

Model Checking

Aleksei Sholokhov, Yuan-Mao Chueh

Agenda

Formal reasoning:

1. Deductive Reasoning (like writing math proof)
2. Model checking (finite; run automatically)

Model Checking:

1. Language used: Temporal logic: ex: CTL* (which contains CTL, LTL) [Alex]
2. Explicit-state \rightarrow state explosion problem [Alex]
 - a. Partial Order Reduction (X)
 - b. BDD-based symbolic model checking (Yuan-Mao)
 - c. SAT/SMT based model checking (X)
 - d. Abstraction (\leftarrow our required reading) (Yuan-Mao)
3. Applications (\leftarrow our optional readings) [Alex] 2-3 pages

Model-Checking Overview

1. Property Specification Language
 - Typically expressed based on a temporal logic
2. Model Specification Language
 - Encoding the system (program, hardware) as a finite-state transition system
3. Verification Procedure
 - Algorithms that does an exhaustive search of the model state space
 - Provides a counterexample if it finds a state that breaks the specification

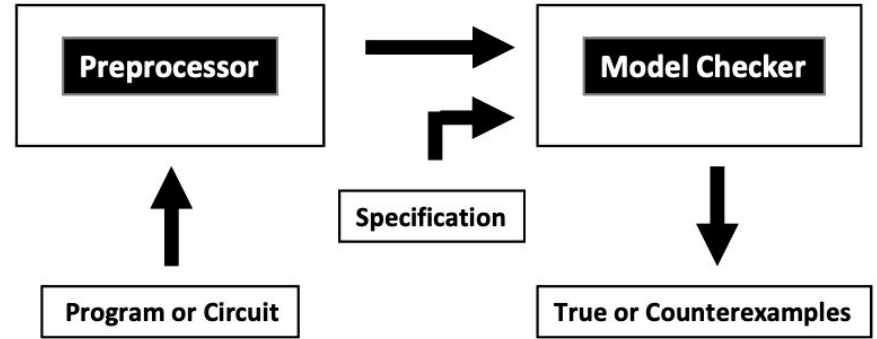


Fig. 4: A Model Checker with Counterexamples

LTL – Linear-time Temporal Logic

Definition 3 (LTL). *Linear temporal logic formulas are of the form $\mathbf{A}\psi$, with ψ given by the grammar:*

$$\psi ::= p \mid \neg\psi \mid \psi \vee \psi \mid \mathbf{X}\psi \mid \psi \mathbf{U} \psi$$

where $p \in AP$.

- Temporal Operators: $\mathbf{X}\varphi \mid \mathbf{G}\varphi \mid \mathbf{F}\varphi \mid \mathbf{R}\varphi \mid \mathbf{W}\varphi \mid \mathbf{M}\varphi$
 - $\mathbf{X}\varphi$ – Next: φ has to hold at the next state (this operator is sometimes noted \mathbf{N} instead of \mathbf{X}).
 - $\mathbf{G}\varphi$ – Globally: φ has to hold on the entire subsequent path.
 - $\mathbf{F}\varphi$ – Finally: φ eventually has to hold (somewhere on the subsequent path).
 - $\psi \mathbf{W} \varphi$ – Weak until: ψ has to hold *at least* until φ ; if φ never becomes true, ψ must remain true forever.
 - $\psi \mathbf{U} \varphi$ – Until: ψ has to hold *at least* until φ becomes true, which must hold at the current or a future position.
 - $\psi \mathbf{R} \varphi$ - Release: φ has to be true until and including the point where ψ first becomes true; if ψ never becomes true, φ must remain true forever
 - $\psi \mathbf{M} \varphi$ - Strong release: φ has to be true until and including the point where ψ first becomes true, which must hold at the current or a future position
- All formulas have an implicit \mathbf{A} in front

Computation Tree Logic (CTL) – branching-time logic

Definition 4 (CTL). *Computation tree logic formulas are inductively defined as follows:*

$$\begin{aligned}\phi &::= p \mid \neg\phi \mid \phi \vee \phi \mid \mathbf{A}\psi \mid \mathbf{E}\psi \quad (\text{state formulas}) \\ \psi &::= \mathbf{X}\phi \mid \mathbf{F}\phi \mid \mathbf{G}\phi \mid \phi \mathbf{U}\phi \quad (\text{path formulas})\end{aligned}$$

where $p \in AP$.

- Each basic temporal (X, F, G, U) operator must be immediately preceded by a path quantifier (A or E)
- Quantifiers over paths
 - **A** Φ – **All**: Φ has to hold on all paths starting from the current state.
 - **E** Φ – **Exists**: there exists at least one path starting from the current state where Φ holds.
- Path-specific quantifiers
 - **X** ϕ – **Next**: ϕ has to hold at the next state (this operator is sometimes noted **N** instead of **X**).
 - **G** ϕ – **Globally**: ϕ has to hold on the entire subsequent path.
 - **F** ϕ – **Finally**: ϕ eventually has to hold (somewhere on the subsequent path).
 - ϕ **U** ψ – **Until**: ϕ has to hold *at least* until at some position ψ holds. This implies that ψ will be verified in the future.
 - ϕ **W** ψ – **Weak until**: ϕ has to hold until ψ holds. The **W** operator is sometimes called "unless".

CTL* – combines state- and path-specific qualifiers

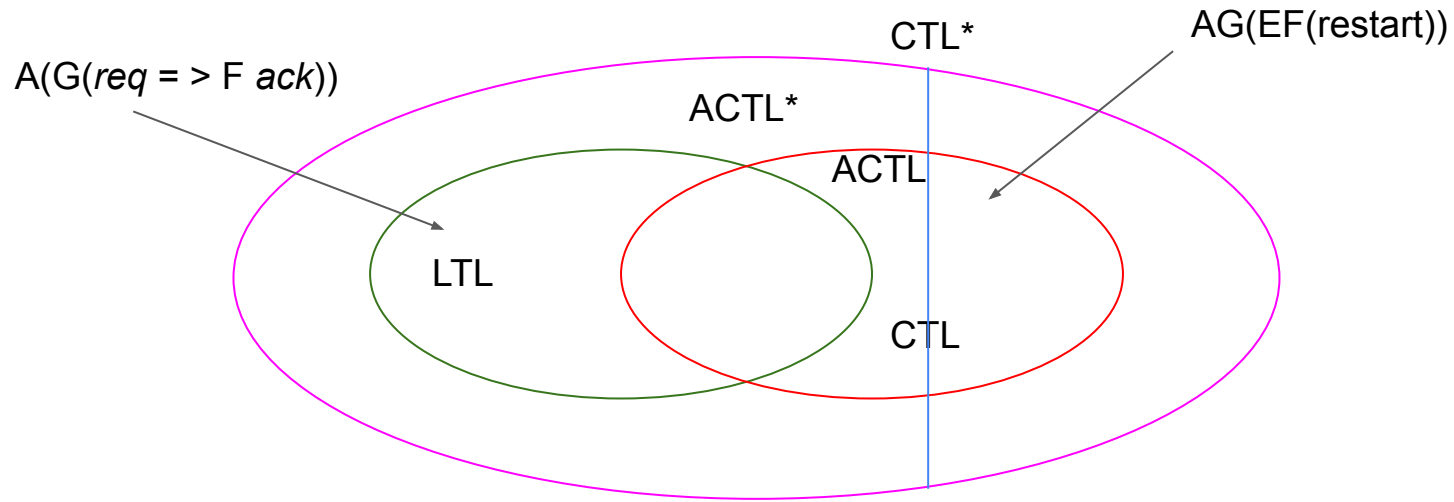
Definition 2 (CTL*). *The syntax of CTL* is given by the grammar:*

$$\begin{aligned}\phi &::= p \mid \neg\phi \mid \phi \vee \phi \mid \mathbf{A}\psi \mid \mathbf{E}\psi && \text{(state formulas)} \\ \psi &::= \phi \mid \neg\psi \mid \psi \vee \psi \mid \mathbf{X}\psi \mid \mathbf{F}\psi \mid \mathbf{G}\psi \mid \psi \mathbf{U} \psi && \text{(path formulas)}\end{aligned}$$

where $p \in AP$.

- ϕ is satisfied with respect to the state: $s \models \phi$
- Ψ is satisfied with respect to the path: $\pi \models \psi$
- ACTL* – CTL* where the A (forall) qualifier is excluded and all formulas are in NNF.
 - Because of the latter, we can not define $E\phi = \neg A\neg\phi$. Thus $ACTL^* \subset CTL^*$

LTL vs CTL vs CTL* vs ACTL*




State Explosion Problem

- Each n -bit number has 2^n states
- k branch conditions give 2^k states
- m asynchronous processes with n states each have m^n states

Superposition of those quickly yields unmanageable number of states

Model Checking

- System is modeled as transition system
 - $M = (S, \rightarrow, L)$ with a set of atoms (p, q, r, \dots : either True or False)
 - S : States
 - \rightarrow : Transitions (rule: transitions are always possible) (paths are infinite)
 - L : which atoms are true in which states
- Problem: Is “ $M, s \models \varphi$ ” true?
- Input:
 - Model $M = (S, \rightarrow, L)$
 - Formula φ in a temporal logic




Model Checking Algorithm (CTL)

- Strategy: Starting from the smallest subformulas and working outward towards φ , label the states of M with the subformulas of φ that are satisfied there.
- Follow the following rules until the whole φ has been considered:
 - \perp : then no states are labelled with \perp .
 - p : then label s with p if $p \in L(s)$.
 - $\psi_1 \wedge \psi_2$: label s with $\psi_1 \wedge \psi_2$ if s is already labelled both with ψ_1 and with ψ_2 .
 - $\neg\psi_1$: label s with $\neg\psi_1$ if s is not already labelled with ψ_1 .
 - $AF \psi_1$:
 - If any state s is labelled with ψ_1 , label it with $AF \psi_1$.
 - Repeat: label any state with $AF \psi_1$ if all successor states are labelled with $AF \psi_1$, until there is no change. This step is illustrated in Figure 3.24.
 - $E[\psi_1 U \psi_2]$:
 - If any state s is labelled with ψ_2 , label it with $E[\psi_1 U \psi_2]$.
 - Repeat: label any state with $E[\psi_1 U \psi_2]$ if it is labelled with ψ_1 and at least one of its successors is labelled with $E[\psi_1 U \psi_2]$, until there is no change. This step is illustrated in Figure 3.25.
 - $EX \psi_1$: label any state with $EX \psi_1$ if one of its successors is labelled with ψ_1 .


Example of CTL algorithm

Q: Is “E [Ψ_1 U Ψ_2]”
true in M?




Example of CTL algorithm

Q: Is “E [ψ_1 U ψ_2]”
true in M?




Example of CTL algorithm

Q: Is “E [ψ_1 U ψ_2]”
true in M?




Example of CTL algorithm

Q: Is “ $E [\Psi_1 U \Psi_2]$ ”
true in M ?



Example of CTL algorithm


Q: Is “ $E [\Psi_1 U \Psi_2]$ ”
true in M ?



Example of CTL algorithm

Q: Is “E [Ψ_1 U Ψ_2]”
true in M?


A: There are 3 states in
which “E [Ψ_1 U Ψ_2]” is
true.



Model Checking Algorithm (LTL)



- Step 1: Construct an automaton $A_{\neg\varphi}$ that accepts formula $\neg\varphi$
- Step 2: Combine $A_{\neg\varphi}$ and model M into a new automaton.
- Step 3: Check if the new automaton accepts any path. If no, $M, s \models \varphi$; if yes, the path is a counterexample.

Our model M:




Example of LTL algorithm


Our model M:




Step 1



Step 2



Step 3



Ways to overcome State Explosion Problem

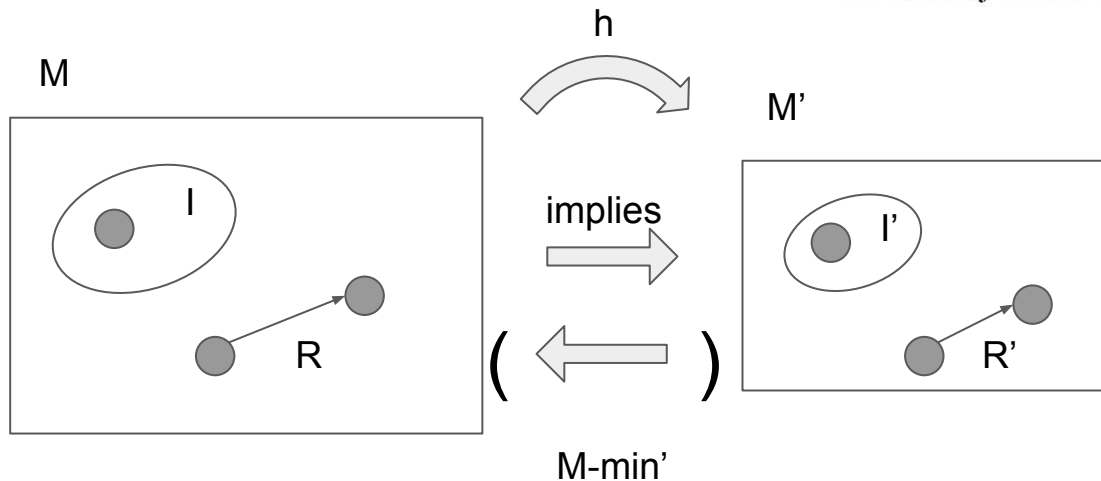
- Abstraction (←- our required reading)
- Partial order reduction
- BDD-based symbolic model checking
- Bounded model checking with SAT / SMT

Abstraction

- Approximation

- M' approximates M
- Use M' to deduce properties of M

- (1) congruence modulo an integer, for dealing with arithmetic operations;
- (2) single bit abstractions, for dealing with bitwise logical operations;
- (3) product abstractions, for combining abstractions such as the above; and
- (4) symbolic abstractions, which is a powerful type of abstraction that allows us to verify an entire class of formulas simultaneously.



Abstraction

- How to produce M-min'
 - Impractical to construct directly from M explicitly (what if b is 64-bit)
 - Solution: Compute it directly from the program text using relational semantics + approximation tricks

0: $p := 0$

1: **while** $b \neq 0$

$p := p \oplus \text{lsb}(b)$

$b := b \gg 1$

endwhile

2: **end**

Symbolic Representation of the program:

$(\text{PC} = 0 \wedge p' = 0 \wedge b' = b \wedge \text{PC}' = 1)$

$\vee (\text{PC} = 1 \wedge b = 0 \wedge p' = p \wedge b' = b \wedge \text{PC}' = 2)$

$\vee (\text{PC} = 1 \wedge b \neq 0 \wedge p' = p \oplus \text{lsb}(b) \wedge b' = b \gg 1 \wedge \text{PC}' = 1)$

$\vee (\text{PC} = 2 \wedge p' = p \wedge b' = b \wedge \text{PC}' = 2).$

Abstraction

- How to produce M-min' from program text
 - Step 1. Derive formula I for initial condition and formula R for the transition relation using *relational semantic*. I and R can represent M .
 - Step 2. Try to compute I -min' and R -min' (which represent M-min') directly from I and R .
 - Step 3. Step 2 is too difficult. Use approximation tricks to derive I -app' and R -app' for M-app' instead. M-app' somehow similar to M-min'.
- Result: we get M-app' instead of M-min'

The tricks:

- (1) If P is a primitive relation, then $\mathcal{A}(P(x_1, \dots, x_m)) = [P](\hat{x}_1, \dots, \hat{x}_m)$ and $\mathcal{A}(\neg P(x_1, \dots, x_m)) = [\neg P](\hat{x}_1, \dots, \hat{x}_m)$.
- (2) $\mathcal{A}(\phi_1 \wedge \phi_2) = \mathcal{A}(\phi_1) \wedge \mathcal{A}(\phi_2)$.
- (3) $\mathcal{A}(\phi_1 \vee \phi_2) = \mathcal{A}(\phi_1) \vee \mathcal{A}(\phi_2)$.
- (4) $\mathcal{A}(\forall x \phi) = \forall \hat{x} \mathcal{A}(\phi)$.
- (5) $\mathcal{A}(\exists x \phi) = \exists \hat{x} \mathcal{A}(\phi)$.

How similar:

- 1.M-min' transition -> M-app' transition
- 2.M-min' initial state -> M-app' initial state

Abstraction

- Now we have $M\text{-app}'$. The paper shows $M\text{-app}'$ also approximates M .
- Main result in the paper:

COROLLARY 5.7. *Assume $M \sqsubseteq_h \hat{M}$, and let ϕ be a $\forall\text{CTL}^*$ formula describing \hat{M} . Then $\hat{M} \models \phi$ implies $M \models \mathcal{E}(\phi)$.*

Definition 5.4. \mathcal{E} is the mapping from formulas describing \hat{M} to formulas describing M that is defined as follows:


- (1) $\mathcal{E}(\text{true}) = \text{true}$. $\mathcal{E}(\text{false}) = \text{false}$. $\mathcal{E}(\hat{v}_i = \hat{d}_i)$ is $\bigvee \{v_i = d_i \mid h_i(d_i) = \hat{d}_i\}$.
 $\mathcal{E}(\hat{v}_i \neq \hat{d}_i) = \neg \mathcal{E}(\hat{v}_i = \hat{d}_i)$.
- (2) If ϕ and ψ are state formulas, then $\mathcal{E}(\phi \wedge \psi) = \mathcal{E}(\phi) \wedge \mathcal{E}(\psi)$, and $\mathcal{E}(\phi \vee \psi) = \mathcal{E}(\phi) \vee \mathcal{E}(\psi)$.
- (3) If ϕ is a path formula, then $\mathcal{E}(\forall(\phi)) = \forall(\mathcal{E}(\phi))$, and $\mathcal{E}(\exists(\phi)) = \exists(\mathcal{E}(\phi))$.
- (4) If ϕ is a path formula that is also a state formula, then $\mathcal{E}(\phi)$ is given by the above rules.
- (5) If ϕ and ψ are path formulas, then $\mathcal{E}(\phi \wedge \psi) = \mathcal{E}(\phi) \wedge \mathcal{E}(\psi)$, and $\mathcal{E}(\phi \vee \psi) = \mathcal{E}(\phi) \vee \mathcal{E}(\psi)$.
- (6) If ϕ and ψ are path formulas, then
 - (a) $\mathcal{E}(\mathbf{X}\phi) = \mathbf{X}\mathcal{E}(\phi)$,
 - (b) $\mathcal{E}(\phi\mathbf{U}\psi) = \mathcal{E}(\phi)\mathbf{U}\mathcal{E}(\psi)$, and
 - (c) $\mathcal{E}(\phi\mathbf{V}\psi) = \mathcal{E}(\phi)\mathbf{V}\mathcal{E}(\psi)$.

Binary Decision Diagram

- Binary Decision Tree - - - some reduction rules - - > Reduced OBDD
- Characteristics of Reduced OBDD
 - Compact representation of boolean functions
 - Canonical: all semantically-equivalent boolean func have exactly the same BDD structures
 - Common operations (+, *, ^) have reasonable complexities (not exponential). The complexities depends on **the size of OBDD**.
 - Size of OBDD critically relies on the variable order. Worst case can be exponential. In some cases we can have only worst case (ex: integer multiplication function).

How is BDD useful?

- State space and transition relations in model M can be represented as Reduced OBDD.



States:

$$x_1 \cdot \bar{x}_2 + \bar{x}_1 \cdot x_2 + \bar{x}_1 \cdot \bar{x}_2$$

Transitions:

x_1	x_2	x'_1	x'_2	\rightarrow
0	0	0	0	1
0	0	0	1	0
0	0	1	0	1
0	0	1	1	0
0	1	0	0	1
0	1	0	1	0
0	1	1	0	0
0	1	1	1	0
1	0	0	0	0
1	0	0	1	1
1	0	1	0	0
1	0	1	1	0
1	1	0	0	0
1	1	0	1	0
1	1	1	0	0
1	1	1	1	0

$$f \rightarrow \stackrel{\text{def}}{=} \bar{x}_1 \cdot \bar{x}_2 \cdot \bar{x}'_1 \cdot \bar{x}'_2 + \bar{x}_1 \cdot \bar{x}_2 \cdot x'_1 \cdot \bar{x}'_2 + x_1 \cdot \bar{x}_2 \cdot \bar{x}'_1 \cdot x'_2 + \bar{x}_1 \cdot x_2 \cdot \bar{x}'_1 \cdot \bar{x}'_2.$$

Application 1: Model Checking of Linux TCP (2004)

- 50k lines of code
- Size of the system state – 250 KBs ($\sim 2^{2048000}$ states)
 - The observable universe has $\sim 2^{273}$ atoms


CMC System:

- Runs two Linux Kernels in parallel (two TCP peers)
- Containerised TCP code via an interface
- Compresses states efficiently to deal with state explosion
- Attempts to visit as many states as possible before running out of resources
- Checks for memory leaks, resource leaks, and protocol conformance

Results: Found **4 bugs** in implementation

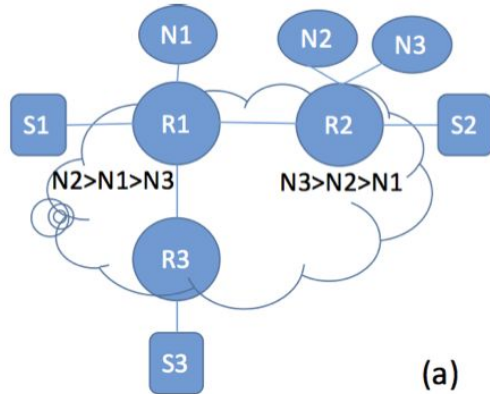
Application 2: Network Configuration Verification

- Most of the network outages happen due to misconfiguration
- We need tools that could verify all data planes for a given configuration

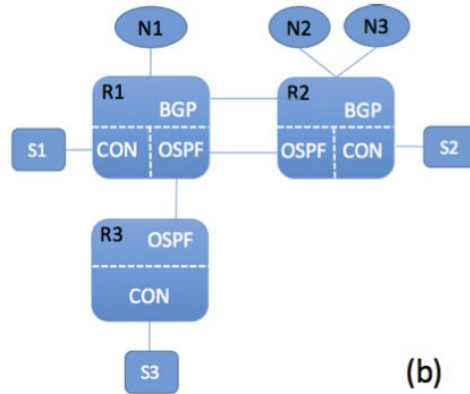


Minesweeper

- Reasoning about **networks as graphs**, not as paths
- Combinatorial search (**formal logic**) instead of message construction
- BGP as a **stable path problem**
- **Multiple optimizations** to scale to size of real networks





(a)



(b)

Stable Path Problem

- Graph (V, E) with a special node 0 that every other node is trying to reach.
- Paths to 0 from v_k : $P = (v_k, v_{k-1}, \dots, 0)$
- Path value $L(P)$
- Stable path assignment $s(v) = P$ if P maximizes the value
- A stable path problem is solvable if every node can have a stable assignment



Results

- Created formal system F that embeds network configuration and constraints
- Type of constraints supported:
 - Reachability and isolation
 - Waypoints, path length, equal paths, disjoint paths
 - Identifying forwarding loops and black holes
 - Load balancing, fault tolerance,
 - Full and partial equivalence
 - Many more
- Testing:
 - Applied to 152 real network, found **120 violations** of must-hold properties
 - Including one that possesses significant security threat