Timed and Hybrid Systems 2009-2014

Oded Maler

CNRS - VERIMAG
Grenoble, France

November 2014

Introduction

- Our focus: **model-based analysis of systems** in the large sense, not attached too strongly to a specific application domain (but motivated and inspired by some)
- Phenomena that can be modeled as complex **dynamical systems** of different types
- For which we develop and adapt analysis techniques originating from algorithmic verification
- **Hybrid** Systems:
- Analysis of systems that admit **numerical state variables**: differential equations, discrete-time systems, hybrid automata, programs (in principle)
- **Timed** systems:
- Analysis of discrete system where **quantitative timing information** (execution time, delay) is represented explicitly

Human Resources

- **Relatively permanent**:
- Oded Maler (DR1 CNRS), Thao Dang (DR2 CNRS)
  Goran Frehse (MdC UJF), Olivier Lebeltel (IR CNRS)
- **Post-docs**: Stefano Minopoli, Eduardo Carrilho
- **PhD students**:
- Irini-Eleftheria Mens, Jan Lanik, Abhinav Srivastav, Thomas Ferrere, Dogan Ulus, Tommaso Dreossi, Alexandre Rocca

Past Members

- **Long-term Visitors**: Adam Halasz
- **Post-docs**: Alexandre Donzé, Scott Cotton, Piotr Niemczyk
- **Engineers and interns**: Noa Shalev, Rodolfo Ripado, Gabriel Vincent, Brian Vautier, Subhankar Mukherjee, Rajat Kateja, Manish Goyal, Ioannis Galanomatis, Poorna Alamanda
- **PhD Students**:
  - Julien Legriel, (10/2011)
  - Rajarshi Ray, (06/2012)
  - Jean-Francois Kempf (10/2012)
  - Selma Saidi, (10/2012)
  - Romain Testylier (10/2012)
  - Pranav Tendulkar (10/2014)

Team's "Philosophy"

- Focus on the following aspects:
- **Modeling**: how to model **new phenomena** mathematically, with algorithmic analysis in mind, not bound to the **current tradition** and practice of the (academic or industrial) domain
- **Complexity**: how to scale up beyond **toy problems**
- This is done first on clean mathematical models, neglecting many details of real world which are important but sometime premature
- Only after/if something significant has been demonstrated we move to the details (externally usable tools, full development chain, modeling all aspects of an application)
- Like any other approach it has pros and cons

Hybrid Automata

- Systems that can switch between several continuous modes due to external or internal events
- For example: opening/closing of valves

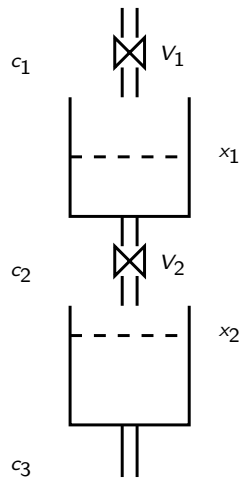

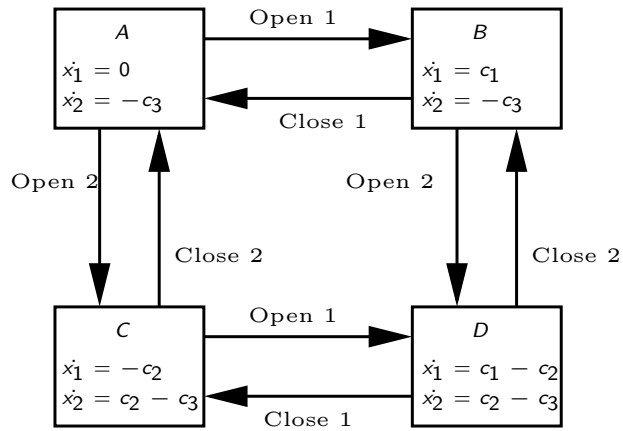

- Hybrid systems, even with trivial continuous dynamics, are very difficult to analyze

Hybrid at a Glance

- Our team is among the world wide creators and leaders of the hybrid systems domain (also known as cyber-physical systems)
- Steering committee of the international conference HSCC, participation in European (Multiform) and national (Malthy, Compacs) projects, dissemination
- Extension of algorithmic verification: computing reachable sets
- Simulation-based techniques, test-generation, state-space and parameter-space exploration, monitoring
- Tool development and integration
- Applications domains: control systems, analog and mixed-signal circuits, systems biology, validation of numerical software

Computing Reachable Sets

- Extension of model-checking to continuous and hybrid systems
- Compute **all** behaviors of a continuous/hybrid system under **all** choices of initial conditions, external disturbances, parameters and transition scheduling; exhaustive simulation
- Set integration: combination of numerical analysis, computational geometry and graph algorithms

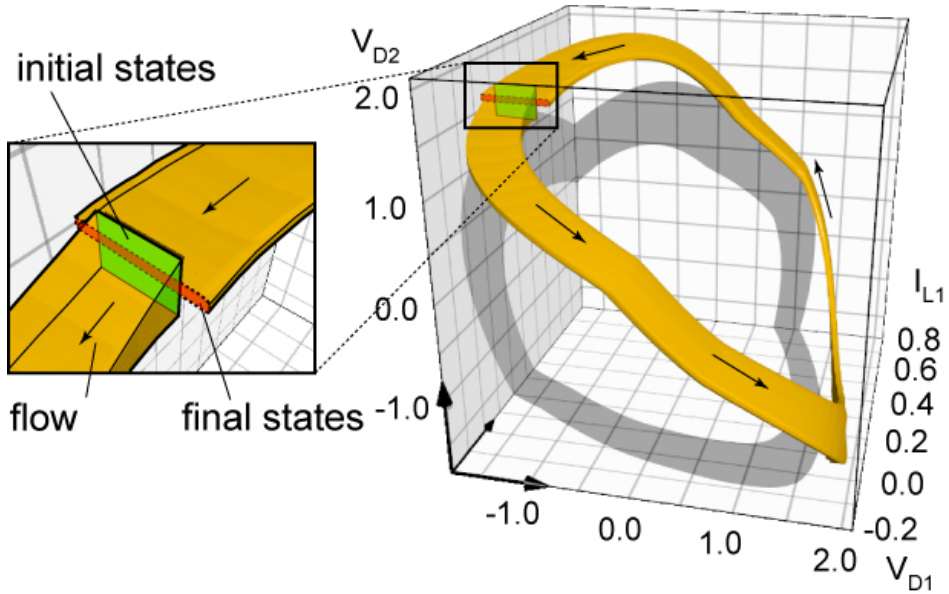

Computing Reachable Sets: Linear Systems

- Systems defined by linear and piecewise-linear differential equations
- During previous period we had a breakthrough in the size of systems that can be treated (Le Guernic thesis 2009)
- Symbolic representation by **support functions** (Girard, LeGuernic 2009) we could increase the dimensionality of systems that can be treated from 10 to several hundreds
- How to consolidate these results developed within a thesis ?

SpaceEx: the State-Space Explorer I

- Under the direction of Goran Frehse we developed a more mature tool implementing these algorithms and much more
- Many "small" details that can be ignored in a scientific publication have to be treated if the tool is to be robust
- Exmaple: splitting reachable sets when the intersection with transition guards happens in several steps (Frehse, Le Guernic, Kateja 2013)
- Usability: a graphical user interface, a model editor, a simulator and the capability to import a sub-class of Simulink models

SpaceEx: the State-Space Explorer II

- **SpaceEx** became the reference tool in the domain with 208 citations since its announcement in 2011
- It has 247 registered users coming from 140 institutions with 10% from industry
- Researchers in other universities use the platform to test new algorithms and teach cyber-physical systems
- The tool is at the center of a new H2020 project with strong industrial participation (Bosch, Esterel, ..)
- Supported by three consecutive Carnot projects and the possibility of a start-up is investigated

SpaceEx: the State-Space Explorer III

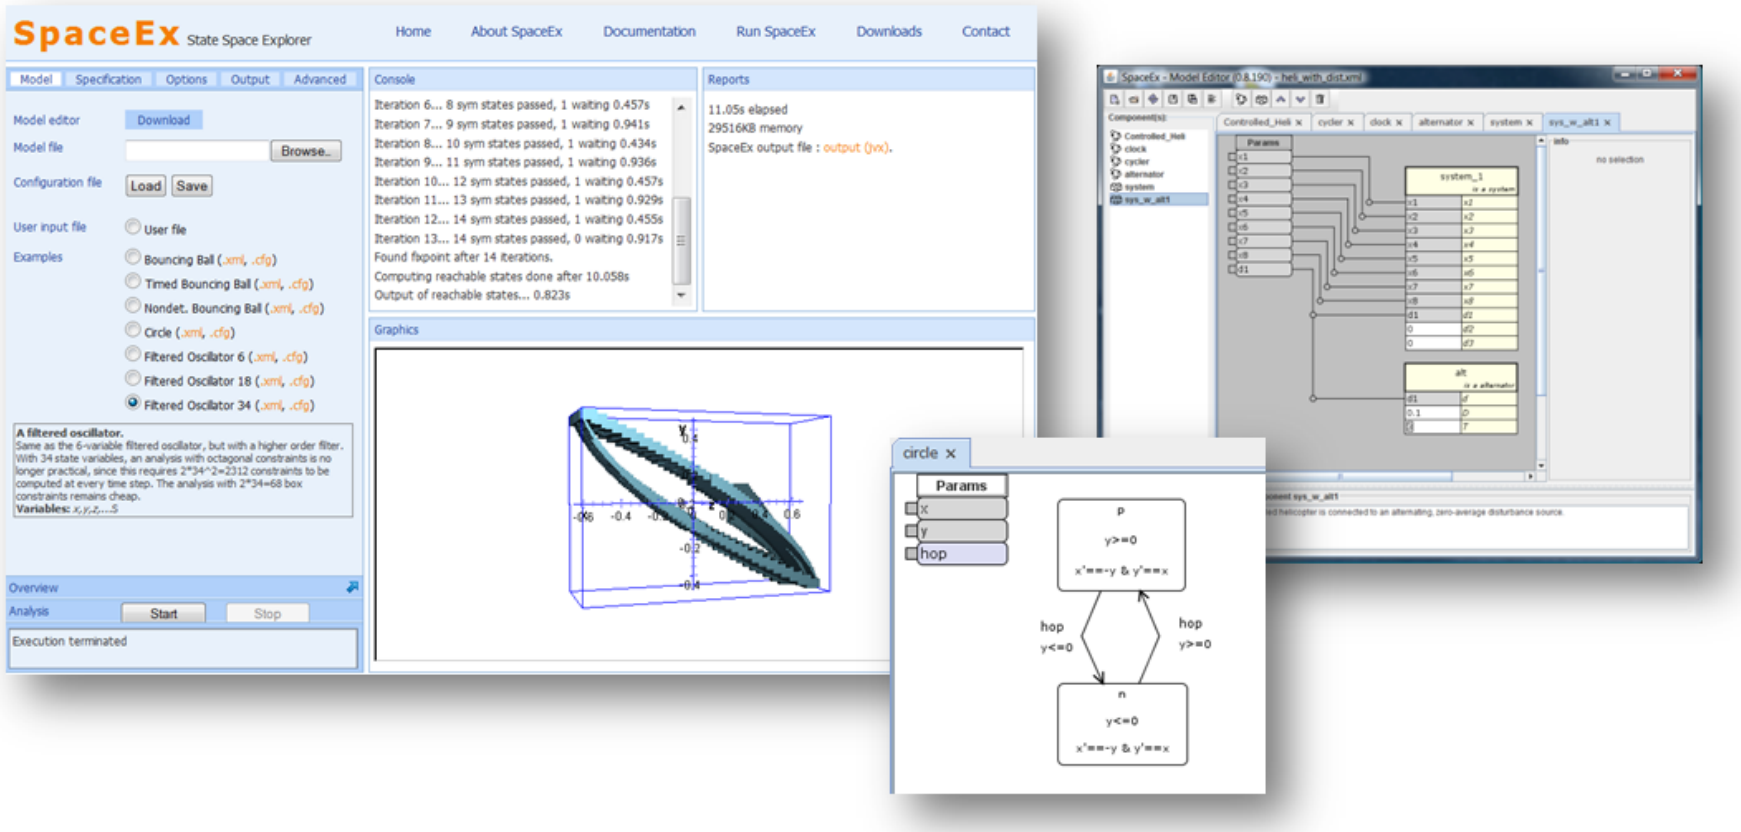

Computing Reachable Sets: Nonlinear Systems I

- Analyzing nonlinear systems is a major challenge in many domains including electrical circuits and biochemical reactions
- More difficult because the nonlinear dynamics does not preserve convexity
- Subject to two theses supervised by Thao Dang (R. Testylier, T. Dreossi)
- Used a variety of methods, the first class specialized to polynomial dynamics
- Using Bernstein expansion of polynomials (Dang, Testylier 2012), applied recently to parameter synthesis of biological models (Dreossi, Dang 2014)

Computing Reachable Sets: Nonlinear Systems II

- Hybridization: a general technique for approximating nonlinear by piecewise-linear systems and then using linear reachability in each linearization domain
- Dynamic hybridization: nonlinear biological models with more than 10 variables

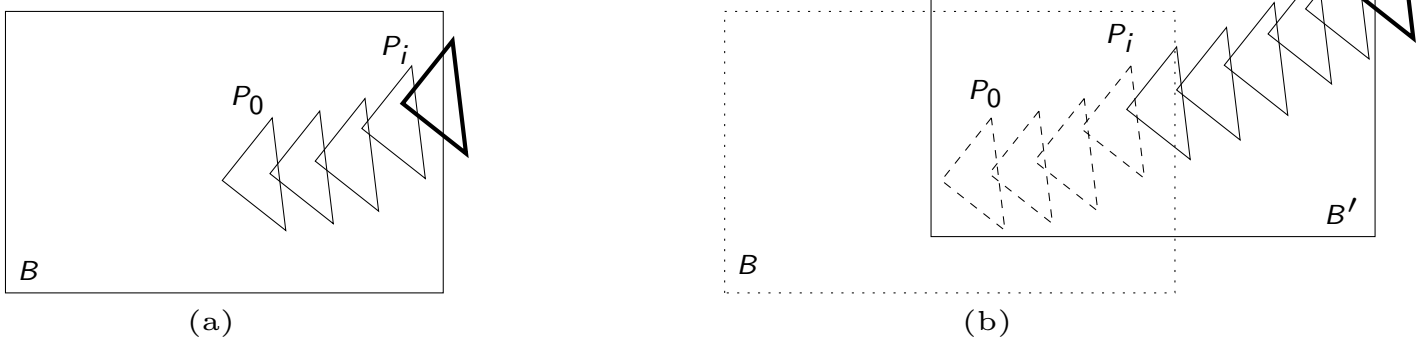

- Implemented in a publicly available tool NLTOOLBOX

Simulation-based Verification

- Verification is somewhat romantic, simulation will always remain the major validation method
- How to improve its effectiveness and rigor?
- Techniques developed by Alexandre Donze, can explore by simulation the parameter-space of a system
- It can approximate the boundary between parameter-values that yield some quantitative-qualitative behavior and those that do not
- Can be applied to systems that can be simulated even if they are not linear (or not even mathematical)
- Scales well with the dimension of the state-space

The Breach Toolboox

- Parameter-space exploration for arbitrary continuous dynamical systems relative to properties expressed in **signal temporal logic** (STL)
- Applied to embedded control systems, analog circuits, biochemical reactions

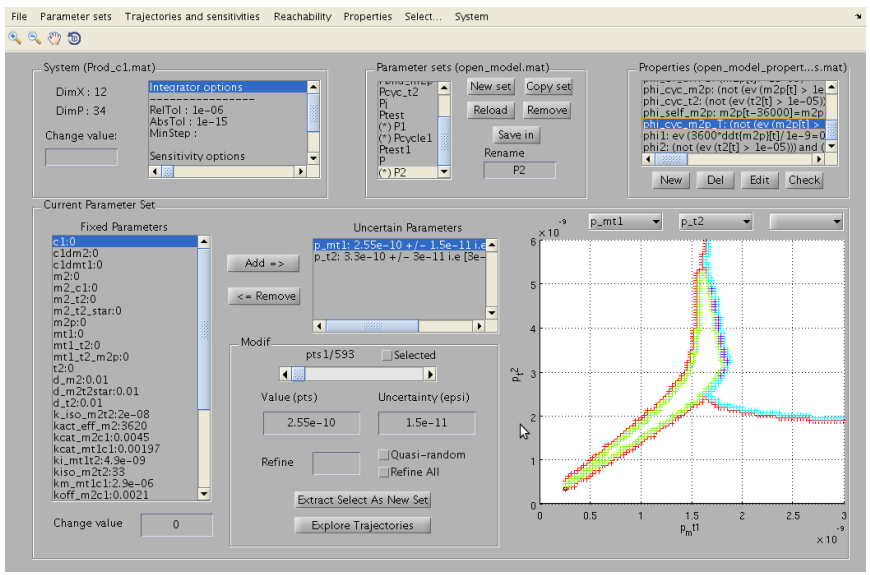

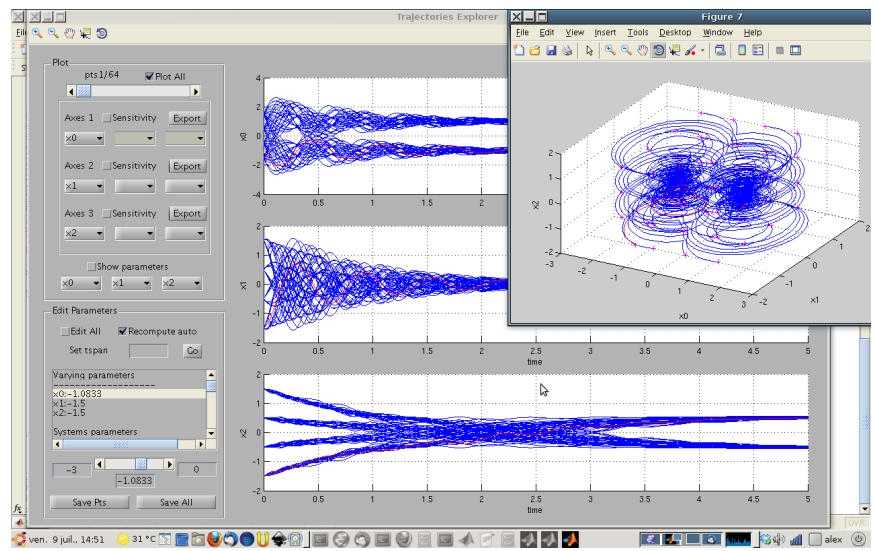

Test Generation I

- How to generate dynamic input stimuli that yield trajectories that cover nicely the reachable state space?
- RRT technique from **robotic motion planning**: biased random search using statistical coverage measures
- HTG tool (Dang, Nahhal 2009)

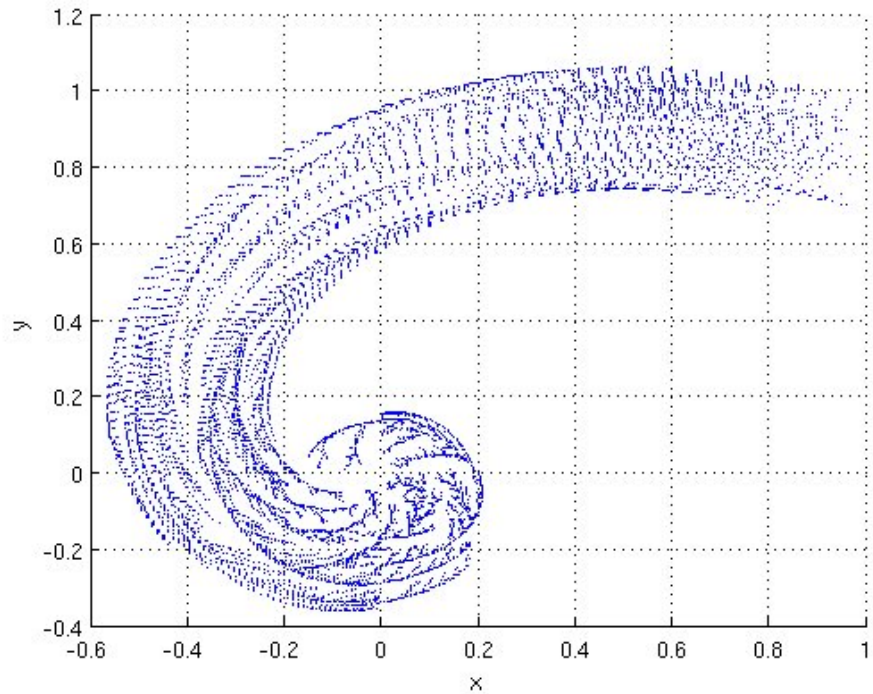

Test Generation II

- Has been applied to SPICE netlists (transistor-level simulation)
- Extended to systems with partial observability (Dang, Shalev 2012)
- Guidance toward the falsification of temporal properties with application to biology (Dang, Dreossi 2013)
- Used in our projects with Toyota and United Technologies

Monitoring Temporal Properties of Continuous Signals

- Monitoring: lightweight (runtime) verification: checking property satisfaction by **individual** behaviors
- In previous period we developed AMT (analog monitoring tool) for **signal temporal logic**
- Automatic derivation of temporal testers, liberate designers from the need to observe simulation traces

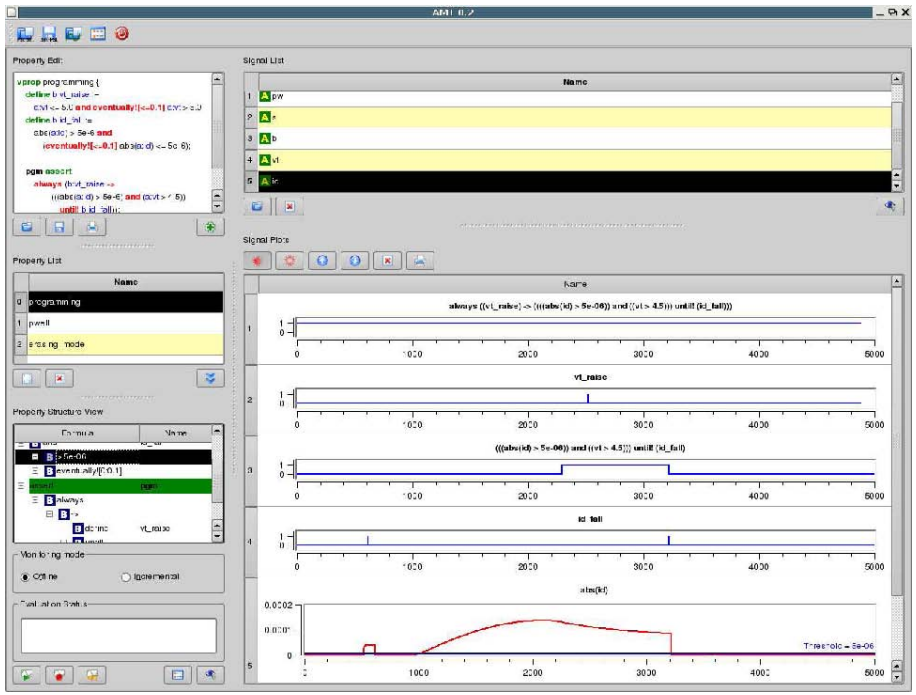

Example: Specifying Stabilization in Temporal Logic

- A **water-level controller** for a **nuclear plant** should maintain a variable y around a fixed level despite external disturbances
- We want y to stay always in the interval $[-30, 30]$ except, possibly, for an initialization period of duration 300
- If y goes outside the interval $[-0.5, 0.5]$, it should return to it within 150 time units and stay there for at least 20 time units
- The property is expressed as

$$\Box_{[300,2500]}((|y| \leq 30) \wedge ((|y| > 0.5) \Rightarrow \Diamond_{[0,150]} \Box_{[0,20]} (|y| \leq 0.5)))$$

Monitoring Stabilization

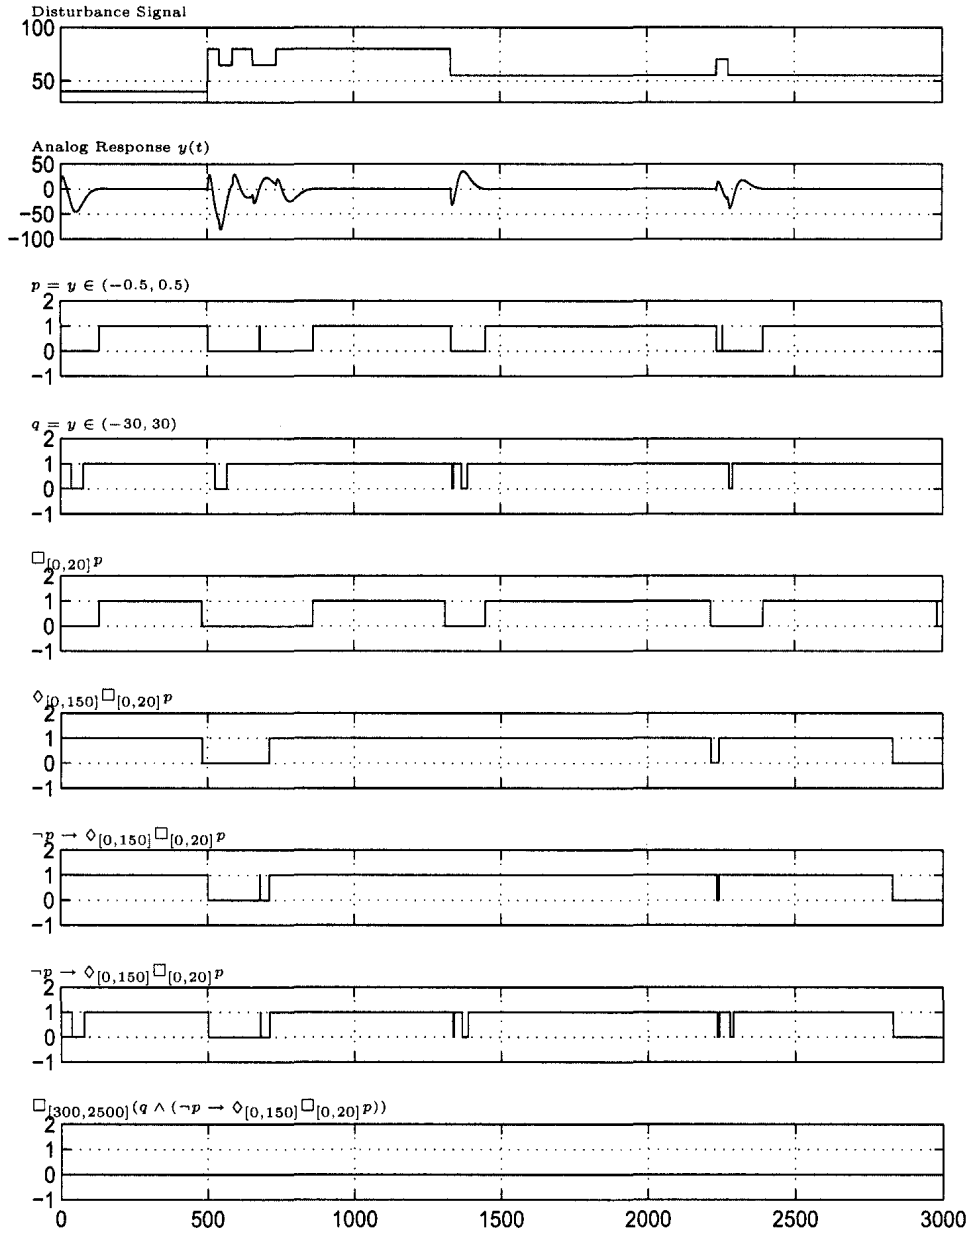

State of the Art

- Surprisingly this relatively simple work (compared to heroic verification efforts), immediately attracted industrial interest. It underlies a CIFRE thesis with Mentor
- We developed numerous extensions:
- Quantitative semantics: not only yes/no but how much (robustness), and an efficient algorithm to compute it (Donze, Ferrere, Maler 2013)
- Inverse problem: how to compute parameters in the formula that render it satisfied by a given set of simulation traces (parametric identification)
- Adding sliding window Fourier operators to specify music and other properties that combine time and frequency domains
- New monitoring/pattern matching algorithm for timed regular expressions (Asarin, Ferrere, Ulus, Maler 2014)
- Underlies a new EDA project with AIT and Easii-ic

Timed Systems: Motivation

- The level of abstraction captured by timed automata is extremely important
- Modeling a piece of hardware or software (or physics or wetware) as a process that **takes some time to complete** has a huge advantage over a detailed model (software, cycle-accurate, SPICE, proteins)

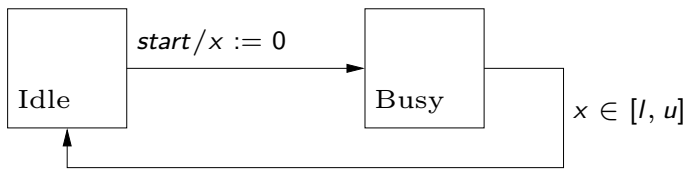

- It allows to do **fast** simulation for performance evaluation and design-space exploration
- It would help the world if timed automata tools could scale beyond toy problems
- Initiated the FORMATS workshop (steering committee)

Fighting the Clock Explosion

- We spend more than a decade trying to scale up the size of timed automata that can be handled by verification tools (Kronos, Uppaal, IF) beyond toy problems
- Covering all possible consequences of timing uncertainty is important for the safety-critical part of the embedded market, but is very difficult
- It is too strong and too weak for "best effort" systems
- Too strong because we do not care about rare events
- Too weak because it cannot give **average** case performance
- Set-theoretic non determinism vs. probability

The ATHOLE Project

- This French regional project with ST, CEA and Thales, brought us closer to surface of the earth
- The issue: performance evaluation and optimization for running application on an embedded multi-core architecture
- Insight: our real contribution is not in exhaustive verification but in high-level **modeling** - in contrast with the overly detailed models used by developers
- We bring some more quantitative abstract thinking to software and hardware engineering
- Below we summarize some results from this project (4 PhD theses)

Multi-criteria Optimization

- Systems are evaluated according to various criteria - cost, performance, consumption, ...
- The optimum concept applies only to one-dimensional domains and functions (linear orders)
- In partial orders there is no unique optimum but a set of **Pareto** solutions: they cannot be improved in one criterion without being worsened in the other
- They represent the trade-offs between conflicting criteria
- We developed a general methodology for approximating the Pareto front using an SMT solver, a **multi-dimensional** generalization of **binary search** (Thesis J. Legriel, 2011)
- Developed an alternative technique based on stochastic local search (Legriel, Cotton) and working on a general scheme (ongoing thesis of A. Srivastav)

Application to Deployment

- Extensive application of these ideas for deploying (mapping and scheduling) streaming applications on multi-cores (thesis of P. Tendulkar, 2014)
- Applications expressed as split-join graphs
- Architectures with shared (Tilera) and distributed (Kalray) memory
- Show trade-offs between latency, power consumption and memory in the
- New symmetry breaking results to reduce search by the SMT solver (with P. Poplavko)
- Extensive experimental evaluation
- Automatic sizing and shaping of DMA blocks (Thesis of S. Saidi, 2012)

Performance Evaluation and Design-Space Exploration I

- We developed a performance exploration tool (DESPEX) based on high-level models
- It has an input language to define all the components that influence performance:
- **Application**: task graphs annotated with workloads and size of data
- **Architecture**: simple models of processors, interconnect, memories and their features: speed, latency, bandwidth
- **Environment**: input generators that produce streams of jobs to be execute according to some constraints: periodic, jitter, bounded uncertainty, bounded variability
- **Deployment**: mapping and scheduling policies

Performance Evaluation and Design-Space Exploration II

- All these are compiled into timed automata used for three types of analysis:
- 1) Standard timing verification by reachability analysis of timed automata
- 2) Monte-Carlo simulation interpreting timing uncertainty probabilistically
- 3) Piecewise-analytic computation of expected performance
- Thesis of JF Kempf (2012)

Additional Research Results

- Duration probabilistic automata: time automata with durations distributed uniformly in bounded intervals: new average case analysis and optimal synthesis algorithms (Kempf, Bozga, Maler 2011, 2013)
- Results on entropy of timed languages (Asarin, Degorre 2009)
- Learning over large alphabets (Mens, Maler 2014)
- ...

Impact

- Publications 2009-2014: 82, 5 per permanent per year
- Major ongoing projects: Malthy, Compacs, Eqinocs, Cadmidia (ANR), Toyota, UTRC (industrial)
- Thesis defended 2010-2014: 6
- Organization and leadership of conferences (HSCC, FORMATS, CAV 2009) workshops (FAC, HSB, SynCoP, ARCH)
- Dissemination by numerous survey and tutorial articles

Major Industrial Collaborations

- ST: multi-core (ATHOLE project, ended) and analog CAD (Nano2017 submissions)
- Mentor Graphics: CIFRE on measurement and assertions in circuit simulators
- ATRENTA: CIFRE on power reduction in SoC
- Toyota (USA): application of our technology to find bad behaviors in automotive models
- United Technology Research Center (Ireland): application of our technology to find bad behaviors in HVAC models