Light-weight hybrid model checking facilitating online prediction of temporal properties

Gerald Sauter[1], Henning Dierks[2], Martin Fränzle[1], and Michael R. Hansen[3]

[1] Carl von Ossietzky Universität, Oldenburg, Germany
{fraenzle|sauter}@informatik.uni-oldenburg.de
[2] Hochschule für Angewandte Wissenschaften, Hamburg, Germany
henning.dierks@haw-hamburg.de
[3] Technical University of Denmark, Kgs. Lyngby, Denmark
mrh@imm.dtu.dk

We address the question of system-in-the-loop monitoring of dynamic systems with a moderately fast sampling rate. The objective is, given a partial time series of (possibly inexact) data obtained from online measurements of observable variables of the system, to safely predict the truth value of a Linear-time Temporal Logic (LTL) formula including future modalities, in real-time while the system is evolving. We therefore develop an online monitoring tool providing early warning based on safe extrapolation from system history, where the history is recorded as a time series of observations and, in order to achieve sound predictions, the future is safely approximated using differential equations and other constraints pertinent to the system dynamics. The systems considered are heterogeneous systems from the automotive domain, featuring a seamless integration of cognitive models for human behavior, traffic models, model-based designs of driver assistance systems, and models describing the dynamic behavior of the car.

Tool support for online monitoring of hybrid systems with respect to requirements expressed using temporal logics has previously been studied in [NM07], where a property based tool (named ATM) for monitoring analogue systems is described. The properties are expressed as formulae of STL/PSL, an extension of MITL [AFH96] and STL [MN04]. A serious limitation of AST, however, is that formulae are just interpreted over a history of past events, forcing online monitoring to fall behind real-time. Our work can be seen as an extension of that in the sense that we interpret formulae over infinite extensions of the monitored traces, where the future events are safely approximated using laws describing the dynamics of the system. To arrive at useful safe over-approximations of the possible infinite extensions of the sampled finite trace, interval constraint propagation [BG06] is used to efficiently estimate ranges of observations in the bounded future based on ordinary differential equations (ODE) governing parts of the system dynamics, thereby exploiting a mean-value form of the ODEs akin to that used in the hybrid-system model-checker HSolver [RS05]. An offline preprocessing step is used to permit efficient online computation of estimates covering the unbounded future.

Following the intuition that dynamical systems gradually turn their state to the good or bad, we let our monitor output a quantitative figure in the form of a numerical interval providing a safe estimate of how severe a property violation can be, instead of a qualitative yes/no answer. To this end, we resort to a *robust semantics* for LTL, where the meaning of a formula is the degree of accuracy within which it is satisfied. This notion of robustness originates from [Rat00], where it is applied to arithmetic first-order constraints, and has been lifted to the linear-time temporal logics Duration Calculus [FH05] and LTL [FLS08]. We adopt the latter and extend it to an interval-valued interpretation over interval-valued traces, thus being able to accommodate inexact measurements as well as not fully determined extrapolated values in the semantics.

1 Semantics

Reflecting the fact that in online monitoring, traces are obtained by measurements which may be inexact, we associate intervals to the individual measurements on the observable variables. Given a finite set of observables $\mathcal{O}$ (with typical elements $x, y, z, \dots$), we assume that measurements for such an observable $x \in \mathcal{O}$ range over a non-empty interval $[l_x, u_x] \subseteq \mathbb{R}$. We call a mapping

$V : \mathcal{O} \longrightarrow 2^{\mathbb{R}}$, where for each observable $x \in \mathcal{O}$ it holds that $V(x)$ is a non-empty interval and $V(x) \subseteq [l_x, u_x]$ an *interval valuation* of $\mathcal{O}$. Let $\mathcal{V}$ denote the set of all interval valuations of $\mathcal{O}$. We call a mapping $T : \mathbb{N} \longrightarrow \mathcal{V}$ an *interval trace* of $\mathcal{O}$. Hence, an interval trace describes the evolution of the observables over discrete time where for each time instance and each observable an interval of values is given. An interval valuation can be conceived as a (convex) set of standard valuations. Based on interval traces we define a semantics of LTL that robustly interprets LTL formulae over such traces and assigns an interval encoding the range of possible truth values as well as the robustness of these truth values. The *quantitative interval semantics* $T[\![\varphi]\!] : \mathbb{N} \longrightarrow 2^{\mathbb{R}}$ for an LTL formula φ, of an interval trace T, is inductively defined as follows:

$$\begin{array}{ll} T[\![x > k]\!]\, i = T(i)(x) - [k,k] & T[\![\neg\varphi]\!]\, i = -T[\![\varphi]\!]\, i \\ T[\![\varphi \wedge \psi]\!]\, i = \min(T[\![\varphi]\!]\, i, T[\![\psi]\!]\, i) & T[\![\bigcirc\varphi]\!]\, i = T[\![\varphi]\!]\, (i+1) \\ T[\![\Box\varphi]\!]\, i = \inf_{j \geq i} T[\![\varphi]\!]\, j & T[\![\varphi \mathsf{U} \psi]\!]\, i = \sup_{j \geq i} \min(T[\![\psi]\!]\, j, \min_{k=i}^{j-1} T[\![\varphi]\!]\, k) \end{array}$$

Here, i denotes the temporal position in the trace at which the formula is evaluated and $-$, min, sup, etc., denote the interval liftings of the corresponding base operations.

The idea of this definition is that the interval $T[\![\varphi]\!]\, i$ encodes two kinds of information. First, it encodes the possible truth values arising from the contained point-valued valuations in the standard interpretation of LTL: If the supremum of $T[\![\varphi]\!]\, i$ is greater than 0, then T contains a classical (i.e., point-valued) valuation such that the LTL formula φ becomes true in the standard semantics of LTL. Analogously, a negative infimum encodes that T contains a (standard) valuation for which φ is false. The second kind of information encoded is a measure of robustness: If $m > 0$ is the supremum of $T[\![\varphi]\!]\, i$ this means that there is a (standard) valuation T_s contained in T such that T_s at time instant i satisfies φ and does also satisfy all formulae φ' which are structurally equal to φ, yet differ in the constants in the formulae, which may vary by up to $\pm m$. If $m < 0$ is the infimum the valuation is respectively used.

The pragmatics of using this semantics in online monitoring is as follows. It permits the direct accommodation of inexact measurements and the capturing and continuously reporting the time-wise monotonic increase in information about the requirements monitored: For all interesting formulae, i.e. all formulae being neither tautologic nor unsatisfiable, the "truth interval" reported by the monitor starts with a large interval covering both positive and negative values. In the course of monitoring, this truth interval will be monotonically refined with every new measurement until it finally may fall completely into the positive range (or, alternatively, into the negative), which implies that every possible continuation of the observed time series robustly satisfies (violates, resp.) the formula with a robustness margin at least equalling the infimum (supremum, resp.) of the reported interval.

2 Online monitoring with prediction

Aiming at reliable prediction of truth values of temporal formulae based on finite time series of measurements and on known constraints on the system dynamics, we provide mechanisms to achieve a *finite representations of the infinite interval traces*, to *evaluate these under the terms of the interval semantics* and to *refine these by the means of known facts about the dynamics and upon the appearance of new measurements*. Thereby a decomposition into an offline phase, which precomputes all computationally expensive entities, in particular forms of branching, and an online part which exploits these precomputed properties in an analysis based on computationally inexpensive interval constraint propagation (ICP), is obtained.

The finite representation of infinite traces is based on extending a finite interval trace $tr : \{0, \ldots, n\} \longrightarrow V$ by an extra *collecting interval* $v_\infty \in V$ representing the indefinite future from time instant $n+1$ onward. This is achieved by employing overapproximation, i.e., an infinite trace T is represented by (t, v) iff $\forall i \leq n\, \forall o \in \mathcal{O} : T(i)(o) \subseteq t(i)(o)$ and $\forall i > n\, \forall o \in \mathcal{O} : T(i)(o) \subseteq v(o)$. Evaluation of the interval semantics on such a finite representation (t, v) is straightforward, even in case of unbounded future temporal operators. For the sake of efficiency, as required in real-time online monitoring, the actual implementation does however require extensive use of dynamic programming.

During online monitoring, only a finite trace prefix $tr' \in V^*$ is known from measurements. In order to extend it to infinity, as required due to unbounded future modalities, we may append some concrete interval observations beyond the length of tr' and finally have to add a collecting node, thus arriving at a representation (tr, v) with tr being a suffix of tr'. In order to make the representation (tr, v) as refined as possible, we exploit known facts about the system dynamics to extrapolate the measurements from tr' into the future. To obtain the suffix tr, where the length of the extension is heuristically determined based on the bounds in the bounded future modalities in the formula monitored, we, online, apply ICP to Ratschan's mean-value form constraint pertinent to the ODEs of car dynamics and other continuous components [RS05]. An online computation of the collecting interval, however, is infeasible as it entails an unbounded reachability analysis. To this end, we employ an offline precomputation of reachable states based on a predicate abstraction. Prior to monitoring, we use PHAVer [Fre05] to precompute the state sets reachable from the partitions of the phase space defined by the predicates in the formula. During online monitoring, the collecting node v is generated by intersecting the union of the (precomputed) reachable sets with the last interval valuation in tr.

Refining an interval time series upon appearance of new measurements, finally, is achieved by simply intersecting the interval valuation present in the trace with the new measurement, which may itself be interval valued due to quantization and measurement error. Please note that this together with ICP imposes a filtering on the measurements, as extremal measured values may be inconsistent with previous measurements and system dynamics such that they are pruned off by intersecting the expectation already present in the interval trace with the actual measurement.

References

[AFH96] R. Alur, T. Feder, and T.A. Henzinger. The benefits of relaxing punctuality. *Journal of the ACM*, 43:116–146, 1996.

[BG06] F. Benhamou and L. Granvilliers. Continuous and interval constraints. In F. Rossi, P. van Beek, and T. Walsh, editors, *Handbook of Constraint Programming*, Foundations of Artificial Intelligence, chapter 16, pages 571–603. Elsevier, Amsterdam, 2006.

[FH05] M. Fränzle and M. R. Hansen. A robust interpretation of duration calculus. In Dang Van Hung and Martin Wirsing, editors, *Theoretical Aspects of Computing - ICTAC 2005*, volume 3722 of *LNCS*, pages 257–271. Springer-Verlag, 2005.

[FLS08] M. Faella, A. Legay, and M.I.A. Stoelinga. Model checking quantitative linear time logic. In A. Aldini and C.Baier, editors, *Proceedings of the Sixth Workshop on Quantitative Aspects of Programming Languages (QAPL'08)*, ENTCS, 2008. to appear.

[Fre05] Goran Frehse. Phaver: Algorithmic verification of hybrid systems past hytech. In Morari and Thiele [MT05], pages 258–273.

[MN04] Oded Maler and Dejan Nickovic. Monitoring temporal properties of continuous signals. In Yassine Lakhnech and Sergio Yovine, editors, *Proceedings of the Joint International Conferences on Formal Modelling and Analysis of Timed Systems, (FORMATS'04) and Formal Techniques in Real-Time and Fault-Tolerant Systems (FTRTFT'04)*, volume 3253 of *Lecture Notes in Computer Science*, pages 152–166. Springer-Verlag, September 2004.

[MT05] M. Morari and L. Thiele, editors. *Hybrid Systems: Computation and Control*, volume 3414 of *LNCS*. Springer, 2005.

[NM07] Dejan Nickovic and Oded Maler. AMT: A property-based monitoring tool for analog systems. In Jean-Francois Raskin and P. S. Thiagarajan, editors, *FORMATS*, volume 4763 of *Lecture Notes in Computer Science*, pages 304–319. Springer, 2007.

[Rat00] Stefan Ratschan. Uncertainty propagation in heterogeneous algebras for approximate quantified constraint solving. *Journal of Universal Computer Science*, 6(9), 2000.

[RS05] Stefan Ratschan and Zhikun She. Safety verification of hybrid systems by constraint propagation based abstraction refinement. In Morari and Thiele [MT05], pages 573–589.