

Second Workshop on Proactive Failure Avoidance, Recovery, and Maintenance (PFARM)

Mirosław Malek and Felix Salfner
Humboldt-Universität zu Berlin, Germany
{salfner\malek}@informatik.hu-berlin.de

Kishor S. Trivedi
Duke University, Durham, USA
kst@ee.duke.edu

Abstract

Proactive approaches to failure avoidance, recovery and maintenance have recently attracted increased interest among researchers and practitioners from various areas of dependable system design and operation. This first workshop provided a stimulating, and fruitful forum to foster collaboration among researchers working on proactive fault management, to discuss ideas, exchange experiences and to find new answers to the overall challenge of significantly improving system dependability in contemporary computing and communication systems.

1. Themes, Goals and Objectives

Over the last decade, research on dependable computing has undergone a shift from reactive towards proactive-reactive and purely proactive methods: In classical fault tolerance a system reacts to errors or component failures in order to prevent them from turning into system failures, and maintenance follows fixed, time-based plans. However, due to an ever increasing system complexity, use of commercial-off-the-shelf components, virtualization, ongoing system patches and updates and dynamicity the effectiveness of such approaches has diminished and some of them have become difficult to apply. Therefore, a new area in dependability research has emerged focusing on proactive approaches that start acting before a problem arises in order to increase time-to-failure and/or reduce time-to-repair. These techniques frequently build on the anticipation of upcoming problems based on runtime monitoring. Industry and academia use several terms for such techniques, each focusing on different aspects, including self-* computing, autonomic computing, proactive fault management, trustworthy computing, failure prediction or forecasting, software rejuvenation, or preventive/proactive maintenance. It is the goal of this workshop to increase collaboration among re-

searchers from various communities all over the world working on the topic of PFARM. We want to provide a stimulating, and fruitful forum to foster collaboration among researchers working on similar topics, to discuss ideas, exchange experiences and to find new answers to the overall challenge of improving system dependability in contemporary computing and communication systems by an order of magnitude or more. More specifically, these techniques usually build on

- runtime monitoring to permanently measure the system's state
- algorithms to anticipate upcoming problems
- predictive diagnosis focused on the root cause of the upcoming problem
- decision and scheduling techniques in order to decide which method should be applied and when
- efficient and effective techniques to avoid upcoming failures, minimize their negative effects or improve system dependability by maintenance

An increased interest in proactive failure avoidance, recovery and maintenance techniques can be observed from the rising number of papers published in major general dependability conferences such as DSN, SRDS, PRDC and EDCC. It was the goal of this workshop to bring together researchers from all over the world to form a community within the community of dependable computing. In order to achieve this, the workshop provided a stimulating, and fruitful forum to foster collaboration among researchers working on similar topics, to discuss ideas, exchange experiences and to find new answers to the overall challenge of significantly improving system dependability in contemporary computing and communication systems. In fact, there is some evidence that with comprehensive use of proactive fault management methods the dependability in the server clusters environments (e.g., cloud computing) can be increased by an order of magnitude or more.

2. Main Challenges

Several topics were addressed from theory and practice to inspire research, design techniques and implementation from both industry and academia. Topics included, but were not limited to:

- **Runtime dependability assessment** and evaluation (reliability, availability, etc.)
- **Runtime monitoring** for online fault detection and diagnosis, including monitoring data processing
- **Prediction methods** to anticipate failures, resource exhaustion or other critical situations in complex systems, distributed systems, adaptive or peer-to-peer networks.
- **Predictive diagnosis and fault location** as well as root-cause analysis.
- **Online recovery, updates and upgrades**, non-intrusive hardware installation and software deployment
- **Proactive maintenance strategies** (short-term as well as long-term)
- **Optimal decision algorithms** and policies to manage and schedule the application of actions
- **Downtime minimization or avoidance mechanisms** such as preventive failover, state-clean up, proactive reconfiguration, failure-prevention driven load balancing, prediction-driven restarts, rejuvenation, adaptive checkpointing, or other prediction-driven enhancements of traditional repair methods
- **Proactive fault management and maintenance techniques** such as monitoring-based replacement, configuration and management of computer systems and components
- **Dependability evaluation** including models to assess the impact on metrics such as availability, reliability, security, performability, survivability and user-oriented metrics such as service availability, downtime, quality-of-service and quality-of-experience.
- Case-studies, applications, experiments, experience reports

3. The PFARM Karaoke

At the first edition of the PFARM workshop in 2009 a karaoke session on the biggest challenges in PFARM was held. We list the questions and challenges that were proposed and do hope that some answers to these questions and challenges can be found in the 2010 PFARM Workshop.

- Can uncertainty in measured data lead to bad behavior?
- Investigating the relation between faults and misconfigurations
- Measuring system reliability against potential faults
- Improvement and exploitation of existing monitoring infrastructures for proactive failure recovery / maintenance
- How do I proactively compose my systems or services in a highly dynamic environment (such as the Web)?
- Requirements of monitoring infrastructures for proactive actions (metrics ...)
- Relationships between prediction and monitoring. Should we focus on a "design for prediction" methodology?
- How to factor human factors / whole ecosystem analysis
- Diagnosis of problems that have never occurred before
- Blackbox systems
- Dependability assessment of off-the-shelf / open-source component-based complex systems
- Dependability assessment of dynamic distributed systems
- How can we model the network accurately?
- Deal with diversity of equipment, protocols, etc.
- Increased diversity & complexity of corporate environments
- Rapid growth brought about by mergers, acquisitions, integration of new areas
- In proactive recovery, how can we choose the "best" frequency of recoveries, in order to cause only minimal performance penalty, but also to recover from as many faults as possible?
- If we use modeling to choose the best frequencies of recoveries, how do we model malicious intruders?
- Automation scares operators! How to build trust of operators on automating recovery / configuration and reactive tools?
- Benchmark for PFARM problems
- Monitoring: How, what, when?
- Proactive reliable systems design; how can proactive policies be integrated in the development cycle?
- Methods to estimate the costs of introducing proactive policies versus reactive policies. How can we assess such costs and compose them with reactive policies' costs?

- Testing of proactive reliable systems: How can we test the introduced mechanisms for proactive reliability assurance?
- Diversity injection
- How do I know that the (service) components I am using from other people are reliable?
- How to process / select relevant data from all the data available?
- How to predict and diagnose problems never seen before?
- How to deal with multiple administrative domains (i.e., not complete visibility) in diagnosis?
- How to diagnose transient failures?
- Field data sources: How can we achieve an effective source of failure data?

4. Program Committee

We would like to thank Program Committee Members for their support and their insightful reviews.

- Andrea Bondavalli, University of Florence, Italy
- Tadashi Dohi, Hiroshima University, Japan
- Michael Grottke, University of Erlangen-Nuremberg, Germany
- Matti Hiltunen, AT&T Labs Research, USA
- Hermann Kopetz, TU Vienna, Austria
- Michael R. Lyu, Chinese University of Hong Kong
- Aad van Moorsel, University of Newcastle, UK
- Takashi Nanya, University of Tokyo, Japan
- András Pataricza, Budapest University of Technology and Economics, Hungary
- Manfred Reitenspiess, Fujitsu Siemens Computers
- Lisa Spainhower, IBM
- Kalyan Vaidyanathan, Sun Microsystems
- Ricardo Vilalta, University of Houston, USA

5. Program

This year we have received 15 contributions of which eight, after a careful reviewing process, were selected. The papers cover the entire spectrum of the proactive approaches ranging from adaptive monitoring to proactive recovery such as rejuvenation.

Design and Theory

- Aspect Oriented Software Fault Tolerance and Analytically Redundant Design Framework
K. Hameed, R. Williams and J. Smith
- A Translation of State Machines to Temporal Fault Trees
N. Mahmud, Y. Papadopoulos and M. Walker

Monitoring and Alerting

- Fast Entropy Based Alert Detection in Super Computer Logs
A. Makanju, N. Zincir-Heywood and E. Milios
- Qualitative Performance Control in Supervised IT Infrastructures
G. J. Paljak, Z. Egel, D. Toth, I. Kocsis, T. Kovacs-hazy and A. Pataricza
- Adaptive Monitoring in Microkernel OSs
D. Cotroneo, D. Di Leo and R. Natella

Modeling for Proactive Fault Management

- Hybrid, Recursive, Nested Monitoring of Control Systems Using Petri Nets and Particle Filters
L. Zouaghi, A. Wagner and E. Badreddin
- Rejuvenation with Workload Migration
V. Mendiratta and R. Hanmer

Virtualization

- CacheMind: Fast Performance Recovery Using a Virtual Machine Monitor
K. Kourai
- Guest Talk: Error Detection and Resolution in the Freescale Virtualized Embedded Environment
V. Sethi

PFARM Panel (chair: *Lisa Spainhower*)