Project no.: **ICT-FP7-STREP-214755**

Project full title: **Quantitative System Properties in Model-Driven Design**

Project Acronym: **QUASIMODO**

Deliverable no.: D3.4

Title of Deliverable: Synthesizing controllers with bounded resources

Contractual Date of Delivery to the CEC:	**Month 24**
Actual Date of Delivery to the CEC:	**Month 24 (January 1, 2010)**
Organisation name of lead contractor for this deliverable:	**CNRS**
Author(s):	**Nicolas Markey, Shuhao Li, Jean-François Raskin, Mariëlle Stoelinga**
Participant(s):	**P01 AAU, P02 ESI, P03 CNRS, P06 CFV**
Work package contributing to the deliverable:	**WP 3**
Nature:	**R+P**
Version:	**1.00**
Total number of pages:	**11**
Start date of project:	1 Jan. 2008 **Duration:** 36 month

Project co-funded by the European Commission within the Seventh Framework Programme (2007-2013)

Dissemination Level

PU	Public	X
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	

Abstract:

This deliverable reports on the works carried out inside the QUASIMODO consortium on the synthesis of controllers in a setting where the resources are bounded. This includes many different cases: controllers that work e.g. with a finite amount of memory, controllers that have access to a limited part of the system, or controllers that ensure that the system under supervision preserves the amount of some resource (energy, level of oil in a tank, ...) between some given bounds.

Keyword list: controller synthesis, bounded memory, partial observability, energy constraints

Contents

Introduction

The automatic verification is of utmost importance in the design of reactive and embedded systems. In case the system must react to the uncontrollable actions of the external environment, it amounts to synthesizing a controller, whose role is to keep the system in a safe state. In most cases, the controller will be embedded on the system and may have to work with a bounded amount of resources (*e.g.*, bounded memory), may have access to only part of the system (*e.g.*, imperfect information, partial controllability, ...), or may have to ensure that the system under control uses a bounded amount of resources (*e.g.*, energy).

Our work relates to all three issues:

- first, we developped an extension of ATL, a temporal logic which can express controllability properties. Roughly, ATL extends CTL by replacing the usual (existential and universal) path quantifiers with *strategy quantifiers*, which may be used to express that *there is a strategy for the controller to keep the system in a safe state*. Our extension can express the extra requirement that the strategy should only use a limitted amount of memory on the history of the computation.

- second, we designed compositional and symbolic algorithms for synthesizing controllers under imperfect information and partial controllability.

- finally, we present *energy games*, *i.e.*, games played on a weighted graph where the weights represent energy consumption or storage, and where the aim is to never run out of energy. These games have been defined within the QUASIMODO consortium in order to model the HYDAC case study. The algorithmic study of the numerous variants of this problem is still in progress, and we report here our preliminary results.

1 Bounded-memory strategies

1.1 ATL with strategy contexts and bounded memory

1.1.1 Participants

- Thomas Brihaye, CFV, University of Mons, Belgium
- Arnaud Da Costa, Nicolas Markey, CNRS/LSV, ENS Cachan, France
- François Laroussinie, CNRS/LIAFA, University Paris 7, France

1.1.2 Contributions

Over the last ten years, ATL [AHK02] has been proposed as a new flavor of temporal logics for specifying and verifying properties in multi-agent systems (modeled as Concurrent Game Structures (CGS)), in which several agents can concurrently act upon the behaviour of the system. In these models, it is not only interesting to know if something can or will happen, as is expressed in CTL or LTL, but also if some agent(s) can control the evolution of the system in order to enforce a given property, whatever the other agents do. ATL can express this kind of properties thanks to its quantifier over strategies.

We extend ATL in two directions [BDLM09]: first, while classical ATL strategy quantifiers drop strategies introduced by earlier quantifiers in the evaluation of the formula, our logic keep "executing" those strategies. To achieve this idea, we naturally adapt the semantics of ATL in order to interpret a formula within a stategy context. Our new strategy quantifier can for instance express that "agent A has a strategy s.t. (1) Agent B always has a strategy (given that of A) to enforce some property ϕ_B and (2) Agent C always has a strategy (under the same strategy of A) to enforce ϕ_C". This allows to really express collaboration between different agents.

While ATL can be proved to admit memoryless strategies, it is not the case for our logic ATL$_{sc}$: in general, when the above property holds true, any winning strategy for agent A will use memory. Our second extension is quantitative: it decorates strategy quantifiers with the maximal amount of memory allowed for the strategy. Memory here is measured in terms of the number of states of an auxiliary automaton for each agent, which she can manage as she wants.

We studied the expressiveness of our logic, and proved that it has many surprising properties: first, our logic is more powerful than the usual *alternating bisimulation*. As a consequence, it is strictly more expressive than the usual ATL. Moreover, it can express many interesting properties, such as the one above, and also several kinds of equilibria (*e.g.* Nash equilibria). Also, bounds on the amount of memory cannot be expressed in ATL or ATL*.

On the algorithmic side, we prove that our logic is decidable: in the most general case, the algorithm involves alternating parity tree automata, having non-elementary complexity. In the case where all the strategy quantifiers have bounded memory, the algorithm runs in exponential space.

References

[AHK02] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, September 2002.

[BDLM09] Thomas Brihaye, Arnaud Da Costa, François Laroussinie, and Nicolas Markey. ATL with strategy contexts and bounded memory. In Sergei N. Artemov and Anil Nerode, editors, *Proceedings of the Symposium on Logical Foundations of Computer Science (LFCS'09)*, volume 5407 of *Lecture Notes in Computer Science*, pages 92–106, Deerfield Beach, FL, USA, January 2009. Springer.

2 Games with imperfect information

2.1 Compositional Control Synthesis for Partially Observable Systems

2.1.1 Participants

- Wouter Kuijper and Jaco C. van de Pol, ESI, University of Twente, the Netherlands.

2.1.2 Contribution

The main difficulty that any effective procedure for controller synthesis must face is that the uncontrolled state space generated by the plant description is typically large. This is mainly due to concurrency in the model, which is a central issue also in model checking. However, for synthesis the problem is amplified by two additional, complicating factors. First, we typically see a higher degree of non-determinism because *a priori* no control constraints are given. Second, it is often the case that the state of the plant P is only partially observable for the controller C. Resolving this may incur another exponential blowup. On instances, this blowup may be avoided by using smart, symbolic methods.

In [KvdP09], we present a compositional method for deriving control constraints on a network of interconnected, partially observable and partially controllable plant components. The constraint derivation method works in conjunction with an antichain-based, symbolic algorithm for computing weakest strategies in safety games of imperfect information. We demonstrate how the technique allows a reactive controller to be synthesized in an incremental manner, exploiting locality and independence in the problem specification.

The paper focuses on the compositional synthesis of a reactive controller under the assumption of partial observability. Our main contributions are a compositional framework for describing control synthesis problems as a network of interconnected, partially controllable, partially observable plant components, and a compositional method for synthesizing safety controllers over such a plant model. We believe there are at least two novel aspects to our approach. First, there is the combination of imperfect information with compositionality. In particular, we make sure that context assumptions take into account partial observability of the components. Second, our framework ensures context assumptions gradually shift in the direction of control constraints as the scope widens. In this way we avoid, to some extent, unrealistic assumptions, and generally obtain a less permissive context. Note that the size of the context assumptions is an important factor in the efficiency of assume-guarantee based methods.

References

[KvdP09] W. Kuijper and J. C. van de Pol. Compositional control synthesis for partially observable systems. In M. Bravetti and G. Zavattaro, editors, *Proc. of CONCUR'09, Concurrency Theory*, volume 5710 of *Lecture Notes in Computer Science*, pages 431–447. Springer, 2009.

2.2 Computing Weakest Strategies for Safety Games of Imperfect Information

2.2.1 Participants

- Wouter Kuijper and Jaco C. van de Pol, ESI, University of Twente, the Netherlands.

2.2.2 Contribution

CEDAR (Counter Example Driven Antichain Refinement) is a new symbolic algorithm for computing weakest strategies for safety games of imperfect information. The algorithm computes a fixed point over the lattice of contravariant antichains. Here contravariant antichains are antichains over pairs consisting of an information set and an allow set representing the associated move. In [KvdP09], we demonstrate how the richer structure of contravariant antichains for representing antitone functions, as opposed to standard antichains for representing sets of downward closed sets, allows CEDAR to apply a significantly less complex controllable predecessor step than previous algorithms.

In our approach we limit the scope to safety objectives which allows for a symbolic algorithm that can compute weakest strategies. This is complementary to the approach of [CDH09]. Our approach is useful for cases where (1) everything that one wants to synthesize is expressible as a safety property (*e.g.* hard timeliness constraints for instance are expressible as a safety property) and/or (2) the result must be reusable and amenable to further analysis/composition/optimization. It is especially in case (2) where weakest strategies really shine. Since a weakest strategy for a given game subsumes all possible safe strategies it is useful as a safety monitor (*i.e*, for supervising software or users that cannot be completely guaranteed safe). Computing the weakest strategy may also form part of a preprocessing step for generating a safe input graph to a second synthesis procedure that can optimize some performance measure that is not expressible as a safety property. Finally, weakest safety strategies are useful in a compositional setting where the behaviour of the context is not known beforehand so that a most general solution is necessary in order not to exclude possible safe compositions with a concrete context. Our main contribution is a new algorithm named CEDAR (Counter Example Driven Antichain Refinement). In a nutshell, the algorithm computes a fixed point over an enriched form of antichains which we call contravariant antichains. Contravariant Antichains enjoy most of the properties of normal antichains, but they can represent knowledge based strategies, which are antitone functions from information sets to allow sets, rather than just sets of downward closed information sets. This additional structure allows us to symbolically compute, not just the set of winning initial information sets, but the entire weakest knowledge based strategy. As a second contribution our approach permits to significantly simplify the controllable predecessor step. In contrast to [CDH09] we only treat a single counterexample observation to the observation-closedness condition for the contravariant antichain, as opposed to treating all counterexample observations at every iteration.

References

[CDH09] Krishnendu Chatterjee, Laurent Doyen, and Thomas A. Henzinger. Probabilistic weighted automata. In Mario Bravetti and Gianluigi Zavattaro, editors, *Proc. of CONCUR'09, Concurrency Theory*, volume 5710 of *Lecture Notes in Computer Science*, pages 244–258. Springer, 2009.

[KvdP09] W. Kuijper and J. C. van de Pol. Computing weakest strategies for safety games of imperfect information. In S. Kowalewski and A. Philippou, editors, *Proc. of TACAS'09, Tools and Algorithms for the Construction and Analysis of Systems*, volume 5505 of *Lecture Notes in Computer Science*, pages 92–106. Springer, 2009.

3 Games with energy constraints

3.1 Timed automata and games with energy constraints

3.1.1 Participants

- Patricia Bouyer, Nicolas Markey, CNRS/LSV, ENS Cachan, France;
- Uli Fahrenberg, Kim G. Larsen, Jiří Srba, CISS, Aalborg University, Denmark.

3.1.2 Contribution

Priced timed games are timed games whose locations and edges are decorated with prices: prices in locations are *prices per time unit spent in the location*, while prices on edges are paid each time the transition is fired. As opposed to hybrid systems, prices cannot serve in guards on transitions. In the 1-player case (*i.e.*, in priced timed automata), this restriction makes model checking and optimization decidable [BFH$^+$01, ALP01].

The Hydac case study led us to consider an alternative semantics for priced timed games (and automata): in this case study, a pump must be turned on and off regularly in order to maintain the level of oil in a tank between two bounds. This can be modelled by priced timed automata under *energy constraints*: the accumulated price since the beginning of the run must be kept between a lower and an upper bound. The name *energy constraint* comes from the analogy with batteries, where energy can be consummed and regain, with the aim to never run out of energy.

Our preliminary results are as follows: the very bad news is that, even with only one clock, the problem of controlling the system in order to keep the accumulated price between a lower- and an upper-bound is undecidable [BFL$^+$08]. On the other hand, in the untimed case, we proved that the problem is solvable in exponential time.

In order to have more intuitions on the game problem, we also studied the 1-player case (model-checking): for this restricted case, we proved that the 1-clock case with only a lower-bound constraint is already very intricate [BFLM09]: we designed an algorithm computing the optimal accumulated cost that can be achieved when reaching a final location, under a lower-bound constraint along the path, and depending on the initial credit. We also extended the model of prices to also handle "exponential prices", *i.e.*, prices that follow a differential equation of the form $\frac{dp}{dt} = k \cdot p$ where k is the "rate" of the location where time is elapsed. We also gave an algorithm for optimizing the final credit under lower-bound constraint in this case, under the restriction that discrete prices on transitions must be nonpositive.

In this setting, many problems remain open: our next aim is to tackle the game problem with only lower-bound constraints. Another important question w.r.t. priced timed games is the following: all the undecidability proofs rely on an encoding of a two-counter machines, with the values of both counters being encoded in the value of clocks or costs. This requires infinite precision. Studying this problem in the presence of guard enlargement (or clock drifts) would be especially interesting. Finally, a more precise study of the algorithms for solving the above problems would be worth it. Our first results are presented in the next section.

References

[ALP01] Rajeev Alur, Salvatore La Torre, and George J. Pappas. Optimal paths in weighted timed automata. In Maria Domenica Di Benedetto and Alberto L. Sangiovani-Vincentelli, editors, *Proceedings of the 4th International Workshop on Hybrid Systems: Computation and Control (HSCC'01)*, volume 2034 of *Lecture Notes in Computer Science*, pages 49–62. Springer-Verlag, March 2001.

[BFH+01] Gerd Behrmann, Ansgar Fehnker, Thomas Hune, Kim Gulstrand Larsen, Paul Pettersson, Judi Romijn, and Frits Vaandrager. Minimum-cost reachability for priced timed automata. In Maria Domenica Di Benedetto and Alberto L. Sangiovani-Vincentelli, editors, *Proceedings of the 4th International Workshop on Hybrid Systems: Computation and Control (HSCC'01)*, volume 2034 of *Lecture Notes in Computer Science*, pages 147–161. Springer-Verlag, March 2001.

[BFL+08] Patricia Bouyer, Uli Fahrenberg, Kim G. Larsen, Nicolas Markey, and Jiří Srba. Infinite runs in weighted timed automata with energy constraints. In Franck Cassez and Claude Jard, editors, *Proceedings of the 6th International Conference on Formal Modelling and Analysis of Timed Systems (FORMATS'08)*, volume 5215 of *Lecture Notes in Computer Science*, pages 33–47, Saint-Malo, France, September 2008. Springer.

[BFLM09] Patricia Bouyer, Uli Fahrenberg, Kim G. Larsen, and Nicolas Markey. Timed automata with observers under energy constraints. Submitted, October 2009.

3.2 Pseudo-polynomial algorithms for energy games and mean-payoff games

3.2.1 Participants

- Laurent Doyen, Rafaella Gentilini, and Jean-Franois Raskin, CFV, Brussels, Belgium

3.2.2 Contribution

We show that (untimed) energy games with lower-bound constraints can be solved elegantly and efficiently using a notion of progress measure[DGR09]. Progress measures for weighted graphs are functions that impose local conditions to ensure global properties of the graph, in this case that all cycles are nonnegative. We show how to transfer this notion from graphs to games, and we provide an efficient fixpoint algorithm to synthesize a progress measure when it exists.

We end up with an algorithm solving energy games in time $O(|E| \cdot |M_G|)$, where E is the size of graph and M_G is the sum of the negative weights in the automaton. We also proved that lower-bound energy games are logspace reducible to mean-payoff games (games where the aim is to maintain the mean accumulated price above 0; these games are fundamental in theoretical computer science, as they have been proven reducible in polynomial time to μ-calculus model checking or to parity games). Quite surprisingly, when adapted to the setting of mean-payoff

games, our algorithm for energy games has better complexity than the best-known algorithms so far: for instance, our new algorithm for mean payoff games runs in deterministic time $O(|E| \cdot |V| \cdot W)$ where W is the maximal absolute price, while the best known deterministic algorithm so far would run in time $\Theta(|E| \cdot |V|^2 \cdot W)$ or $O(|E| \cdot |V| \cdot 2^{|V|})$ [ZP96, LP07].

References

[DGR09] Laurent Doyen, Rafaella Gentilini, and Jean-Franois Raskin. Faster pseudo-polynomial algorithms for mean-payoff games. Submitted, June 2009.

[LP07] Yuri M. Lifshits and Dmitri S. Pavlov. Potential theory for mean payoff games. *Journal of Mathematical Science*, 145(3):4967–4974, September 2007.

[ZP96] Uri Zwick and Mike Paterson. The complexity of mean payoff games on graphs. *Theoretical Computer Science*, 158(1-2):343–359, May 1996.

3.3 Verification and controller synthesis for resource-constrained real-time systems: case study of an autonomous truck

3.3.1 Participants

- Kim G. Larsen, Shuhao Li, CISS, Aalborg University, Denmark;
- Paul Petterson, Mälardalen University, Sweden.

3.3.2 Contribution

We model a real-time embedded system as a network of interacting timed game automata, where the system masters the controllable actions, and the (hostile) environment masters the uncontrollable actions. We propose to model the resources that are of continuous nature (*e.g.*, memory, network bandwidth, energy) by using piecewise constant integer functions and integer arithmetics. The resource status may affect the behaviors of the system, both functionality-wise and timing-wise. We verify a number of properties which involve resource constraints in the timed model checker Uppaal. After specifying some resource-constrained reachability or safety control objectives, we check whether the control problems are solvable in the timed game solver Uppaal-Tiga, and if yes, we synthesize winning strategies for the system. The suggested approach is tried out on a case study of a battery-powered autonomous truck. Experimental results indicate that the method is effective and computationally feasible.