27th Ada-Europe International Conference on Reliable Software Technologies (AEiC 2023)

13-16 June 2023, Lisbon, Portugal

BOOKLET OF PRESENTATIONS

http://www.ada-europe.org/conference2023

In cooperation with

acm In-Cooperation

SIGAda

SIGBED

SIGPLAN

Ada Resource ARA Association

Ciências ULisboa

About this Booklet

This booklet contains short summaries of all the presentations included in the conference core program. The booklet groups presentations by session, prefixing their title with their type: ‘IP’ for industrial presentations, ‘RP’ for research presentations, ‘WiP’ for work-in-progress presentations (which are accompanied by posters exposed during the refreshment breaks).

The proceedings of the research papers will appear in a dedicated Special Issue of Elsevier’s Journal on Systems Architecture.
The proceedings of the industrial papers will appear in forthcoming issues of Ada-Europe’s Ada User Journal, along with papers drawn from the Work-in-Progress presentations.

We invite you to use this booklet as “navigational tool” throughout the conference program and its blank spaces provided in it, as a pad for your notes and commentaries. Enjoy the conference!

Core Conference Program

	Morning	Before Lunch	After Lunch	Afternoon
Wednesday, June 14th	Keynote Talk	Session 1: *Verification and Validation 1*	Session 2: *Advanced Systems*	Session 3: *Reliability and Performance*
Thursday, June 15th	Panel Discussion	Session 4: *Verification and Validation 2*	Session 5: *Reliable Programming*	Session 6: *Real-Time Systems*

Conference Sponsors

AdaCore

gmv INNOVATING SOLUTIONS

TABLE OF CONTENTS

SESSION 1: VERIFICATION AND VALIDATION 1 4

SESSION 2: ADVANCED SYSTEMS 6

SESSION 3: RELIABILITY AND PERFORMANCE 9

SESSION 4: VERIFICATION AND VALIDATION 2 11

SESSION 5: RELIABLE PROGRAMMING 13

SESSION 6: REAL-TIME SYSTEMS 15

Session 1: Verification and Validation 1

RP: Systematic Review on Contract-based Safety Assurance and Guidance for Future Research

Samina Kanwala, Vrije Universiteit Amsterdam, Netherlands (*speaker*)
Faiz Ul Muramb, Department of Computer Science and Media Technology, Linnaeus University, Sweden
Muhammad Atif Javedc, Amaris Consulting, Goteborg, Sweden

Abstract

The safety requirements are often described via specifications called contracts. To verify that the system fulfills certain safety requirements, for instance, in the assume-guarantee contract specification, the key safety indicators are organized, so that if certain assumptions hold then the respective behaviour/properties are guaranteed. Safety contracts provide a means of exposing potential incompatibilities early in the development process, selecting components to reuse, certifying systems as well as identifying uncertainty sources during the operational phase. There exist several studies on contract-based safety assurance, however, there is not any systematic study in this field. For this, a first Systematic Literature Review (SLR) is carried out to obtain an overview of the various contract-based assurance concepts, problems, proposed solutions, and their usefulness. In our study, the identification and selection of the primary studies were based on a well-planned search strategy. The search process identified a total of 2881 studies published between 1969 and 2021, out of which 66 studies are selected through a multi-stage process according to our predefined SLR protocol. This SLR aims to highlight the state-of-the-art of contract-based safety assurance and identify potential gaps for future research. Based on research topics in selected studies, we identified the following main categories: contract type, hazard and safety analysis techniques, compliance with standard requirements, development stage, domain, level of automation, type of study and evaluation, and tool support. The findings of the systematic review not just highlight that the contracts are even more important for advanced safety-critical systems, but also strategies to exploit their full potential should be considered in future studies. The suggestions revealed for future research include the usage of contracts for adapting new behaviour, defining system boundaries, interacting with other systems, managing risk during operation, dynamic/runtime safety assurance, and integration of safety with security.

RP: Compositional Verification of Embedded Real-Time Systems

Mohammed Aristide Foughali, IRIF, Université Paris Cité, France (*speaker*)
Pierre-Emmanuel Hladik, Nantes Université, École Centrale Nantes, Nantes, France
Alexander Zuepke, Technical University of Munich, Garching, Germany

Abstract

In an embedded real-time system (ERTS),real-time tasks (software) are typically executed on a multicore shared-memory platform (hardware).The number of cores is usually small, contrasted with a larger number of complex tasks that share data to collaborate. Since most ERTSs are safety-critical, it is crucial to rigorously verify their software against various real-time requirements under the actual hardware constraints(concurrent access to data, number of cores).Both the real-time systems and the formal methods communities provide elegant techniques to realize such verification, which nevertheless face major challenges. For instance, model checking (formal methods) suffers from the state-space explosion problem, whereas schedulability analysis (real-time systems)is pessimistic and restricted to simple task models and schedulability properties. In this paper, we propose a scalable and generic approach to formally verify ERTSs. The core contribution is enabling, through joining the forces of both communities, compositional verification to tame the state-space size. To that end, we formalize a realistic ERTS model where tasks are complex with an arbitrary number of jobs and critical sections, then show that compositional verification of such model is possible, using a hybrid approach (from both communities),under the state-of-the-art partitioned fixed-priority (P-FP)with limited preemption scheduling algorithm. The approach consists of the following steps, given the above ERTS model and scheduling algorithm. First, we compute fine-grained data sharing overheads for each critical section that

reads or writes some data from the shared memory. Second, we generalize an algorithm that, aware of the data sharing overheads, computes an affinity (task-core allocation)guaranteeing the schedulability of hard-real-time (HRT) tasks. Third, we devise a timed automata (TA) model of the ERTS, that takes into account the affinity, the data sharing overheads and the scheduling algorithm, on which we demonstrate that various properties can be verified compositionally, i.e., on a subset of cores instead of the whole ERTS, therefore reducing the state-space size. In particular, we enable the scalable computation of tight worst-case response times (WCRTs)and other tight bounds separating events on different cores, thus overcoming the pessimism of schedulability analysis techniques. We fully automate our approach and show its benefits on three real-world complex ERTSs, namely two autonomous robots and an automotive case study from the WATERS 2017 industrial challenge.

SESSION 2: ADVANCED SYSTEMS

WiP: VR-based Teleoperation of Autonomous Ground Vehicles for Operation Recovery

Georg Jäger (*speaker*), Gero Licht, Norman Seyffer, Stefan Reitmann, TU Bergakademie Freiberg, Germany

Abstract

While research will enable deployment of autonomous system in harsh and inaccessible environments, their operation may be interrupted due to unforeseen situations. A possibility to recover operation nonetheless is to employ teleoperation. However, what requirements and criteria need to be fulfilled by such a system when deployed in safety-critical operation scenarios? How can a timely and safely operation recovery be ensured? The present work aims to report our progress on developing a research platform for addressing these and similar questions.

WiP: Denoising Autoencoder-based Defensive Distillation as an Adversarial Robustness Algorithm

Bakary Badjie (*speaker*), José Cecílio, António Casimiro, LASIGE, FCUL, Portugal

Abstract

Adversarial attacks significantly threaten the robustness of deep neural networks (DNNs). Despite the multiple defensive methods employed, they are nevertheless vulnerable to poison attacks, where attackers meddle with the initial training data. In order to defend DNNs against such adversarial attacks, this work proposes a novel method that combines the defensive distillation mechanism with a denoising autoencoder (DAE). This technique tries to lower the sensitivity of the distilled model to poison attacks by spotting and reconstructing poisonous adversarial inputs in the training data. We added carefully created adversarial samples to the initial training data to assess the proposed method's performance. Our experimental findings demonstrate that our method successfully identified and reconstructed the poisonous inputs while also considering enhancing the DNN's resilience. The proposed approach provides a potent and robust defence mechanism for DNNs in various applications where data poisoning attacks are a concern. Thus, the defensive distillation technique's limitation posed by poisonous adversarial attacks is overcome.

WiP: Software-based security approach for networked embedded devices

José Ferreira (*speaker*), Alan Oliveira, André Souto, José Cecílio, LASIGE, FCUL, Portugal

Abstract

As the Internet of Things (IoT) continues to expand, data security has become increasingly important for ensuring privacy and safety, especially given the sensitive and, sometimes, critical nature of the data handled by IoT devices.
There exist hardware-based trusted execution environments used to protect data, but they are not compatible with low-cost devices that lack hardware-assisted security features.

The research in this paper presents software-based protection and encryption mechanisms explicitly designed for embedded devices. The proposed architecture is designed to work with low-cost, low-end devices without requiring the usual changes on the underlying hardware. It protects against hardware attacks and supports runtime updates, enabling devices to write data in protected memory. The proposed solution is an alternative data security approach for low-cost IoT devices without compromising performance or functionality. Our work underscores the importance of developing secure and cost-effective solutions for protecting data in the context of IoT.

WiP: Cooperative Autonomous Driving in Simulation

Gonçalo Costa (*speaker*), José Cecílio, António Casimiro, LASIGE, FCUL, Portugal

Abstract

Autonomous driving is an area that has been growing in recent years. However, cars are unprepared to cooperate with others nearby, wasting resources and computational power. Thus, cooperative autonomous driving reveals its importance in the future. In this work-in-progress paper, we define, implement and test an architecture for a simulation environment where cooperative autonomous driving protocols can be tested. Additionally, a Manoeuvre Negotiation Protocol is implemented. This protocol will make an existing autonomous driving (AD) stack more resilient in real driving scenarios, improving its robustness and safety.

WiP: Exploring Trade-offs in Explainable AI

Dene Brown (*speaker*), Glenn Hawe, Ulster University, UK

Abstract

Machine learning (ML) plays an increasing role in the technology that we interact with. However, many of the ML models that control these systems are classified as 'black box', meaning that their complexity is too great to understand how their decision making is achieved. This has led to the need for explainable AI (xAI) methods [4], such as LIME (Local Interpretable Model-agnostic Explanations) [2], which aim to present the reasoning behind a ML model's predictions in terms that is understandable to humans.

Evaluating the performance of an xAI method is challenging, as no single criterion effectively captures when one explanation is better than another, in all circumstances. A recent review by [1] identified twelve important criteria for explanations to satisfy. These have been dubbed "Co-12" criteria (because they all have names beginning with "co"). Examples include "compactness" (simpler explanations are better) and "completeness" (explanations should provide sufficient detail, and not omit anything important). Finding "perfect" explanations that optimize all twelve criteria simultaneously is generally not possible, as different criteria may naturally compete with one another, e.g. improving an explanation in terms of completeness may well come at the cost of compactness. Generating explanations is therefore inherently multi-objective.

The aim of this work is to identify trade-offs between a number of the Co-12 properties for xAI. In particular, we aim to modify an existing xAI approach [3] to generate a Pareto-optimal set of explanations that trade off different Co-12 criteria, rather than return a single explanation. This enables the user to select the explanation that best suits their needs in the application.

WiP: Cataloging Prompt Patterns to Enhance the Discipline of Prompt Engineering

Douglas C. Schmidt (*speaker*), Jesse Spencer-Smith, Quchen Fu, Jules White, Vanderbilt University, USA

Abstract

The rapid advent of Large language models (LLMs), such as ChatGPT, are disrupting a number of domains, ranging from education to medicine and software engineering. LLMs rely on "prompts", which are natural language statements given to the LLM to query and program its capabilities. This paper provides several contributions to research on LLMs. First, discusses the importance of codifying "prompt patterns" to enable a more disciplined and repeatable means of interacting with and evaluating LLMs than current ad hoc discussions based on individual examples. Second, it provides examples of prompt patterns that improve human interaction with LLMs in the context of software engineering, as well as other domains. We contend that prompt patterns play an essential role in providing the foundation for prompt engineering.

WiP: Achieving Crash Fault Tolerance In Autonomous Vehicle Autopilot Software Stacks Through Safety-Critical Module Rejuvenation

Federico Lucchetti, University of Luxembourg, Luxembourg

Abstract

Autonomous driving vehicles (ADV), have been in recent years, victims of their own success. Through their use of increasingly sophisticated sensor modalities and deep learning capabilities, ADVs have not only learned how to probe their chaotic environment with higher granularity coupled with smooth trajectory execution but also inherited all the vulnerabilities that were hiding behind these new features. Ensuring the safety of ADVs is crucial since a simple fault along their underlying autopilot software stack can lead to catastrophic accidents with the loss of human lives. Therefore we propose a crash-fault tolerant scheme that can be triggered whenever a crash fault of the safety critical submodules of the autopilot software stack is detected, which executes an emergency trajectory and effectively steers the car into a safe spot where the autopilot can be rejuvenated. We implement and evaluate the efficacy of this recovery scheme in the Apollo ADV software stack in conjunction with the SVL simulator.

IP: Safety-critical software in the EROSS+ on-orbit servicing project

Kristoffer Nyborg Gregertsen, SINTEF Digital, Norway (*speaker*)
Vincent Dubanchet, Thales Alenia Space, France
Carolina Pinto Dos Santos Serra, GMV SkySoft, Portugal
Juan Antonio Bejar Romero, GMV, Spain

Abstract

This industrial presentation will give an update on the EROSS+ on-orbit servicing project in the H2020 programme, with emphasis on the development and maturation process of safety-critical software.

SESSION 3: RELIABILITY AND PERFORMANCE

WiP: Exploiting container based microservices for reliable smart mobility applications

Paolo Ferrari, Emiliano Sisinni, Massimiliano Gaffurini (*speaker*), University of Brescia, Italy

Abstract

Modern smart city must intelligently manage resources to become sustainable without sacrificing the quality of life and the needs of its citizens. Smart mobility is emerging as a key concept because of its impact on research on advanced technology, infrastructures and solutions, IoT services and devices, and people. Safety and availability are mandatory, forcing the design of new reliable services for localization, health monitoring of the user, maintenance of vehicle, and protection of the environment. This paper proposes a container-based microservice approach to the edge computing in IoT smart mobility scenarios. Since smart mobility backend must manage a large heterogeneity of applications, the proposed approach outperforms classical solutions (based on "monolithic hardware+software" devices), from the point of view of flexibility, upgradability, security, scalability and reliability. A demo prototype based on industry-grade hardware and Docker has been realized and multiple implementations of the same services have been executed in parallel, showing strong independence between them. Moreover, average delays of less than 10 ms are obtained, confirming the usability in several smart mobility (and smart city) applications.

IP: Towards Reliable Distributed Edge-Cloud Applications

Michael Pressler (*speaker*), Dirk Ziegenbein, Arne Hamann, Robert Bosch GmbH, Germany

Abstract

We will present our first version of the Silverline framework for distributed CPS applications. We consider an Edge-Cloud Continuum with a heterogeneous portfolio of platforms, ranging from datacenter-style servers to microprocessor- and microcontroller-based solutions with diverse capabilities in computing, memory, I/O, etc. Real-time/safety-criticality applications from various CPS domains are executed distributed across an Edge-Cloud continuum. We make use of the technological trend in virtualization is towards lightweight, bytecode-based sandbox solutions that run on small microcontrollers as well as in data centers.

RP: Monintainer: An orchestration-independent extensible container-based monitoring solution for large clusters

Miguel Correia, Wellington Oliveira, José Cecílio (*speaker*), LASIGE, FCUL, Portugal

Abstract

Container virtualization has recently gained popularity due to its low performance and resource allocation overhead. The rise of this technology can be attributed to the advancement of cloud computing and the adoption of micro-services architecture. These new approaches offer a more efficient and fine-grained system design through the benefits of containerization, such as isolation, portability, and improved performance. However, container-based systems have created new challenges in monitoring due to their automated flexibility, ephemerality, and the increasing number of containers in a system. So there is a practical need for effective monitoring and performance management tools. This paper analyses the key performance metrics for machine, container and application services, including CPU usage, memory usage, disk usage, and network usage. Furthermore, we review several widespread tools for collecting and monitoring these metrics and present the Monintainer tool. It is a solution designed to monitor entire container-based systems, from applications to their underlying infrastructure, allowing users to better understand their systems' behaviour in run-time. The tool's results can aid container-based systems' design, implementation and optimization.

RP: Analyzing the performance of persistent storage for fault-tolerant stateful fog applications

Zeinab Bakhshi (*speaker*), Guillermo Rodriguez-Navas, Hans Hansson, Mälardalen University, Sweden

Abstract

In this paper, we analyze the scalability and performance of a persistent, fault-tolerant storage approach that provides data availability and consistency in a distributed container-based architecture with intended use in industrial control applications. We use simulation to evaluate the performance of this storage system in terms of scalability and failures. As the industrial applications considered have timing constraints, the simulation results show that for certain failure patterns, it is possible to determine whether the storage solution can meet critical deadlines. The presented approach is applicable for evaluating timing constraints also of other container-based critical applications that require persistent storage.

SESSION 4: VERIFICATION AND VALIDATION 2

WiP: Symbolic Assurance Refinement for CPS

Dionisio de Niz (*speaker*), Lutz Wrage, Carnegie Mellon University, USA

Abstract

In this paper we present an analysis contract approach that takes advantage of efficient domain-specific analysis algorithms, enable incremental analysis of architectural model refinements, and implement assume-guarantee reasoning in symbolic domains in SMT.

WiP: Towards a Methodology to Design Provably Secure Cyber-Physical Systems

Felipe Lisboa Malaquias (*speaker*), Georgios Giantamidis, Stylianos Basagiannis, Collins Aerospace, Ireland
Simone Fulvio Rolllini, Collins Aerospace, Italy
Isaac Amundson, Collins Aerospace, USA

Abstract

The inordinate financial cost of mitigating post-production cybersecurity vulnerabilities in cyber-physical systems (CPS) is forcing the industry to rethink systems design cycles: greater attention is being given to the design phase -- with the goal of reducing the attack surface of systems at an early stage (i.e., before silicon tape out). Fortunately, formal methods have advanced to the point that they can address such needs and contribute towards achieving security certification. However, new methods and tools focusing on industrial scalability and usability for systems engineers are required. In this ongoing research paper, we describe a framework that will help systems engineers to: a) design cyber-assured CPS using a Model Based Engineering (MBE) approach; b) formally map security requirements to different hardware and software blocks in the model; and c) formally verify security requirements. Based on the nature of each requirement, our framework collects formal correctness evidence from different tools: while high-level architectural properties are suitable for a contract- or ontology-based reasoning, more complex properties with rich semantics require the use of model checking or theorem proving techniques.

IP: Safety of the Intended Functionality Concept Integration into a Validation Tool Suite

Víctor J. Expósito Jiménez, Helmut Martin, Bernhard Winkler, Joaquim M. Castella Triginer (*speaker*), Virtual Vehicle Research GmbH, Austria
Heiko Scharke, Hannes Schneider, AVL List GmbH, Austria

Abstract

Nowadays, the introduction of more complex Advanced Driver Assistance Systems (ADAS) as well as Automated Driving (AD) force the industry to shift to a more scenario-based validation from the standard used technology-based approach. In line with this new focus, support for Safety of the Intended Functionality (SOTIF) has to be integrated into existing tools. Unlike functional safety or cybersecurity, which covers failures and malfunctions, and external attacks respectively, SOTIF focuses on the technical shortcomings and human misuses and any hazardous behaviour they could include in the system. Technical shortcomings could be related to environment sensing technologies. For example, lidar sensors could produce ghost reflections in a heavy snow scenario. Therefore, a heavy snow scenario could be identified as a potential triggering condition. Triggering conditions are the specific situations that could activate a performance insufficiency to finally result into a hazardous scenario, and it is one of the main topics that SOTIF introduces. The main goal is to reduce the number of unknown hazardous scenarios as much as possible by defining SOTIF measures.

The integration of the SOTIF concept is achieved by incorporating the identified potential triggering conditions. As a first step, an extensive list of them has been investigated based on current state-of-the-art and available standards. They are then parametrised by using an existing system ontology, which is used to model the scenarios. Once one or more potential triggering conditions are selected in the scenario, the scenario constraints (e.g., limited visibility, reduced friction...) associated to each potential triggering condition are included in the generated test cases. The resulting metrics of the matrix test cases show the impact of the selected potential triggering conditions on the function, which are compared with the nominal performance of the function (i.e., no potential triggering conditions included) to determine them not longer as potential but triggering conditions for the function, and to identify the thresholds at which they are relevant to impact and effect on the function output.

The validation tool suite in which the concept has been integrated is the AVL SCENIUS, which was developed with the scenario-based validation approach in mind. The complete validation process is covered, from scenario design to scenario management, test case generation, test allocation, and result reporting. The suite is based on three main tools modules. First, the scenario designer allows the user to handle all aspects of the scenario including the parametrisation. It fully supports ASAM OpenScenario and OpenDrive. All scenarios are immediately verified for standard conformity as well as by the enhanced data and logic checks. Then, the user could manage all stored scenarios in the scenario data manager. All the elements relevant for the sufficient description of scenarios such as road content, traffic content, and other environmental data are managed and stored in a central database. Finally, the test case generator provides the user the possibility of defining test orders in a simulation or transfer to another different execution environment. The implemented smart testing algorithms enable the automatic reduction of the vast amount of test cases and parameter variations.

In addition to the main benefits provided by the tool suite such as time-cost saving, efficiency, fast integration and traceability; the inclusion of the SOTIF concept extends and improves the identification and validation of both unknown and known hazardous scenarios of a ADAS/AD function to obtain a more precise safety argumentation.

IP: Application of a method for evaluation of software used in naval nuclear means

Eduardo Bezerra, Marinha do Brasil, Brazil

Abstract

This article presents the summary of the practical application of a software evaluation method, developed to be used in software used in naval means with nuclear propulsion, in an integration system of sensors for navigation.

Session 5: Reliable Programming

WiP: Automatic test value generation for Ada

Léo Creuse (*speaker*), Matthieu Eyraud, Viviane Garèse, AdaCore, France

Abstract

This article introduces novel tools to automatically generate pertinent Ada values in order to produce higher quality tests for Ada subprograms. This development is being led by AdaCore, conjointly with Thales Research and Technologies, as part of the RAPID initiative [1]. A first tool will generate valid Ada values based on a structural analysis of the types of the parameters of the subprogram under test following various customizable strategies. Those values will then be filtered in order to satisfy the specifications of the subprogram, and new coverage criteria for executable specifications will be used to assess the relevance of the generated testsuite. This first set of values will then be used as seeds both for a fuzzing process, and a symbolic execution campaign, from which values of interest will be then extracted. This integrated process will enable users to generate a high value starting test corpus, which can then be expanded upon by domain-specific tests.

WiP: Mechanization of the Ravenscar profile in Coq

Jerome Hugues, CMU/SEI, USA

Abstract

The Ravenscar profile has been added to the Ada language as part of the 2005 revision. It is a pragmatic definition of concurrency patterns for real-time systems for mono-core processor. In this paper, we report on an ongoing effort to mechanize a toy language derived from IMP that embeds the Ravenscar profile constructs as they are specified in the Ada Reference Manual. We define the denotational semantics of the language and derive two interpreters for the language.

WiP: A Real-Time Parallel Programming Approach for Rust

Hugo Silva, Tiago Carvalho (*speaker*), Luis Miguel Pinho, Polytechnic Institute of Porto – ISEP, Portugal

Abstract

The development of real-time systems is one of the areas with the highest relevance in computer science, and the number of critical systems has increased significantly. These systems considers several applications running concurrently, and inside each of those applications code might be parallelized to improve their performance and control the priority of each parallelizable task. Several efforts have been done in different programming languages to provide real-time systems with parallel programming models, whether by code extensions or annotations, or with specific features in the actual language core.

Rust is a recent programming language that have quickly grown in potential and already with a large community, being continuously formed. The language is a good candidate in terms of both real-time systems and parallel programming. However, there is a lack of work that joins these two important concepts in an efficient and reliable way.

In this work we aim to design and provide a framework for real-time parallel systems. We conduct a study over the existing work in other programming languages and aim to bring their advantages and useful programming models into the Rust programming language, in the format of a real-time parallel programming library.

IP: Security Hardening Ada Programs through Innovative Fuzz Testing

Paul Butcher, AdaCore, UK

Abstract

This article describes a talk on the research and development undertaken by AdaCore into fuzz testing Ada programs. It introduces the basic concepts of fuzz testing before describing the barriers to adoption, regulatory objectives, advanced coverage findings techniques and the critical components needed to make up an effective fuzz testing solution for applications developed in Ada. This article was submitted as an outline for a talk on the industrial-track for the Ada-Europe International Conference on Reliable Software Technologies 2023.

IP: An Update On the Tasking Profiles in Ada 2022

Patrick Rogers, AdaCore, USA

Abstract

Ada 2005 formally defined the Ravenscar profile, a tasking subset supporting the maximal application and run-time implementation simplicity required by the sort of stringent analyses expected for certification by, for example, DO-178.

In 2016, AdaCore developed and deployed a new tasking profile based directly on the Ravenscar profile. The new profile relaxes certain restrictions in Ravenscar in order to increase the expressive power of the language subset without sacrificing predictability or (relative) efficiency.

We presented the AdaCore profile in 2017 at Ada Europe [1], including the justification, additional constructs allowed, execution-time costs, and resulting schedulability analysis.

The new profile is included in the Ada 2022 draft revision, with some additional enhancements and a new name. In addition, Ada 2022 makes some updates to the Ravenscar profile.

In this presentation we first provide a brief overview of the Ravenscar profile, including the justification. We then use that foundation to explore the new profile, including the enhancements in the Ada 2022 draft revision and the rationale for the limits of the extensions. We conclude with an identification of the users expected to continue using the Ravenscar profile, and those who might benefit from using the new profile.

IP: Ada on a New Embedded Target

Dylan Eskew, Western Washington University, USA

Abstract

The Ada programming language has many attractive qualities to it, one of which is its usefulness on bare-board embedded platforms via the Ravenscar profile tasking library. In order for these Ada features to make it to these embedded platforms, an adaptation of the Ada runtime library has to made for the specific device. While runtime libraries already exist for many popular platforms such as the STM32 or Raspberry Pi microprocessors, there are times when a developer might need or want to develop a library for a new platform. In an effort to explore embedded systems deeper through the implementation of a new Ada bare-board runtime, a new goal emerged as the process of development posed unexpected challenges.

Session 6: Real-Time Systems

WiP: Worst Case Execution Time Estimation of Multicore and GPU Software: A Pedestrian Detection Use Case

Van Rodiguez-Ferrandez (*speaker*), UPC and BSC, Spain
Alvaro Jover-Alvarez, BSC, Spain
Matina Maria Trompouki, Leonidas Kosmidis, BSC and UPC, Spain
Francisco J Cazorla, BSC, Spain

Abstract

Worst Case Execution Time estimation of software running on parallel platforms is a challenging task, due to resource interference of other tasks and the complexity of the underlying CPU and GPU hardware architectures. In this work in progress paper, we employ Measurement Based Probabilistic Timing Analysis (MBPTA), which is capable of managing complex architectures such as multicores. We enable its use by software randomisation, which we show for the first time that is also possible on GPUs. We demonstrate our method on a pedestrian detection use case on an embedded multicore and GPU platform for the automotive domain, the NVIDIA Xavier.

WiP: A POSIX/RTEMS monitoring tool and a benchmark to detect real-time scheduling anomalies

Blandine Djika Mezatio (*speaker*), Georges Kouamou, University of Yaoundé 1, Cameroon
Frank Singhoff, Alain Plantec, Lab-STICC/Université de Bretagne Occidentale, France

Abstract

This article deals with scheduling anomalies in real-time systems.
We present MONANO, a POSIX user-level library allowing applications to dynamically detect a pre-identified set of real-time scheduling anomalies. This library is based on the modelling of runtime constraints. We present also a benchmark to evaluate our approach. The benchmark is composed of several programs implementing the most frequent real-time scheduling anomalies.

RP: Time-Predictable Task-to-Thread Mapping in Multi-Core Processors

Mohammad Samadi (*speaker*), ISEP, IPP, Portugal
Sara Royuela, BSC, Spain
Luis Miguel Pinho, Tiago Carvalho, ISEP, IPP, Portugal
Eduardo Quiñones, BSC, Spain

Abstract

The performance of time-predictable systems can be improved in multi-core processors using parallel programming models (e.g., OpenMP). However, schedulability analysis of parallel applications is a big challenge due to their sophisticated structure. The common drawbacks of current task-to-thread mapping approaches in OpenMP are that they (i) utilize a global queue in the mapping process, which may increase contention, (ii) do not apply heuristic techniques, which may reduce the predictability and performance of the system, and (iii) use basic analytical techniques, which may cause notable pessimism in the temporal conditions. Accordingly, this paper proposes a task-to-thread mapping method in multi-core processors based on the OpenMP framework. The mapping process is carried out through two phases: allocation and dispatching. Each thread has an allocation queue in order to minimize contention, and the allocation and dispatching processes are performed using several heuristic algorithms to enhance predictability. In the allocation phase, each task-part from the OpenMP DAG is allocated to one of the allocation queues, which includes both sibling and child task-parts. A suitable thread (i.e., allocation queue) is selected using one of the suggested heuristic allocation algorithms. In the dispatching phase, when a thread is idle, a task-part is selected from its allocation queue using one of the suggested heuristic dispatching algorithms and then dispatched to and executed by the thread. The performance of the proposed method is evaluated under different conditions (e.g., varying the number of tasks and the number of threads) in terms of application response time and overhead of the mapping process. The simulation results show that the proposed method surpasses the other methods, especially in the scenario that includes overhead of the mapping. In addition, a prototype implementation of the main heuristics is evaluated using two kernels from real-world applications, showing that the methods work better than LLVM's default scheduler in most of the configurations.

RP: Fine-grained adaptive parallelism for automotive systems through AMALTHEA and OpenMP

Adrian Munera, Sara Royuela (*speaker*), BSC, Spain
Michael Pressler, Harald Mackamul, Dirk Ziegenbein, Robert Bosch GmbH, Germany
Eduardo Quiñones, BSC, Spain

Abstract

The software development complexity of automotive systems has significantly increased during the last decade due to the latest Advanced Driving Assistance System (ADAS) functionalities. To effectively address this complexity, domain specific modeling languages (DSMLs) like AUTOSAR or an opensource system performance model for AUTOSAR-aligned systems, APP4MC, have become a common trend in the automotive industry. DSMLs allow for easily capturing the functional and non-functional requirements of the system without needing to master low level details of the programming model or the processor architecture. Unfortunately, current DSMLs do not support the parallel programming models, like OpenMP and CUDA, that are used to exploit parallel heterogeneous architectures featuring acceleration devices such as GPUs and FPGAs required. These architectures are however essential to cope with the performance needs of ADAS. This exposes a gap between the DSMLs used by automotive designers to enhance software productivity and leverage verification and validation processes, and the parallel processor architectures used in this domain. This paper presents a complete framework to safely exploit the inherent parallelism exposed by the AMALTHEA system description, supported in APP4MC, by: (1) automatically transforming the high-level design into the OpenMP parallel programming model targeting both host and accelerator parallelism, and (2) using compiler analysis techniques to prove the correctness of the model transformed to OpenMP code. The paper contributes also with (3) an analysis of the parallel execution model allowed by the AMALTHEA DSML and that of OpenMP, and (4) a performance plus productivity evaluation of the proposed framework on real automotive systems executed on an embedded GPU-based processor architecture.

ORGANIZERS

Conference & Program Chair
António Casimiro
University of Lisbon, Portugal

Journal-track Chair
Elena Troubitsyna
KTH Royal Institute of Technology, Sweden

Industrial-track Co-Chairs
Alexandre Skrzyniarz
Thales, France

Sara Royuela
Barcelona Supercomputing Center, Spain

Work-in-Progress-track Co-Chairs
Björn Andersson
Software Engineering Institute - Carnegie Mellon University, USA

José Cecílio
University of Lisbon, Portugal

Tutorial & Education Chair
Luis Miguel Pinho
ISEP, Portugal

Workshop Chair
Frank Singhoff
University of Brest, France

Sponsorship Chair
Ahlan Marriott
White Elephant GmbH, Switzerland

Publicity Chair
Dirk Craeynest
Ada-Belgium & KU Leuven, Belgium

Local Chair
Madalena Almeida
Abreu Travel Agency, Portugal

Web Master
Hai Nam Tran
University of Brest, France

JOURNAL-TRACK COMMITTEE

Mario Aldea Rivas (University of Cantabria, Spain), Matthias Becker (KTH Royal Institute of Technology, Sweden), Bernd Burgstaller (Yonsei University, South Korea), Daniela Cancila (CEA LIST, France), António Casimiro (University of Lisbon, Portugal), Xiaotian Dai (University of York, England), Juan A. de la Puente (Polytechnic University of Madrid, Spain), Barbara Gallina (Mälardalen University, Sweden), J. Javier Gutiérrez (University of Cantabria, Spain), Jérôme Hugues (Software Engineering Institute, Carnegie Mellon University, USA), Patricia López Martínez (University of Cantabria, Spain), Linas Laibinis (University of Vilnius, Lithuania), Alejandro Mosteo (CUD Zaragoza, Spain), Kristoffer Nyborg Gregertsen (SINTEF Digital, Norway), Laurent Pautet (Telecom ParisTech, France), Luís Miguel Pinho (CISTER/ISEP, Portugal), Sara Royuela (Barcelona Supercomputing Center, Spain), José Ruiz (AdaCore, France), Sergio Sáez (Valencia Polytechnic University, Spain), Frank Singhoff (University of Brest, France), Tucker Taft (AdaCore, USA), Santiago Urueña (GMV, Spain), Tullio Vardanega (University of Padua, Italy)

INDUSTRIAL-TRACK COMMITTEE

Alessandro Biondi (Scuola Superiore Sant'Anna, Italy), Raúl de la Cruz (Collins Aerospace, Ireland), Claire Dross (AdaCore, France), Emmanuel Ledinot (Thales Research and Technology, France), José Neves (GMV, Portugal), Maxime Perrotin (ESA/ESTEC, France), José Prado (Capgemini Engineering, Portugal), Michael Pressler (Bosch, Germany), Helder Silva (Edisoft, Portugal), Hugo Torres Vieira (Evidence Srl, Italy)

WORK-IN-PROGRESS-TRACK COMMITTEE

Patricia Balbastre Betoret (Valencia Polytechnic University, Spain), Kalinka Branco (University of São Paulo, Brazil), Leandro Buss Becker (University of Manchester, UK), Tiago Carvalho (ISEP, Portugal), Li-Pin Chang (National Yang Ming Chiao Tung University, Taiwan), João Carlos Cunha (Polytechnic of Coimbra, Portugal), Catherine Dezan (University of Brest, France), Alwyn Goodloe (NASA Langley, USA), Jérémie Guiochet (LAAS-CNRS, France), Marine Kadar (CEA-Leti, France), Robert Kaiser (RheinMain University of Applied Sciences, Germany), Federico Lucchetti (University of Luxembourg, Luxembourg), Ruben Martins (CMU, USA), Alan Oliveira (University of Lisbon, Portugal), Kristin Yvonne Rozier (Iowa State University, USA), Douglas Schmidt (Vanderbilt University, USA)