Algorithmic Analysis of a Model Checking Partial Order Reduction Scheme Using Alloy

Stephen F. Siegel

Verified Software Laboratory Department of Computer and Information Sciences University of Delaware Newark, DE, USA

> Subodh Sharma's Class Anywhere December 14, 2020

> > (日) (四) (王) (王) (王) (王)

background: model checking

- background: model checking
- on-the-fly ample set POR
 - Peled, CAV 1994; FMSD 1996
 - Combining partial order reductions with on-the-fly model-checking

▲□▶▲□▶▲□▶▲□▶ □ クタウ

- Holzmann & Peled, FORTE 1994
 - An improvement in formal verification

- background: model checking
- on-the-fly ample set POR
 - Peled, CAV 1994; FMSD 1996
 - Combining partial order reductions with on-the-fly model-checking

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ - 三 - つくぐ

- Holzmann & Peled, FORTE 1994
 - An improvement in formal verification
- ► ~2009: gap in proof

- background: model checking
- on-the-fly ample set POR
 - Peled, CAV 1994; FMSD 1996
 - Combining partial order reductions with on-the-fly model-checking
 - Holzmann & Peled, FORTE 1994
 - An improvement in formal verification
- \blacktriangleright \sim 2009: gap in proof

fix the proof? find a counterexample?



- background: model checking
- on-the-fly ample set POR
 - Peled, CAV 1994; FMSD 1996
 - Combining partial order reductions with on-the-fly model-checking
 - Holzmann & Peled, FORTE 1994
 - An improvement in formal verification
- ► ~2009: gap in proof



▲□▶ ▲□▶ ▲□▶ ▲□▶ □ ● ● ●

- fix the proof? find a counterexample?
- heard talk by Pamela Zave on using Alloy to find defect in Chord

- background: model checking
- on-the-fly ample set POR
 - Peled, CAV 1994; FMSD 1996
 - Combining partial order reductions with on-the-fly model-checking
 - Holzmann & Peled, FORTE 1994
 - An improvement in formal verification
- \blacktriangleright \sim 2009: gap in proof



- fix the proof? find a counterexample?
- heard talk by Pamela Zave on using Alloy to find defect in Chord
- used Alloy to model on-the-fly ample set POR

- background: model checking
- on-the-fly ample set POR
 - Peled, CAV 1994; FMSD 1996
 - Combining partial order reductions with on-the-fly model-checking
 - Holzmann & Peled, FORTE 1994
 - An improvement in formal verification
- \blacktriangleright \sim 2009: gap in proof



- fix the proof? find a counterexample?
- heard talk by Pamela Zave on using Alloy to find defect in Chord
- used Alloy to model on-the-fly ample set POR
- counterexample! (in theory and in Spin)

- background: model checking
- on-the-fly ample set POR
 - Peled, CAV 1994; FMSD 1996
 - Combining partial order reductions with on-the-fly model-checking
 - Holzmann & Peled, FORTE 1994
 - An improvement in formal verification
- \blacktriangleright \sim 2009: gap in proof



- fix the proof? find a counterexample?
- heard talk by Pamela Zave on using Alloy to find defect in Chord
- used Alloy to model on-the-fly ample set POR
- counterexample! (in theory and in Spin)

🕨 a fix

- > an automated approach to the verification of software and hardware systems
 - integrated circuits
 - network protocols
 - software: programs written in C, Java, machine language, ...

- an automated approach to the verification of software and hardware systems
 - integrated circuits
 - network protocols
 - software: programs written in C, Java, machine language, ...
- particularly effective for concurrent systems
 - multithreaded programs, ...

- > an automated approach to the verification of software and hardware systems
 - integrated circuits
 - network protocols
 - software: programs written in C, Java, machine language, ...
- particularly effective for concurrent systems
 - multithreaded programs, ...
- many robust, widely-used model checking tools
 - SPIN has its own input language, PROMELA, highly optimized
 - ▶ Divine (C++), CIVL (C+...), NuSMV, Java PathFinder (Java), ...

1. represent the system/program as a finite-state transition system

- 1. represent the system/program as a finite-state transition system
- 2. specify correctness properties you wish to check
 - example "at most one thread is in the critical section at any time"
 - temporal logic, automata, ...

- 1. represent the system/program as a finite-state transition system
- 2. specify correctness properties you wish to check
 - example "at most one thread is in the critical section at any time"
 - temporal logic, automata, ...
- 3. apply automated algorithmic (graph, automata, ...) techniques to check that all executions of the program satisfy the properties

- typically: explore the reachable states of the program
- simultaneously updating a property automaton
- report a violation as soon as an accepting cycle is detected

Example: mutual exclusion: Promela program

```
bool lock = false;
int ncrit = 0:
active [2] proctype thread() {
  do
  :: atomic { !lock -> lock = true }
     ncrit++;
     // ...this is the critical section...
     ncrit--;
     lock = false
  od
}
```

Example: mutual exclusion: state-transition system



Example: mutual exclusion: state-transition system



- an execution specifies a sequence of states
- a property specifies which state sequences are OK

- an execution specifies a sequence of states
- ► a property specifies which state sequences are OK
- > properties may be specified in natural language, using LTL formulas, or automata

- an execution specifies a sequence of states
- a property specifies which state sequences are OK
- > properties may be specified in natural language, using LTL formulas, or automata
- natural language: "ncrit is always at most 1"

- an execution specifies a sequence of states
- a property specifies which state sequences are OK
- > properties may be specified in natural language, using LTL formulas, or automata

- natural language: "ncrit is always at most 1"
- ▶ linear temporal logic (LTL): $G(ncrit \leq 1)$

- an execution specifies a sequence of states
- a property specifies which state sequences are OK
- properties may be specified in natural language, using LTL formulas, or automata
- natural language: "ncrit is always at most 1"
- ▶ linear temporal logic (LTL): $G(ncrit \leq 1)$
 - **G** p = "henceforth p" ("from now on", "always")
 - ▶ $\mathbf{F} p =$ "eventually p"
 - ▶ $p \mathbf{U} q = "p$ until q"

- an execution specifies a sequence of states
- a property specifies which state sequences are OK
- properties may be specified in natural language, using LTL formulas, or automata

- natural language: "ncrit is always at most 1"
- ▶ linear temporal logic (LTL): $G(ncrit \leq 1)$
 - **G** p = "henceforth p" ("from now on", "always")
 - **F** p = "eventually p"
 - ▶ $p \mathbf{U} q = "p$ until q"
- Büchi automaton (negated) . . .

Büchi automata

A Büchi automaton is a tuple

$$\mathcal{B} = \langle S, \Delta, \Sigma, \delta, F \rangle,$$

where

- 1. S is a finite set of *automaton states*
- 2. $\Delta \subseteq S$ is the set of initial states
- 3. Σ is a finite set called the *alphabet*
- 4. $\delta \subseteq S \times \Sigma \times S$ is the transition relation
- 5. $F \subseteq S$ is the set of *accepting states*.

Büchi automata

A Büchi automaton is a tuple

$$\mathcal{B} = \langle S, \Delta, \Sigma, \delta, F \rangle,$$

where

- 1. S is a finite set of *automaton states*
- 2. $\Delta \subseteq S$ is the set of *initial states*
- 3. Σ is a finite set called the *alphabet*
- 4. $\delta \subseteq S \times \Sigma \times S$ is the transition relation
- 5. $F \subseteq S$ is the set of *accepting states*.

The language of \mathcal{B} , denoted $\mathcal{L}(B)$, is the set of all $\xi \in \Sigma^{\omega}$ generated by infinite paths in \mathcal{B} that pass through an accepting state infinitely often.

(日)

8 S.F.Siegel & Subodh's Class & Algorithmic Analysis of POR Scheme Using Alloy



イロト イポト イヨト イヨト

= nar





"after some b, a occurs before c"



= nar

Using Büchi automata to specify program properties

10 S.F.Siegel & Subodh's Class & Algorithmic Analysis of POR Scheme Using Alloy

<□> <圖> < E> < E> E のQ@

Using Büchi automata to specify program properties

- first, identify the set AP of atomic propositions
 - an atomic proposition is either true or false in each state
 - example: $p = (\texttt{ncrit} \le 1)$

Using Büchi automata to specify program properties

- first, identify the set AP of atomic propositions
 - an atomic proposition is either true or false in each state
 - example: $p = (\texttt{ncrit} \le 1)$
- the alphabet Σ of the Büchi Automaton is 2^{AP}
 - an element of Σ is a set of atomic propositions
 - interpretation: the set of propositions which are true in that state
Using Büchi automata to specify program properties

- first, identify the set AP of atomic propositions
 - an atomic proposition is either true or false in each state
 - example: $p = (\texttt{ncrit} \le 1)$
- \blacktriangleright the alphabet Σ of the Büchi Automaton is $2^{\sf AP}$
 - an element of Σ is a set of atomic propositions
 - interpretation: the set of propositions which are true in that state

• example:
$$AP = \{p\}$$

• $\Sigma = \{\{p\}, \emptyset\}$
• $\{p\}: p \text{ is true}$

•
$$\emptyset$$
: p is false

Using Büchi automata to specify program properties

- first, identify the set AP of atomic propositions
 - an atomic proposition is either true or false in each state
 - example: $p = (\texttt{ncrit} \le 1)$
- the alphabet Σ of the Büchi Automaton is 2^{AP}
 - an element of Σ is a set of atomic propositions
 - interpretation: the set of propositions which are true in that state
- example: $AP = \{p\}$
 - $\blacktriangleright \Sigma = \{\{p\}, \emptyset\}$
 - $\{p\}: p \text{ is true}$
 - \emptyset : p is false
- \blacktriangleright each execution yields a trace, which is an infinite word in Σ

Using Büchi automata to specify program properties

- first, identify the set AP of atomic propositions
 - an atomic proposition is either true or false in each state
 - example: $p = (\texttt{ncrit} \le 1)$
- \blacktriangleright the alphabet Σ of the Büchi Automaton is $2^{\sf AP}$
 - an element of Σ is a set of atomic propositions
 - interpretation: the set of propositions which are true in that state
- example: $AP = \{p\}$
 - $\blacktriangleright \Sigma = \{\{p\}, \emptyset\}$
 - $\{p\}: p \text{ is true}$
 - \emptyset : p is false
- \blacktriangleright each execution yields a trace, which is an infinite word in Σ
- the BA specifies which infinite words violate the property
- example: a trace violating "always $\texttt{ncrit} \leq 1$ ":



(日)

Automated technique: compute the sycnchronous product



11 S.F.Siegel & Subodh's Class & Algorithmic Analysis of POR Scheme Using Alloy

▲□▶▲圖▶▲≣▶▲≣▶ ▲国▼▲○▲

Automated technique: compute the sycnchronous product





Automated technique: compute the sycnchronous product



▶ in the synchronous product, no accepting cycle is reached – try it in SPIN!

12 S.F.Siegel & Subodh's Class & Algorithmic Analysis of POR Scheme Using Alloy

- POR is a model checking optimization
 - POR can dramatically reduce the number of states that need to be explored in order to verify certain properties
 - the yes/no result should always be the same whether or not POR is used

- POR is a model checking optimization
 - POR can dramatically reduce the number of states that need to be explored in order to verify certain properties
 - the yes/no result should always be the same whether or not POR is used
- POR can only be applied when the property is stutter invariant
 - let $\mathcal{L} \subset \Sigma^{\omega}$ (a set of infinite words over Σ)
 - \mathcal{L} is stutter invariant if for all $a_1, a_2, \ldots \in \Sigma$, $i_1, i_2, \ldots \geq 1$:

$$a_1 a_2 \dots \in \mathcal{L} \iff a_1^{i_1} a_2^{i_2} \dots \in \mathcal{L}$$

- POR is a model checking optimization
 - POR can dramatically reduce the number of states that need to be explored in order to verify certain properties
 - the yes/no result should always be the same whether or not POR is used
- POR can only be applied when the property is stutter invariant
 - let $\mathcal{L} \subset \Sigma^{\omega}$ (a set of infinite words over Σ)
 - \mathcal{L} is stutter invariant if for all $a_1, a_2, \ldots \in \Sigma$, $i_1, i_2, \ldots \geq 1$:

$$a_1 a_2 \dots \in \mathcal{L} \iff a_1^{i_1} a_2^{i_2} \dots \in \mathcal{L}$$

- hence the property is invariant under repeating an entry (stuttering)
- > and invariant under removing stutters of an entry (but leave at least one occurrence)

- POR is a model checking optimization
 - POR can dramatically reduce the number of states that need to be explored in order to verify certain properties
 - the yes/no result should always be the same whether or not POR is used
- POR can only be applied when the property is stutter invariant
 - let $\mathcal{L} \subset \Sigma^{\omega}$ (a set of infinite words over Σ)
 - \mathcal{L} is stutter invariant if for all $a_1, a_2, \ldots \in \Sigma$, $i_1, i_2, \ldots \geq 1$:

$$a_1 a_2 \dots \in \mathcal{L} \iff a_1^{i_1} a_2^{i_2} \dots \in \mathcal{L}$$

- hence the property is invariant under repeating an entry (stuttering)
- > and invariant under removing stutters of an entry (but leave at least one occurrence)
- all properties we have seen so far are stutter invariant

- POR is a model checking optimization
 - POR can dramatically reduce the number of states that need to be explored in order to verify certain properties
 - the yes/no result should always be the same whether or not POR is used
- POR can only be applied when the property is stutter invariant
 - let $\mathcal{L} \subset \Sigma^{\omega}$ (a set of infinite words over Σ)
 - \mathcal{L} is stutter invariant if for all $a_1, a_2, \ldots \in \Sigma$, $i_1, i_2, \ldots \geq 1$:

$$a_1 a_2 \dots \in \mathcal{L} \iff a_1^{i_1} a_2^{i_2} \dots \in \mathcal{L}$$

- hence the property is invariant under repeating an entry (stuttering)
- > and invariant under removing stutters of an entry (but leave at least one occurrence)
- all properties we have seen so far are stutter invariant
- ▶ all properties expressible with LTL operators F, G, U are stutter invariant

- POR is a model checking optimization
 - POR can dramatically reduce the number of states that need to be explored in order to verify certain properties
 - the yes/no result should always be the same whether or not POR is used
- POR can only be applied when the property is stutter invariant
 - let $\mathcal{L} \subset \Sigma^{\omega}$ (a set of infinite words over Σ)
 - \mathcal{L} is stutter invariant if for all $a_1, a_2, \ldots \in \Sigma$, $i_1, i_2, \ldots \geq 1$:

$$a_1 a_2 \dots \in \mathcal{L} \iff a_1^{i_1} a_2^{i_2} \dots \in \mathcal{L}$$

- hence the property is invariant under repeating an entry (stuttering)
- > and invariant under removing stutters of an entry (but leave at least one occurrence)
- all properties we have seen so far are stutter invariant
- ▶ all properties expressible with LTL operators F, G, U are stutter invariant
- almost any property you would ever want to express is stutter invariant

Partial Order Reduction: intuition

14 S.F.Siegel & Subodh's Class & Algorithmic Analysis of POR Scheme Using Alloy

<□> <圖> < E> < E> E のQ@

Partial Order Reduction: intuition

- ▶ in a parallel program, an action α on local variables commutes with all actions β from other threads
 - the order of the two actions doesn't matter
 - you end up in the same state either way
 - \blacktriangleright if the local action α can not change the value of any atomic proposition...
 - ... you only need to explore one of the two orders!

Partial Order Reduction: intuition

 α

 L_1

- \blacktriangleright in a parallel program, an action α on local variables commutes with all actions β from other threads
 - the order of the two actions doesn't matter
 - you end up in the same state either way
 - if the local action α can not change the value of any atomic proposition...
 - ... you only need to explore one of the two orders!
 - one path yields trace $\cdots L_1 L_1 L_2 \cdots$
 - ▶ other path yields trace $\cdots L_1L_2L_2\cdots$
 - stutter equivalent!
 - either both paths satisfy the property or both violate

 L_2

- A program consists of
 - \blacktriangleright a set Q of states
 - \blacktriangleright an initial state $\iota \in Q$
 - \blacktriangleright a set T of operations
 - ▶ each $\alpha \in T$ is a (deterministic) function $\alpha : en_{\alpha} \rightarrow Q$, for some $en_{\alpha} \subseteq Q$

- A program consists of
 - \blacktriangleright a set Q of states
 - \blacktriangleright an initial state $\iota \in Q$
 - \blacktriangleright a set T of operations
 - ▶ each $\alpha \in T$ is a (deterministic) function $\alpha : en_{\alpha} \rightarrow Q$, for some $en_{\alpha} \subseteq Q$



- A program consists of
 - \blacktriangleright a set Q of states
 - \blacktriangleright an initial state $\iota \in Q$
 - ▶ a set T of operations
 - ▶ each $\alpha \in T$ is a (deterministic) function $\alpha : en_{\alpha} \rightarrow Q$, for some $en_{\alpha} \subseteq Q$



▶ for $q \in Q$, define $en(q) = \{\alpha \in T \mid q \in en_{\alpha}\}$

Program Traces

- Iet AP be a set of atomic propositions
- ▶ an interpretation is a function $L: Q \to \Sigma = 2^{\mathsf{AP}}$

Program Traces

- Iet AP be a set of atomic propositions
- ▶ an interpretation is a function $L: Q \to \Sigma = 2^{\mathsf{AP}}$
- $\mathcal{L}(P) \subseteq \Sigma^{\omega}$ is the set of infinite traces

 $L(q_0)L(q_1)\cdots$

where $\iota = q_0 \rightarrow q_1 \rightarrow \cdots$ is an infinite initial path in P

Program Traces

- Iet AP be a set of atomic propositions
- ▶ an interpretation is a function $L: Q \to \Sigma = 2^{\mathsf{AP}}$
- $\mathcal{L}(P) \subseteq \Sigma^{\omega}$ is the set of infinite traces

$$L(q_0)L(q_1)\cdots$$

where $\iota = q_0 \rightarrow q_1 \rightarrow \cdots$ is an infinite initial path in P



Properties

 \blacktriangleright a property is a subset of Σ^{ω}

Properties

- \blacktriangleright a property is a subset of Σ^{ω}
- ▶ a Büchi automaton \mathcal{B} with alphabet Σ specifies the property $\Sigma^{\omega} \setminus \mathcal{L}(\mathcal{B})$
 - ▶ the program violates the property if $\mathcal{L}(P) \cap \mathcal{L}(\mathcal{B}) \neq \varnothing$

Properties

- \blacktriangleright a property is a subset of Σ^{ω}
- ▶ a Büchi automaton \mathcal{B} with alphabet Σ specifies the property $\Sigma^{\omega} \setminus \mathcal{L}(\mathcal{B})$
 - ▶ the program violates the property if $\mathcal{L}(P) \cap \mathcal{L}(\mathcal{B}) \neq \emptyset$



• violation: $\varnothing \varnothing \varnothing \{p\}^{\omega} \in \mathcal{L}(P) \cap \mathcal{L}(\mathcal{B})$

Given program $P = (Q, T, \iota)$ and Büchi automaton \mathcal{B} with states S: The product automaton $P \otimes \mathcal{B}$ is the Büchi automaton

▶ alphabet: $T \times \Sigma$, states: $Q \times S$, (q, s) accepting $\Leftrightarrow s$ accepting

Given program $P = (Q, T, \iota)$ and Büchi automaton \mathcal{B} with states S: The product automaton $P \otimes \mathcal{B}$ is the Büchi automaton

- ▶ alphabet: $T \times \Sigma$, states: $Q \times S$, (q, s) accepting $\Leftrightarrow s$ accepting
- transitions:

Given program $P = (Q, T, \iota)$ and Büchi automaton \mathcal{B} with states S: The product automaton $P \otimes \mathcal{B}$ is the Büchi automaton

- ▶ alphabet: $T \times \Sigma$, states: $Q \times S$, (q, s) accepting $\Leftrightarrow s$ accepting
- transitions:

 $\mathcal{L}(P \otimes \mathcal{B}) = \mathcal{L}(P) \cap \mathcal{L}(\mathcal{B})$

Given program $P = (Q, T, \iota)$ and Büchi automaton \mathcal{B} with states S: The product automaton $P \otimes \mathcal{B}$ is the Büchi automaton

- ▶ alphabet: $T \times \Sigma$, states: $Q \times S$, (q, s) accepting $\Leftrightarrow s$ accepting
- transitions:



18 S.F.Siegel & Subodh's Class & Algorithmic Analysis of POR Scheme Using Alloy

- $\blacktriangleright \mathcal{A} = P \otimes \mathcal{B}$
- ▶ explore a sub-automaton $\mathcal{A}' \subseteq \mathcal{A}$ such that $\mathcal{L}(\mathcal{A}') = \emptyset \Leftrightarrow \mathcal{L}(A) = \emptyset$

- $\blacktriangleright \mathcal{A} = P \otimes \mathcal{B}$
- ▶ explore a sub-automaton $\mathcal{A}' \subseteq \mathcal{A}$ such that $\mathcal{L}(\mathcal{A}') = \varnothing \Leftrightarrow \mathcal{L}(A) = \varnothing$
- ▶ "offline" POR: $A' = P' \otimes B$. "on-the-fly" POR: not necessarily

- $\blacktriangleright \mathcal{A} = P \otimes \mathcal{B}$
- ▶ explore a sub-automaton $\mathcal{A}' \subseteq \mathcal{A}$ such that $\mathcal{L}(\mathcal{A}') = \varnothing \Leftrightarrow \mathcal{L}(A) = \varnothing$
- "offline" POR: $\mathcal{A}' = P' \otimes \mathcal{B}$. "on-the-fly" POR: not necessarily
- on-the-fly ample set POR
 - $\blacktriangleright \ \, {\rm given} \ \, (q,s) \in Q \times S \text{, specify } {\rm amp}(q,s) \subseteq {\rm en}(q)$

 $\blacktriangleright \mathcal{A} = P \otimes \mathcal{B}$

- ▶ explore a sub-automaton $\mathcal{A}' \subseteq \mathcal{A}$ such that $\mathcal{L}(\mathcal{A}') = \varnothing \Leftrightarrow \mathcal{L}(A) = \varnothing$
- "offline" POR: $\mathcal{A}' = P' \otimes \mathcal{B}$. "on-the-fly" POR: not necessarily
- on-the-fly ample set POR

Dependence and visibility

Definition

Two operations $\alpha,\beta\in T$ are independent if for

- all $q \in en_{\alpha} \cap en_{\beta}$:
 - 1. $\alpha(q) \in \mathsf{en}_{\beta}$
 - 2. $\beta(q) \in \mathsf{en}_{\alpha}$ and
 - **3**. $\alpha(\beta(q)) = \beta(\alpha(q)).$



Definition

An operation $\alpha \in T$ is invisible if, for all $q \in en_{\alpha}$: $L(q) = L(\alpha(q)).$



Conditions on the ample sets
C0 For all $q \in Q$, $s \in S$: $en(q) \neq \emptyset \implies amp(q, s) \neq \emptyset$

- $\textbf{C0} \ \text{For all} \ q \in Q \text{, } s \in S \text{: } \textbf{en}(q) \neq \emptyset \implies \texttt{amp}(q,s) \neq \emptyset$
- **C1** For all $q \in Q$, $s \in S$, and paths in P starting from q:
 - \blacktriangleright no operation dependent on an operation in $\mathsf{amp}(q,s)$ can occur before some operation in $\mathsf{amp}(q,s)$ occurs

 $\textbf{C0} \ \text{For all } q \in Q \text{, } s \in S \text{: } \textbf{en}(q) \neq \emptyset \implies \texttt{amp}(q,s) \neq \emptyset$

C1 For all $q \in Q$, $s \in S$, and paths in P starting from q:

 \blacktriangleright no operation dependent on an operation in $\mathsf{amp}(q,s)$ can occur before some operation in $\mathsf{amp}(q,s)$ occurs

C2 For all $q \in Q$, $s \in S$:

• if $amp(q, s) \neq en(q)$ then $\forall \alpha \in amp(q, s)$, α is invisible

 $\textbf{C0} \ \text{For all } q \in Q \text{, } s \in S \text{: } \textbf{en}(q) \neq \emptyset \implies \texttt{amp}(q,s) \neq \emptyset$

C1 For all $q \in Q$, $s \in S$, and paths in P starting from q:

 \blacktriangleright no operation dependent on an operation in $\mathsf{amp}(q,s)$ can occur before some operation in $\mathsf{amp}(q,s)$ occurs

C2 For all $q \in Q$, $s \in S$:

- ▶ if $amp(q, s) \neq en(q)$ then $\forall \alpha \in amp(q, s)$, α is invisible
- C3 [weaker version] In any cycle in \mathcal{A}' :
 - \blacktriangleright there is a transition from $\langle q,s\rangle$ to $\langle q',s'\rangle$ for which ${\rm amp}(q,s')={\rm en}(q)$

 $\textbf{C0} \ \text{For all } q \in Q \text{, } s \in S \text{: } \textbf{en}(q) \neq \emptyset \implies \texttt{amp}(q,s) \neq \emptyset$

C1 For all $q \in Q$, $s \in S$, and paths in P starting from q:

 \blacktriangleright no operation dependent on an operation in $\mathsf{amp}(q,s)$ can occur before some operation in $\mathsf{amp}(q,s)$ occurs

C2 For all $q \in Q$, $s \in S$:

- ▶ if $amp(q, s) \neq en(q)$ then $\forall \alpha \in amp(q, s)$, α is invisible
- C3 [weaker version] In any cycle in \mathcal{A}' :
 - \blacktriangleright there is a transition from $\langle q,s\rangle$ to $\langle q',s'\rangle$ for which ${\rm amp}(q,s')={\rm en}(q)$

Combined Theorem:

If $\mathcal{L}(\mathcal{B})$ is stutter-invariant and **C0–C3** hold, then $\mathcal{L}(A') = \emptyset \Leftrightarrow \mathcal{L}(A) = \emptyset$.

Counterexample to combined theorem



- **C0** For all $q \in Q$, $s \in S$: $en(q) \neq \emptyset \implies amp(q, s) \neq \emptyset$
- **C1** For all $q \in Q$, $s \in S$, and paths in P starting from q:
 - \blacktriangleright no operation dependent on an operation in $\mathsf{amp}(q,s)$ can occur before some operation in $\mathsf{amp}(q,s)$ occurs
- **C2** For all $q \in Q$, $s \in S$: if $amp(q, s) \neq en(q)$ then $\forall \alpha \in amp(q, s)$, α is invisible
- C3 In any cycle in \mathcal{A}' : there is a transition from $\langle q, s \rangle$ to $\langle q', s' \rangle$ for which $\operatorname{amp}(q, s') = \operatorname{en}(q)$

Spin

```
bit p = 0;
active proctype p0() { p=1 }
active proctype p1() { bit x=0; do :: x=0 od }
never {
   B0: do :: !p :: !p -> break od
   accept_B1: do :: p od
}
```



Spin

. . .

```
bit p = 0;
active proctype p0() { p=1 }
active proctype p1() { bit x=0; do :: x=0 od }
never {
   B0: do :: !p :: !p -> break od
   accept_B1: do :: p od
}
```



\$ spin -a test1.pml; cc -o pan -DNOREDUCE pan.c; ./pan -a -n
pan:1: acceptance cycle (at depth 2)

Spin

. . .

```
bit p = 0;
active proctype p0() { p=1 }
active proctype p1() { bit x=0; do :: x=0 od }
never {
 BO: do :: !p :: !p -> break od
  accept_B1: do :: p od
}
```



```
$ spin -a test1.pml; cc -o pan -DNOREDUCE pan.c; ./pan -a -n
pan:1: acceptance cycle (at depth 2)
```

```
$ spin -a test1.pml; cc -o pan pan.c; ./pan -a -n
warning: for p.o. reduction to be valid the never claim must be stutter-invariant
(never claims generated from LTL formulae are stutter-invariant)
+ Partial Order Reduction
State-vector 36 byte, depth reached 2, errors: 0
       3 states, stored
       2 states, matched
       5 transitions (= stored+matched)
```

S.F.Siegel & Subodh's Class & Algorithmic Analysis of POR Scheme Using Alloy 23

Alloy

- specification language based on relational/first order logic
- a model specifies "signatures" (sets), relations, constraints



Alloy

- specification language based on relational/first order logic
- a model specifies "signatures" (sets), relations, constraints

The Alloy Analyzer



given

- an Alloy model
- an assertion
- an upper bound on the size of each signature

produces either

- "assertion may hold"
 - it holds within the bounds, or
 - "assertion is violated"





- constructed Alloy model of on-the-fly ample-set POR
- 71 lines of Alloy code (reproduced in paper in entirety)

- constructed Alloy model of on-the-fly ample-set POR
- 71 lines of Alloy code (reproduced in paper in entirety)
- modeled in complete generality
 - program, operation, independence, invisiblity, Büchi automaton, ample set, product automaton, language emptiness

 $\blacktriangleright \text{ asserted } \mathcal{L}(A') = \varnothing \Leftrightarrow \mathcal{L}(A) = \varnothing$

- constructed Alloy model of on-the-fly ample-set POR
- > 71 lines of Alloy code (reproduced in paper in entirety)
- modeled in complete generality
 - program, operation, independence, invisibility, Büchi automaton, ample set, product automaton, language emptiness
- $\blacktriangleright \text{ asserted } \mathcal{L}(A') = \varnothing \Leftrightarrow \mathcal{L}(A) = \varnothing$
- could not find a way to model stutter-invariance in Alloy
 - resorted to specifying concrete Buchi automata for most experiments

- constructed Alloy model of on-the-fly ample-set POR
- 71 lines of Alloy code (reproduced in paper in entirety)
- modeled in complete generality
 - program, operation, independence, invisibility, Büchi automaton, ample set, product automaton, language emptiness
- $\blacktriangleright \text{ asserted } \mathcal{L}(A') = \varnothing \Leftrightarrow \mathcal{L}(A) = \varnothing$
- could not find a way to model stutter-invariance in Alloy
 - resorted to specifying concrete Buchi automata for most experiments
- counterexample shown earlier
 - I specified the Büchi automaton
 - Alloy found the (minimal) program and ample sets

- constructed Alloy model of on-the-fly ample-set POR
- ▶ 71 lines of Alloy code (reproduced in paper in entirety)
- modeled in complete generality
 - program, operation, independence, invisibility, Büchi automaton, ample set, product automaton, language emptiness
- $\blacktriangleright \text{ asserted } \mathcal{L}(A') = \varnothing \Leftrightarrow \mathcal{L}(A) = \varnothing$
- could not find a way to model stutter-invariance in Alloy
 - resorted to specifying concrete Buchi automata for most experiments
- counterexample shown earlier
 - I specified the Büchi automaton
 - Alloy found the (minimal) program and ample sets

- make the ample set a function of the source product state
 - not of the intermediate state

make the ample set a function of the source product state

not of the intermediate state

Define \mathcal{A}' :

 $\alpha \in \mathsf{amp}(q,s)$

make the ample set a function of the source product state

not of the intermediate state

Define \mathcal{A}' :

 $\alpha \in \mathsf{amp}(q,s)$

C0–C2: (same) **C3**: In any cycle in \mathcal{A}' : \blacktriangleright there is a state $\langle q, s \rangle$ for which $\operatorname{amp}(q, s) = \operatorname{en}(q)$

make the ample set a function of the source product state

not of the intermediate state

Define A':

 $\alpha \in \mathsf{amp}(q,s)$

C0–C2: (same) **C3**: In any cycle in \mathcal{A}' : • there is a state $\langle q, s \rangle$ for which amp(q, s) = en(q)

minor changes to the Alloy model

q

make the ample set a function of the source product state

not of the intermediate state

Define \mathcal{A}' :

 $\alpha \in \mathsf{amp}(q,s)$

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ - 三 - のへで

C0–C2: (same) **C3**: In any cycle in \mathcal{A}' :

▶ there is a state $\langle q, s \rangle$ for which amp(q, s) = en(q)

minor changes to the Alloy model

Alloy Analyzer: previous Büchi automaton now OK! S.F.Siegel & Subodh's Class & Algorithmic Analysis of POR Scheme Using Alloy

New counterexample



Note: only A0 and A2 are not fully enabled S.F.Siegel \diamond Subodh's Class \diamond Algorithmic Analysis of POR Scheme Using Alloy

▲□▶▲圖▶▲≣▶▲≣▶ ≣ のQで

- define a "stutter-invariant normal form" for Büchi automata
- every BA with a SI language can be transformed into SI normal form
- similar ideas in literature
 - Etessami 1999 (Muller automata)
 - Holzmann & Kupferman 1997 (Büchi transition system)

- define a "stutter-invariant normal form" for Büchi automata
- every BA with a SI language can be transformed into SI normal form
- similar ideas in literature
 - Etessami 1999 (Muller automata)
 - Holzmann & Kupferman 1997 (Büchi transition system)

Consequence of normal form:

1. If $s_1 \xrightarrow{a} s_2 \xrightarrow{b} s_3$ is a path in \mathcal{B} , then for some $s'_2 \in S$, $s_1 \xrightarrow{a} s_2 \xrightarrow{a} s'_2 \xrightarrow{b} s_3$ is a path in \mathcal{B} .

2. If $s_1 \xrightarrow{a} s_2 \xrightarrow{a} s_3 \xrightarrow{b} s_4$ is a path in \mathcal{B} , then $s_1 \xrightarrow{a} s_2 \xrightarrow{b} s_4$ is a path in \mathcal{B} . Moreover, if s_3 is accepting, then s_2 is accepting.

- define a "stutter-invariant normal form" for Büchi automata
- every BA with a SI language can be transformed into SI normal form
- similar ideas in literature
 - Etessami 1999 (Muller automata)
 - Holzmann & Kupferman 1997 (Büchi transition system)

Consequence of normal form:

- 1. If $s_1 \xrightarrow{a} s_2 \xrightarrow{b} s_3$ is a path in \mathcal{B} , then for some $s'_2 \in S$, $s_1 \xrightarrow{a} s_2 \xrightarrow{a} s'_2 \xrightarrow{b} s_3$ is a path in \mathcal{B} .
- 2. If $s_1 \xrightarrow{a} s_2 \xrightarrow{a} s_3 \xrightarrow{b} s_4$ is a path in \mathcal{B} , then $s_1 \xrightarrow{a} s_2 \xrightarrow{b} s_4$ is a path in \mathcal{B} . Moreover, if s_3 is accepting, then s_2 is accepting.
- this is all you need to make on-the-fly POR work

A Fix: Stutter-Invariant Normal Form

Theorem

If \mathcal{B} is in SI normal form and CO–C3 holds, then $\mathcal{L}(A') = \emptyset \Leftrightarrow \mathcal{L}(A) = \emptyset$.

- ► a (LATEX) proof is in the paper
- proof goes through for original and variant schemes
- verified for Alloy for small scope [thanks to reviewers]
 - $\blacktriangleright \leq 3$ symbols in alphabet Σ
 - \blacktriangleright ≤ 4 states in the Büchi automaton
 - \blacktriangleright ≤ 6 transitions in Büchi automaton
 - \blacktriangleright ≤ 3 operations in program
 - ► ≤ 4 program states
 - \blacktriangleright ≤ 16 product states
 - time: 2,265 seconds

- 1. a theorem combining on-the-fly model checking and POR is incorrect
 - a theoretical counterexample
 - ▶ for same reason, Spin can declare an erroneous program correct

- 1. a theorem combining on-the-fly model checking and POR is incorrect
 - a theoretical counterexample
 - for same reason, Spin can declare an erroneous program correct
- 2. but can be fixed by restricting to certain normalized Büchi automata
 - what other classes of Büchi automata are OK?

- 1. a theorem combining on-the-fly model checking and POR is incorrect
 - a theoretical counterexample
 - for same reason, Spin can declare an erroneous program correct
- 2. but can be fixed by restricting to certain normalized Büchi automata
 - what other classes of Büchi automata are OK?
- 3. Alloy is an effective tool for checking POR algorithms
 - models can be constructed with little effort
 - can find counterexamples quickly
 - can verify correctness within small scopes

- 1. a theorem combining on-the-fly model checking and POR is incorrect
 - a theoretical counterexample
 - for same reason, Spin can declare an erroneous program correct
- 2. but can be fixed by restricting to certain normalized Büchi automata
 - what other classes of Büchi automata are OK?
- 3. Alloy is an effective tool for checking POR algorithms
 - models can be constructed with little effort
 - can find counterexamples quickly
 - can verify correctness within small scopes
- 4. but even better would be mechanized proofs of correctness
 - see Brunner & Lammich, 2018
 - Formal verification of an executable LTL model checker with partial order reduction, J. Automated Reasoning

- 1. a theorem combining on-the-fly model checking and POR is incorrect
 - a theoretical counterexample
 - for same reason, Spin can declare an erroneous program correct
- 2. but can be fixed by restricting to certain normalized Büchi automata
 - what other classes of Büchi automata are OK?
- 3. Alloy is an effective tool for checking POR algorithms
 - models can be constructed with little effort.
 - can find counterexamples quickly
 - can verify correctness within small scopes
- 4. but even better would be mechanized proofs of correctness
 - see Brunner & Lammich, 2018
 - Formal verification of an executable LTL model checker with partial order reduction. J. Automated Reasoning
- 5. perhaps the best approach
 - use Alloy to prototype (quick iterations)
 - once Alloy says everything is OK, invest in a full mechanized proof (日)

S.F.Siegel & Subodh's Class & Algorithmic Analysis of POR Scheme Using Alloy 30