Sophia Antipolis, France
April 16-17, 2018

IJLTS

MoE Trustworthy International Joint Lab
Advisory Committee
and Workshop

An event organized by :

Inria — Université Nice Sophia Antipolis — i3S sophia antipolis — cnrs

Membre de UNIVERSITÉ CÔTE D'AZUR

TUESDAY APRIL 17TH ---- AGENDA

INRIA Kahn building, room 1-2

9:00: Allocution of David Simplot, Director of INRIA Sophia-Antipolis center.

9:20: Session 1 (Chair E. Madelaine)

- Pr. HE Jifeng

Algebra of Probabilistic Programs

- Robert de Simone, DR, PhD

Formal concurrent Models of Communication and Computation facing Parallel Programming Models

- Pr. Frederic Vivien

Checkpointing Workflows for Fail-Stop Errors

We consider the problem of orchestrating the execution of workflow applications structured as Directed Acyclic Graphs (DAGs) on parallel computing platforms that are subject to fail-stop failures. The objective is to minimize the expected overall execution time, or makespan. A solution to this problem consists of a schedule of the workflow tasks on the available processors and of a decision of which application data to checkpoint to stable storage, so as to mitigate the impact of processor failures. We revisit classical mapping heuristics such as HEFT and MinMin and complement them with several checkpointing strategies. The objective is to derive an efficient trade-off between checkpointing every task (CkptAll), which is an overkill when failures are rare events, and checkpointing no task (CkptNone), which induces dramatic re-execution overhead even when only a few failures strike during execution. Contrarily to previous work, our approach applies to arbitrary workflows, not just special classes of dependence graphs such as M-SPGs (Minimal Series-Parallel Graphs).

10:50: coffee break

11:10: Session 2 (Chair: Min Zhang)

- Eric Madelaine, CR, PhD

Using an SMT engine to generate Symbolic Automata: a case-study

Abstract:

We present a symbolic and hierarchical model called “pNet” expressing the behaviour of open concurrent systems. The pNet model is endowed with a symbolic operational semantics in terms of so-called “Open Automata”. This allows to check properties of such systems in a compositional manner. We implemented an algorithm computing this semantics, building predicates expressing the synchronization conditions between the events of the pNet sub-systems. Checking such predicates requires symbolic reasoning over first order logics, but also over application-specific data. We use the Z3 SMT engine to check satisfiability of the predicates, and prune the open automaton of its unsatisfiable transitions. As an industrial oriented use-case, we use so-called "architectures" for BIP systems, that have been used in the framework of an ESA project and to specify the control software of a nanosatellite at the EPFL Space Engineering Center. We use pNets to encode a BIP architecture extended with explicit data, and compute its open automaton semantics.

- Pr. DENG Yuxin

Logical Characterizations of Probabilistic Bisimilarity

Abstract:

Larsen and Skou initiated the study of probabilistic bisimilarity and its characterisation in terms of tests. Later on, van Breugel et al. showed that, for labelled Markov processes with continuous state spaces, probabilistic bisimilarity nicely coincides with a simple notion of testing equivalence. Their proof employs advanced machinery from topology. In the discrete case of finite-state reactive probabilistic processes, we prove that coincidence result with an elementary and more accessible proof.

12:10: lunch at INRIA

AFTERNOON

14:00: Collaborations with Université Côte d'Azur
By Pr. Jean-Christophe MARTIN, Vice-president for international Affairs, UCA

- Presentation of UCA
- Memorandum of Understanding between ECNU and UCA
- Group photo

14:30: Session 3 (Chair: Deng Yuxin)

- Pr. Marcello Bonsangue

Android, Java and Python libraries can be wrong ... but they can be fixed, if wanted!

Abstract:
Timsort is a hybrid sorting algorithm used as the default in Java, in Android, and Python.. Given the popularity of these platforms, this means that the number of computers, cloud services and mobile phones that use TimSort for sorting is well into the billions. In this talk I will introduce the algorithm, and show that it is formally wrong, in the sense that it is possible to find an input which eventually causes TimSort to crash. I will discuss how TimSort can be fixed and the impact of the result in the programming language community.

- Pr. Frederic Mallet

Model-Based System Engineering for Cyber-Physical Systems

Abstract:
Cyber-Physical Systems (CPSs) are networks of heterogeneous embedded systems immersed within a physical environment. Several ad-hoc frameworks and mathematical models have been studied to deal with challenging issues raised by CPSs. We explore a more standard-based approach that relies on UML/SysML/MARTE to capture different aspects of CPSs, including structure, system behaviors, clock constraints, and non-functional properties. The novelty of our work lies in the use of logical clocks and MARTE/CCSL to drive and coordinate different models. Meanwhile, to capture stochastic behaviors of CPSs, we propose an extension of CCSL, called pCCSL, where logical clocks are adorned with stochastic properties. The stochastic information is used to capture the uncertain behavior of the environment. Possible variants are explored using Statistical Model Checking (SMC) via a transformation from the MARTE/pCCSL models into Stochastic Hybrid Automata. The whole process is illustrated through a case study of energy-aware building in which the system is modeled by UML/MARTE/pCCSL and different variants are explored through SMC to help expose the best alternative solutions.

15:30: coffee break

15:50: Session 3 (chair: Frederic Mallet):

- Pr. CHEN Mingsong

Efficient Resource Constrained Scheduling using Parallel Branch-and-Bound Pruning Techniques

Abstract:
Branch-and-bound (B&B) approaches are widely investigated in resource constrained scheduling (RCS). However, due to the lack of approaches that can generate a tight schedule at the beginning of the search and prune the fruitless solutions during the search, B&B approaches usually start with a large initial search space, which makes the following search of an optimal schedule time-consuming. To address this problem, in this talk i will introduce novel parallel B&B approaches that can drastically reduce the overall RCS time. Our approach makes three major contributions: i) it proposes various partial-search heuristics that can quickly find a tight schedule to compact the initial search space; ii) it presents a two-phase search framework that supports the efficient parallel search of an optimal schedule; iii) it investigates various bound sharing and speculation techniques among collaborative tasks to further improve the parallel search performance at different search phases. Experimental results based on well-established benchmarks demonstrate the efficacy of our proposed approach.

- Ass. Pr. ZHANG Min

BIAN: Accelerating Backbone Computing Using Assumptions

Abstract:
Backbone literals of a satisfiable propositional formula are always assigned to 1 in every model of Partial backbone literals are able to increase the performance of SAT-based model enumeration applications (e.g. test case generation) and SAT-based applications (e.g. model checking). We propose a tool BIAN that is able to compute backbone literals efficiently with two main strategies: generating initial assumptions and using neighbor information. Experiments show that BIAN outperforms minibones (state-of-art) using the formulas from SAT competitions (from 2011 to 2017) as benchmarks. Within 3600 seconds, BIAN (178) is able to compute 9% more formulas then minibones(170). For the formulas that are solved by both tools, BIAN (107129 seconds) saved 16% seconds then minibones (126036 seconds).

- Ass. Pr. Julien DeAntoni

Making Concurrency and Time Explicit in the Definition of Operational Semantics

Abstract:
We will briefly present how we defined concurrent and timed operational semantics based on a formal model of time and rewriting rules. Based on such definition, we were able to generate a dedicated interpreter (providing debugging/model checking facilities) and to define the synchronization between heterogeneous model executions, based on coordination patterns.

17:20: Closing.