Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction

Introduction to Model Checking

Fabio Somenzi

Department of Electrical, Computer, and Energy Engineering University of Colorado at Boulder

July 25, 2009

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Outline						

◆□> ◆□> ◆三> ◆三> ・三 のへで



- 2 Modeling Systems and Properties
- ③ Specification Mechanisms
- 4 CTL Model Checking
- 5 Model Checking LTL and CTL*
- 6 SAT-Based Model Checking

7 Abstraction



• Two thirds of ASIC budget goes into verification

- Dynamic verification has improved, but ...
- The verification crisis only got worse over the last decade
- Over 60% of IC designs requires a second *spin*
- Bugs that go undetected may end up costing hundreds of millions of dollars

▲ロ ▶ ▲周 ▶ ▲ 国 ▶ ▲ 国 ▶ ● の Q @

• The FDIV bug costed Intel over \$ 500M

• Security and dependability are increasingly important



- Two thirds of ASIC budget goes into verification
- Dynamic verification has improved, but ...
- The verification crisis only got worse over the last decade
- Over 60% of IC designs requires a second spin
- Bugs that go undetected may end up costing hundreds of millions of dollars

• The FDIV bug costed Intel over \$ 500M

• Security and dependability are increasingly important



- Two thirds of ASIC budget goes into verification
- Dynamic verification has improved, but ...
- The verification crisis only got worse over the last decade
- Over 60% of IC designs requires a second spin
- Bugs that go undetected may end up costing hundreds of millions of dollars

• The FDIV bug costed Intel over \$ 500M

• Security and dependability are increasingly important



- Two thirds of ASIC budget goes into verification
- Dynamic verification has improved, but ...
- The verification crisis only got worse over the last decade
- Over 60% of IC designs requires a second spin
- Bugs that go undetected may end up costing hundreds of millions of dollars

- The FDIV bug costed Intel over \$ 500M
- Security and dependability are increasingly important



- Two thirds of ASIC budget goes into verification
- Dynamic verification has improved, but ...
- The verification crisis only got worse over the last decade
- Over 60% of IC designs requires a second spin
- Bugs that go undetected may end up costing hundreds of millions of dollars

- The FDIV bug costed Intel over \$ 500M
- Security and dependability are increasingly important



- Two thirds of ASIC budget goes into verification
- Dynamic verification has improved, but
- The verification crisis only got worse over the last decade
- Over 60% of IC designs requires a second *spin*
- Bugs that go undetected may end up costing hundreds of millions of dollars

- The FDIV bug costed Intel over \$ 500M
- Security and dependability are increasingly important

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Binary-G	iray-Bina	ary				



◆□ > ◆□ > ◆ Ξ > ◆ Ξ > → Ξ = ∽ ۹ < ↔

 $\mathsf{AG}(p \leftrightarrow z)$: invariably (AG) p and z have the same value

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Binary-G	iray-Bina	ary				



 $AG(p \leftrightarrow z)$: invariably (AG) p and z have the same value

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Binary-G	ray-Bina	ary				



◆□▶ ◆□▶ ◆三▶ ◆三▶ 三三 - のへで

 $AG(p \leftrightarrow z)$: invariably (AG) p and z have the same value

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction				
Verilog D	Verilog Description									

```
module gray (clock, i, z);
     input clock, i;
     output z;
     reg p, q, r;
     wire w:
     always @ (posedge clock) begin
          \mathbf{r} = \mathbf{z}:
          q = p;
          \mathbf{p} = \mathbf{i};
     end
     assign w = p^q, z = w^r;
endmodule // gray
```

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
States ar	nd Trans	sitions				



 $(p \leftrightarrow z) \rightarrow AG(p \leftrightarrow z)$ holds of all initial states $(q \leftrightarrow r) \rightarrow AG(p \leftrightarrow z)$ holds of all initial states $AG((p \leftrightarrow z) \leftrightarrow (q \leftrightarrow r))$ holds of all initial states

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
States ar	nd Trans	sitions				



 $(p \leftrightarrow z) \rightarrow AG(p \leftrightarrow z)$ holds of all initial states $(q \leftrightarrow r) \rightarrow AG(p \leftrightarrow z)$ holds of all initial states $AG((p \leftrightarrow z) \leftrightarrow (q \leftrightarrow r))$ holds of all initial states

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
States ar	nd Trans	sitions				



 $(p \leftrightarrow z) \rightarrow AG(p \leftrightarrow z)$ holds of all initial states $(q \leftrightarrow r) \rightarrow AG(p \leftrightarrow z)$ holds of all initial states $AG((p \leftrightarrow z) \leftrightarrow (q \leftrightarrow r))$ holds of all initial states

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
States ar	nd Trans	sitions				



 $(p \leftrightarrow z) \rightarrow AG(p \leftrightarrow z)$ holds of all initial states $(q \leftrightarrow r) \rightarrow AG(p \leftrightarrow z)$ holds of all initial states $AG((p \leftrightarrow z) \leftrightarrow (q \leftrightarrow r))$ holds of all initial states

◆□ > ◆□ > ◆臣 > ◆臣 > ○臣 ○ のへ⊙





 $\mathsf{EF}(p \leftrightarrow z)$ "*p* and *z* may become equal"

does not hold of all initial states

◆□▶ ◆□▶ ◆三▶ ◆三▶ 三三 - のへで

Introduction

Modeling

Specifications

CTL

LTL and CTL*

T Abstraction

The CGW Puzzle



rlrl: boat and goat on right bank; cabbage and wolf on left

 $E \neg$ yellow U cyan

there is a path to a cyan state not going through any yellow states



 μZ . cyan \vee (\neg yellow \wedge EX Z)





 μZ . cyan \vee (\neg yellow \wedge EX Z)





 μZ . cyan \vee (\neg yellow \wedge EX Z)





 μZ . cyan \vee (\neg yellow \wedge EX Z)





 μZ . cyan \vee (\neg yellow \wedge EX Z)





 μZ . cyan \vee (\neg yellow \wedge EX Z)





Compute

 μZ . cyan \vee (\neg yellow \wedge EX Z)





Compute

 μZ . cyan \vee (\neg yellow \wedge EX Z)





Compute

 μZ . cyan \vee (\neg yellow \wedge EX Z)





- Check whether the given finite-state system is a model for a property
- That is, check whether the computations of the system satisfy the property

• The check is based on exploring the states of the system

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
The Cub	e Puzzle					



$$\mathsf{AG}(\mathsf{position} = \mathsf{center} \to \bigvee_{0 \le i < 27} \neg \mathsf{visited}_i)$$

<□ > < @ > < E > < E > E のQ @

There are 3.46 $\cdot\,10^7$ reachable states out of 4.29 $\cdot\,10^9$

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
The Cub	e Puzzle					



$\mathsf{AG}(\mathsf{position} = \mathsf{center} \to \bigvee_{0 \le i < 27} \neg \mathsf{visited}_i)$

◆□ > ◆□ > ◆ Ξ > ◆ Ξ > → Ξ = ∽ 9 < @

There are 3.46 $\cdot\,10^7$ reachable states out of 4.29 $\cdot\,10^9$

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
The Cube Puzzle						



$$\mathsf{AG}(\mathsf{position} = \mathsf{center} \to \bigvee_{0 \le i < 27} \neg \mathsf{visited}_i)$$

There are $3.46\cdot 10^7$ reachable states out of $4.29\cdot 10^9$





Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Further	Up					





- Galaxies in the (observable) universe: $2.5 \cdot 10^{11}$
- Stars in the Milky Way: $4 \cdot 10^{11}$
- Stars in the universe: 10²³
- Average number of neutrinos per cubic meter: $3.3\cdot 10^8$
- Neutrinos in the universe: 10⁹³ (wild guess)
- A small sequential circuit may have more than 10^{100} states $(10^{100} = 1 \text{ googol})$



- Galaxies in the (observable) universe: $2.5 \cdot 10^{11}$
- Stars in the Milky Way: $4 \cdot 10^{11}$
- Stars in the universe: 10²³
- Average number of neutrinos per cubic meter: 3.3 · 10⁸
- Neutrinos in the universe: 10⁹³ (wild guess)
- A small sequential circuit may have more than 10^{100} states $(10^{100} = 1 \text{ googol})$

▲ロト ▲帰 ト ▲ ヨ ト ▲ ヨ ト ● ● ● ● ● ●



- Galaxies in the (observable) universe: $2.5 \cdot 10^{11}$
- Stars in the Milky Way: 4 · 10¹¹
- Stars in the universe: 10²³
- Average number of neutrinos per cubic meter: 3.3 · 10⁸
- Neutrinos in the universe: 10⁹³ (wild guess)
- A small sequential circuit may have more than 10^{100} states $(10^{100} = 1 \text{ googol})$

▲ロト ▲帰 ト ▲ ヨ ト ▲ ヨ ト ● ● ● ● ● ●


• Symbolic Model Checking

- Represent sets (of states and transitions) by their characteristic functions
 - $\{00, 01, 10\} \longrightarrow \neg x_1 \lor \neg x_2$
 - $x_1 \vee x_{100}$ represents $3 \cdot 2^{98}$ elements
- BDDs and CNF popular representations
- Do not enumerate the elements of the sets
 - Manipulate the characteristic functions instead

Abstraction

• Infer properties of the concrete model from the analysis of a simplified abstract model

◆□▶ ◆□▶ ◆□▶ ◆□▶ ▲□ ◆ ○ ◆



• Symbolic Model Checking

- Represent sets (of states and transitions) by their characteristic functions
 - $\{00, 01, 10\} \longrightarrow \neg x_1 \lor \neg x_2$
 - $x_1 \vee x_{100}$ represents $3 \cdot 2^{98}$ elements
- BDDs and CNF popular representations
- Do not enumerate the elements of the sets
 - Manipulate the characteristic functions instead
- Abstraction
 - Infer properties of the concrete model from the analysis of a simplified abstract model

◆□▶ ◆□▶ ◆□▶ ◆□▶ ▲□ ◆ ○ ◆





◆□ > ◆□ > ◆豆 > ◆豆 > 「豆 」 のへで

```
#include <iostream>
int main()
{
   std::cout << "Hello, world!" << std::endl;
   return 0;
}</pre>
```





900

- 2

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Peterson	Mutex	Algorithm				

```
always @ (posedge clock) begin
    self = select:
    case (pc[self])
    L0: if (!pause) pc[self] = L1;
    L1: begin intr[self] = 1; pc[self] = L2; end
    L2: begin turn = "self; pc[self] = L3; end
    L3: if (!intr[~self] || turn == self) pc[self] = L4;
    L4: if (!pause) pc[self] = L5; // critical
    L5: begin intr[self] = 0; pc[self] = L0; end
    endcase
```

end



Properties for Mutex Algorithm

- Mutual exclusion
 - AG \neg (pc[0] = L4 \land pc[1] = L4)
- Absence of starvation
 - $AG(pc[0] = L1 \rightarrow AF pc[0] = L4)$
 - $AG(pc[1] = L1 \rightarrow AFpc[1] = L4)$
- Temporal logic operators
 - AG: invariably, AF: inevitably

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Nondeter	minism					

```
always @ (posedge clock) begin
    self = select;
    case (pc[self])
    L0: if (!pause) pc[self] = L1;
    L1: begin intr[self] = 1; pc[self] = L2; end
    L2: begin turn = "self; pc[self] = L3; end
    L3: if (!intr[~self] || turn == self) pc[self] = L4;
    L4: if (!pause) pc[self] = L5; // critical
    L5: begin intr[self] = 0; pc[self] = L0; end
    endcase
```

end



• Fair scheduling

• $GF self = 1 \land GF self = 0$

• Neither process dwells forever in the critical section

▲ロ ▶ ▲周 ▶ ▲ 国 ▶ ▲ 国 ▶ ● の Q @

• $G F pc[0] \neq L4 \land G F pc[1] \neq L4$

- Temporal logic operators
 - GF: infinitely often

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Outline						



- 2 Modeling Systems and Properties
- ③ Specification Mechanisms
- 4 CTL Model Checking
- 5 Model Checking LTL and CTL*
- 6 SAT-Based Model Checking

7 Abstraction



- Finite transition systems without inputs
- $\langle S, T, S_0, A, L \rangle$
 - S: finite set of states
 - $T \subseteq S \times S$: transition relation
 - $S_0 \subseteq S$: set of initial states
 - A: set of atomic propositions
 - $L: S \rightarrow 2^A$: labeling function







- Complex systems are composed of several modules
- Each module is described as a finite state machine (FSM)
- The overall Kripke structure is obtained as the product of the FSMs

- State explosion!
- The product can be either synchronous or asynchronous (interleaving)

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Specifica	ations					

• Properties are sets of behaviors

- Various specification mechanisms are in use: Temporal logics and automata are popular
- The examples we have seen are formulae of the temporal logic CTL

◆□ > ◆□ > ◆豆 > ◆豆 > 「豆 」 のへで



- Properties are sets of behaviors
- Various specification mechanisms are in use: Temporal logics and automata are popular
- The examples we have seen are formulae of the temporal logic CTL



- Properties are sets of behaviors
- Various specification mechanisms are in use: Temporal logics and automata are popular
- The examples we have seen are formulae of the temporal logic CTL



- Properties are sets of behaviors
- Various specification mechanisms are in use: Temporal logics and automata are popular
- The examples we have seen are formulae of the temporal logic CTL



- Temporal logics add temporal operators and path quantifiers to standard (e.g., propositional) logics
- Temporal operators allow one to conveniently describe the order of occurrence of events and other statements involving time without explicitly mentioning time

• Expressiveness of propositional temporal logic in between those of propositional logic and predicate logic



- Temporal logics add temporal operators and path quantifiers to standard (e.g., propositional) logics
- Temporal operators allow one to conveniently describe the order of occurrence of events and other statements involving time without explicitly mentioning time

• Expressiveness of propositional temporal logic in between those of propositional logic and predicate logic

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Operato	rs and Q	uantifiers				

- G φ : φ holds globally $(\Box \varphi)$
- F φ : φ holds eventually ($\Diamond \varphi$)
- $\psi \cup \varphi$: ψ holds until φ holds
- $\psi \operatorname{\mathsf{R}} \varphi$: ψ releases φ
- X φ : φ holds at the next state ($\bigcirc \varphi$)
- E: along at least one path
- A: along all paths
- In CTL all temporal operators are immediately preceeded by a path quantifier

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Operato	rs and Q	uantifiers				

- G φ : φ holds globally $(\Box \varphi)$
- F φ : φ holds eventually ($\Diamond \varphi$)
- $\psi \cup \varphi$: ψ holds until φ holds
- $\psi \operatorname{\mathsf{R}} \varphi$: ψ releases φ
- X φ : φ holds at the next state ($\bigcirc \varphi$)
- E: along at least one path
- A: along all paths
- In CTL all temporal operators are immediately preceded by a path quantifier



Branching time logics reason about computation trees





Linear time logics reason about sets of computation paths





- A linear-time temporal logic formula can contain only one path quantifier at the beginning
 - A F G *p*
 - $A(F p \rightarrow F q)$
 - EGF p (existential linear-time formula)
- A branching-time formula may have multiple quantifiers

- AF AG p
- AG EF p
- EG EF *p*





• A F G φ holds in this structure

• Infinite paths eventually dwell in either a or c where φ holds

• AF AG φ does not hold in this structure

• As long as a run dwells in a, it can always go to a state where φ does not hold

(日) (字) (日) (日) (日)

• No CTL formula exists that is equivalent to A F G φ





- A F G φ holds in this structure
 - Infinite paths eventually dwell in either a or c where φ holds
- AF AG φ does not hold in this structure
 - As long as a run dwells in a, it can always go to a state where φ does not hold

(日) (字) (日) (日) (日)

• No CTL formula exists that is equivalent to A F G φ





- A F G φ holds in this structure
 - Infinite paths eventually dwell in either a or c where φ holds
- AF AG φ does not hold in this structure
 - As long as a run dwells in a, it can always go to a state where φ does not hold

(日) (字) (日) (日) (日)

• No CTL formula exists that is equivalent to A F G φ





◆□▶ ◆□▶ ◆ □▶ ★ □▶ = □ の < @

Linear-time properties cannot distinguish these two



- Linear time properties cannot distinguish two structures if they are trace (language) equivalent
- Branching time properties cannot distinguish two structures if they are bisimilar

• Bisimilarity is stronger than trace equivalence (see previous example)



- Linear time properties cannot distinguish two structures if they are trace (language) equivalent
- Branching time properties cannot distinguish two structures if they are bisimilar

• Bisimilarity is stronger than trace equivalence (see previous example)





◆□▶ ◆□▶ ◆三▶ ◆三▶ 三三 - のへで

p and p' are bisimilar



 A relation B among the states of two Kripke structures K and K' with A = A' is a bisimulation relation if (p, p') ∈ B implies

◆□▶ ◆□▶ ◆□▶ ◆□▶ ▲□ ◆ ○ ◆

- L(p) = L(p')
- $(p,q) \in T \rightarrow \exists q'.(p',q') \in T' \land (q,q') \in B$
- $(p',q') \in T' \rightarrow \exists q.(p,q) \in T \land (q,q') \in B$



- Two Kripke structures K and K' are bisimulation equivalent (K ≡ K') if there is a bisimulation relation between their states such that every initial state of one structure is bisimilar to some initial state of the other structure
- Two structures are bisimulation equivalent iff they satisfy the same branching-time properties

• CTL properties suffice (Browne et al. [1988])



- Two Kripke structures K and K' are bisimulation equivalent (K ≡ K') if there is a bisimulation relation between their states such that every initial state of one structure is bisimilar to some initial state of the other structure
- Two structures are bisimulation equivalent iff they satisfy the same branching-time properties

• CTL properties suffice (Browne et al. [1988])



- Two Kripke structures K and K' are bisimulation equivalent (K ≡ K') if there is a bisimulation relation between their states such that every initial state of one structure is bisimilar to some initial state of the other structure
- Two structures are bisimulation equivalent iff they satisfy the same branching-time properties

• CTL properties suffice (Browne et al. [1988])





◆□ > ◆□ > ◆豆 > ◆豆 > 「豆 」 のへで

p is simulated by p'



A relation Σ among the states of two Kripke structures K and K' with A ⊆ A' is a simulation relation if (p, p') ∈ Σ implies
L(p) = L'(p') ∩ A
(p,q) ∈ T → ∃q'.(p',q') ∈ T' ∧ (q,q') ∈ Σ
We say that p is simulated by p'

◆□▶ ◆□▶ ◆□▶ ◆□▶ ▲□ ◆ ○ ◆
Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Example	of Simu	lation				





- Kripke structure K' simulates structure K (written K ≤ K') if there exists a simulation relation Σ ⊆ S × S' such that for every initial state s of K there is an initial state s' of K' such that (s, s') ∈ Σ
- If K ≤ K' and φ is a universal branching-time formula over A, then K' ⊨ φ implies K ⊨ φ.
- Two Kripke structures K and K' are simulation equivalent $(K \sim K')$ iff $K \preceq K'$ and $K' \preceq K$
- Two structures are simulation equivalent iff they satisfy the same branching-time universal properties



- Kripke structure K' simulates structure K (written K ≤ K') if there exists a simulation relation Σ ⊆ S × S' such that for every initial state s of K there is an initial state s' of K' such that (s, s') ∈ Σ
- If K ≤ K' and φ is a universal branching-time formula over A, then K' ⊨ φ implies K ⊨ φ.
- Two Kripke structures K and K' are simulation equivalent (K ~ K') iff K ≤ K' and K' ≤ K
- Two structures are simulation equivalent iff they satisfy the same branching-time universal properties



- Kripke structure K' simulates structure K (written K ≤ K') if there exists a simulation relation Σ ⊆ S × S' such that for every initial state s of K there is an initial state s' of K' such that (s, s') ∈ Σ
- If K ≤ K' and φ is a universal branching-time formula over A, then K' ⊨ φ implies K ⊨ φ.
- Two Kripke structures K and K' are simulation equivalent $(K \sim K')$ iff $K \preceq K'$ and $K' \preceq K$
- Two structures are simulation equivalent iff they satisfy the same branching-time universal properties



- Kripke structure K' simulates structure K (written K ≤ K') if there exists a simulation relation Σ ⊆ S × S' such that for every initial state s of K there is an initial state s' of K' such that (s, s') ∈ Σ
- If K ≤ K' and φ is a universal branching-time formula over A, then K' ⊨ φ implies K ⊨ φ.
- Two Kripke structures K and K' are simulation equivalent $(K \sim K')$ iff $K \preceq K'$ and $K' \preceq K$
- Two structures are simulation equivalent iff they satisfy the same branching-time universal properties



- Branching time is more powerful, but also trickier
 - Two vending machine models
 - A F G φ vs. AF AG φ
- Linear time is more suitable for compositional verification and Bounded Model Checking

• Counterexample generation simpler for linear time



- Branching time is more powerful, but also trickier
 - Two vending machine models
 - A F G φ vs. AF AG φ
- Linear time is more suitable for compositional verification and Bounded Model Checking

• Counterexample generation simpler for linear time



- Branching time is more powerful, but also trickier
 - Two vending machine models
 - A F G φ vs. AF AG φ
- Linear time is more suitable for compositional verification and Bounded Model Checking

• Counterexample generation simpler for linear time

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Safety						

- A safety property describes something bad that should not happen
- AG \neg (grant₀ \land grant₁): requestors 0 and 1 should not be granted access to the shared resource simultaneously

<日 > < 同 > < 目 > < 目 > < 目 > < 目 > < 0 < 0</p>

• AG \neg (door = open \land engine = running)

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Liveness						

• A liveness property describes something good that should happen

◆□▶ ◆□▶ ◆ □▶ ★ □▶ = = - のへぐ

 \bullet AG(req \rightarrow Fack): requests should be acknowledged



- The definitions given so far may be confusing
- $AG(\texttt{command} = \texttt{stop} \rightarrow X \texttt{state} = \texttt{halt})$
- $AG(\text{command} = \text{stop} \rightarrow F \text{state} = \text{halt})$
- $AG(req \rightarrow Fack)$



- A linear-time property is a set of (linear) traces or infinite sequences over the atomic propositions
- A property φ is a safety property if every trace not in φ has a prefix that cannot be extended to a trace in φ
- Intuitively, the prefix includes the bad event that causes the property to fail
- $AG(p \rightarrow Xq)$ is a safety property
 - A counterexample includes a state where *p* holds followed by a state where ¬*q* holds
 - The prefix that includes these two states cannot be extended to an infinite path satisfying the property



- A linear-time property is a set of (linear) traces or infinite sequences over the atomic propositions
- A property φ is a safety property if every trace not in φ has a prefix that cannot be extended to a trace in φ
- Intuitively, the prefix includes the bad event that causes the property to fail
- $AG(p \rightarrow Xq)$ is a safety property
 - A counterexample includes a state where *p* holds followed by a state where ¬*q* holds
 - The prefix that includes these two states cannot be extended to an infinite path satisfying the property

<日 > < 同 > < 目 > < 目 > < 目 > < 目 > < 0 < 0</p>



- A linear-time property φ is a liveness property if every finite sequence over the atomic propositions can be extended to an infinite sequence in φ
- Intuitively, finite prefixes do not affect the fulfillment of the eventualities that characterize liveness properties
- $AG(p \rightarrow Fq)$ is a liveness property
 - One can add a state where q holds to any finite prefix



- A linear-time property φ is a liveness property if every finite sequence over the atomic propositions can be extended to an infinite sequence in φ
- Intuitively, finite prefixes do not affect the fulfillment of the eventualities that characterize liveness properties
- $AG(p \rightarrow Fq)$ is a liveness property
 - One can add a state where q holds to any finite prefix



- $A(p \cup q)$ is neither a safety property nor a liveness property
 - A trace satisfying G(p ∧ ¬q) is a counterexample with no prefix that cannot be extended
 - A trace satisfying ¬q U(¬p¬q) is a counterexample with a finite prefix that cannot be extended

◆□▶ ◆□▶ ◆□▶ ◆□▶ ▲□ ◆ ○ ◆

• Every property can be written as the intersection of a safety property and a liveness property

• $A(p \cup q) = A((q \cap (p \vee q)) \wedge F q)$



- $A(p \cup q)$ is neither a safety property nor a liveness property
 - A trace satisfying G(p ∧ ¬q) is a counterexample with no prefix that cannot be extended
 - A trace satisfying ¬q U(¬p¬q) is a counterexample with a finite prefix that cannot be extended

◆□▶ ◆□▶ ◆□▶ ◆□▶ ▲□ ◆ ○ ◆

- Every property can be written as the intersection of a safety property and a liveness property
 - $A(p \cup q) = A((q R(p \lor q)) \land Fq)$

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Classific	ation of	Properties				

- Time structure: branching, linear
- Bisimilarity, trace equivalence
- Safety, liveness
- Existential, universal
- Tense: future, past
 - X versus Y (neXt vs. Yesterday)
 - See (Laroussinie and Schnoebelen [2000])

◆□ > ◆□ > ◆臣 > ◆臣 > ○臣 ○ のへ⊙

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Outline						

◆□▶ ◆□▶ ◆□▶ ◆□▶ ▲□ ◆ ○ ◆



- 2 Modeling Systems and Properties
- Specification Mechanisms
- 4 CTL Model Checking
- 5 Model Checking LTL and CTL*
- 6 SAT-Based Model Checking

7 Abstraction

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
CTL*						

• CTL* is a powerful branching-time temporal logic

- We use a subset of the operators (X and U) and the E quantifier to define the syntax
- The remaining operators are defined as abbreviations
- We need both state formulae and path formulae to recursively define the logic

◆□▶ ◆□▶ ◆□▶ ◆□▶ ▲□ ◆ ○ ◆

- The state formulae give CTL*
- Path formulae can only appear as subformulae



- CTL* is a powerful branching-time temporal logic
- We use a subset of the operators (X and U) and the E quantifier to define the syntax
- The remaining operators are defined as abbreviations
- We need both state formulae and path formulae to recursively define the logic

- The state formulae give CTL*
- Path formulae can only appear as subformulae



- CTL* is a powerful branching-time temporal logic
- We use a subset of the operators (X and U) and the E quantifier to define the syntax
- The remaining operators are defined as abbreviations
- We need both state formulae and path formulae to recursively define the logic

- The state formulae give CTL*
- Path formulae can only appear as subformulae

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
CTL*						

- CTL* is a powerful branching-time temporal logic
- We use a subset of the operators (X and U) and the E quantifier to define the syntax
- The remaining operators are defined as abbreviations
- We need both state formulae and path formulae to recursively define the logic

<日 > < 同 > < 目 > < 目 > < 目 > < 目 > < 0 < 0</p>

- The state formulae give CTL*
- Path formulae can only appear as subformulae



- An atomic proposition is a state formula
- A state formula is also a path formula
- If φ and ψ are state formulae, so are $\neg\varphi$ and $\varphi\wedge\psi$
- If φ is a path formula, $E \varphi$ is a state formula
- If φ and ψ are path formulae, so are $\neg\varphi$ and $\varphi\wedge\psi$
- If φ and ψ are path formulae, so are X φ and ψ U φ

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
CTL* AI	bbreviati	ons				

◆□ > ◆□ > ◆ Ξ > ◆ Ξ > → Ξ = ∽ ۹ < ↔

•
$$\varphi \lor \psi = \neg (\neg \varphi \land \neg \psi)$$

• true =
$$\varphi \lor \neg \varphi$$

• false =
$$\varphi \land \neg \varphi$$

•
$$\psi \mathsf{R} \varphi = \neg (\neg \psi \mathsf{U} \neg \varphi)$$

• F
$$\varphi = \mathsf{true}\,\mathsf{U}\,\varphi$$

•
$$\mathsf{G} \, \varphi = \mathsf{false} \, \mathsf{R} \, \varphi$$

• A
$$\varphi = \neg E \neg \varphi$$



- $\neg \mathsf{A}\mathsf{G}\mathsf{F}\varphi = \mathsf{E}\mathsf{F}\mathsf{G}\neg\varphi$
- Negations can be pushed "down" toward the atomic propositions

▲ロト ▲帰 ト ▲ ヨ ト ▲ ヨ ト ● ● ● ● ● ●

• The basic rules are (beside DeMorgan)

•
$$\psi \mathsf{R} \varphi = \neg (\neg \psi \mathsf{U} \neg \varphi)$$

• A
$$\varphi = \neg E \neg \varphi$$

 $\bullet\,$ From the first we get ${\rm G}\,\varphi=\neg\,{\rm F}\,\neg\varphi$



• The semantics of CTL* formulae are defined with respect to a Kripke structure *K*

- If formula φ holds of state s (path π) of K, we write $K, s \models \varphi$ ($K, \pi \models \varphi$)
- The double turnstile is read "models"
- K is omitted when no ambiguity arises
- π^i is π without the first *i* states



• The semantics of CTL* formulae are defined with respect to a Kripke structure *K*

- If formula φ holds of state s (path π) of K, we write $K, s \models \varphi$ ($K, \pi \models \varphi$)
- The double turnstile is read "models"
- K is omitted when no ambiguity arises
- π^i is π without the first *i* states



• The semantics of CTL* formulae are defined with respect to a Kripke structure *K*

- If formula φ holds of state s (path π) of K, we write
 K, s ⊨ φ (K, π ⊨ φ)
- The double turnstile is read "models"
- K is omitted when no ambiguity arises
- π^i is π without the first *i* states

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
CTL* Se	mantics					

•
$$s \models p, p \in A$$
 iff $p \in L(s)$
• $s \models \neg \varphi$ iff $s \not\models \varphi$
• $s \models \varphi \land \psi$ iff $s \models \varphi$ and $s \models \psi$
• $s \models E\varphi$ iff $\exists \pi$ from s such that $\pi \models \varphi$
• $\pi \models \neg \varphi$ iff $\pi \not\models \varphi$
• $\pi \models \varphi \land \psi$ iff $\pi \models \varphi$ and $\pi \models \psi$
• $\pi \models X\varphi$ iff $\pi^1 \models \varphi$
• $\pi \models \psi \cup \varphi$ iff $\exists i \ge 0 . \pi^i \models \varphi$ and $0 \le j < i \rightarrow \pi^j \models \psi$

◆□ ◆ ▲□ ◆ ▲□ ◆ ▲□ ◆ ▲□ ◆

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Semantic	s of X a	ind U				



◆□ > ◆□ > ◆ Ξ > ◆ Ξ > → Ξ = ∽ ۹ < ↔

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
CTL						

 Computational Tree Logic is a branching time fragment of CTL*

<日 > < 同 > < 目 > < 目 > < 目 > < 目 > < 0 < 0</p>

- In CTL every temporal operator must be immediately preceded by a path quantifier
 - AF AG φ is a CTL formula
 - A F G φ is not a CTL formula
- Model checking easy relative to CTL*

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
CTL						

 Computational Tree Logic is a branching time fragment of CTL*

▲ロ ▶ ▲周 ▶ ▲ 国 ▶ ▲ 国 ▶ ● の Q @

- In CTL every temporal operator must be immediately preceded by a path quantifier
 - AF AG φ is a CTL formula
 - A F G φ is not a CTL formula
- Model checking easy relative to CTL*

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
LTL						

- Linear Temporal Logic is a linear-time fragment of CTL*
- In LTL there can be only one quantifier at the beginning of the formula
- The quantifier is usually A, in which case it is usually omitted
 FGφ means AFGφ

▲□▶ ▲□▶ ▲□▶ ▲□▶ □ のQ@

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
LTL						

- Linear Temporal Logic is a linear-time fragment of CTL*
- In LTL there can be only one quantifier at the beginning of the formula
- The quantifier is usually A, in which case it is usually omitted

<日 > < 目 > < 目 > < 目 > < 目 > < 目 > < 目 > < 0 < 0</p>

• FG φ means AFG φ

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction

Omega-Automata

- Properties can be described by automata that take the computation of the system as input and either accept it or reject it
- For computations that finish regular automata suffice
- For non-terminating computations we need ω -automata

<日 > < 同 > < 目 > < 目 > < 目 > < 目 > < 0 < 0</p>
Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction

Omega-Automata

- Properties can be described by automata that take the computation of the system as input and either accept it or reject it
- For computations that finish regular automata suffice
- For non-terminating computations we need ω -automata

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction

Omega-Automata

- Properties can be described by automata that take the computation of the system as input and either accept it or reject it
- For computations that finish regular automata suffice
- $\bullet\,$ For non-terminating computations we need $\omega\textsc{-}automata$

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction

Automaton for Safety Property

• $AG(\varphi \rightarrow X \neg \psi)$ is negated to get the formula $EF(\varphi \land X \psi)$



ヘロン 人間 とくほど 人間 とう

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction

Why Omega Automata?

- For safety properties we can "stretch" regular automata because it suffices to find the non-extensible prefixes of the counterexamples
- For liveness properties we need to address the fact that the computations we consider do not terminate

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction

Büchi Automata

- Büchi automata are similar to (nondeterministic) regular automata except for the acceptance conditions
- A Büchi automaton accepts an infinite sequence if there is a run of it that goes through an accepting state infinitely often

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction

Büchi Automaton to Model Check AFG φ

 $\bullet\,$ We negate the property to get EGF $\neg\varphi$



◆□▶ ◆□▶ ◆□▶ ◆□▶ ▲□ ◆ ○ ◆

• A computation accepted by the automaton is a counterexample to A F G φ



- Omega-automata describe linear-time properties
- They are more expressive than LTL



◆□ > ◆□ > ◆豆 > ◆豆 > 「豆 」 のへで

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Fairness						

- Fairness constraints are used to
 - Model features of the environment
 - Mitigate the effects of simplifications of the model
- They instruct the model checker to disregard certain (unfair) computations



- Property: The combination lock will open unless we keep making mistakes
- Fairness constraint: provided we never give up
- This constraint implies that the knob is turned infinitely often

◆□ > ◆□ > ◆豆 > ◆豆 > 「豆 」 のへで



- Property: The combination lock will open unless we keep making mistakes
- Fairness constraint: provided we never give up
- This constraint implies that the knob is turned infinitely often

	Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction	
Modeling the Environment								

- Property: The combination lock will open unless we keep making mistakes
- Fairness constraint: provided we never give up
- This constraint implies that the knob is turned infinitely often

◆□▶ ◆□▶ ◆三▶ ◆三▶ ○三 のへぐ



• Property: All requests are eventually acknowledged

- Fairness constraints: provided all grantees eventually relinquish the shared resource
- These constraints imply that each requestor is not using the shared resource infinitely often

▲ロ ▶ ▲周 ▶ ▲ 国 ▶ ▲ 国 ▶ ● の Q @



- Property: All requests are eventually acknowledged
- Fairness constraints: provided all grantees eventually relinquish the shared resource
- These constraints imply that each requestor is not using the shared resource infinitely often

▲ロ ▶ ▲周 ▶ ▲ 国 ▶ ▲ 国 ▶ ● の Q @

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction		
Mitigating Abstraction								

- Property: All requests are eventually acknowledged
- Fairness constraints: provided all grantees eventually relinquish the shared resource
- These constraints imply that each requestor is not using the shared resource infinitely often

◆□ > ◆□ > ◆豆 > ◆豆 > 「豆 」 のへで

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction	
Büchi Fairness Constraints							

- A Büchi fairness constraint is a set of states that a fair run must intersect infinitely often
- LTL can express fairness constraints
 - A(G F $\psi\to\varphi)$ says that φ must hold only along paths where ψ holds infinitely often

▲ロ ▶ ▲周 ▶ ▲ 国 ▶ ▲ 国 ▶ ● の Q @

- CTL cannot express fairness constraints
 - They must be specified separately

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Outline						

◆□▶ ◆□▶ ◆□▶ ◆□▶ ▲□ ◆ ○ ◆



- 2 Modeling Systems and Properties
- 3 Specification Mechanisms
- 4 CTL Model Checking
- 5 Model Checking LTL and CTL*
- 6 SAT-Based Model Checking

7 Abstraction



◆□ > ◆□ > ◆豆 > ◆豆 > 「豆 」 のへで

- We find $\llbracket \varphi \rrbracket_K$, the set of all states *s* such that $K, s \models \varphi$
- We then check whether $S_0 \subseteq \llbracket \varphi \rrbracket_K$
- When no confusion arises we write simply φ for $\llbracket \varphi \rrbracket_K$



◆□ > ◆□ > ◆豆 > ◆豆 > 「豆 」 のへで

- We find $\llbracket \varphi \rrbracket_{\mathcal{K}}$, the set of all states s such that $\mathcal{K}, s \models \varphi$
- We then check whether $S_0 \subseteq \llbracket \varphi \rrbracket_K$
- When no confusion arises we write simply φ for $[\![\varphi]\!]_K$



- We find $\llbracket \varphi \rrbracket_{\mathcal{K}}$, the set of all states s such that $\mathcal{K}, s \models \varphi$
- We then check whether $S_0 \subseteq \llbracket \varphi \rrbracket_K$
- When no confusion arises we write simply φ for $\llbracket \varphi \rrbracket_K$



- We find $\llbracket \varphi \rrbracket_{\mathcal{K}}$, the set of all states s such that $\mathcal{K}, s \models \varphi$
- We then check whether $S_0 \subseteq \llbracket \varphi \rrbracket_K$
- When no confusion arises we write simply φ for $\llbracket \varphi \rrbracket_K$



- Work bottom-up on the parse graph of the CTL formula
- Annotate every node with the satisfying set of the subformula rooted at the node
- The computation at each node of the parse graph depends on its label

▲ロ ▶ ▲周 ▶ ▲ 国 ▶ ▲ 国 ▶ ● の Q @









▲ロト ▲園 ト ▲ 国 ト ▲ 国 ト ▲ 回 ト ④ ヘ () ()















• Atomic proposition *p*: return set of states with *p* in their labels

- $\neg \varphi$: return the complement of $\llbracket \varphi \rrbracket_K$
- $\varphi \land \psi$: return $\llbracket \varphi \rrbracket_{\mathcal{K}} \cap \llbracket \psi \rrbracket_{\mathcal{K}}$
- EX φ: return the set of predecessors in K of the states in [[φ]]_K
- E ψ U φ: return the set of states with paths to states in [[φ]]_K that are entirely contained in [[ψ]]_K (except possibly for the last state of each path)

◆□▶ ◆□▶ ◆□▶ ◆□▶ ▲□ ◆ ○ ◆



- Atomic proposition *p*: return set of states with *p* in their labels
- $\neg \varphi$: return the complement of $\llbracket \varphi \rrbracket_K$
- $\varphi \wedge \psi$: return $\llbracket \varphi \rrbracket_{\mathcal{K}} \cap \llbracket \psi \rrbracket_{\mathcal{K}}$
- EX φ : return the set of predecessors in K of the states in $\llbracket \varphi \rrbracket_K$
- $E \psi \cup \varphi$: return the set of states with paths to states in $\llbracket \varphi \rrbracket_{\mathcal{K}}$ that are entirely contained in $\llbracket \psi \rrbracket_{\mathcal{K}}$ (except possibly for the last state of each path)

◆□▶ ◆□▶ ◆□▶ ◆□▶ ▲□ ◆ ○ ◆



- Atomic proposition *p*: return set of states with *p* in their labels
- $\neg \varphi$: return the complement of $\llbracket \varphi \rrbracket_K$
- $\varphi \wedge \psi$: return $\llbracket \varphi \rrbracket_K \cap \llbracket \psi \rrbracket_K$
- EX φ: return the set of predecessors in K of the states in [[φ]]_K
- E ψ U φ: return the set of states with paths to states in [[φ]]_K that are entirely contained in [[ψ]]_K (except possibly for the last state of each path)

◆□▶ ◆□▶ ◆□▶ ◆□▶ ▲□ ◆ ○ ◆



- Atomic proposition *p*: return set of states with *p* in their labels
- $\neg \varphi$: return the complement of $\llbracket \varphi \rrbracket_K$
- $\varphi \wedge \psi$: return $\llbracket \varphi \rrbracket_K \cap \llbracket \psi \rrbracket_K$
- EX φ: return the set of predecessors in K of the states in [[φ]]_K
- E ψ U φ: return the set of states with paths to states in [[φ]]_K that are entirely contained in [[ψ]]_K (except possibly for the last state of each path)



- Atomic proposition *p*: return set of states with *p* in their labels
- $\neg \varphi$: return the complement of $\llbracket \varphi \rrbracket_K$
- $\varphi \wedge \psi$: return $\llbracket \varphi \rrbracket_K \cap \llbracket \psi \rrbracket_K$
- EX φ: return the set of predecessors in K of the states in [[φ]]_K
- E ψ U φ: return the set of states with paths to states in [[φ]]_K that are entirely contained in [[ψ]]_K (except possibly for the last state of each path)



- To fix ideas, suppose Kripke structures are represented by adjacency lists and sets by bit vectors
- Boolean operations on sets can be performed in linear time
- Computing EU amounts to reachability in a graph
- Computing EG amounts to finding the strongly connected components (SCCs) of a subgraph of *K*
- Both reachability and SCC computation are based on depth-first search and take time linear in the size of the Kripke structure

▲ロ ▶ ▲周 ▶ ▲ 国 ▶ ▲ 国 ▶ ● の Q @

• In practice sets are often represented by hash tables



- To fix ideas, suppose Kripke structures are represented by adjacency lists and sets by bit vectors
- Boolean operations on sets can be performed in linear time
- Computing EU amounts to reachability in a graph
- Computing EG amounts to finding the strongly connected components (SCCs) of a subgraph of *K*
- Both reachability and SCC computation are based on depth-first search and take time linear in the size of the Kripke structure

▲ロト ▲帰 ト ▲ ヨ ト ▲ ヨ ト ● ● ● ● ● ●

• In practice sets are often represented by hash tables

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction		
Explicit Model Checking								

- To fix ideas, suppose Kripke structures are represented by adjacency lists and sets by bit vectors
- Boolean operations on sets can be performed in linear time
- Computing EU amounts to reachability in a graph
- Computing EG amounts to finding the strongly connected components (SCCs) of a subgraph of *K*
- Both reachability and SCC computation are based on depth-first search and take time linear in the size of the Kripke structure

<日 > < 同 > < 目 > < 目 > < 目 > < 目 > < 0 < 0</p>

• In practice sets are often represented by hash tables



Complexity of Explicit CTL MC

- Considering fairness constraints does not change the complexity of the algorithm
 - It suffices to discard SCCs that do not intersect all fair sets

▲ロト ▲帰 ト ▲ ヨ ト ▲ ヨ ト ● ● ● ● ● ●

• CTL model checking is linear in the size of the formula and the size of the structure



- Linear time complexity is great unless your system has 10⁵⁰ states
 - Number of states grows exponentially with number of state variables

▲ロト ▲帰 ト ▲ ヨ ト ▲ ヨ ト ● ● ● ● ● ●

- Explicit model checking limited to a few billion states
- Symbolic model checking can do much more (though not uniformly)


- Linear time complexity is great unless your system has 10⁵⁰ states
 - Number of states grows exponentially with number of state variables

- Explicit model checking limited to a few billion states
- Symbolic model checking can do much more (though not uniformly)



- Linear time complexity is great unless your system has 10⁵⁰ states
 - Number of states grows exponentially with number of state variables

- Explicit model checking limited to a few billion states
- Symbolic model checking can do much more (though not uniformly)

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Characte	ristic Fu	nctions				

- States are encoded as binary strings of length *n* (the number of binary state variables)
- We often drop the " χ " from χ_V when there is no ambiguity

◆□ ▶ ◆□ ▶ ◆□ ▶ ◆□ ▶ ◆□ ◆ ○ ◆

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Characte	ristic Fu	nctions				

- States are encoded as binary strings of length *n* (the number of binary state variables)
- We often drop the " χ " from χ_V when there is no ambiguity

◆□ ▶ ◆□ ▶ ◆□ ▶ ◆□ ▶ ◆□ ◆ ○ ◆

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Symbolic	Represe	entation				



• $T(x_1, x_0, y_1, y_0) = (\neg x_1 \land \neg x_0 \land \neg y_1 \land y_0) \lor (\neg x_1 \land x_0 \land \neg y_0) \lor (x_1 \land \neg x_0 \land \neg y_0)$

◆□ > ◆□ > ◆豆 > ◆豆 > ・ 亘 ・ 今 Q @

• $S_0(x_1, x_0) = \neg x_1 \land \neg x_0$

•
$$p(x_1, x_0) = x_1 \wedge \neg x_0$$

•
$$q(x_1, x_0) = \neg x_1 \wedge x_0$$

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Characte	ristic Fu	nctions				

• Let V contain all states with either the first or the last bit set to 1

$$\chi_V = x_1 \vee x_n$$

◆□▶ ◆□▶ ◆三▶ ◆三▶ ○三 のへぐ

- Set V has $3 \cdot 2^{n-2}$ elements
- Great, but let us not get carried away, because it is not possible to find a representation that is compact for most functions

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Charac	cteristic Fu	nctions				

• Let V contain all states with either the first or the last bit set to 1

 $\chi_V = x_1 \vee x_n$

- Set V has $3 \cdot 2^{n-2}$ elements
- Great, but let us not get carried away, because it is not possible to find a representation that is compact for most functions

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Character	istic Fun	ictions				

• Let V contain all states with either the first or the last bit set to 1

$$\chi_V = x_1 \vee x_n$$

◆□▶ ◆□▶ ◆三▶ ◆三▶ ○三 のへぐ

- Set V has $3 \cdot 2^{n-2}$ elements
- Great, but let us not get carried away, because it is not possible to find a representation that is compact for most functions



- Symbolic model checking enumerates states implicitly
- No explicit loop on the states or the transitions is used
- The cost of implicit enumeration is more affected by the size of the representation than by the cardinality of the set



- Symbolic model checking enumerates states implicitly
- No explicit loop on the states or the transitions is used
- The cost of implicit enumeration is more affected by the size of the representation than by the cardinality of the set



- Let x(y) be the vector of current (next) state variables
- The set of predecessors of the states in V is given by

 $\exists y . T(x,y) \land V(y)$

▲ロ ▶ ▲周 ▶ ▲ 国 ▶ ▲ 国 ▶ ● の Q @

• No loops over states and transitions



- Let x(y) be the vector of current (next) state variables
- The set of predecessors of the states in V is given by

 $\exists y . T(x,y) \land V(y)$

▲ロ ▶ ▲周 ▶ ▲ 国 ▶ ▲ 国 ▶ ● の Q @

No loops over states and transitions

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Symbolic	Model	Checking				

- Boolean connectives give no difficulties
- Complementation turns into negation
- Union becomes disjunction and intersection becomes conjunction
- $\bullet\,$ We have seen how to deal with EX φ
- For EU and EG we use a fixpoint characterization

• A fixpoint x of f is such that f(x) = x



- The satisfying sets of EU and EG computations are fixpoints of monotonic functions over 2^S
 - $\mathsf{E} \psi \mathsf{U} \varphi = \varphi \lor [\psi \land \mathsf{EX} \mathsf{E} \psi \mathsf{U} \varphi]$
 - EG $\varphi = \varphi \wedge \operatorname{\mathsf{EX}}\operatorname{\mathsf{EG}}\varphi$
- Specifically, E ψ U φ is the least fixpoint and EG φ is the greatest fixpoint. This is written

•
$$\mathsf{E} \psi \mathsf{U} \varphi = \mu Z \cdot \varphi \lor [\psi \land \mathsf{EX} Z]$$

• EG
$$\varphi = \nu Z \cdot \varphi \wedge \mathsf{EX} Z$$

• μ -calculus notation



- The satisfying sets of EU and EG computations are fixpoints of monotonic functions over 2^S
 - $\mathsf{E} \psi \mathsf{U} \varphi = \varphi \lor [\psi \land \mathsf{EX} \mathsf{E} \psi \mathsf{U} \varphi]$
 - EG $\varphi = \varphi \wedge \text{EX EG } \varphi$
- Specifically, E ψ U φ is the least fixpoint and EG φ is the greatest fixpoint. This is written

•
$$\mathsf{E} \psi \mathsf{U} \varphi = \mu Z \cdot \varphi \vee [\psi \land \mathsf{EX} Z]$$

• EG
$$\varphi = \nu Z \cdot \varphi \wedge \mathsf{EX} Z$$

• μ -calculus notation



- The satisfying sets of EU and EG computations are fixpoints of monotonic functions over 2^S
 - $\mathsf{E} \psi \mathsf{U} \varphi = \varphi \lor [\psi \land \mathsf{EX} \mathsf{E} \psi \mathsf{U} \varphi]$
 - EG $\varphi = \varphi \wedge \text{EX EG } \varphi$
- Specifically, E ψ U φ is the least fixpoint and EG φ is the greatest fixpoint. This is written

•
$$\mathsf{E} \psi \mathsf{U} \varphi = \mu Z \cdot \varphi \vee [\psi \land \mathsf{EX} Z]$$

• EG
$$\varphi = \nu Z \cdot \varphi \wedge \text{EX } Z$$

• μ -calculus notation



$$x \leq y \to f(x) \leq f(y)$$

• A monotonic function over a finite lattice has a least fixpoint that can be computed as the limit of the sequence

 $0, f(0), f(f(0)), f(f(f(0))), \dots$

• For greatest fixpoints

 $1, f(1), f(f(1)), f(f(f(1))), \ldots$

▲ロ ▶ ▲周 ▶ ▲ 国 ▶ ▲ 国 ▶ ● の Q @



$$x \leq y \to f(x) \leq f(y)$$

• A monotonic function over a finite lattice has a least fixpoint that can be computed as the limit of the sequence

 $0, f(0), f(f(0)), f(f(f(0))), \ldots$

• For greatest fixpoints

 $1, f(1), f(f(1)), f(f(f(1))), \ldots$

▲ロ ▶ ▲周 ▶ ▲ 国 ▶ ▲ 国 ▶ ● の Q @



$$x \leq y \to f(x) \leq f(y)$$

• A monotonic function over a finite lattice has a least fixpoint that can be computed as the limit of the sequence

 $0, f(0), f(f(0)), f(f(f(0))), \ldots$

• For greatest fixpoints

$1, f(1), f(f(1)), f(f(f(1))), \ldots$

<日 > < 同 > < 目 > < 目 > < 目 > < 目 > < 0 < 0</p>



$$x \leq y \to f(x) \leq f(y)$$

• A monotonic function over a finite lattice has a least fixpoint that can be computed as the limit of the sequence

 $0, f(0), f(f(0)), f(f(f(0))), \ldots$

• For greatest fixpoints

```
1, f(1), f(f(1)), f(f(f(1))), \ldots
```

▲ロ ▶ ▲周 ▶ ▲ 国 ▶ ▲ 国 ▶ ● の Q @

Introduction Modeling Specifications CTL LTL and CTL* SAT Abstraction

Computing E ψ U φ and E ψ R φ

$$Z = \emptyset;$$

$$\zeta = S;$$

while
$$(Z \neq \zeta)$$
 {
 $\zeta = Z$;
 $Z = \varphi \lor (\psi \land \mathsf{EX} Z)$;
}

$$Z = S;$$

$$\zeta = \emptyset;$$

while
$$(Z \neq \zeta)$$
 {
 $\zeta = Z$;
 $Z = \varphi \land (\psi \lor \mathsf{EX} Z)$;
}

◆ロト ◆聞 ▶ ◆臣 ▶ ◆臣 ▶ ○ 臣 ○ の Q @









Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Computi	ng E ψ L	Jφ				



Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Computi	ng E ψ L	Jφ				







▲□▶▲圖▶▲圖▶▲圖▶ 圖 のQ@

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Computi	ng EG $arphi$					



- ▲日 > ▲ 国 > ▲ 国 > ▲ 国 > ● ④ ● ●

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Comput	ing EG $arphi$					



- ▲日 > ▲ 国 > ▲ 国 > ▲ 国 > ● ④ ● ●





▲□▶▲圖▶▲圖▶▲圖▶ 圖 のQ@







- The algorithm we have presented is based on backward search in the Kripke structure (EX computes predecessors)
- Part of CTL can be model checked using forward search (using EY that computes successors)
- Forward search popular for invariants (AG p)
- Forward search is also used to find states reachable from initial states and prune backward search



- In checking K ⊨ EF p we can stop the computation of the least fixpoint as soon as all initial states are acquired
- This is an instance of on-the-fly model checking
- In general, on-the-fly model checking stops as soon as the answer is known



- ECTL is a fragment of CTL that consists of existential properties
- Negation can only occur in front of atomic propositions
- The path quantifier A cannot be used
- ACTL consists of the negation of ECTL formulae (universal properties)



- ECTL is a fragment of CTL that consists of existential properties
- Negation can only occur in front of atomic propositions
- The path quantifier A cannot be used
- ACTL consists of the negation of ECTL formulae (universal properties)



- $\mathsf{EF}(\varphi \land \mathsf{EG} \neg \psi) \mathsf{ECTL}$
- $AG(\varphi \rightarrow AF\psi) ACTL$
 - Negation of the previous formula
- $AG(AF \varphi \rightarrow AF \psi)$ mixed
- $EF(EG \varphi \lor EG \psi) ECTL$



- If a property fails it is useful to show an execution of the system that is a counterexample
- For a property that holds, it may be useful to show an execution of the system that is a witness
- A witness to $K, s \models \varphi$ is a counterexample to $K, s \models \neg \varphi$ and vice versa


• Witness for the ECTL property

 $\mathsf{EF}(\varphi \wedge \mathsf{EG} \neg \psi)$



◆□ > ◆□ > ◆臣 > ◆臣 > ○臣 ○ のへで



- \bullet One witness computation does not suffice to show that AG φ holds in a structure with multiple paths
- One counterexample computation does not suffice to show that ${\rm EF}\,\neg\varphi$ fails
- Complete witnesses can be found for ECTL formulae and complete counterexamples for ACTL formulae



- \bullet One witness computation does not suffice to show that AG φ holds in a structure with multiple paths
- One counterexample computation does not suffice to show that ${\rm EF}\,\neg\varphi$ fails
- Complete witnesses can be found for ECTL formulae and complete counterexamples for ACTL formulae



- \bullet One witness computation does not suffice to show that AG φ holds in a structure with multiple paths
- One counterexample computation does not suffice to show that EF $\neg \varphi$ fails
- Complete witnesses can be found for ECTL formulae and complete counterexamples for ACTL formulae



• Consider a witness for the ECTL formula

 $\mathsf{EF}(\mathsf{EG}\,\varphi\wedge\mathsf{EG}\,\psi)$





- Fairness constraints are given as sets of states that must be traversed infinitely often (Büchi fairness)
- In CTL the fair sets are the satisfying sets of CTL formulae (distinct from the formulae for the properties to be verified)

- Quantifiers are restricted to fair paths
 - Paths that intersect all the fair sets infinitely often



- Fairness constraints are given as sets of states that must be traversed infinitely often (Büchi fairness)
- In CTL the fair sets are the satisfying sets of CTL formulae (distinct from the formulae for the properties to be verified)

- Quantifiers are restricted to fair paths
 - Paths that intersect all the fair sets infinitely often



- Fairness constraints are given as sets of states that must be traversed infinitely often (Büchi fairness)
- In CTL the fair sets are the satisfying sets of CTL formulae (distinct from the formulae for the properties to be verified)

- Quantifiers are restricted to fair paths
 - Paths that intersect all the fair sets infinitely often



- The fair states are those along fair paths
- They are computed by the μ -calculus formula

 νZ . EX[E Z U($Z \wedge c$)]

▲ロ ▶ ▲周 ▶ ▲ 国 ▶ ▲ 国 ▶ ● の Q @

in the case of one fairness constraint \boldsymbol{c}





- For EX and EU, we append witness to EG true to finite witness
- For EG φ , while tracing a loop in a fair SCC we need to insure that all fair sets are visited

• Finding a shortest witness is hard, but heuristics work reasonably well

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Outline						

◆□ > ◆□ > ◆豆 > ◆豆 > 「豆 」 のへで



- 2 Modeling Systems and Properties
- 3 Specification Mechanisms
- 4 CTL Model Checking
- 5 Model Checking LTL and CTL*
- 6 SAT-Based Model Checking

Abstraction



• LTL operators can be characterized in terms of fixpoints like their CTL counterparts

•
$$\psi \cup \varphi = \varphi \vee [\psi \land \mathsf{X}(\psi \cup \varphi)]$$

•
$$\psi \mathsf{R} \varphi = \varphi \land [\psi \lor \mathsf{X}(\psi \mathsf{R} \varphi)]$$

 The problem is that we have sets of paths instead of sets of states

◆□ > ◆□ > ◆豆 > ◆豆 > 「豆 」 のへで

• We take a different approach



- Negate the given formula φ to get $\neg \varphi$
- Build a Büchi automaton A_{¬φ} that accepts exactly the computations that model ¬φ
- Compose $A_{\neg \varphi}$ with the system to be verified
- Check whether there is a fair path in the composition

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
From Fo	rmula to	Automato	on			

$$\psi \, \mathsf{U} \, \varphi = \varphi \vee [\psi \wedge \mathsf{X}(\psi \, \mathsf{U} \, \varphi)]$$



◆□▶ ▲□▶ ▲目▶ ▲目▶ ▲□▶

Introduction Modeling Specifications CTL LTL and CTL* SAT Abstraction $Model Checking AG(p \rightarrow Fq)$

- A counterexample to AG(p → F q) is an execution in which p happens, but q does not follow it
- Compose automaton for all counterexamples with model and check for accepting path



◆□▶ ◆□▶ ◆臣▶ ◆臣▶ 三臣 - のへ⊙



- If formula is propositional, compute its satisfying set directly, otherwise
- There is a subformula that is an LTL formula E ψ : find its satisfying set with LTL model checking algorithm and use it as a new atomic proposition
- Repeat until the satisfying set of the root of the parse tree is found

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Outline						

◆□▶ ◆□▶ ◆三▶ ◆三▶ ○三 ○○○



- 2 Modeling Systems and Properties
- ③ Specification Mechanisms
- 4 CTL Model Checking
- 5 Model Checking LTL and CTL*
- 6 SAT-Based Model Checking

7 Abstraction

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Bounded	Model	Checking				

• Checks for a counterexample to a property of a model

• We assume finite state and LTL

- Encodes the property checking problem as propositional satisfiability (SAT)
- Constructs a propositional formula that is satisfiable iff there exits a length-k counterexample, e.g.,

$$I(s_0) \wedge \bigwedge_{0 \leq i < k} T(s_i, s_{i+1}) \wedge \neg p(s_k)$$

- If no counterexample is found, BMC increases k until
 - a counterexample is found,
 - the search becomes intractable, or
 - k reaches a certain bound

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Bounded	Model	Checking				

- Checks for a counterexample to a property of a model
 - We assume finite state and LTL
- Encodes the property checking problem as propositional satisfiability (SAT)
- Constructs a propositional formula that is satisfiable iff there exits a length-k counterexample, e.g.,

$$I(s_0) \wedge \bigwedge_{0 \leq i < k} T(s_i, s_{i+1}) \wedge \neg p(s_k)$$

- If no counterexample is found, BMC increases k until
 - a counterexample is found,
 - the search becomes intractable, or
 - k reaches a certain bound

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Bounded	Model	Checking				

- Checks for a counterexample to a property of a model
 - We assume finite state and LTL
- Encodes the property checking problem as propositional satisfiability (SAT)
- Constructs a propositional formula that is satisfiable iff there exits a length-k counterexample, e.g.,

$$I(s_0) \wedge \bigwedge_{0 \leq i < k} T(s_i, s_{i+1}) \wedge \neg p(s_k)$$

- If no counterexample is found, BMC increases k until
 - a counterexample is found,
 - the search becomes intractable, or
 - k reaches a certain bound

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Bounded	Model (Checking				

- Checks for a counterexample to a property of a model
 - We assume finite state and LTL
- Encodes the property checking problem as propositional satisfiability (SAT)
- Constructs a propositional formula that is satisfiable iff there exits a length-k counterexample, e.g.,

$$I(s_0) \wedge \bigwedge_{0 \leq i < k} T(s_i, s_{i+1}) \wedge \neg p(s_k)$$

- If no counterexample is found, BMC increases k until
 - a counterexample is found,
 - the search becomes intractable, or
 - k reaches a certain bound



- Sometime, proving the absence of an error is more useful than finding one
 - When checking an abstract model for a universal property
 - Finding an error in the abstract model does not imply its existence in the original model
 - However, proving that the property passes in the abstract model guarantee the absence of errors in the original model
- BMC efficiency reduces as the length of the counterexample increases
 - It may be more efficient to prove the property and stop early than keep searching for a counterexample



- The original BMC algorithm (Biere et al. [1999]), although complete in theory, is limited in practice to falsification of LTL properties
- BMC can prove that an LTL property ψ passes on a model \mathcal{M} only if a bound, κ , is known such that:
 - if no counterexample of length up to κ is found, then $\mathcal{M} \models \mathsf{A}\,\psi$
- Several methods exist to compute a suitable κ
- The optimum value of κ , however, is usually very expensive to obtain
 - Finding it is at least as hard as checking whether $M \models A \psi$ (Clarke et al. [2004])



- The original BMC algorithm (Biere et al. [1999]), although complete in theory, is limited in practice to falsification of LTL properties
- BMC can prove that an LTL property ψ passes on a model \mathcal{M} only if a bound, κ , is known such that:
 - if no counterexample of length up to ${\color{black}\kappa}$ is found, then $\mathcal{M}\models \mathsf{A}\,\psi$
- Several methods exist to compute a suitable κ
- The optimum value of κ , however, is usually very expensive to obtain
 - Finding it is at least as hard as checking whether $M \models A \psi$ (Clarke et al. [2004])



- The original BMC algorithm (Biere et al. [1999]), although complete in theory, is limited in practice to falsification of LTL properties
- BMC can prove that an LTL property ψ passes on a model \mathcal{M} only if a bound, κ , is known such that:
 - if no counterexample of length up to ${\color{black}\kappa}$ is found, then $\mathcal{M}\models \mathsf{A}\,\psi$
- Several methods exist to compute a suitable κ
- The optimum value of κ , however, is usually very expensive to obtain
 - Finding it is at least as hard as checking whether $M \models A \psi$ (Clarke et al. [2004])

<日 > < 同 > < 目 > < 目 > < 目 > < 目 > < 0 < 0</p>



- The original BMC algorithm (Biere et al. [1999]), although complete in theory, is limited in practice to falsification of LTL properties
- BMC can prove that an LTL property ψ passes on a model \mathcal{M} only if a bound, κ , is known such that:
 - if no counterexample of length up to ${\color{black}\kappa}$ is found, then $\mathcal{M}\models \mathsf{A}\,\psi$
- Several methods exist to compute a suitable κ
- The optimum value of κ , however, is usually very expensive to obtain
 - Finding it is at least as hard as checking whether $M \models A \psi$ (Clarke et al. [2004])

<日 > < 同 > < 目 > < 目 > < 目 > < 目 > < 0 < 0</p>

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Interpola	ation					

 $I(s_0) \land T(s_0, s_1) \land T(s_1, s_2) \land \dots \land T(s_{k-1}, s_k) \land \neg p(s_k)$ is unsatisfiable

- Let $F_1 = I(s_0) \land T(s_0, s_1)$ and $F_2 = T(s_1, s_2) \land \dots \land T(s_{k-1}, s_k) \land \neg p(s_k)$
- Then $F_1(s_0, s_1) \wedge F_2(s_1, \dots, s_k)$ is unsatisfiable
- Interpolant $\mathcal{I}_1(s_1)$ (McMillan [2003]) is such that
 - $F_1(s_0, s_1) \rightarrow \mathcal{I}_1(s_1)$
 - $\mathcal{I}_1(s_1) \wedge F_2(s_1, \dots, s_k)$ is unsatisfiable
- $\mathcal{I}_1(s_1)$ can be computed in linear time from a resolution proof that $F_1(s_0, s_1) \wedge F_2(s_1, \dots, s_k)$ is unsatisfiable
- $\exists s_0 \, . \, I(s_0) \wedge T(s_0, s_1)$ is the strongest interpolant
 - set of states reachable from $I(s_0)$ in one step

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Interpola	ation					

 $I(s_0) \wedge T(s_0, s_1) \wedge T(s_1, s_2) \wedge \cdots \wedge T(s_{k-1}, s_k) \wedge \neg p(s_k)$ is unsatisfiable

- Let $F_1 = I(s_0) \wedge T(s_0, s_1)$ and $F_2 = T(s_1, s_2) \wedge \cdots \wedge T(s_{k-1}, s_k) \wedge \neg p(s_k)$
- Then $F_1(s_0, s_1) \wedge F_2(s_1, \dots, s_k)$ is unsatisfiable
- Interpolant $\mathcal{I}_1(s_1)$ (McMillan [2003]) is such that
 - $F_1(s_0, s_1) \rightarrow \mathcal{I}_1(s_1)$
 - $\mathcal{I}_1(s_1) \wedge F_2(s_1, \dots, s_k)$ is unsatisfiable
- $\mathcal{I}_1(s_1)$ can be computed in linear time from a resolution proof that $F_1(s_0, s_1) \wedge F_2(s_1, \dots, s_k)$ is unsatisfiable

- $\exists s_0 \, . \, I(s_0) \wedge T(s_0, s_1)$ is the strongest interpolant
 - set of states reachable from $I(s_0)$ in one step

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Interpola	ation					

 $I(s_0) \land T(s_0, s_1) \land T(s_1, s_2) \land \dots \land T(s_{k-1}, s_k) \land \neg p(s_k)$ is unsatisfiable

- Let $F_1 = I(s_0) \land T(s_0, s_1)$ and $F_2 = T(s_1, s_2) \land \cdots \land T(s_{k-1}, s_k) \land \neg p(s_k)$
- Then $F_1(s_0,s_1) \wedge F_2(s_1,\ldots,s_k)$ is unsatisfiable
- Interpolant $\mathcal{I}_1(s_1)$ (McMillan [2003]) is such that
 - $F_1(s_0, s_1) \rightarrow \mathcal{I}_1(s_1)$
 - $\mathcal{I}_1(s_1) \wedge F_2(s_1, \dots, s_k)$ is unsatisfiable
- $\mathcal{I}_1(s_1)$ can be computed in linear time from a resolution proof that $F_1(s_0, s_1) \wedge F_2(s_1, \dots, s_k)$ is unsatisfiable

- $\exists s_0 \, . \, l(s_0) \wedge T(s_0, s_1)$ is the strongest interpolant
 - set of states reachable from $I(s_0)$ in one step

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Interpola	tion					

 $I(s_0) \land T(s_0, s_1) \land T(s_1, s_2) \land \dots \land T(s_{k-1}, s_k) \land \neg p(s_k)$ is unsatisfiable

- Let $F_1 = I(s_0) \land T(s_0, s_1)$ and $F_2 = T(s_1, s_2) \land \cdots \land T(s_{k-1}, s_k) \land \neg p(s_k)$
- Then $F_1(s_0, s_1) \wedge F_2(s_1, \dots, s_k)$ is unsatisfiable
- Interpolant $\mathcal{I}_1(s_1)$ (McMillan [2003]) is such that
 - $F_1(s_0, s_1) \rightarrow \mathcal{I}_1(s_1)$
 - $\mathcal{I}_1(s_1) \wedge F_2(s_1, \dots, s_k)$ is unsatisfiable
- $\mathcal{I}_1(s_1)$ can be computed in linear time from a resolution proof that $F_1(s_0, s_1) \wedge F_2(s_1, \dots, s_k)$ is unsatisfiable

- $\exists s_0 . I(s_0) \land T(s_0, s_1)$ is the strongest interpolant
 - set of states reachable from $I(s_0)$ in one step

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Interpola	ation					

 $I(s_0) \land T(s_0, s_1) \land T(s_1, s_2) \land \dots \land T(s_{k-1}, s_k) \land \neg p(s_k)$ is unsatisfiable

- Let $F_1 = I(s_0) \land T(s_0, s_1)$ and $F_2 = T(s_1, s_2) \land \cdots \land T(s_{k-1}, s_k) \land \neg p(s_k)$
- Then $F_1(s_0,s_1) \wedge F_2(s_1,\ldots,s_k)$ is unsatisfiable
- Interpolant $\mathcal{I}_1(s_1)$ (McMillan [2003]) is such that
 - $F_1(s_0, s_1) \rightarrow \mathcal{I}_1(s_1)$
 - $\mathcal{I}_1(s_1) \wedge F_2(s_1, \dots, s_k)$ is unsatisfiable
- $\mathcal{I}_1(s_1)$ can be computed in linear time from a resolution proof that $F_1(s_0, s_1) \wedge F_2(s_1, \ldots, s_k)$ is unsatisfiable

<日 > < 同 > < 目 > < 目 > < 目 > < 目 > < 0 < 0</p>

- $\exists s_0 \, . \, I(s_0) \wedge T(s_0, s_1)$ is the strongest interpolant
 - set of states reachable from $I(s_0)$ in one step

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Interpola	ation					

 $I(s_0) \land T(s_0, s_1) \land T(s_1, s_2) \land \dots \land T(s_{k-1}, s_k) \land \neg p(s_k)$ is unsatisfiable

- Let $F_1 = I(s_0) \land T(s_0, s_1)$ and $F_2 = T(s_1, s_2) \land \cdots \land T(s_{k-1}, s_k) \land \neg p(s_k)$
- Then $F_1(s_0,s_1) \wedge F_2(s_1,\ldots,s_k)$ is unsatisfiable
- Interpolant $\mathcal{I}_1(s_1)$ (McMillan [2003]) is such that
 - $F_1(s_0, s_1) \rightarrow \mathcal{I}_1(s_1)$
 - $\mathcal{I}_1(s_1) \wedge F_2(s_1, \dots, s_k)$ is unsatisfiable
- $\mathcal{I}_1(s_1)$ can be computed in linear time from a resolution proof that $F_1(s_0, s_1) \wedge F_2(s_1, \ldots, s_k)$ is unsatisfiable
- $\exists s_0 \, . \, I(s_0) \wedge T(s_0, s_1)$ is the strongest interpolant
 - set of states reachable from $I(s_0)$ in one step



- $\mathcal{I}_1(s_1)$ is a superset of the states reachable in one step such that no member state has a path of length k-1 to a bad state
- Replace $I(s_0)$ with $I(s_0) \lor \mathcal{I}_1(s_0)$ and repeat
 - If formula still unsatisfiable, interpolant \$\mathcal{I}_2(s_1)\$ is a superset of states reachable in one or two steps such that no member state has a path of length \$k-1\$ to a bad state

 A converging sequence of interpolants means that no states satisfying ¬p (bad states) are reachable



- $\mathcal{I}_1(s_1)$ is a superset of the states reachable in one step such that no member state has a path of length k-1 to a bad state
- Replace $I(s_0)$ with $I(s_0) \lor \mathcal{I}_1(s_0)$ and repeat
 - If formula still unsatisfiable, interpolant I₂(s₁) is a superset of states reachable in one or two steps such that no member state has a path of length k − 1 to a bad state

 A converging sequence of interpolants means that no states satisfying ¬p (bad states) are reachable



- $\mathcal{I}_1(s_1)$ is a superset of the states reachable in one step such that no member state has a path of length k-1 to a bad state
- Replace $I(s_0)$ with $I(s_0) \lor \mathcal{I}_1(s_0)$ and repeat
 - If formula still unsatisfiable, interpolant I₂(s₁) is a superset of states reachable in one or two steps such that no member state has a path of length k − 1 to a bad state

 A converging sequence of interpolants means that no states satisfying ¬p (bad states) are reachable
Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Outline						

◆□> ◆□> ◆三> ◆三> ・三 のへで



- 2 Modeling Systems and Properties
- ③ Specification Mechanisms
- 4 CTL Model Checking
- 5 Model Checking LTL and CTL*
- 6 SAT-Based Model Checking

Abstraction



- In use since the early days of simulation
 - What if we write to a memory when the address contains Xs?
- Still indispensable, but
 - May be laborious
 - Comes with no implied warranty





- In use since the early days of simulation
 - What if we write to a memory when the address contains Xs?
- Still indispensable, but
 - May be laborious
 - Comes with no implied warranty





- Given concrete model C and property φ
- Derive somehow an abstract model A from C
 - $\bullet\,$ Abstraction should preserve detail relevant to φ
- Check whether $A \models \varphi$
- If $A \models \varphi$ confidence in correctness of C is increased
- If $A \not\models \varphi$ check counterexample
 - Is failure due to a real bug or an artifact of abstraction?

• Either fix bug or refine abstraction



- Comes with warranty
- Simulation Relations



- bisimilarity preserves all of μ -calculus
- simulation equivalence preserves linear-time properties

(日)

3

• simulation preserves passing universal properties



- Comes with warranty
- Simulation Relations



- bisimilarity preserves all of μ -calculus
- simulation equivalence preserves linear-time properties

・ロト ・ 同ト ・ ヨト ・ ヨト

3

• simulation preserves passing universal properties



- Computing simulation relations is expensive
- Avoid direct computation
 - Cone of Influence (COI) Reduction and slicing yield a bisimilar model

◆□▶ ◆□▶ ◆□▶ ◆□▶ ▲□ ◆ ○ ◆

- The symmetry quotient induced by an invariance group is bisimilar to the original model
- Predicate abstraction leads to bisimulation or simulation
- Freeing some components leads to simulation
- Simulation is sufficient for language containment



Rule of thumb: The more you want to preserve, the less you can abstract

- Bisimulation does not allow much abstraction
- Property driven abstraction
 - focus on preserving only the property of interest
 - Start with a coarse abstraction
 - Let the property guide the initial abstraction and the refinement

▲ロ ▶ ▲周 ▶ ▲ 国 ▶ ▲ 国 ▶ ● の Q @



- What is not in the abstract model is ignored
- Refinement brings in something that was previously ignored





- What is not in the abstract model is ignored
- Refinement brings in something that was previously ignored



◆□ > ◆□ > ◆豆 > ◆豆 > 「豆 = つへぐ



- What is not in the abstract model is ignored
- Refinement brings in something that was previously ignored



◆□ > ◆□ > ◆豆 > ◆豆 > 「豆 = つへぐ

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Abstract	ion Refi	nement				

- For universal properties:
 - Choose initial abstract model
 - while the property fails
 - if concrete model also fails, report failure

- refine
- report success
- For mixed properties, use two abstractions
- If property fails on abstract model, an abstract counterexample is produced



- A cube is divided into 27 equal smaller cubes (blocks hereafter). Is it possible to trace a path from any point on the surface to the center that always travels parallel to some side of the cube and visits each block exactly once?
- If we number the smaller cubes consecutively, top-down, front to back, and left to right, we find that there are 14 odd-numbered blocks and 13 even-numbered blocks
- Since all adjacent blocks have different parity, we cannot find a path of length 27 ending with an even-numbered block.
- Since the center block is numbered 14, the answer is "no"
- The same argument applies whenever the side of the cube is divided into an odd number of parts, so that the number of blocks is (2n+1)³ for some n > 0



- A cube is divided into 27 equal smaller cubes (blocks hereafter). Is it possible to trace a path from any point on the surface to the center that always travels parallel to some side of the cube and visits each block exactly once?
- If we number the smaller cubes consecutively, top-down, front to back, and left to right, we find that there are 14 odd-numbered blocks and 13 even-numbered blocks
- Since all adjacent blocks have different parity, we cannot find a path of length 27 ending with an even-numbered block.
- Since the center block is numbered 14, the answer is "no"
- The same argument applies whenever the side of the cube is divided into an odd number of parts, so that the number of blocks is (2n+1)³ for some n > 0



- A cube is divided into 27 equal smaller cubes (blocks hereafter). Is it possible to trace a path from any point on the surface to the center that always travels parallel to some side of the cube and visits each block exactly once?
- If we number the smaller cubes consecutively, top-down, front to back, and left to right, we find that there are 14 odd-numbered blocks and 13 even-numbered blocks
- Since all adjacent blocks have different parity, we cannot find a path of length 27 ending with an even-numbered block.
- Since the center block is numbered 14, the answer is "no"
- The same argument applies whenever the side of the cube is divided into an odd number of parts, so that the number of blocks is (2n + 1)³ for some n > 0



- A cube is divided into 27 equal smaller cubes (blocks hereafter). Is it possible to trace a path from any point on the surface to the center that always travels parallel to some side of the cube and visits each block exactly once?
- If we number the smaller cubes consecutively, top-down, front to back, and left to right, we find that there are 14 odd-numbered blocks and 13 even-numbered blocks
- Since all adjacent blocks have different parity, we cannot find a path of length 27 ending with an even-numbered block.
- Since the center block is numbered 14, the answer is "no"
- The same argument applies whenever the side of the cube is divided into an odd number of parts, so that the number of blocks is (2n + 1)³ for some n > 0



- A cube is divided into 27 equal smaller cubes (blocks hereafter). Is it possible to trace a path from any point on the surface to the center that always travels parallel to some side of the cube and visits each block exactly once?
- If we number the smaller cubes consecutively, top-down, front to back, and left to right, we find that there are 14 odd-numbered blocks and 13 even-numbered blocks
- Since all adjacent blocks have different parity, we cannot find a path of length 27 ending with an even-numbered block.
- Since the center block is numbered 14, the answer is "no"
- The same argument applies whenever the side of the cube is divided into an odd number of parts, so that the number of blocks is $(2n + 1)^3$ for some n > 0

Introduction Modeling Specifications CTL LTL and CTL* SAT Abstraction
Model Checking the Cube

We encode the problem with two state variables:

$$posn \in \{1, \dots, 27\}$$
 and $visited \subseteq \{1, \dots, 27\}$
 $S = \{1, \dots, 27\} \times 2^{\{1, \dots, 27\}}$
 $S_0 = \{(p, \{p\}) \mid p \in \{1, \dots, 27\} \setminus \{14\}\}$

Let adj(p) return the set of positions adjacent to block p

 $T = \{((p, v), (p', v')) \mid (p, v) \in S \land p' \in \operatorname{adj}(p) \setminus v \land v' = v \setminus \{p'\}\}$

Finally, the property:

$$\varphi = \mathsf{AG}(\operatorname{posn} = 14 \to \operatorname{visited} \neq \{1, \dots, 27\})$$
.

In Verilog, we need 32 binary variables. Out of the 2^{32} states, 3.46426 \cdot 10⁷ are reachable and φ is proved in about one minute Introduction Modeling Specifications CTL LTL and CTL* SAT Abstraction
Model Checking the Cube

We encode the problem with two state variables:

$$ext{posn} \in \{1, \dots, 27\}$$
 and $ext{visited} \subseteq \{1, \dots, 27\}$
 $S = \{1, \dots, 27\} imes 2^{\{1, \dots, 27\}}$
 $S_0 = \{(p, \{p\}) \mid p \in \{1, \dots, 27\} \setminus \{14\}\}$

Let adj(p) return the set of positions adjacent to block p

$$\mathcal{T} = \{((p, v), (p', v')) \mid (p, v) \in S \land p' \in \operatorname{adj}(p) \setminus v \land v' = v \setminus \{p'\}\}$$

Finally, the property:

$$\varphi = AG(posn = 14 \rightarrow visited \neq \{1, \dots, 27\})$$
.

In Verilog, we need 32 binary variables. Out of the 2^{32} states, 3.46426 \cdot 10⁷ are reachable and φ is proved in about one minute Introduction Modeling Specifications CTL LTL and CTL* SAT Abstraction
Model Checking the Cube

We encode the problem with two state variables:

$$ext{posn} \in \{1, \dots, 27\}$$
 and $ext{visited} \subseteq \{1, \dots, 27\}$
 $S = \{1, \dots, 27\} imes 2^{\{1, \dots, 27\}}$
 $S_0 = \{(p, \{p\}) \mid p \in \{1, \dots, 27\} \setminus \{14\}\}$

Let adj(p) return the set of positions adjacent to block p

$$T = \{((p, v), (p', v')) \mid (p, v) \in S \land p' \in \operatorname{adj}(p) \setminus v \land v' = v \setminus \{p'\}\}$$

Finally, the property:

$$\varphi = \mathsf{AG}(\operatorname{posn} = 14 \to \operatorname{visited} \neq \{1, \dots, 27\})$$
.

In Verilog, we need 32 binary variables. Out of the 2^{32} states, 3.46426 \cdot 10⁷ are reachable and φ is proved in about one minute Introduction Modeling Specifications CTL LTL and CTL* SAT Abstraction

Augmenting the Transition Relation

- We can do significantly better by keeping the same state encoding, initial states and property, but augmenting the transition relation
- Let opp(p) returns the set of positions in the cube that have parity opposite to block p. The transition relation can then be defined as before, but with opp(p) replacing adj(p)
- The new relation has many more transitions. For instance, it is now possible to go from Block 1 to Block 6
- The number of reachable states correspondingly increases to $5.41574\cdot 10^8$, but the model checking time decreases to less than a second
- Since φ is universal, the fact that it passes on the model with the augmented transition relation guarantees that it passes also on the original model

 Introduction
 Modeling
 Specifications
 CTL
 LTL and CTL*
 SAT
 Abstraction

 Augmenting the Transition Relation
 Feature
 Feature

- We can do significantly better by keeping the same state
 - encoding, initial states and property, but augmenting the transition relation
 - Let opp(p) returns the set of positions in the cube that have parity opposite to block p. The transition relation can then be defined as before, but with opp(p) replacing adj(p)
 - The new relation has many more transitions. For instance, it is now possible to go from Block 1 to Block 6
 - The number of reachable states correspondingly increases to $5.41574\cdot 10^8$, but the model checking time decreases to less than a second
 - Since φ is universal, the fact that it passes on the model with the augmented transition relation guarantees that it passes also on the original model

 Introduction
 Modeling
 Specifications
 CTL
 LTL and CTL*
 SAT
 Abstraction

 Augmenting the Transition Relation
 Feature
 Feature

- We can do significantly better by keeping the same state encoding, initial states and property, but augmenting the transition relation
 - Let opp(p) returns the set of positions in the cube that have parity opposite to block p. The transition relation can then be defined as before, but with opp(p) replacing adj(p)
 - The new relation has many more transitions. For instance, it is now possible to go from Block 1 to Block 6
 - The number of reachable states correspondingly increases to $5.41574 \cdot 10^8$, but the model checking time decreases to less than a second
 - Since φ is universal, the fact that it passes on the model with the augmented transition relation guarantees that it passes also on the original model

Introduction Modeling Specifications CTL LTL and CTL* SAT Abstraction

- Augmenting the Transition Relation
 - We can do significantly better by keeping the same state encoding, initial states and property, but augmenting the transition relation
 - Let opp(p) returns the set of positions in the cube that have parity opposite to block p. The transition relation can then be defined as before, but with opp(p) replacing adj(p)
 - The new relation has many more transitions. For instance, it is now possible to go from Block 1 to Block 6
 - The number of reachable states correspondingly increases to $5.41574\cdot 10^8$, but the model checking time decreases to less than a second
 - Since φ is universal, the fact that it passes on the model with the augmented transition relation guarantees that it passes also on the original model

Specifications LTL and CTL* Modeling Abstraction

- Augmenting the Transition Relation
 - We can do significantly better by keeping the same state encoding, initial states and property, but augmenting the transition relation
 - Let opp(p) returns the set of positions in the cube that have parity opposite to block p. The transition relation can then be defined as before, but with opp(p) replacing adj(p)
 - The new relation has many more transitions. For instance, it is now possible to go from Block 1 to Block 6
 - The number of reachable states correspondingly increases to $5.41574 \cdot 10^8$, but the model checking time decreases to less than a second
 - Since φ is universal, the fact that it passes on the model with the augmented transition relation guarantees that it passes also on the original model



The proof that there is no path refers to the parity of the block numbers and the total number of blocks of each parity.

parity
$$\in \{0,1\}$$
, visited $0 \in \{1,\ldots,14\}$, visited $E \in \{1,\ldots,13\}$
 $S = \{0,1\} \times \{1,\ldots,14\} \times \{1,\ldots,13\}$
 $S_0 = \{(0,0,1),(1,1,0)\}$

The transition relation is

 $\{((0, v_O, v_E), (1, v_O, v'_E)) \mid (0, v_O, v_E) \in S \land v_E < 13 \land v'_E = v_E + 1\} \cup \{((1, v_O, v_E), (0, v'_O, v_E)) \mid (1, v_O, v_E) \in S \land v_O < 14 \land v'_O = v_O + 1\}$

Transitions are possible to states of opposite parity with one of the two counters incremented (if it has not saturated).

◆□▶ ◆□▶ ◆□▶ ◆□▶ ▲□ ◆ ○ ◆



The proof that there is no path refers to the parity of the block numbers and the total number of blocks of each parity.

parity
$$\in \{0,1\}$$
, visited $0 \in \{1,\ldots,14\}$, visited $E \in \{1,\ldots,13\}$
 $S = \{0,1\} \times \{1,\ldots,14\} \times \{1,\ldots,13\}$
 $S_0 = \{(0,0,1),(1,1,0)\}$

The transition relation is

 $\{ ((0, v_O, v_E), (1, v_O, v'_E)) \mid (0, v_O, v_E) \in S \land v_E < 13 \land v'_E = v_E + 1 \} \cup \\ \{ ((1, v_O, v_E), (0, v'_O, v_E)) \mid (1, v_O, v_E) \in S \land v_O < 14 \land v'_O = v_O + 1 \}$

Transitions are possible to states of opposite parity with one of the two counters incremented (if it has not saturated).

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Changi	ng the St	ate Space				

The property needs to be changed

$$arphi_1 = \mathsf{AG}(\texttt{parity} = 0
ightarrow (\texttt{visited0}
eq 14 \lor \texttt{visitedE}
eq 13))$$

Since parity = 0 corresponds to posn $\in \{2, 4, \dots, 26\}$, our new property actually corresponds to

$$\varphi_2 = \mathsf{AG}(\texttt{posn} \in \{2, 4, \dots, 26\} \rightarrow \texttt{visited} \neq \{1, \dots, 27\})$$

If φ_2 is satisfied on the original model, φ is satisfied as well. The number of binary variables has been reduced from 32 to 9, there are only 53 reachable states, and verification takes negligible time. Interestingly, the following property also holds:

$$arphi_1' = \mathsf{AG}(\texttt{parity} = 0 o \texttt{visited0}
eq 14)$$

◆□▶ ◆□▶ ◆□▶ ◆□▶ ▲□ ◆ ○ ◆

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Changing	the Sta	ate Space				

The property needs to be changed

$$arphi_1 = \mathsf{AG}(\texttt{parity} = 0
ightarrow (\texttt{visited0}
eq 14 \lor \texttt{visitedE}
eq 13))$$

Since parity=0 corresponds to $\texttt{posn}\in\{2,4,\ldots,26\},$ our new property actually corresponds to

$$\varphi_2 = \mathsf{AG}(\texttt{posn} \in \{2, 4, \dots, 26\} \rightarrow \texttt{visited} \neq \{1, \dots, 27\})$$

If φ_2 is satisfied on the original model, φ is satisfied as well. The number of binary variables has been reduced from 32 to 9, there are only 53 reachable states, and verification takes negligible time. Interestingly, the following property also holds:

$$arphi_1' = \mathsf{AG}(\mathsf{parity} = 0 o \mathsf{visited0}
eq 14)$$

◆□▶ ◆□▶ ◆□▶ ◆□▶ ▲□ ◆ ○ ◆

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Changing	the Stat	e Space				

The property needs to be changed

$$arphi_1 = \mathsf{AG}(\texttt{parity} = 0
ightarrow (\texttt{visited0}
eq 14 \lor \texttt{visitedE}
eq 13))$$

Since parity=0 corresponds to $\texttt{posn}\in\{2,4,\ldots,26\},$ our new property actually corresponds to

$$\varphi_2 = \mathsf{AG}(\texttt{posn} \in \{2, 4, \dots, 26\} \rightarrow \texttt{visited} \neq \{1, \dots, 27\})$$

If φ_2 is satisfied on the original model, φ is satisfied as well. The number of binary variables has been reduced from 32 to 9, there are only 53 reachable states, and verification takes negligible time. Interestingly, the following property also holds:

$$arphi_1' = \mathsf{AG}(\texttt{parity} = 0 o \texttt{visited0}
eq 14)$$

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Another	Twist					

In the reachable states of our model, the values of visitedO and visitedE never differ by more than one in absolute value. If we let

```
diff = visitedE - visitedO
```

then we can choose the following encoding:

parity $\in \{0,1\}$ and diff $\in \{-1,0,+1\}$

Each state is a pair consisting of parity and difference:

 $S = \{0, 1\} imes \{-1, 0, +1\}$ $S_0 = \{(0, +1), (1, -1)\}$

The transition relation is

 $\{((0,d),(1,d-1)) \mid d \in \{0,+1\}\} \cup \{((1,d),(0,d+1)) \mid d \in \{-1,0\}\}$

Introduction	Modeling	Specifications	CTL	LTL and CTL*	SAT	Abstraction
Another	Twist					

In the reachable states of our model, the values of visitedO and visitedE never differ by more than one in absolute value. If we let

```
diff = visitedE - visitedO
```

then we can choose the following encoding:

parity $\in \{0,1\}$ and diff $\in \{-1,0,+1\}$

Each state is a pair consisting of parity and difference:

$$S = \{0, 1\} \times \{-1, 0, +1\}$$

 $S_0 = \{(0, +1), (1, -1)\}$

The transition relation is

 $\{((0,d),(1,d{-}1)) \mid d \in \{0,+1\}\} \cup \{((1,d),(0,d{+}1)) \mid d \in \{-1,0\}\}$



Once again, we need to rewrite the property to be checked because the state encoding has changed:

$$arphi_3 = \mathsf{AG}(\texttt{parity} = 0
ightarrow \texttt{diff}
eq -1)$$

Since diff $\neq -1$ implies visitedO $\neq 14 \lor$ visitedE $\neq 13$, we have further strengthened our property.

Since φ_3 holds in our new model, as long as satisfaction of φ_3 in the new, 3-bit model implies satisfaction of φ_3 in the previos, 9-bit model, then we can conclude that φ holds in the original model. There are four reachable states in the 3-bit model, and φ_3 is proved in almost no time.



Once again, we need to rewrite the property to be checked because the state encoding has changed:

$$arphi_3 = \mathsf{AG}(\texttt{parity} = 0
ightarrow \texttt{diff}
eq -1)$$

Since diff $\neq -1$ implies visitedO $\neq 14 \lor$ visitedE $\neq 13$, we have further strengthened our property.

Since φ_3 holds in our new model, as long as satisfaction of φ_3 in the new, 3-bit model implies satisfaction of φ_3 in the previos, 9-bit model, then we can conclude that φ holds in the original model.

◆□▶ ◆□▶ ◆□▶ ◆□▶ ▲□ ◆ ○ ◆

proved in almost no time.



Once again, we need to rewrite the property to be checked because the state encoding has changed:

$$arphi_3 = \mathsf{AG}(\texttt{parity} = 0
ightarrow \texttt{diff}
eq -1)$$

Since diff $\neq -1$ implies visitedO $\neq 14 \lor$ visitedE $\neq 13$, we have further strengthened our property.

Since φ_3 holds in our new model, as long as satisfaction of φ_3 in the new, 3-bit model implies satisfaction of φ_3 in the previos, 9-bit model, then we can conclude that φ holds in the original model. There are four reachable states in the 3-bit model, and φ_3 is proved in almost no time.

◆□▶ ◆□▶ ◆□▶ ◆□▶ ▲□ ◆ ○ ◆



If we use the following encoding:

$$\begin{split} \text{diff} &= -1 \rightarrow \text{diff}_1 = 0 \land \text{diff}_0 = 0\\ \text{diff} &= 0 \rightarrow \text{diff}_1 = 0 \land \text{diff}_0 = 1\\ \text{diff} &= +1 \rightarrow \text{diff}_1 = 1 \land \text{diff}_0 = 0 \end{split}$$

then, over all four reachable states, we have

$$\texttt{diff}_1 \leftrightarrow \texttt{parity} = 0 \land \texttt{diff}_0 = 0$$

We can replace one state bit with a combinational function of the other two. In the resulting model, φ_3 becomes

$$AG(parity = 0 \rightarrow (parity = 0 \lor diff_0 = 1))$$

◆□ > ◆□ > ◆豆 > ◆豆 > 「豆 」 のへで

which simplifies to AG true; this property trivially holds.


If we use the following encoding:

$$\begin{split} \text{diff} &= -1 \rightarrow \text{diff}_1 = 0 \land \text{diff}_0 = 0\\ \text{diff} &= 0 \rightarrow \text{diff}_1 = 0 \land \text{diff}_0 = 1\\ \text{diff} &= +1 \rightarrow \text{diff}_1 = 1 \land \text{diff}_0 = 0 \end{split}$$

then, over all four reachable states, we have

$$diff_1 \leftrightarrow parity = 0 \land diff_0 = 0$$

We can replace one state bit with a combinational function of the other two. In the resulting model, φ_3 becomes

$$\mathsf{AG}(\mathsf{parity} = 0 \rightarrow (\mathsf{parity} = 0 \lor \mathsf{diff}_0 = 1))$$

<日 > < 同 > < 目 > < 目 > < 目 > < 目 > < 0 < 0</p>

which simplifies to AG true; this property trivially holds.