emergency situations. The interactions between these two systems need to be carefully analyzed to understand and characterize their (inter)dependencies, that is how the state of each infrastructure influences or is correlated to the state of the others.

This presentation focuses on a model-based framework for quantitative analysis of the propagation and impact of malfunctions in the infrastructures composing EPS, mainly developed in the context of the EU project CRUTIAL. The framework has been implemented using the Stochastic Activity Networks formalism and has been extensively applied to concrete case studies, to illustrate a representative set of different typologies of analyses for understanding and assessing the impact of interdependencies in EPS. The obtained results can be fruitfully exploited to guide the set-up of proper countermeasures to lessen the EPS vulnerabilities revealed through the analysis.

Recent extensions of the framework to consider more complex grid and information control organizations are also briefly discussed.

Keywords: Critical Infrastructures, Electric Power Systems, Dependability analysis, State-based stochastic assessment

Power and Performance: Green ICT R+D problems and opportunities

Carlos Juiz (University of the Balearic Islands, ES)

Current and future Power Management technologies will affect all system components. For example, slowing components down, such as Dynamic Voltage and Frequency Scaling of processor and memory, slowing down the rotational speed of magnetic disks, speed control of data communication and network switches, resource deactivation such as turning off individual cores, disabling some cache lines and turning off specific memory banks, turning off disks (used in mobility for years), etc., may produce some unexpected interactions. The interaction of Power Management could be not optimal either due to the overhead produced to control on/off systems or due to the reduction of speedup ir order to save energy consumption. Interaction of power management with system resilience is likely to be negative: power management will create temperature variations stressing the chips and the continuous change in the disk rotational speed or the frequent on/off at disks will produce reliability problems to the mechanical components of I/O. Even the power installation of datacenters will suffer from this enormous potential variations in energy consumption. Other interactions could be produced by software. Virtualization helps to reduce servers and storage, so that power seems to be reduced but components may be exhausted. Virtualization helps to increase reliability, but the consequence should be performance

overhead and consolidation overload. Software bugs and power management are also interacting with the system in ways that were not tested before. All these and other problems may produce resilience and performance problems that are also opportunities of research and development at Green ICT.

Keywords: Virtualization, dependability, quality-of-service, service-oriented systems

Towards Self-Aware Dependability Management in Virtualized Service Infrastructures

Samuel Kounev (KIT - Karlsruhe Institute of Technology, DE)

Virtualization, the major enabling technology for Cloud Computing, comes at the cost of increased system complexity and dynamicity making it hard to provide dependability guarantees. The increased dynamicity is caused by the introduction of virtual resources and the lack of direct control over the underlying physical hardware. The increased complexity is caused by the interactions between the applications and workloads sharing the physical infrastructure. The inability to predict such interactions and adapt the system accordingly makes it hard to provide dependability guarantees in terms of availability and responsiveness as well as resilience to attacks and operational failures. Moreover, the consolidation of workloads translates into higher utilization of physical resources which makes the system much more vulnerable to threats resulting from unforeseen load fluctuations, hardware failures or network attacks. Thus, virtualization introduces new sources of failure and threats degrading the dependability of virtualized service infrastructures. In this talk, we discuss the above challenges and present an approach and a long-term research agenda aiming to provide the basis for building next generation self-aware virtualized platforms and services. The latter will be self-aware in the sense that they will be aware of changes that occur in their environment and will be able to predict the effect of such changes on their dependability. Moreover, they will automatically adapt to ensure that system resources are utilized efficiently and dependability requirements are continuously satisfied. This vision will be realized using online models generated dynamically from the evolving system configuration and exploiting these models for autonomic performance and resource management.

Keywords: Virtualization, dependability, quality-of-service, service-oriented systems

Advanced Data Processing Methods for Dependability Benchmarking and Log Analysis

Andras Pataricza (Budapest Univ. of Technology & Economics, HU)

A collection of data and signal processing problems related to dependability monitoring has been presented and discussed.

Keywords: Dependability, monitoring, control theory, data and signal processing

Stochastic Fault-Models for Fault-Injection

Philipp Reinecke (FU Berlin, DE)

Traditional fault-injection is based on injecting a specific fault (e.g. stuck-at-X) and studying the reaction of the system to that fault. This technique is very successful at analysing the effectiveness of fault-detectors and fault-tolerance mechanisms. Fault-injection using stochastic fault-models, on the other hand, reproduces the occurrence of a fault by stochastic models.

While this technique does not allow a detailed study of the reaction of the system to the occurrence of a fault, it enables the study of QoS impairments that may be observed in typical usage scenarios with faults. Furthermore, stochastic fault-models can be used in test-beds, simulation, and analytical approaches, thus allowing comprehensive studies. This talk introduces the basics of fault-injection using stochastic fault-models, provides a classification of fault-models, and gives a short survey of fault-models present in the literature.

Keywords: Fault-injection, stochastic fault-models

Joint work of: Reinecke, Philipp; Wolter, Katinka

Software Reliability Assessment What was achieved? What has changed? What is still challenging?

Francesca Saglietti (Universität Erlangen-Nürnberg, DE)

The talk focused on state-of-the-art techniques and on novel research results related to the problem of assessing software reliability.

Compared with the past (low complexity, low real-time relevance) where classical solutions like proofs of correctness may have been considered as adequate, it is felt that today's verification and validation activities additionally require high testing coverage measures and a systematic evaluation of testing experience.

Both objectives are still seriously challenging the software engineering community by practicability constraints: the former (achieving high coverage by code

or model-based testing) is severely limited w.r.t. test case generation, the latter (statistical analysis of test samples) is often jeopardized by prohibitive effort.

In view of these serious restrictions, the talk suggested two emerging approaches aimed at increasing practicality be heuristics resp. by analysis of large amounts of operational data.

In the first case, evolutionary techniques are applied to generate sets of test cases achieving both maximum coverage and minimum size. Their implementation resulted in tools supporting both unit testing (coverage of classical control and data flow at code level) and integration testing (coverage of pre-defined model-based component interactions).

In the second case, the enormous effort usually required by statistical testing can be considerably reduced by taking into account experience gained during operation. The analysis and extraction of relevant operational data for the purpose of reliability evaluation was illustrated by means of an automotive case study centered on a software-based gearbox controller.

Keywords: Software reliability, test case generation, operational experience, statistical testing

Performance modelling of security protocols

Nigel Thomas (Newcastle University, GB)

Security adds an overhead to the actions performed. Different protocols, algorithms or parameters add different overheads. Hence to optimise the performance we can alter the choice of protocol, algorithm or parameters.

This might also lead to a change in security properties; hence a security vs performance tradeoff. In this talk I describe some work we have done in exploring this trade-off in different scenarios.

Keywords: Performance modelling, security protocols, stochastic process algebra

From Performance to Resilience Benchmarking:Past, Present & Future

Marco Vieira (University of Coimbra, PT)

The work on benchmarking has started long ago. Ranging from simple benchmarks that target a very specific system or component to very complex benchmarks for complex infrastructures, benchmarks have contributed to improve successive generations of systems. This presentation provides an overview of the state-of-the-art on benchmarking and defines a set of research needs and challenges that have to be addressed for the establishment of real resilience benchmarks.

Keywords: Benchmarking, Performance, Dependabilty, Security, Resilience

Trust Economics Framework

Aad van Moorsel (Newcastle University, GB)

This presentation introduces the Trust Economics Framework for objective decision-making in information security. The methodology takes into account human factors as well as economic/business factors. In the first step, stakeholders are interviewed for their concerns and priorities and a technology study is carried out. The second step produces a mathematical model of the system under study, including human and economic factors and incentives. The final stage is a discussion among stakeholders to come to a decision on the security investment question.

Keywords: Security, trust, human factors, economics

Summary

Katinka Wolter (FU Berlin, DE)

The seminar has provided ample opportunity for discussing challenges in a long known, but still changing and evolving field with various new open problems. The discussions have led to a book outline which Springer Verlag has agreed to publish. The book will cover known techniques for the assessment and evaluation of resilience, demonstrate those in case studies and explore into new and evolving applications of computing that demand for resilience as well as new problems in the analysis as well as solution approaches.