Lecture 11
Counterexamples + Bisimulation

Dr. Dave Parker

Department of Computer Science
University of Oxford

Overview

- Counterexamples
  - non-probabilistic model checking
  - counterexamples for PCTL + DTMCs
  - computing smallest counterexamples

- Bisimulation
  - bisimulation equivalences: DTMCs, CTMCs
  - preservation of logics: PCTL, CSL
  - bisimulation minimisation

Non probabilistic counterexamples

- Counterexamples (for non-probabilistic model checking)
  - generated when model checking a (universal) property fails
  - trace through model illustrating why property does not hold
  - major advantage of the model checking approach
  - bug finding vs. verification
- Example:
  - CTL property AG ¬err
  - (or equivalently, ¬EF err)
  - ("an error state is never reached")
  - counterexample is a finite trace to a state satisfying err
  - alternatively, this is a witness to the satisfaction of formula EF err

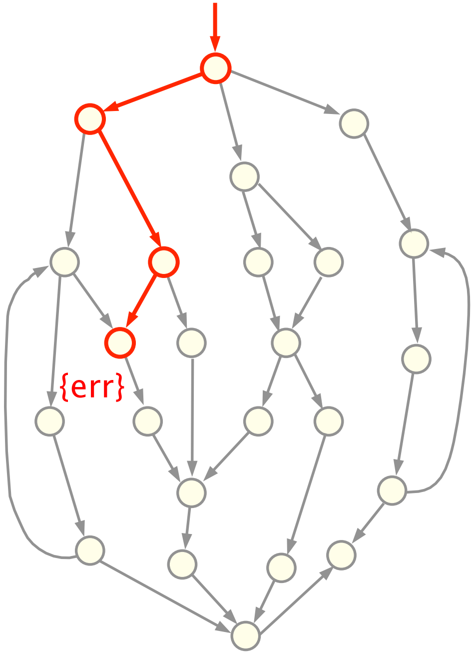

Counterexamples for DTMCs?

- PCTL example: $P_{<0.01}$ [F err]
  - "the probability of reaching an error state is less than 0.01"
  - what is a counterexample for $s \not\models P_{<0.01}$ [F err] ?
  - not necessarily illustrated by a single trace to an err state
  - in fact, "counterexample" is a set of paths satisfying F err whose combined measure is greater than or equal to 0.01
- Alternative approach to "debugging" seen so far:
  - probabilistic model checker provides actual probabilities
  - e.g. queries of the form $P_{=?}$ [F err]
  - anomalous behaviour identified by examining trends
  - e.g. $P_{=?}$ [$F^{\leq T}$ err] for T=0,…,100
- This lecture: DTMC counterexamples in style of [HK07]
  - also some work done on CTMC/MDP counterexamples

DTMC notation

- DTMC: $D = (S, s_{init}, \mathbf{P}, L)$
- $Path(s)$ = set of all infinite paths starting in state s
- $Pr_s : \Sigma_{Path(s)} \rightarrow [0,1]$ = probability measure over infinite paths
  - where $\Sigma_{Path(s)}$ is the σ-algebra on Path(s)
  - defined in terms of probabilities for finite paths
- $\mathbf{P}_s(\omega)$ = probability for finite path $\omega = ss_1 \ldots s_n$
  - $\mathbf{P}_s(s) = 1$
  - $\mathbf{P}_s(ss_1 \ldots s_n) = \mathbf{P}(s,s_1) \cdot \mathbf{P}(s_1,s_2) \cdot \ldots \cdot \mathbf{P}(s_{n-1},s_n)$
  - extend notation to sets: $\mathbf{P}_s(C)$ for set of finite paths C
  - $\mathbf{P}_s$ extends uniquely to Pr_s
- $Path(s, \psi) = \{ \omega \in Path(s) \mid \omega \vDash \psi \}$
  - $Prob(s, \psi) = Pr_s(Path(s, \psi))$
- $Path_{fin}(s, \psi)$ = set of finite paths from s satisfying ψ

Counterexamples for DTMCs

- Consider PCTL properties of the form:
  - $P_{\leq p}$ [Φ_1 $U^{\leq k}$ Φ_2], where $k \in \mathbb{N} \cup \{\infty\}$
  - i.e. bounded or unbounded until formulae with closed upper probability bounds

- Refutation:
  - $s \not\models P_{\leq p}$ [Φ_1 $U^{\leq k}$ Φ_2]
  - $\Leftrightarrow Pr_s(Path(s, \Phi_1\ U^{\leq k}\ \Phi_2)) > p$
  - i.e. total probability mass of Φ_1 $U^{\leq k}$ Φ_2 paths exceeds p

- Since the property is an until formula
  - this is evidenced by a set of finite paths

Counterexamples for DTMCs

- A counterexample for $P_{\leq p}$ [Φ_1 $U^{\leq k}$ Φ_2] in state s is:
  - a set C of finite paths such that $C \subseteq Path_{fin}(s, \psi)$ and $P_s(C) > p$

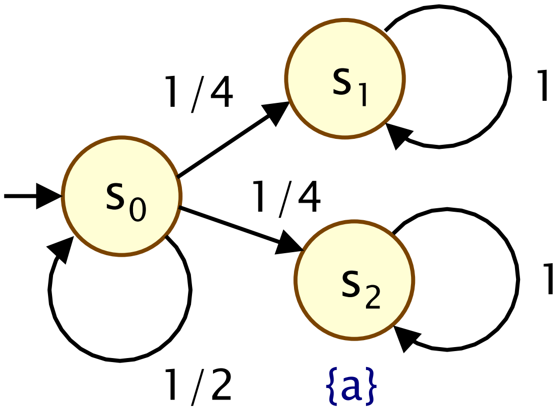

- Example
  - Consider the PCTL formula:
  - $P_{\leq 0.3}$ [F a]
  - This is not satisfied in s_0
  - $Prob(s_0, F\ a) = 1/4+1/8+1/16+\ldots = 1/2$
  - A counterexample: $C = \{ s_0s_2, s_0s_0s_2 \}$
  - $P_{s0}(C) = 1/4 + (1/2)(1/4) = 3/8 = 0.375$

Finiteness of counterexamples

- There is always a finite counterexample for:
  - $s \not\models P_{\leq p} [\ \Phi_1\ U^{\leq k}\ \Phi_2\ ]$
- On the other hand, consider this DTMC:
  - and the PCTL formula:
  - $P_{<1/2} [\ F\ a\ ]$
  - $Prob(s_0, F\ a) = 1/4+1/8+1/16+\ldots$
    $= 1/2$
  - $s_0 \not\models P_{<1/2} [\ F\ a\ ]$
  - counterexample would require infinite set of paths
  - $\{\ (s_0)^i s_2\ \}_{i\in\mathbb{N}}$

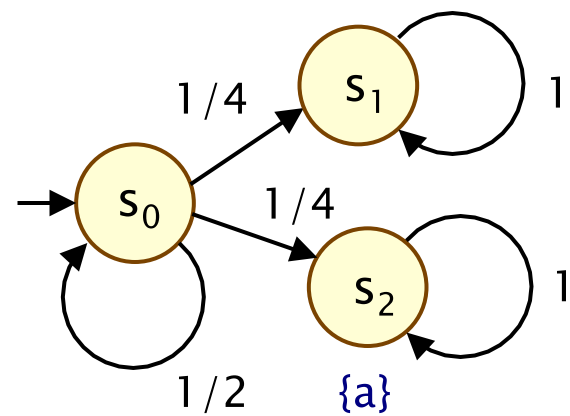

Counterexamples for DTMCs

- Aim: counterexamples should be succinct, comprehensible
- Set of all counterexamples:
  - $CX_p(s,\psi)$ = set of all counterexamples for $P_{\leq p}$ [ψ] in state s
- Minimal counterexample
  - counterexample C with $|C| \leq |C'|$ for all $C' \in CX_p(s,\psi)$
- "Smallest" counterexample
  - minimal counterexample C with $\mathbf{P}(C) \geq \mathbf{P}(C')$ for all minimal $C' \in CX_p(s,\psi)$
  - reduces to finding…
- Strongest (most probable) evidence
  - finite path ω in $Path_{fin}(s, \psi)$ such that $\mathbf{P}(\omega) \geq \mathbf{P}(\omega')$ for all $\omega' \in Path_{fin}(s, \psi)$
  - i.e. contributes most to violation of PCTL formula

Example

- PCTL formula: $P_{\leq 1/2}$ [F b]
  - $s_0 \not\models P_{\leq 1/2}$ [F b]
  - since Prob(s_0, F b) = 0.9

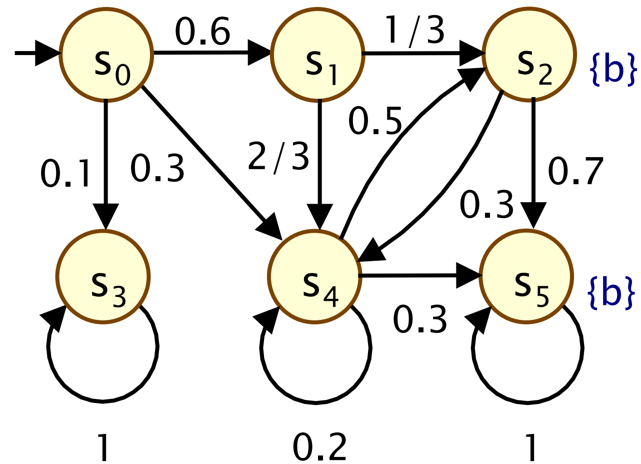

- Counterexamples:
  - $C_1 = \{ s_0s_1s_2, s_0s_1s_4s_2, s_0s_1s_4s_5, s_0s_4s_2 \}$
    - $P_{s0}(C_1) = 0.2+0.2+0.12+0.15 = 0.67$ (not minimal)
  - $C_2 = \{ s_0s_1s_2, s_0s_1s_4s_2, s_0s_1s_4s_5 \}$
    - $P_{s0}(C_2) = 0.2+0.2+0.12 = 0.52$ (not "smallest")
  - $C_3 = \{ s_0s_1s_2, s_0s_1s_4s_2, s_0s_4s_2 \}$
    - $P_{s0}(C_3) = 0.2+0.2+0.15 = 0.55$ ("smallest")

Weighted digraphs

- A weighted directed graph is a tuple $G = (V, E, w)$ where:
  - V is a set of vertices
  - $E \subseteq V \times V$ is a set of edges
  - $w : E \to \mathbb{R}_{\geq 0}$ is a weight function

- Finite path ω in G
  - is a sequence of vertices $v_0 v_1 v_2 \ldots v_n$ such that $(v_i, v_{i+1}) \in E \ \forall i \geq 0$
  - the distance of $\omega = v_0 v_1 v_2 \ldots v_n$ is: $\Sigma_{i=0\ldots n-1} w(v_i, v_{i+1})$

- Shortest path problem
  - given a weighted digraph, find a path between two vertices v_1 and v_2 with the smallest distance
  - i.e. a path ω s.t. $d(\omega) \leq d(\omega')$ for all other such paths ω'

Finding strongest evidences

- Reduction to graph problem…
- Step 1: Adapt the DTMC
  - make states satisfying $\neg\Phi_1 \wedge \neg\Phi_2$ absorbing
    - (i.e. replace all outgoing transitions with a single self-loop)
  - add an extra state t and replace all transitions from any Φ_2 state with a single transition to t (with probability 1)
- Step 2: Convert new DTMC into a weighted digraph
  - for the (adapted) DTMC $D = (S, s_{init}, \mathbf{P}, L)$:
  - corresponding graph is $G_D = (V, E, w)$ where:
  - $V = S$ and $E = \{ (s,s') \in S \times S \mid \mathbf{P}(s,s') > 0 \}$
  - $w(s,s') = \log(1 / \mathbf{P}(s,s'))$
- Key idea: for any two paths ω and ω’ in D (and in G_D)
  - $\mathbf{P}_s(\omega') \geq \mathbf{P}_s(\omega)$ if and only if $d(\omega') \leq d(\omega)$

Example…

- PCTL formula: $P_{\leq 1/2}$ [F b]

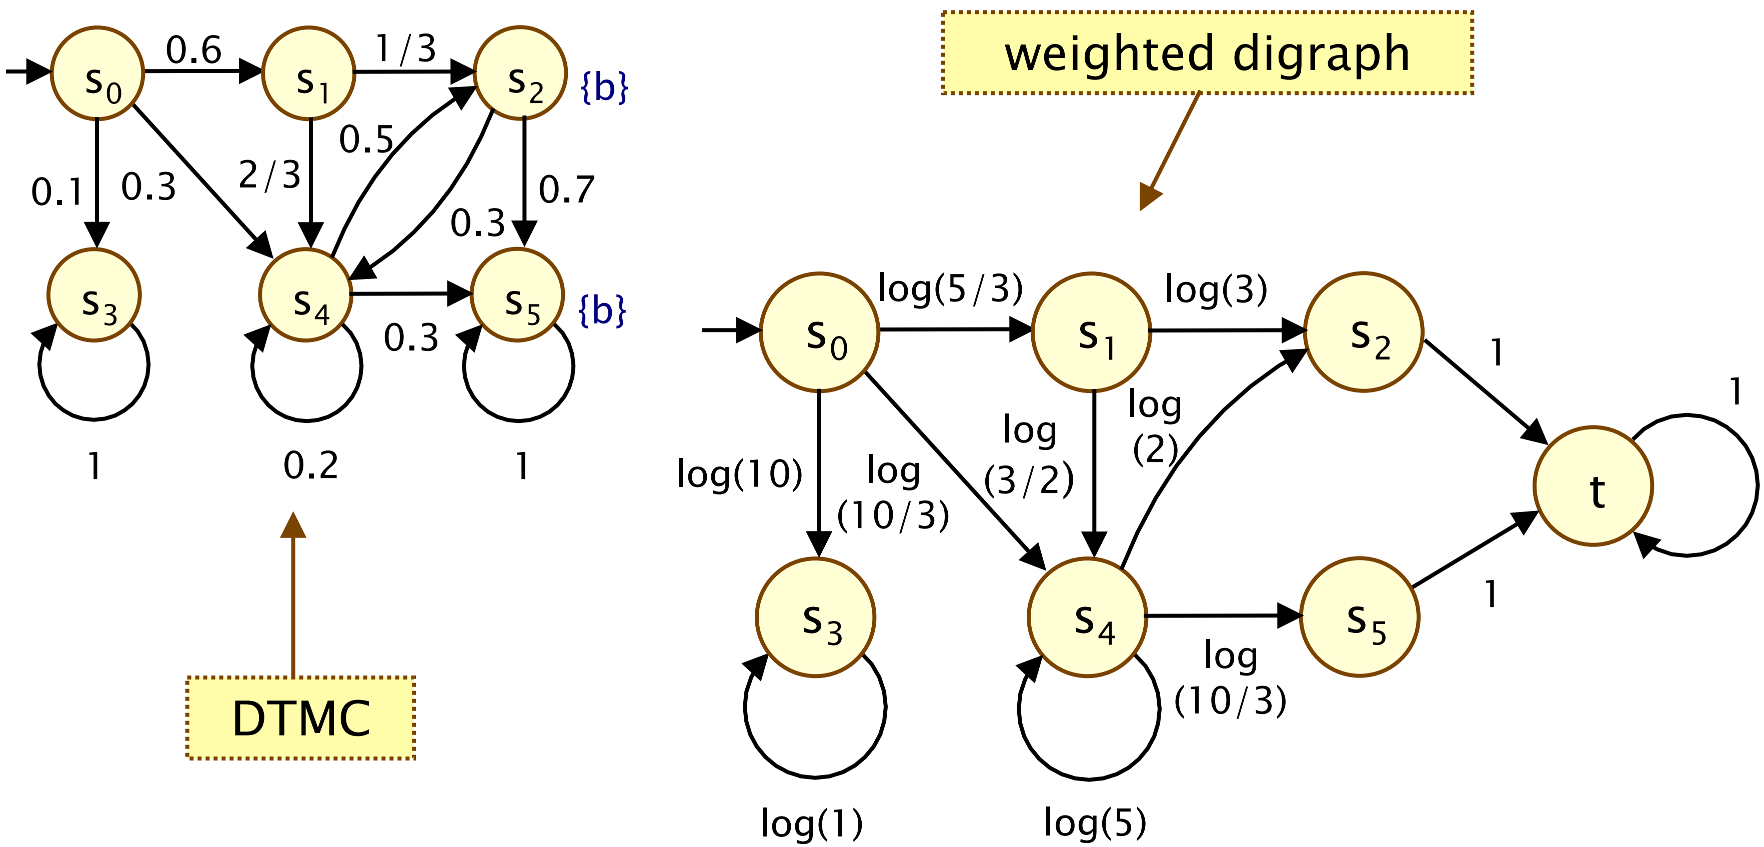

Finding strongest evidences

- To find strongest evidence in DTMC D
  - analyse corresponding digraph
- For unbounded until formula $P_{\leq p}$ [Φ_1 U Φ_2]
  - solve shortest path problem in digraph (target t)
  - polynomial time algorithms exist
    - e.g. Dijsktra's algorithm can be implemented in $O(|E|+|V|\cdot\log|V|)$
- For bounded until formula $P_{\leq p}$ [Φ_1 $U^{\leq k}$ Φ_2]
  - solve special case of the constrained shortest path problem
  - also solvable in polynomial time
- Generation of smallest counterexamples
  - based on computation of k shortest paths
  - k can be computed on the fly

Other cases

- Lower bounds on probabilities
  - i.e. $s \nvDash P_{\geq p} [\ \Phi_1\ U^{\leq k}\ \Phi_2\ ]$
  - negate until formula to reverse probability bound
  - solvable with BSCC computation + probabilistic reachability
  - for details, see [HK07]

- Continuous-time Markov chains
  - these techniques can be extended to CTMCs and CSL [HK07b]
  - naïve approach: apply DTMC techniques to uniformised DTMC
  - modifications required to get smaller counterexamples
  - another possibility: directed search based techniques [AHL05]

Bisimulation

- Identifies models with the same branching structure
  - i.e. the same stepwise behaviour
  - each model can simulate the actions of the other
  - guarantees that models satisfy many of the same properties

- Uses of bisimulation:
  - show equivalence between a model and its specification
  - state space reduction: bisimulation minimisation

- Formally, bisimulation is an equivalence relation over states
  - bisimilar states must have identical labelling and identical stepwise behaviour

Equivalence relations

- Let R be a relation over some set S
  - i.e. $R \subseteq S \times S$
  - we write $s_1\ R\ s_2$ as shorthand for $(s_1,s_2) \in R$
- R is an equivalence relation iff:
  - R is reflexive, i.e. $s\ R\ s$
  - R is symmetric, i.e. if $s_1\ R\ s_2$ then $s_2\ R\ s_1$
  - R is transitive, i.e. if $s_1\ R\ s_2$ and $s_2\ R\ s_3$ then $s_1\ R\ s_3$
- R partitions S:
  - equivalence classes: $[s]_R = \{ s' \in S \mid s'\ R\ s \}$
  - the quotient of S under R is denoted $S/R = \{ [s]_R \mid s \in S \}$

Bisimulation on DTMCs

- Consider a DTMC $D = (S, s_{init}, \mathbf{P}, L)$
- Some notation:
  - $\mathbf{P}(s,T) = \Sigma_{s' \in T} \mathbf{P}(s,s')$ for $T \subseteq S$

- An equivalence relation R on S is a probabilistic bisimulation on D if and only if for all s_1 R s_2:
  - $L(s_1) = L(s_2)$
  - $\mathbf{P}(s_1, T) = \mathbf{P}(s_2, T)$ for all $T \in S/R$ (i.e. for all equivalence classes of R)

- States s_1 and s_2 are bisimulation-equivalent (or bisimilar)
  - if there exists a probabilistic bisimulation R on D with s_1 R s_2
  - denoted $s_1 \sim s_2$

Simple example

- Bisimulation relation ~
- Quotient of S under ~
  - { $\{s_1\}$, $\{u_1, u_2\}$, $\{v_1, v_2\}$ }
- Bisimilar states:
  - $u_1 \sim u_2$
  - $v_1 \sim v_2$

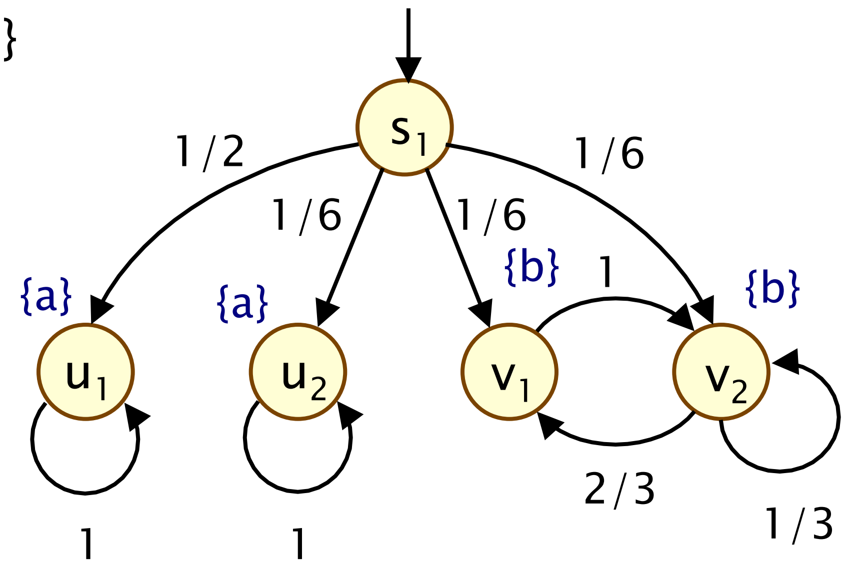

Bisimulation on DTMCs

- Bisimulation between DTMCs D_1 and D_2
  - $D_1 \sim D_2$ if they have bisimilar initial states
- Formally:
  - state labellings for D_1 and D_2 over same set of atomic prop.s
  - bisimulation relation is over disjoint union of D_1 and D_2

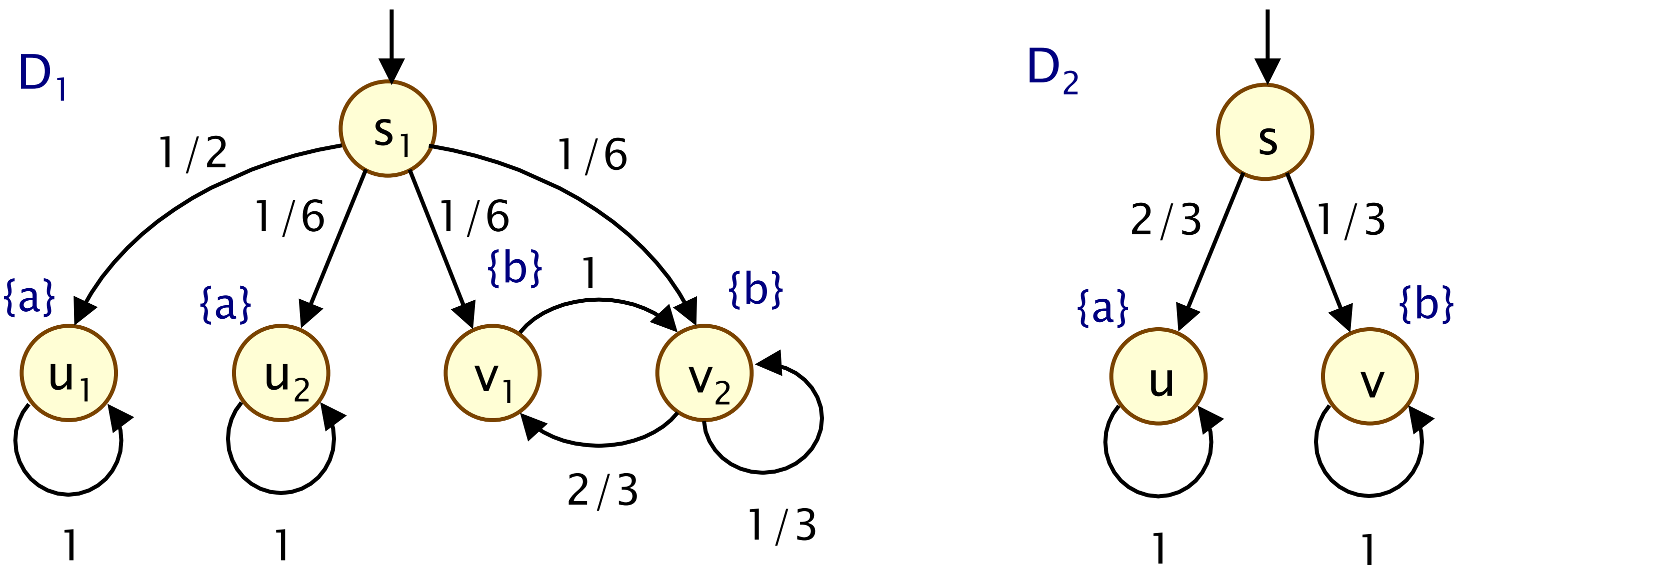

Simple example

- Bisimilar states:
  - $u_1 \sim u_2 \sim u$
  - $v_1 \sim v_2 \sim v$
  - $s_1 \sim s$

Bisimilar DTMCs: $D_1 \sim D_2$

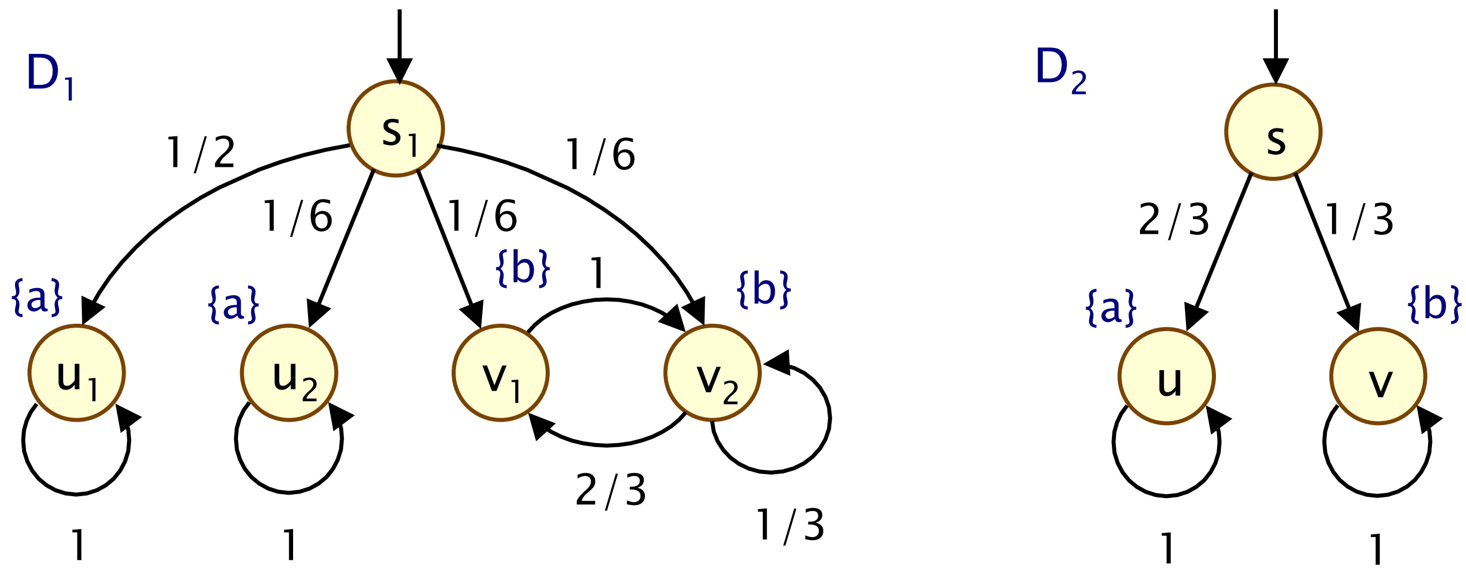

Quotient DTMC

- For a DTMC $D = (S, s_{init}, \mathbf{P}, L)$ and probabilistic bisimulation ~

- Quotient DTMC is
  - $D/\sim \, = (S', s'_{init}, \mathbf{P}', L')$

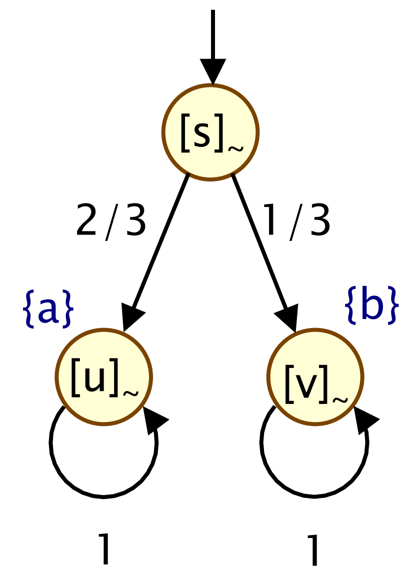

- where:
  - $S' = S/\sim \, = \{ [s]_\sim \mid s \in S \}$
  - $s'_{init} = [s_{init}]_\sim$
  - $\mathbf{P}'([s]_\sim, [s']_\sim) = \mathbf{P}(s, [s']_\sim)$
  - $L'([s]_\sim) = L(s)$

well defined since bisimulation ensures $\mathbf{P}(s, [s']_\sim)$ same for all s in $[s]_\sim$

Bisimulation and PCTL

- Probabilistic bisimulation preserves all PCTL formulae

- For all states s and s':

$$s \sim s'$$
$$\Leftrightarrow$$
$$\text{for all PCTL formulae } \Phi,\ s \vDash \Phi \text{ if and only if } s' \vDash \Phi$$

- Note also:
  - every pair of non-bisimilar states can be distinguished with some PCTL formula
  - ~ is the coarsest relation with this property
  - in fact, bisimulation also preserves all PCTL* formulae

CTMC bisimulation

- Check equivalence of rates, not probabilities…

- An equivalence relation R on S is a probabilistic bisimulation on CTMC $C=(S,s_{init},\mathbf{R},L)$ if and only if for all s_1 R s_2:
  - $L(s_1) = L(s_2)$
  - $\mathbf{R}(s_1, T) = \mathbf{R}(s_2, T)$ for all classes T in S/R

- Alternatively, check:
  - $L(s_1) = L(s_2)$, $\mathbf{P}^{emb(C)}(s_1, T) = \mathbf{P}^{emb(C)}(s_2, T)$, $E(s_1) = E(s_2)$

- Bisimulation on CTMCs preserves CSL
  - (see [BHHK03] and also [DP03])

Bisimulation minimisation

- More efficient to perform PCTL/CSL model checking on the quotient DTMC/CTMC
  - assuming quotient model can be constructed efficiently
  - (see [KKZJ07] for experimental results on this)

- Bisimulation minimisation
  - algorithm to construct quotient model
  - based on partition refinement
  - repeated splitting of an initially coarse partition
  - final partition is coarsest bisimulation wrt. initial partition
  - (optimisations/variants possible by changing initial partition)
  - complexity: $O(|\mathbf{P}| \cdot \log|S| + |AP| \cdot |S|)$ [DHS'03]
    - assuming suitable data structure used (splay trees)

Bisimulation minimisation

- 1. Start with initial partition
  - say $\Pi = \{ \{ s \in S \mid L(s)=lab \} \mid lab \in 2^{AP} \}$
- 2. Find a splitter $T \in \Pi$ for some block $B \in \Pi$
  - a splitter T is a block such that probability of going to T differs for some states in block B
  - i.e. $\exists s,s' \in B \,.\, \mathbf{P}(s,T) \neq \mathbf{P}(s',T)$
- 3. Split B into sub-blocks
  - such that $\mathbf{P}(s,T)$ is the same for all states in each sub-block
- 4. Repeat steps 2/3 until no more splitters exist
  - i.e. no change to partition Π

replace **P** with **R** for CTMCs

CTMC example

- Consider model checking $P_{\sim p}$ [$F^{[0,t]}$ a] on this CTMC:

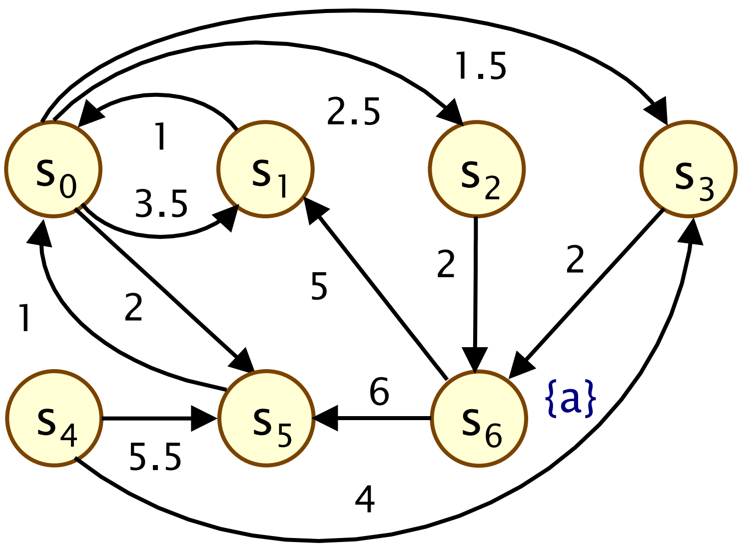

Minimisation:

Π_0: $B_1=\{s_0,s_1,s_2,s_3,s_4,s_5\}$, $B_2=\{s_6\}$

B_2 is a splitter for B_1

(since e.g. $\mathbf{R}(s_1,B_2)=0\neq2=\mathbf{R}(s_2,B_2)$)

Π_1: $B_1=\{s_0,s_1,s_4,s_5\}$, $B_2=\{s_6\}$, $B_3=\{s_2,s_3\}$

B_3 is a splitter for B_1

(since e.g. $\mathbf{R}(s_1,B_3)=0\neq4=\mathbf{R}(s_0,B_3)$)

Π_2: $B_1=\{s_1,s_5\}$, $B_2=\{s_6\}$, $B_3=\{s_2,s_3\}$, $B_4=\{s_0,s_4\}$

No more splitters…

$S/\sim\ = \{ \{s_1,s_5\}, \{s_6\}, \{s_2,s_3\}, \{s_0,s_4\} \}$

CTMC example…

C

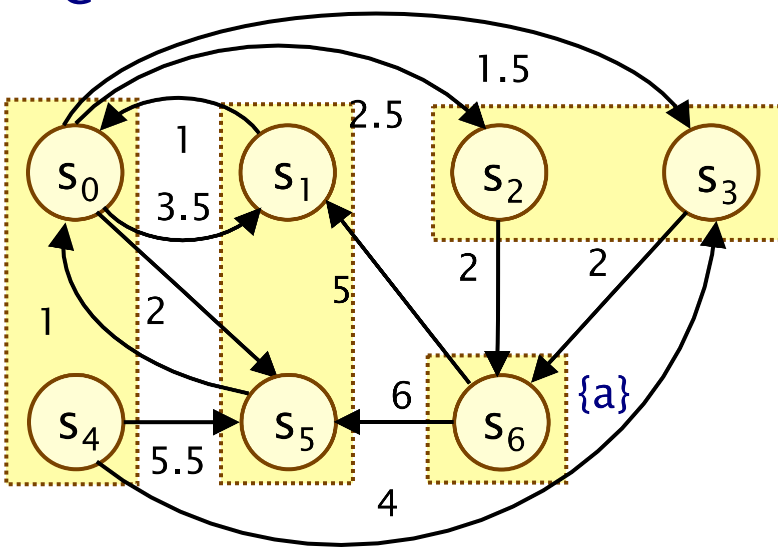

$S/\sim = \{ \{s_1,s_5\}, \{s_6\}, \{s_2,s_3\}, \{s_0,s_4\} \}$

C/~

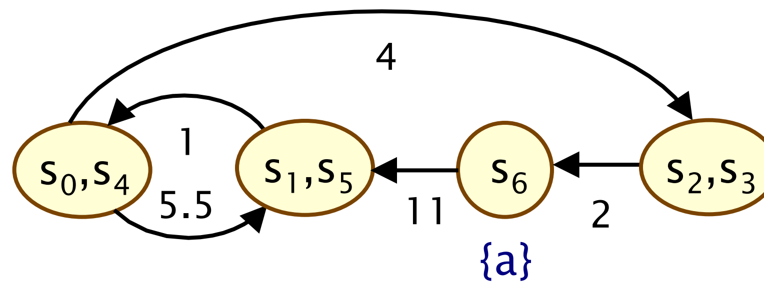

$$Prob^{C}(s_0, F^{[0,t]}\ a) \;=\; Prob^{C/\sim}(\{s_0,s_4\}, F^{[0,t]}\ a)$$

Summing up…

- Counterexamples
  - essential ingredient of non-probabilistic model checking
  - counterexamples for PCTL + DTMCs
    - finite set of paths showing $\not\models P_{\leq p} [\ \Phi_1\ U^{\leq k}\ \Phi_2\ ]$
  - computing smallest counterexamples
    - reduction to well-known graph problems

- Bisimulation
  - relates states/Markov chains with identical labelling and identical stepwise behaviour
  - preserves PCTL, CSL, …
  - bisimulation minimisation: automated reduction to quotient model