

Assumption-Based Runtime Verification of Infinite-State Systems

(RV 2021 regular paper)

Alessandro Cimatti¹

Chun Tian^{1,2}

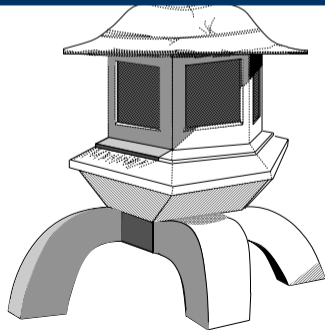
Stefano Tonetta¹

{cimatti, ctian, tonettas}@fbk.eu

¹Fondazione Bruno Kessler, Italy

²University of Trento, Italy

October 12, 2021



Outline

Introduction

Preliminaries

Basic Algorithms


Optimized Algorithms

Experimental Evaluation

Conclusion


Runtime Verification (RV)

- Runtime monitors synthesized from a specification (or property);
- The monitor returns verdicts indicating if the input trace violates the property or not.



Assumption-Based Runtime Verification (ABRV)

- Runtime monitors may be synthesized from a system model (assumptions);
- The ABRV-LTL semantics (as monitor verdicts) is based on LTL_3 semantics, adding one more verdict: \times (*error or out-of-model*);
- Partially observable traces are naturally supported;
- The monitors can be reset, to evaluate the specification at later positions of the trace.



ABRV of Infinite-State Systems

In 2019, the authors have given BDD-based algorithms^a and tools^b for ABRV of *finite-state systems*. The present work is about ABRV of *infinite-state systems*.

^aA. Cimatti, C. Tian, and S. Tonetta. *Assumption-Based Runtime Verification with Partial Observability and Resets*. In *LNCS 11757 - Runtime Verification (RV 2019)*, pages 165–184. Springer, 2019

^bA. Cimatti, C. Tian, and S. Tonetta. *NuRV: A nuXmv Extension for Runtime Verification*. In *LNCS 11757 - Runtime Verification (RV 2019)*, pages 382–392. Springer, 2019

Main ideas

- Runtime Verification (RV) reduced to Model Checking (MC);
- First-Order Quantifier Elimination (QE) for forward image computation;
- Fast (incomplete) Bounded Model Checking (BMC) before full MC algorithms;
- Incremental BMC.

Outline

Introduction

Preliminaries

Basic Algorithms

Optimized Algorithms

Experimental Evaluation

Conclusion

Preliminaries (2)

Fair Transition System (FTS)

$$K \doteq \langle V, \Theta, \rho, \mathcal{J} \rangle$$

where $V = \{x_1, \dots, x_n\}$ is a finite set of variables, Θ the *initial condition*, ρ the *transition relation*, and \mathcal{J} a (finite) set of *justice conditions*. (Θ , ρ and each element of \mathcal{J} are quantifier-free \mathcal{T} -formulas.)

Linear Temporal Logic (LTL) or LTL Modulo Theory

$$\varphi ::= \text{true} \mid \alpha \mid \neg\varphi \mid \varphi \vee \varphi \mid \mathbf{X}\varphi \mid \varphi \mathbf{U}\varphi \mid \mathbf{Y}\varphi \mid \varphi \mathbf{S}\varphi$$

where the (quantifier-free) formula α is built by a set of variables V and a first-order signature Σ , and is interpreted according to a Σ -theory \mathcal{T} .

Preliminaries (3, ABRV)

Set of fair paths

$$\mathcal{L}^K(u) \doteq \{ \sigma \in \mathcal{L}(K) \mid \forall i < |u|. \sigma_i \models_{\mathcal{T}} u_i \}$$


ABRV-LTL semantics

$$\llbracket u, i \models \varphi \rrbracket_4^K \doteq \begin{cases} \times, & \text{if } \mathcal{L}^K(u) = \emptyset \\ \top^a, & \text{if } \mathcal{L}^K(u) \neq \emptyset \text{ and } \forall w \in \mathcal{L}^K(u). w, i \models \varphi \\ \perp^a, & \text{if } \mathcal{L}^K(u) \neq \emptyset \text{ and } \forall w \in \mathcal{L}^K(u). w, i \models \neg\varphi \\ ? & \text{otherwise .} \end{cases}$$

ABRV monitor

$$\mathcal{M}_{\varphi}^K(u) \doteq \llbracket u, 0 \models \varphi \rrbracket_4^K .$$

Preliminaries (4, ABRV illustration)



In this case, $\mathcal{M}_\varphi^K(u_0 u_1 u_2 u_3 u_4) = \top^a$.

Outline

Introduction

Preliminaries

Basic Algorithms

Optimized Algorithms

Experimental Evaluation


Conclusion

ABRV reduced to Model Checking

Let c be an integer variable, $S_u = \langle V_k \cup \{c\}, \Theta, \rho, \emptyset \rangle$, where

$$\Theta = (c = 0) \wedge u_0, \quad \rho = \bigwedge_{i=0}^{|u|-1} ((c = i) \rightarrow (c' = i + 1 \wedge u'_{i+1}))$$

$\mathcal{M}_\varphi^K(u)$ can be computed by two MC calls:



$\llbracket K \times S_u \models \varphi \rrbracket$	$\llbracket K \times S_u \models \neg\varphi \rrbracket$	$\mathcal{M}_\varphi^K(u)$
\top	\top	\times
\top	\perp	\top^a
\perp	\top	\perp^a
\perp	\perp	$?$

Model Checking reduced to ABRV

Computing $\llbracket K \models \varphi \rrbracket$ by ABRV monitoring

Let ϵ be an empty trace, by ABRV definition we have

$$\mathcal{M}_{\varphi}^K(\epsilon) = \begin{cases} \top^a, & \text{if } \llbracket K \models \varphi \rrbracket = \top \text{ (and } \llbracket K \models \neg\varphi \rrbracket = \perp), \\ \perp^a, & \text{if } \llbracket K \models \varphi \rrbracket = \perp \text{ (and } \llbracket K \models \neg\varphi \rrbracket = \top), \\ ?, & \text{if } \llbracket K \models \varphi \rrbracket = \llbracket K \models \neg\varphi \rrbracket = \perp \text{ (counterexamples exist on both sides),} \\ \times, & \text{if } \llbracket K \models \varphi \rrbracket = \llbracket K \models \neg\varphi \rrbracket = \top \text{ (i.e. } \mathcal{L}(K) = \emptyset, \text{ an empty model } K). \end{cases}$$

Thus $\llbracket K \models \varphi \rrbracket = \top$ iff $\mathcal{M}_{\varphi}^K(\epsilon) = \top^a$ or \times (usually the model in a MC problem is not empty).

Conclusion

ABRV with infinite-state assumptions is undecidable (because infinite-state MC is known to be undecidable.)

Quantifier Elimination for Computing Belief States

(For an incremental monitoring algorithm ...)

Definition (forward image)

The *forward image* of a set of states $\psi(V)$ on $\rho(V, V')$ is given by

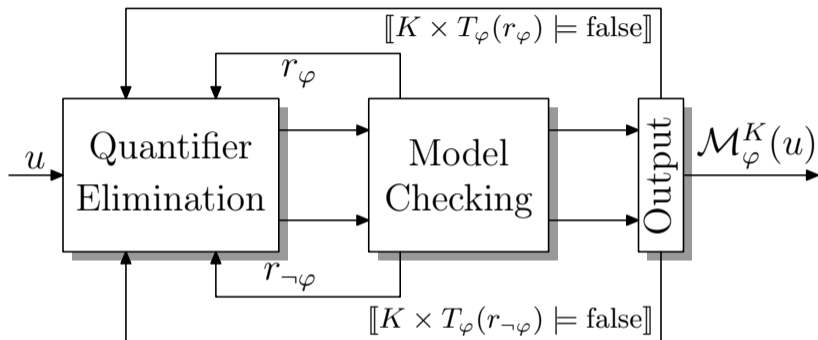
$$\text{fwd}(\psi(V), \rho(V, V'))(V) \doteq (\exists V'. \rho(V, V') \wedge \psi(V'))[V/V']$$

where $[V/V']$ denotes the substitution of (free) variables in V' with the corresponding one in V .

The need of QE procedures

First-order quantifier elimination (QE) can be involved to convert a forward image into an equivalent quantifier-free formula that can be directly sent to SMT solvers, etc.

ABRV reduced to MC and QE (1)



$\neg \llbracket K \times T_\varphi(r_\varphi) \models \text{false} \rrbracket$	$\neg \llbracket K \times T_\varphi(r_{\neg\varphi}) \models \text{false} \rrbracket$	$\mathcal{M}_\varphi^K(\cdot)$
\top	\top	$?$
\top	\perp	\top^a
\perp	\top	\perp^a
\perp	\perp	\times

Outline

Introduction

Preliminaries

Basic Algorithms

Optimized Algorithms

Experimental Evaluation

Conclusion

Basic Optimizations (1)

Basic optimization ideas

- o1 If the monitor has already reached conclusive verdicts (\top^a or \perp^a), then for the runtime verification of the next input state *at most one* MC call is need.
- o2 Before calling model checkers to detect the emptiness of a belief state (w.r.t. fairness), an SMT checking can be done first, to check if the belief state formula can be satisfied or not.
- o3 When `monitor2` is used as online monitor, the same LTL properties are sent to LTL model checkers with different models and are internally translated into equivalent FTS.
- o4 Call the faster but incomplete plain BMC (or any other MC procedure which only detects counterexamples) before calling a unbounded model checker such as IC3IA.

Theorem

Assuming BMC always find the counterexample whenever it exists, IC3_IA is called at most twice in the “online” version of monitor2 with all above optimizations.

Incremental BMC: The Idea

Basic observation

All models used for model checking in (non)emptiness checking $\llbracket K \times T_\varphi(r_\varphi) \models \text{false} \rrbracket$ only differ at the initial condition.

BMC encoding of belief states

The belief states after a sequence of observations $u_0 u_1 \cdots u_n$, denoted by $\text{bs}(u_0 u_1 \cdots u_n)$, can be inductively given by

$$\begin{aligned}\text{bs}(u_0)(V) &= I(V) \wedge u_0(V), \\ \text{bs}(u_0 u_1 \cdots u_{i+1})(V) &= \text{fwd}(\text{bs}(u_0 u_1 \cdots u_i)(V), T(V, V'))(V) \wedge u_{i+1}(V)\end{aligned}$$

Theorem (Equisatisfiability)

When $k > 1$, the SMT formulas $I(V_0) \wedge u_0(V_0) \wedge \bigwedge_{j=0}^{k-1} [T(V_j, V_{j+1}) \wedge u_{j+1}(V_{j+1})]$ and $\text{bs}(u_0 u_1 \cdots u_k)(V)$ are equi-satisfiable.

Incremental BMC: The algorithm (1)

New procedures

- `init_nonemptiness` for creating a persistent SMT solver instance,
- `update_nonemptiness` for checking nonemptiness of belief states after new observation,
- `reset_nonemptiness` for resetting the SMT solver, cleaning up all existing observations.

```
function init_nonemptiness(I, T)
```

```
  e := new BMC solver with initial formula I and transition relation T
```

```
  reset_nonemptiness(e, I)
```

```
  return e
```

```
procedure reset_nonemptiness(e, I)
```

```
  e.problem := I(V0); // the initial formula unrolled at time 0
```

```
  e.observations := []; // an array holding observations
```

```
  e.n := 0; // the number of observations
```

```
  e.map := {}; // a hash map from time to (unused) observations
```

```
  e.k := 0; // the number of unrolled transition relations
```

```
  e.max_k := max_k; // a local copy of max_k
```

Incremental BMC: The algorithm (2)

```
function bmc_monitor( $K \doteq \langle V_K, \Theta_K, \rho_K, \mathcal{J}_K \rangle$ ,  $\varphi$ ,  $u$ ,  $max\_k$ ,  $window\_size$ )
```

```
   $T_\varphi \doteq \langle V_\varphi, \Theta_\varphi, \rho_\varphi, \mathcal{J}_\varphi \rangle := \text{ltl\_translation}(\varphi)$ 
```

```
   $T_{\neg\varphi} \doteq \langle V_\varphi, \Theta_{\neg\varphi}, \rho_\varphi, \mathcal{J}_\varphi \rangle := \text{ltl\_translation}(\neg\varphi)$ 
```

```
   $V := V_K \cup V_\varphi$ 
```

```
   $e_1 := \text{init\_nonemptiness}(\Theta_K \wedge \Theta_\varphi, \rho_K \wedge \rho_\varphi)$ 
```

```
   $e_2 := \text{init\_nonemptiness}(\Theta_K \wedge \Theta_{\neg\varphi}, \rho_K \wedge \rho_\varphi)$ 
```

```
  for  $0 < i < |u|$  do
```

```
     $b_1 := \text{update\_nonemptiness}(e_1, u_i)$ 
```

```
     $b_2 := \text{update\_nonemptiness}(e_2, u_i)$ 
```

```
  if  $b_1 \wedge b_2$  then return ? ;
```

```
  else if  $b_1$  then return  $\top^a$ ;
```

```
  else if  $b_2$  then return  $\perp^a$ ;
```

```
  else return  $\times$ ;
```

```
    // inconclusive
```

```
    // conditionally true
```

```
    // conditionally false
```

```
    // out of model
```

Incremental BMC: The algorithm (3)

```
function update_nonemptiness(e, o)
  e.map[e.n] = o,      e.observations[e.n++] = o;           // store new observation
  for (k, v) : e.map do
    if k ≤ e.k then e.problem := e.problem ∧ v(Vi)
    delete e.map[k] ;
  result := ?
  while e.k ≤ e.max_k and result = ? do
    i := e.k
    if SMT(e.problem) = UNSAT then result := ⊥, break;
    if SMT(e.problem ∧ [[F]]i) = SAT then result = ⊤, break;
    e.problem := e.problem ∧ e.T(Vi, Vi+1)
    if e.map[i + 1] exists then
      e.problem := e.problem ∧ e.map[i + 1](Vi+1),      delete e.map[i + 1]
    e.k++
  e.max_k++;           // increase the search bound for next calls
  if e.k > window_size or result = ? then
    r := compute_belief_states(e),      reset_nonemptiness(e, r)
  if result = ⊤ or result = ⊥ then
    return result
  else
    return ¬IC3_IA(⟨V, r, e.T, JK ∪ Jφ⟩, false)
```

Incremental BMC: The algorithm (4)

```
function compute_belief_states(e)  
  r := e.l(V)  
  for i ← 0 to e.n do  
    if i = 0 then r := r ∧ e.observations[i](V);  
    else  
      └ r := quantifier_elimination(V, r ∧ T(V, V')) ∧ e.observations[i](V)  
  return r
```

Tool Implementation

The RV algorithms presented in this paper have been implemented in NuRV since version 1.6.0 (<https://es.fbk.eu/tools/nurv/>). (Currently we support 5 operating systems, 5 target languages for monitor code generation, and network-based monitoring.)

Outline

Introduction

Preliminaries

Basic Algorithms

Optimized Algorithms

Experimental Evaluation

Conclusion

Performance Tests on Dwyer's LTL patterns (1)

Pattern 49

$\varphi = \mathbf{G}(q \rightarrow (p \rightarrow (\neg r \mathbf{U}(s \wedge \neg r \wedge \mathbf{X}(\neg r \mathbf{U} t)))) \mathbf{U}(r \vee \mathbf{G}(p \rightarrow (s \wedge \mathbf{X} \mathbf{F} t))))$
("s, t responds to p after q until r"),
with $q := (0 \leq i)$, $r := (0.0 \leq x)$, $i \in [-500, 500]$ and $x \in [-0.500, 0.500]$.

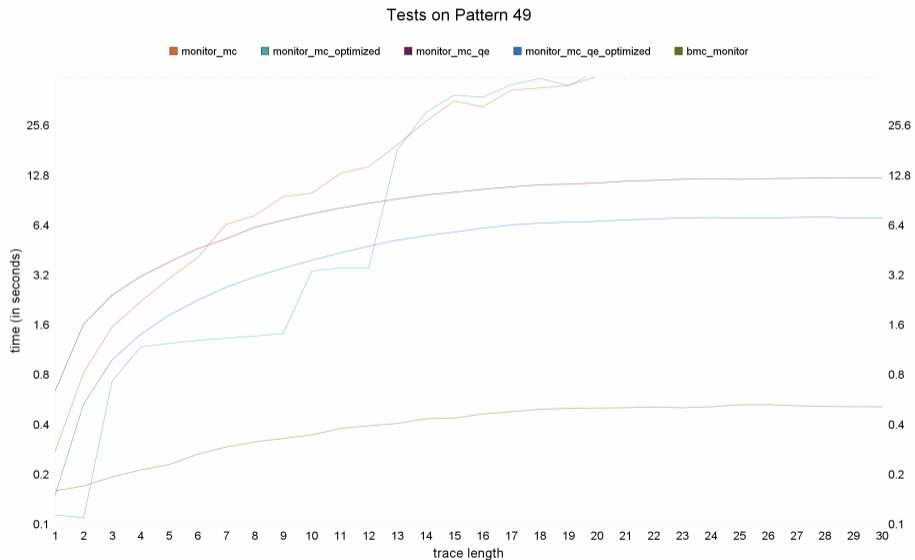
RV assumptions

"The p -transition (i.e., from $\neg p$ to p) happens at most 4 times"

Other settings

The length of input traces increases from 1 to 30.

Performance Tests on Dwyer's LTL patterns (2)



Outline

Introduction

Preliminaries

Basic Algorithms

Optimized Algorithms

Experimental Evaluation

Conclusion

Conclusions

Conclusions

- The ABRV framework has been extended in this paper to assumptions defined as infinite-state system, using existing SMT, MC, QE techniques.
- We start from a trivial reduction from RV to MC, and eventually obtained an highly optimized RV algorithm, based on Incremental BMC. (The final version is hundreds of times faster than the initial one.)

Possible concerns of the present approach

- The use of (slow) SMT solvers in performing (fast) runtime monitoring; (there is a trade-off between the required speed of the monitor and the complexity of the assumptions)
- The boundedness of memory consumptions during runtime monitoring. (unbound in the worst-case but better in practical)

Some Future Directions

- Better performances (including engineering level changes, e.g. 2x faster after paper submission);
- More thoroughly tests on better RV benchmarks;
- More compact presentations of raw formulas (DBM, PPlite, ...);
- Attacking real-time temporal properties with timed assumptions;
- Code generation for infinite-state monitors (e.g. calling external SMT solvers by standard interface like DIMACS).

Preliminaries (1)

Satisfiability Modulo Theory (SMT)

First-order formulas are built as usual by proposition logic connectives, a given set of variables V and a first-order signature Σ , and are interpreted according to a given Σ -theory \mathcal{T} .

SMT is the basis of infinite-state model checking algorithms used in this paper. Note: *Tool implementation and experimental evaluations are based on \mathcal{LRA} (linear arithmetic of reals).*

First-Order Quantifier Elimination (QE)

QE methods convert first-order formulas into \mathcal{T} -equivalent quantifier-free formulas.

Formally speaking, if $\alpha(V_1 \cup V_2)$ is quantifier-free formula (of the theory \mathcal{T}) built by variables from the set $V_1 \cup V_2$, the role of quantifier elimination is to convert the first-order formula $\exists V_1. \alpha(V_1 \cup V_2)$ into an \mathcal{T} -equivalent formula $\beta(V_2)$, where β is quantifier-free and is built by only variables from V_2 .

MathSAT supports two QE procedures: Fourier-Motzkin and Loos-and-Weispfenning (\mathcal{LRA}).

ABRV reduced to Model Checking (2)

function monitor1($K \doteq \langle V_K, \Theta_K, \rho_K, \mathcal{J}_K \rangle, \varphi, u$)

$\Theta := \top, \rho := \top$

if $|u| > 0$ **then**

$\Theta := (c = 0) \wedge u_0$

if $|u| > 1$ **then**

$\rho := \bigwedge_{i=0}^{|u|-1} ((c = i) \rightarrow (c' = i + 1 \wedge u'_{i+1}))$

$S_u := \langle V_K \cup \{c\}, \Theta, \rho, \emptyset \rangle$

$b_1 := \text{model_checking}(K \times S_u, \varphi)$

$b_2 := \text{model_checking}(K \times S_u, \neg\varphi)$

if $b_1 \wedge b_2$ **then return** \times ;

else if b_1 **then return** \top^a ;

else if b_2 **then return** \perp^a ;

else return ?;

// out of model

// conditionally true

// conditionally false

// inconclusive

ABRV reduced to MC and QE (2)

```
function monitor2( $K \doteq \langle V_K, \Theta_K, \rho_K, \mathcal{J}_K \rangle, \varphi, u$ )
   $T_\varphi \doteq \langle V_\varphi, \Theta_\varphi, \rho_\varphi, \mathcal{J}_\varphi \rangle := \text{ltl\_translation}(\varphi)$ 
   $T_{\neg\varphi} \doteq \langle V_\varphi, \Theta_{\neg\varphi}, \rho_\varphi, \mathcal{J}_\varphi \rangle := \text{ltl\_translation}(\neg\varphi)$ 
   $V := V_K \cup V_\varphi$ 
   $\langle r_\varphi, r_{\neg\varphi} \rangle := \langle \Theta_K \wedge \Theta_\varphi, \Theta_K \wedge \Theta_{\neg\varphi} \rangle$ 
  if  $|u| > 0$  then
     $\perp \langle r_\varphi, r_{\neg\varphi} \rangle := \langle r_\varphi \wedge u_0, r_{\neg\varphi} \wedge u_0 \rangle$ 
  for  $1 \leq i < |u|$  do
     $r_\varphi := \text{quantifier\_elimination}(V, \rho_K \wedge \rho_\varphi \wedge r_\varphi) \wedge u_i$ 
     $r_{\neg\varphi} := \text{quantifier\_elimination}(V, \rho_K \wedge \rho_\varphi \wedge r_{\neg\varphi}) \wedge u_i$ 
   $b_1 := \neg\text{model\_checking}(\langle V, r_\varphi, \rho_K \wedge \rho_\varphi, \mathcal{J}_K \cup \mathcal{J}_\varphi \rangle, \text{false})$ 
   $b_2 := \neg\text{model\_checking}(\langle V, r_{\neg\varphi}, \rho_K \wedge \rho_\varphi, \mathcal{J}_K \cup \mathcal{J}_\varphi \rangle, \text{false})$ 
  if  $b_1 \wedge b_2$  then return ?; // inconclusive
  else if  $b_1$  then return  $\top^a$ ; // conditionally true
  else if  $b_2$  then return  $\perp^a$ ; // conditionally false
  else return  $\times$ ; // out of model
```

Basic Optimizations (2)

```
if  $o_3$  then  $F := \text{ltl\_translation}(\bigwedge_{\psi \in \mathcal{J}_K \cup \mathcal{J}_\varphi} \mathbf{GF} \psi) \rightarrow \text{false}$  ;  
function check_nonemptiness( $r$ )  
  if  $o_2 \wedge (\text{SMT}(r) = \text{unsat})$  then return  $\perp$  ;  
  else  
    return  $\neg \text{model\_checking}(\langle V, r, \rho_K \wedge \rho_\varphi, \mathcal{J}_K \cup \mathcal{J}_\varphi \rangle, o_3 ? F : \text{false})$   
function model_checking( $M, \psi$ )  
  if  $o_4$  then  
    if  $\text{BMC}(M, \psi) = \perp$  then return  $\perp$ ; // counterexample found  
    else // max_k reached  
      return  $\text{IC3\_IA}(M, \psi)$   
  else return  $\text{IC3\_IA}(M, \psi)$ ;
```