

A BDD-based BMC Approach for the Verification of Multi-Agent Systems

Andrew V. JONES Alessio LOMUSCIO
{andrewj,alessio}@doc.ic.ac.uk

Department of Computing
Imperial College London

Concurrency, Specification, and Programming
September, 2009

Model Checking vs. Bounded Model Checking

- The majority of “symbolic” model checkers use *binary decision diagrams*
- Current bounded model checkers, either:
 - Require a translation of the problem to SAT, or
 - Aren’t very expressive
- And – they concentrate on “bug hunting”, not verifying correctness

SAT-based BMC and conventional model checking techniques are, at the moment, seen as complementary to each other.

Is it possible to convert an existing model checker to a bounded one?

And ...

How can we distribute the verification process?

What is an agent?

“An **agent** is a computer system that is **situated** in some **environment**, and that is capable of **autonomous action** in this environment in order to meet its design objectives.” – Weiss

Temporal Epistemic Interpreted Systems

- A system is composed of a set of agents $A = \{1, \dots, n\}$ and an environment e .
- Each agent, i , is described by
 - A set of local *states* – \mathcal{L}_i
 - A set of local *actions* – Act_i
 - A local *protocol* function – $\mathcal{P}_i : \mathcal{L}_i \rightarrow 2^{\text{Act}_i}$
 - An *evolution* function – $\tau_i : \mathcal{L}_i \times \text{Act} \rightarrow \mathcal{L}_i$
- Act is the set of *joint actions* –
 $\text{Act} \subseteq \text{Act}_1 \times \dots \times \text{Act}_n \times \text{Act}_e$

Temporal Epistemic Interpreted Systems

- G is the set of *possible global states* – $G \subseteq \mathcal{L}_1 \times \dots \times \mathcal{L}_n \times \mathcal{L}_e$
 - A *global state* $g = (l_1, \dots, l_n, l_e) \in G$ represents a “snapshot” of the system
 - $l_i : G \rightarrow \mathcal{L}_i$ is a *projection* of agent i 's local state from the given global state
 - Each agent i has an *epistemic relation* – $g \sim_i g'$ iff $l_i(g) = l_i(g')$
- T is a *transition relation* for the system – $T \subseteq G \times \text{Act} \times G$
 - $g T g'$ iff there exists actions a_1, \dots, a_n such that for all i , $a_i \in \mathcal{P}_i(l_i(g))$ and $\tau_i(l_i(g), a_1, \dots, a_n) = l_i(g')$

Temporal Epistemic Models

A *model* \mathcal{M}_{IS} is a tuple $(G, \iota, T, \sim_1, \dots, \sim_n, \mathcal{V})$

- $\iota \in G$ is an *initial* global state
- G is the set of *reachable* states accessible from ι via T
- \mathcal{V} is a mapping from global states to propositional variables
(\mathcal{PV}) – $\mathcal{V} : G \rightarrow 2^{\mathcal{PV}}$

A *path* $\pi = (\iota, g_1, \dots)$ is an *infinite* sequence of global states such that $\forall_{k \geq 0} (g_k, g_{k+1}) \in T$.

- $\pi(k)$ is the k^{th} global state of the path π
- $\Pi(g)$ is the set of all paths starting at $g \in G$

- Syntax

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid EX\varphi \mid EG\varphi \mid E[\varphi U\psi] \mid K_i\varphi$$

- Derived Modalities

- $EF\varphi \stackrel{\text{def}}{=} E[\text{true}U\varphi]$
- $AX\varphi \stackrel{\text{def}}{=} \neg EX\neg\varphi$
- $AF\varphi \stackrel{\text{def}}{=} \neg EG\neg\varphi$
- $AG\varphi \stackrel{\text{def}}{=} \neg EF\neg\varphi$
- $\bar{K}_i\varphi \stackrel{\text{def}}{=} \neg K_i\neg\varphi$
- $A[\varphi U\psi]$ – Takes the expected meaning

Epistemic Modalities – K and \overline{K}

$K_i\varphi$ – “agent i knows that φ ”

$\mathcal{M}_{IS}, g \models K_i\varphi$ iff $\forall g' \in G, g \sim_i g'$ implies $\mathcal{M}_{IS}, g' \models \varphi$

$\overline{K}_i\varphi$ – “agent i considers it possible that φ ”

$\mathcal{M}_{IS}, g \models \overline{K}_i\varphi$ iff $\exists g' \in G : g \sim_i g'$ and $\mathcal{M}_{IS}, g' \models \varphi$

Conventional “symbolic” model checkers use a canonical representation called ROBDDs. These can be used to efficiently represent a Boolean function.

Two main stages:

- ① Calculate the entire reachable state space
- ② Recursively evaluate the property using fix point methods

Interpreted Systems Reachable State Space

$$\text{lfp}(Q) = (I(g) \vee \exists g' (T(g, a, g') \wedge Q(g')))$$

MCMAS:

- Symbolic model checker for verifying *certain* aspects of multi-agent systems
- Based on BDDs from the CUDD library
- Based on the “Interpreted Systems Programming Language”

Interpreted Systems Programming Language – ISPL – allows for the definition of multi-agent systems following the interpreted systems formalism. ISPL syntax includes:

- Agent
- Evaluation
- Formulae

The Existential and Universal Fragments

ECTLK

CTLK as before:

- Restricts negation to *only* appear in front of elements of \mathcal{PV}
- Contains \overline{K}_i *not* K_i

ACTLK

A formula is in ACTLK if the negation is ECTLK. . .

$$\text{ACTLK} \subseteq \{\varphi \mid \neg\varphi \in \text{ECTLK}\}$$

Current Approaches to Bounded Model Checking

SAT-based BMC

Idea

- “Unroll” the transition relation k times
- Translate the **negated** property and unrolled model to the Boolean satisfiability problem

Problem

- Requires a SAT solver
- Not straightforward to convert a BDD model checker to a SAT one

BDD-based BMC

Idea

- Represent the “error” as a bad state
- Find all of the reachable states at a depth k
- Check the intersection of the error state with the “fringe”

Problem

How does one specify properties as a *single* error state? Okay for *invariant/safety* properties – but apart from that?

BDD-BMC(ψ : ACTLK FORMULA, \mathcal{I} : INITIAL STATE, Trans :
TRANSITION RELATION) : BOOLEAN

```
1:  $\varphi \leftarrow \neg\psi$  { $\varphi$  : ECLTK FORMULA}
2: Reach  $\leftarrow \mathcal{I}$  {Reach : BDD}
3: while TRUE do
4:   if  $\llbracket \iota \rightarrow \varphi \rrbracket = \text{Reach}$  then
5:     return FALSE {Counterexample found}
6:   end if
7:   Reach  $\leftarrow \text{Reach} \vee (\text{Reach} \wedge \text{Trans})$ 
8:   if Reach Unchanged then
9:     break {Fixed point reached}
10:  end if
11: end while
12: return  $\llbracket \iota \rightarrow \psi \rrbracket = \text{Reach}$ 
```

Current fix point methods work even when using **non-serial** transition relations, but we need a symbolic method for $\overline{K}_i \dots$

We could just use $\overline{K}_i \varphi \stackrel{\text{def}}{=} \neg K_i \neg \varphi \dots$

$\text{SAT}_{\overline{K}}(\varphi : \text{FORMULA}, i : \text{AGENT}) : \textit{set of STATE}$

- 1: $X \leftarrow \text{SAT}_{\text{CTL}\overline{K}}(\varphi)$
- 2: $Y \leftarrow \text{pre}_{\overline{K}}(X, i)$
- 3: **return** Y

```
BDD basic_agent::project_local_state(BDD *state, BDDvector* v)
{
    BDD tmp = bddmgr->bddOne();

    // For all of the state variables before the agent ...
    for (int j = 0; j < get_var_index_start(); j++)
    {
        // ‘and’ them on
        tmp = tmp * (*v)[j];
    }

    // and after the agent ...
    for (int j = get_var_index_end() + 1; j < v->count(); j++)
    {
        // ‘and’ them on
        tmp = tmp * (*v)[j];
    }

    return state->ExistAbstract(tmp);
}
```


Figure: The *simplified* project_local_state method

We restrict verification to invariant properties, i.e. “AG” is the top most connective.


Distributed Approach

- ① *Fixed-Depth BDD-based BMC.*
- ② *Seed State Generation.*
- ③ *Distributed Parallel BDD-based BMC.*

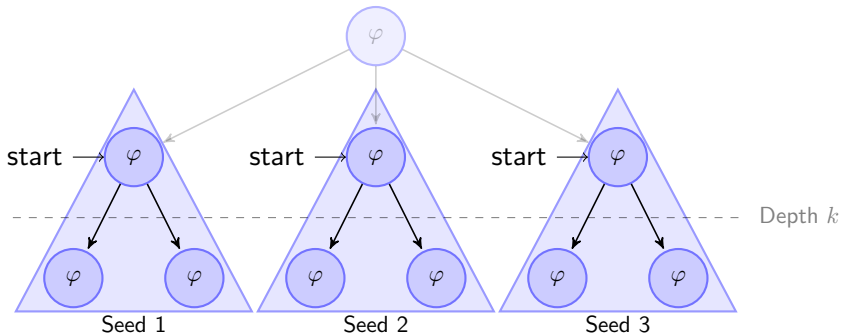
Distributed Verification of Invariant Properties



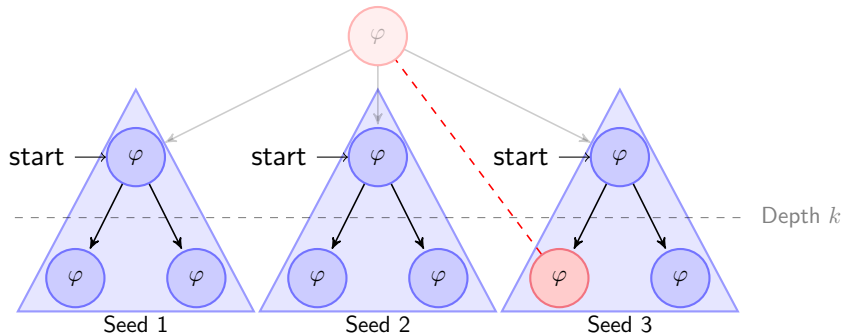
Distributed Verification of Invariant Properties



Distributed Verification of Invariant Properties

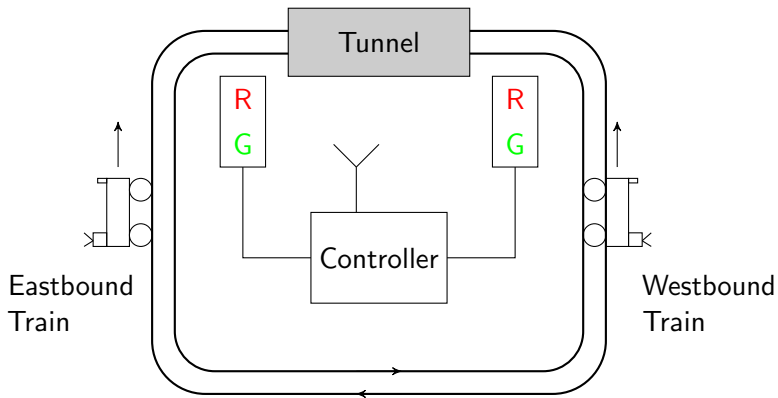


Distributed Verification of Invariant Properties




A Scalable Model

The Train Gate Controller Model




The Faulty Train Gate Controller – Train

Faulty trains now have a **service counter** and a **breaking depth**...



Label	Action
t_1, t_3	SERVICE
t_2	BACK
t_4	SIGNAL
t_5	ENTER
t_6, t_7	BREAK
t_8	LEAVE

The Faulty Train Gate Controller – Controller



Label	Action
c_1	IDLE
c_2	EXIT_TRAIN
c_3	IDLE
c_4	IDLE
c_5	ENTER_TRAIN_T $T \in \{E, W\}$

Example Specifications

φ_{TGC1} : “Trains can’t stay in the tunnel forever”

$AG (AF (\neg \text{TRAINE_IN_TUNNEL}))$

Mutual Exclusion:

φ_{TGC2} : “Two trains never occupy the tunnel at the same time”

$AG (\neg \text{TRAINE_IN_TUNNEL} \vee \neg \text{TRAINW_IN_TUNNEL})$

φ_{TGC3} :

“When a train is in the tunnel it knows that another train is not”

$AG (\text{TRAINE_IN_TUNNEL} \rightarrow K_{\text{TRAINE}} (\neg \text{TRAINW_IN_TUNNEL}))$

Bounded Model Checking Memory Usage

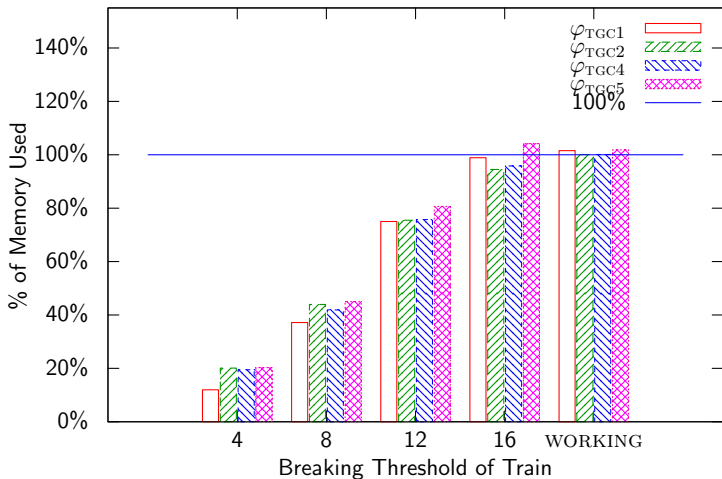


Figure: Memory usage for two trains, with a full service depth of 20 – when checking various formulae

Bounded Model Checking Time Usage




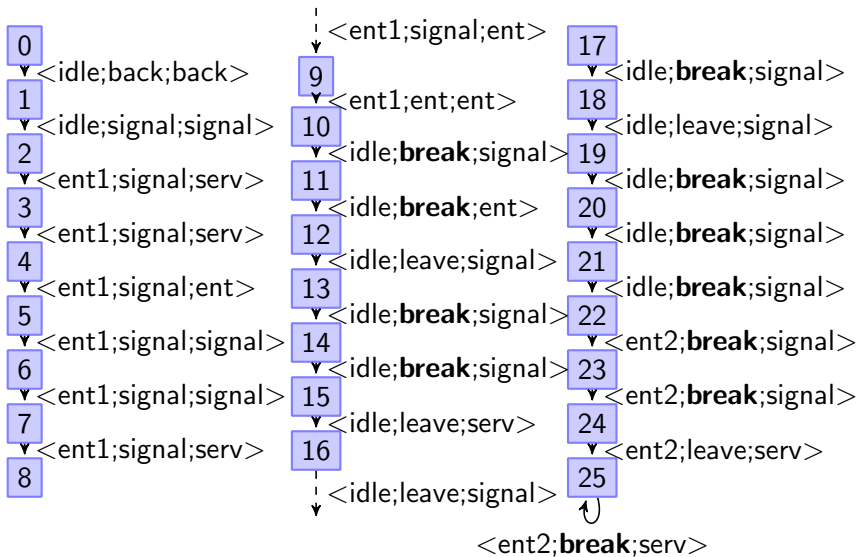
Figure: Time required for two trains, with a full service depth of 20 – when checking various formulae

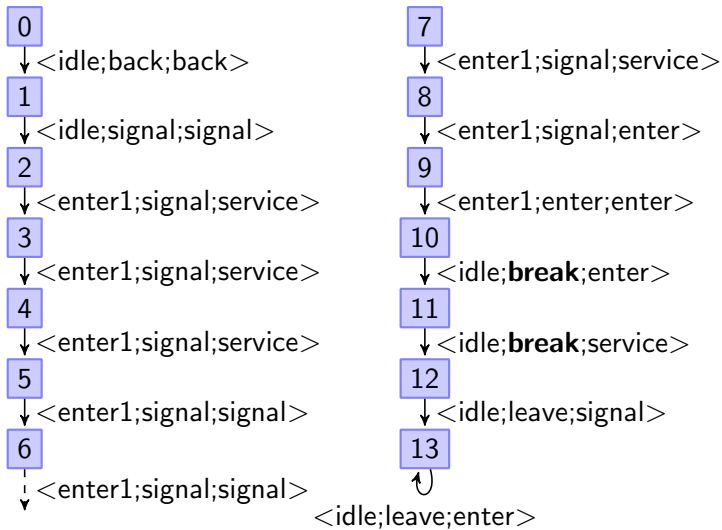
Counterexamples

- Biere *et al.* (1999): bounded model checking “finds counterexamples of minimal length”.

Method	Formula				
	φ_{TGC1}	φ_{TGC2}	φ_{TGC3}	φ_{TGC4}	φ_{TGC5}
Regular	25	17	4	4	12
BMC	13	16	4	4	FAIL

Table: Length of counterexamples generated between BMC and full verification





A Temporal Formula

Given we don't have completeness for epistemic sub-formulae, we need a different formula to verify ...

Temporal-Epistemic

$$AG (\text{TRAINE_IN_TUNNEL} \rightarrow K_{\text{TRAINE}} (\neg \text{TRAINW_IN_TUNNEL}))$$

Temporal-Only

$$AG (\text{TRAINE_IN_TUNNEL} \rightarrow AX (\neg \text{TRAINW_IN_TUNNEL}))$$

Evaluating Distributed Bounded Model Checking

Table: A comparison of seeded BMC vs. BMC for a single master and 3 slaves (seed depth of 4).

Model	Ratio		
	Memory	Time	States
FAULTY	1.8255	3.8130	1.7297
WORKING	0.9500	0.0013	0.0008

Table: Ratios comparing time for seeded BMC vs. BMC, for a varying number of slaves (seed depth of 3).

# Hosts	Ratio
2	0.0033
4	0.0066
6	0.0098
8	0.0131

Conclusion

BDD-based BMC is shown to be effective with variable reordering **disabled**; for satisfiable formulae the overhead imposed is negligible.

Distributed BDD-based BMC:

- Out-performs the sequential approach for falsifiable formulae
- Increasing the number of hosts increases verification efficiency

Further work

- Completeness of temporal-epistemic formulae
- Intersection-based BMC using MCMAS's RedStates
- More Models
- Comparison to VERICS
- Smarter use of CUDD using `Cudd_RecursiveDeref`

Are there any questions?

Backup Slides

SAT requires an encoding of a “back loop”

Semantics of ECTLK


$\mathcal{M}_{IS}, g \models EG\varphi$ iff $\exists \pi \in \Pi(g) \forall_{m \geq 0} \mathcal{M}_{IS}, \pi(m) \models \varphi$

Bounded Semantics of ECLTK

$\mathcal{M}_k, g \models EG\varphi$ iff $\exists \pi \in P_k(\pi(0) = g$ and $\forall_{0 \leq j \leq k} \mathcal{M}_k, \pi(j) \models \varphi$)
... requires $loop(\pi) \neq \emptyset$



$$loop(\pi) = \emptyset$$



$$loop(\pi) = \{2\}$$

BDD approaches require an explicit “error state”

BOUNDEDTRAVERSAL(Trans : Transition Relation, \mathcal{I} : Initial States, Err : Error State, k : Depth)

```
1: Frontier0 ←  $\mathcal{I}$ 
2: for ( $i = 0$ ;  $i < k$ ;  $i++$ ) do
3:   if (Frontier $i$  · Err  $\neq \emptyset$ ) then
4:     return (FAILURE)
5:   end if
6:   Frontier $i+1$  ← IMG(Trans, Frontier $i$ )
7: end for
8: return (PASS)
```

CUDD ExistAbstract

- Quantification – $\mathcal{B}_h = \exists x \mathcal{B}_f$
- Shannon's expression – $h = (\neg x \wedge f|_{x \leftarrow 0}) \vee (x \wedge f|_{x \leftarrow 1})$

$\varphi_{\text{TGC4}}(N)$

$$AG \left(\text{TRAIN}_i\text{-IN_TUNNEL} \rightarrow K_{\text{TRAIN}_i} \left(\bigwedge_{j=1}^{i-1} \left(\neg \text{TRAIN}_j\text{-IN_TUNNEL} \right) \wedge \bigwedge_{j=i+1}^N \left(\neg \text{TRAIN}_j\text{-IN_TUNNEL} \right) \right) \right)$$

$\varphi_{\text{TGC5}}(N)$

$$AG \left(\text{TRAIN}_i\text{-IN_TUNNEL} \rightarrow K_{\text{TRAIN}_i} \left(\bigwedge_{j=1}^{i-1} AX \left(\neg \text{TRAIN}_j\text{-IN_TUNNEL} \right) \wedge \bigwedge_{j=i+1}^N AX \left(\neg \text{TRAIN}_j\text{-IN_TUNNEL} \right) \right) \right)$$