Summary on the

7th Workshop on Adaptive and Reconfigurable Embedded Systems

April 13, 2015, Seattle, USA

http://apres2015.di.fc.ul.pt

In conjunction with

CPSweek 2015

Scope

A system is said to be adaptive if it responds to environmental changes such as hardware/software defects, resource changes, and non-continual feature usage, in ways that extend the area of operation and improve the efficiency in the use of system resources. Adaptivity often incurs overhead in terms of system complexity and resource requirements, but unfortunately the resulting tradeoffs are usually ignored until a very late stage of the system development process. Retrofitting existing prototypes, architectures, middleware, operating systems, and protocols with concepts and means for flexibility such as run-time system reconfiguration or reflexive diagnostics and steering methods, typically leads to disproportionate overhead, unusual tradeoffs, and less satisfactory results. There is a strong need for adaptivity-centered research.

To exploit adaptivity, new specification methods are needed, to define acceptable adaptation ranges which will be explored by the system at run-time to improve a given performance metric. Current operating systems and network protocols are generally not designed to support such flexible requirements nor complementary reflexive mechanisms to help applications adjusting their operation to the current conditions. The same can be said about current fault tolerance mechanisms, which are usually not designed such that they can adapt to different dependability goals during system operation.

Programming such systems also needs adequate middleware layers that provide adequate interfaces for the development of adaptive applications. Building such middleware so that it preserves adaptive properties while providing performance guarantees together with satisfying other usual goals, such as modularity, reusability and scalability, is a challenge still to be conquered.

In general, flexibility and complexity are counterpoised to dependability, but these notions must be reconciled in order to design dependable adaptive systems. This challenge encompasses aspects such as investigation on how adaptivity can be used as a means to achieve improved performance and efficiency without sacrificing dependability (for instance through reconfiguration upon failures) and investigation on how to guarantee that the adaptive mechanisms themselves are dependable, e.g. reliable, available, safe, etc.

The objective of APRES is to bring together experts in the development and use of adaptive and reconfigurable embedded systems and researchers from the embedded systems community at large. Of particular interest are new concepts and ideas for modeling and analyzing tradeoffs of embedded and real-time systems, novel algorithms and mechanisms to realize adaptation and reconfigurability, and experience reports with practical case studies.

Topics

We solicited research contributions on any topics of interest to embedded, real-time and dependable systems research in the areas of systems, languages, software, theory, networking, control and analysis with specific focus on reconfigurability and adaptivity. Topics of particular interest include:

- Capturing and modeling of flexible application and reconfiguration requirements
- Tradeoff analysis and modeling
- Programming-language support for adaptivity
- Middleware support for adaptivity
- Operating system support for adaptivity
- Adaptive fault tolerance mechanisms
- Computation and communication models for adaptivity
- Policies and algorithms for single and multi-resource reconfiguration
- Verification and certification of reconfigurable systems
- Case studies and success stories
- Taxonomies and comparative studies

- Diagnostic and steering of embedded systems
- System architecture and design patterns for adaptivity
- Probabilistic reconfiguration techniques
- Scalability, reusability, and modularity of reconfiguration mechanisms
- Dependability and adaptivity across the architectural levels
- Quality of service management
- Application frameworks for reconfigurable embedded systems

Program Chair

- António Casimiro, University of Lisboa, Portugal

Program Committee

- Carlos Eduardo Pereira, UFRG
- Chang-Gun Lee, Seoul National University
- Frank Singhoff, University of Brest
- Gera Weiss, Ben Gurion University
- Giovani Gracioli, UFSC/LISHA
- Guillermo Rodriguez-Navas, Mälardalen University
- Insik Shin, KAIST
- Jinkyu Lee, Sungkyunkwan University
- Juan Antonio de la Puente, Universidad Politécnica de Madrid
- Kanghee Kim, Soongsil University
- Luca Mottola, Politecnico di Milano (Italy) and SICS Swedish ICT
- Luis Miguel Pinho, CISTER Research Centre/ISEP
- Mário Sousa, University of Porto
- Martina Maggio, Lund University
- Neil Audsley, University of York
- Pau Marti, Technical University of Catalonia
- Paulo Pedreiras, University of Aveiro
- Roman Obermaisser, University of Siegen
- Sebastian Zug, Otto-von-Guericke University Magdeburg
- Thomas Nolte, Mälardalen University

Steering Committee

- Luis Almeida, University of Porto, Portugal
- Karl-Erik Årzén, Lund University, Sweden
- Sebastian Fischmeister, Univ. of Waterloo, Canada
- Insup Lee, University of Pennsylvania, USA
- Julian Proenza, Univ. of the Balearic Islands, Spain

Table of Contents

Keynote Speech

Medical Application Platforms for On-Demand Medical Cyber-Physical Systems

Insup Lee, University of Pennsylvania

Abstract – Medical devices are undergoing significant transformations, embracing the potential of embedded software and network connectivity. Instead of stand-alone devices that can be designed, certified, and used independently of each other for patient treatment, networked medical devices will work as distributed systems that simultaneously monitor and control multiple aspects of the patient's physiology. The combination of embedded software controlling the devices, networking capabilities, and complicated physiological dynamics exhibited by patient bodies makes modern medical device systems a distinct class of cyber-physical systems (CPS). We refer to these as medical cyber-physical systems (MCPS). The goal of MCPS is to improve patient safety and treatment outcomes by leveraging diverse capabilities of individual devices to gain a more detailed and accurate picture of the evolving patient state.

A distinguishing feature of MCPS, compared to other CPS domains such as avionics, is their reliance on a multitude of medical devices, separately developed for specific intended use. This plethora of choices gives a lot of flexibility for the clinical personnel to select a clinical scenario and a set of devices, best suited for treatment of a particular patient. However, traditional medical devices are developed as monolithic stand-alone systems. Due to lack of interoperability between medical devices, clinicians have to carry out manually such coordinated uses in practice. Numerous cases have been reported, where patient safety has been compromised due to an error that could have been prevented or mitigated if these devices were networked and properly coordinated by software that ensure some sort of safety-lock.

There are several on-going efforts to support interoperability between medical devices using open standards. Based on standards, it is envisioned that MCPS would be assembled on-demand since they should be put together for a particular clinical scenario using available medical devices as needed, instead of having dedicated a set of medical devices developed and pre-configured for each clinical scenario. Such an on-demand assembly of MCPS requires a new paradigm to guarantee the safety of clinical applications using the MCPS. In contrast, currently practiced technologies for safety-critical systems assume that the system will be fully designed, manufactured, and tested prior to the customer. The main reason is that safety and effectiveness are emergent system-level properties and whether the properties are satisfied depends on the interactions among the system's components.

This talk will explain why the traditional approach would not scale for on-demand MCPS and then present a newly emerging approach based on medical application platforms. In particular, this talk will describe the needs, challenges, and architecture of medical application platforms.

About the Speaker: Insup Lee is Cecilia Fitler Moore Professor of Computer and Information Science and Director of PRECISE Center, which he co-founded in 2008 at the University of Pennsylvania. His research interests include cyber-physical systems (CPS), real-time systems, embedded systems, high-confidence medical device systems, formal methods and tools, run-time verification, software certification, and trust management. The theme of his research activities has been to assure and improve the correctness, safety, and timeliness of life-critical embedded systems. His papers received the best paper awards in IEEE RTSS 2003, CEAS 2011, IEEE RTSS 2012, and ACM/IEEE ICCPS 2014, and the best student paper in IEEE RTAS 2012. Recently, he has been working in medical cyber-physical systems and security of cyber physical systems.

He has served on many program committees, chaired many international conferences and workshops and served on various steering and advisory committees of technical societies. He has also served on the editorial boards on the several scientific journals, including Journal of ACM, IEEE Transactions on Computers, Formal Methods in System Design, and Real-Time Systems Journal. He is a founding co-Editor-in-Chief of KIISE Journal of Computing Science and Engineering (JCSE). He was a member of Technical Advisory Group (TAG) of President's Council of Advisors on Science and Technology (PCAST) Networking and Information Technology (2006-2007). He is a member of the National Research Council's committee on 21st Century Cyber-Physical Systems Education (Dec 2013-May 2015). He received an appreciation plaque from Ministry of Science, IT and Future Planning, South Korea, for speaking at the Universal Linkage for Top Research Advisor (ULTRA) Program Forum in 2013. He is IEEE fellow and received IEEE TC-RTS Outstanding Technical Achievement and Leadership Award in 2008.

Session 1 - Adaptive mechanisms for distributed embedded systems - Summary

Kiyofumi Tanaka
School of Information Science
Japan Advanced Institute of Science and Technology

This session consists of three presentations which aim to provide new adaptive mechanisms. The first one was "Adaptive environment perception in Cyber-Physical Systems" by Sebastian Zug from Otto-von-Guericke-Universität Magdeburg, Germany, which proposes Adaptive Sensing Controller (ASC) to give optimal application schedules with desirable sensor periods and phase shifts. The optimization is done in two-level analyses; 1st (worst-case static) analysis compares sensor parameters (periods, offsets, delays) and application requirements (# of measurements, quality) based on a worst/best case scenarios. 2nd (dynamic) analysis considers situation-specific properties (phase shift of sensor periods), communication delays and jitter to provide an online optimization. This presentation led to an interesting discussion on consideration of potential communication delay and synchronization problems.

The second presentation was "Adaptive Offloading for Infotainment Systems" by Luis Lino Ferreira from CISTER Research Centre/INESC TEC, Polytechnic Institute of Porto, which supports infotainment applications being dynamically offloaded onto vehicle's infotainment platform by providing QoS according to resource requirements. The offloading decision takes varying execution times into account, where execution times are predicted based on the past information. The offloading principles and scenarios for a user device (Tablet) are described in the paper.

The last presentation was "Towards Adaptive Resource Reservations for Component-Based Distributed Real-Time Systems" by Nima Khalilzad from MRTC/Malardalen University, Sweden, which provides a framework to provide light-weight dynamic resource reservation facilities for component-based distributed systems. According to changes of components' resource demands caused by switching operating modes of tasks, the framework adapts the budgets for resource utilization to the situations. One of discussions was about whether the management should be centralized or distributed.

The aims of the above papers are different, however, they all take dynamic factors in the systems into account and improve efficiency by adapting to them.

Session 2 - Scheduling and replication-based techniques for adaptivity and Hardware-based adaptation of embedded systems – Summary

Sebastian Zug
Otto-von-Guericke Universität Magdeburg

Adaptivity can be applied and implemented on different levels of a CPS. The second session of APRES addressed two aspects in this context: the component configuration in case of error states and the effective scheduling of varying tasks related to power consumption and response time.

One approach for tackling challenges of the last category was given in the presentation of the paper "Digitally Assisted Analog Front-end Power Management Strategy via Dynamic Reconfigurability for Robust Heart Rate Monitoring" by Chengzhi Zong. The paper motivates variable timing for heart beat measurements in order to save energy in wearable measurement units. The main idea was to categorize different intervals in a heartbeat period that have to monitored with different sample rates. As shown at the end of the presentation, the approach provides a significant decreased power consumption. The auditorium was interested in algorithms for signal classification and the benefits of using a powerful DSP for this purpose. A second complex of questions was about the potential level of energy savings related to different heart frequencies.

While the first mentioned paper described an application oriented optimization of the schedule, the second one "Virtual Release Advancing for Earlier Deadlines" was focused on an abstract task scheduling. Kiyofumi Tanaka presented a new approach based on an enhanced Earliest Deadline First (EDF) scheduling mechanism. The main idea is to define a virtual release time for sporadic tasks that should be executed while maintaining schedulability of hard periodic tasks. The virtual release time approach shifts the deployment and the corresponding deadline to the past. Hence, the approach improves the ordering of sporadic tasks to reach a short response. The presenter showed that, for tasks with varying execution times, the technique could outperform comparable approaches especially in case of a high system utilization. The following discussion was focused on specific assumptions of the approach – well known execution times, jitter, etc.

The last paper "Tolerating Partial Failures on IEC 61499 Applications" discussed a framework for building fault-tolerant IEC 61499 application based on distributed replication of function blocks. The authors motivated this approach in two directions, firstly they evaluated the need for flexible replication configurations in industrial application and the secondly, their intuitive development in this context. The presented framework provides the required flexibility to easily perform changes in a production line without emergency breaks. In order to meet this goal the authors proposed a replication mechanism of a specific function block and the needed synchronization and adaptation.

Session 3 - CodeStream: Combining Network Coding and Scheduling for adaptive communications – Summary

Luís Almeida
Instituto de Telecomunicações
Dep. Electrical and Computer Engineering
Faculty of Engineering, University of Porto

The session organized in the scope of the CodeStream project aimed at bringing in new forms of online adaptation and reconfiguration that are possible when using rateless coding techniques in wireless communications. The project itselfs aims at combining rateless coding with scheduling in streaming applications.

Three papers were presented. The first one was an invited submission by Koo and Govindarasu from the Iowa State University in the USA, which explained the basics of Network Coding and how it could be used to minimize transmissions and thus save energy in networked embedded systems. This paper generated an interesting discussion on the potential of Network Coding to allow taking opportunistic advantage of the broadcast nature of the medium in a multi-hop or star topology to send the same information in a more compact form using linear combinations of original packets.

The following two papers were presented by the session organizer and they addressed two different works being carried out within the project. The first of these papers, by Faneca etal, from the Universities of Aveiro and Porto in Portugal, addressed the dissemination of large files among many clients, potentially embedded clients. It used Fountain Codes with a feedback phased at the end of the file transfer and it showed that adapting the feedback phase online according to the actual number of clients in feedback mode can help in reducing the transfer time.

The last paper, by Moreira etal, from the Universities of Porto, Portugal, and Aalborg, Denmark, addressed the case of streaming for large numbers of embedded clients using Network Coding to avoid retransmissions. It showed that it is possible to achieve stochastic worst-case delay guarantees actuating on the specific code structure as well as on the number of packets transmitted per round.

At the end of the session a few comments were made on the potential of rateless coding and its combination with scheduling for efficient and reliable use of the wireless medium in one to many communications.