Probabilistic bisimulation

Lecture #14 of Modeling Concurrent and Probabilistic Systems

Joost-Pieter Katoen

Lehrstuhl 2: Softwaremodeling and Verification

E-mail: `katoen@cs.rwth-aachen.de`

June 25, 2009

Overview Lecture #14

$\Rightarrow$ *Probabilistic bisimulation*

- Bisimulation for labeled transition systems
- Fully probabilistic systems and DTMCs
- Probabilistic bisimulation

Labeled transition system

A *labeled transition system* *LTS* is a quadruple $(S, \mathit{Act}, \longrightarrow, s_0)$ where

- S is a set of states,
- *Act* is a set of actions,
- $\longrightarrow \subseteq S \times \mathit{Act} \times S$ is a transition relation,
- $s_0 \in S$ is the initial state.

S and *Act* are either finite or countably infinite

Notation: $s \xrightarrow{\alpha} s'$ instead of $(s, \alpha, s') \in \longrightarrow$

Strong bisimulation

- Let *LTS* $= (S, \mathit{Act}, \longrightarrow, s_0)$ and R a binary relation on S
- R is a *strong bisimulation* on $S \times S$ whenever for $(s, t) \in R$:

  if $s \stackrel{\alpha}{\longrightarrow} s'$ then there exists $t' \in S$ such that $t \stackrel{\alpha}{\longrightarrow} t'$ and $(s', t') \in R$

  and

  if $t \stackrel{\alpha}{\longrightarrow} t'$ then there exists $s' \in S$ such that $s \stackrel{\alpha}{\longrightarrow} s'$ and $(s', t') \in R$

- s is *strongly bisimilar* to t, notation $s \sim t$, if:

  there exists a strong bisimulation R such that $(s, t) \in R$

property: $\sim$ is an equivalence

Strong bisimulation

$$\begin{array}{ccc} s & \xrightarrow{\alpha} & s' \\ R & & \\ t & & \end{array} \quad \text{can be completed to} \quad \begin{array}{ccc} s & \xrightarrow{\alpha} & s' \\ R & & R \\ t & \xrightarrow{\alpha} & t' \end{array}$$

and

$$\begin{array}{ccc} s & & \\ R & & \\ t & \xrightarrow{\alpha} & t' \end{array} \quad \text{can be completed to} \quad \begin{array}{ccc} s & \xrightarrow{\alpha} & s' \\ R & & R \\ t & \xrightarrow{\alpha} & t' \end{array}$$

Are these transition systems strongly bisimilar?

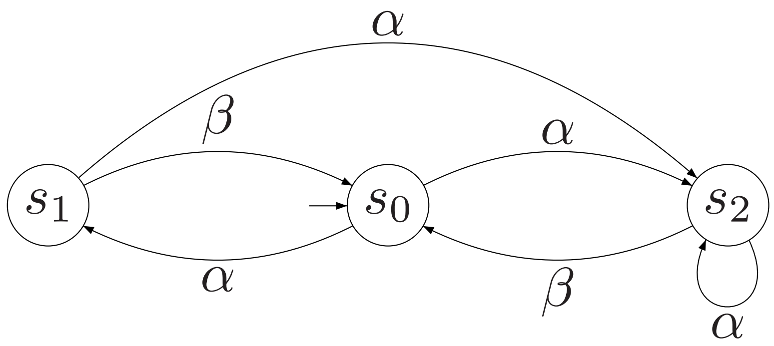

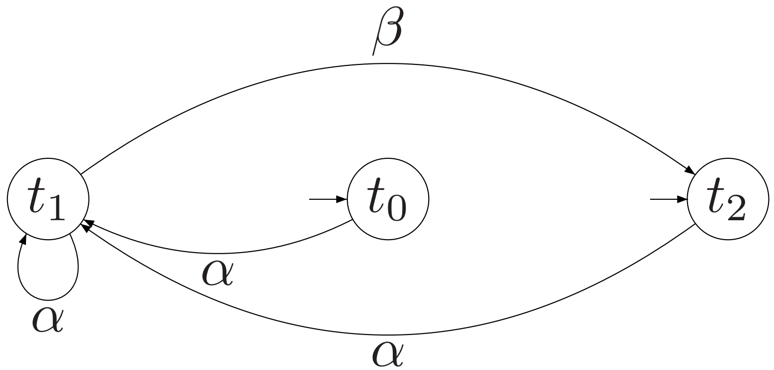

Yes, they are!

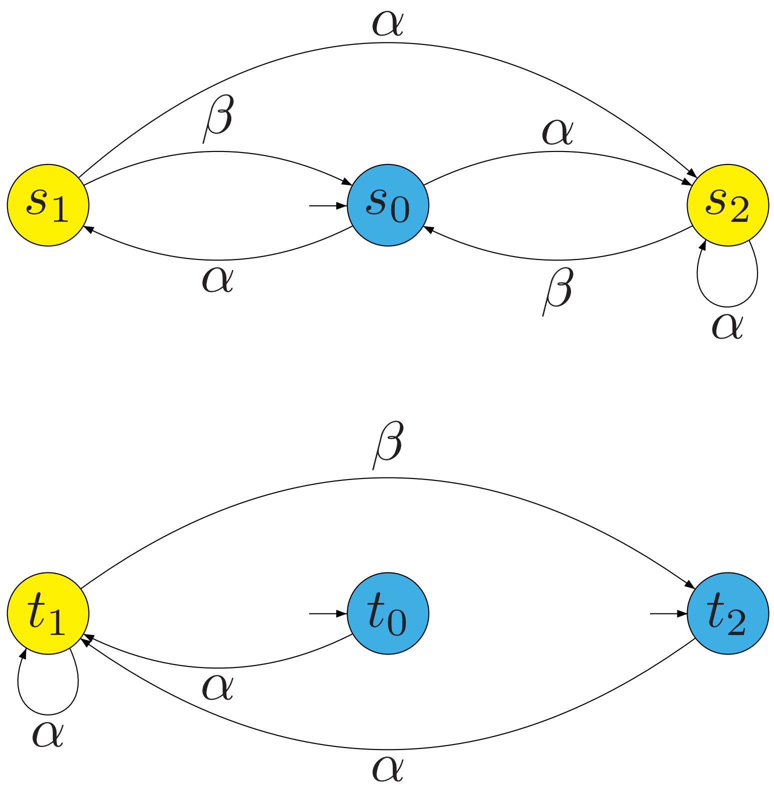

Quotient LTS under $\sim$

For *LTS* $= (S, \textit{Act}, \longrightarrow, s_0)$ and strong bisimulation $\sim \subseteq S \times S$ let

$$\textit{LTS}/\!\sim \; = \; (S', \textit{Act}, \longrightarrow', s_0'), \qquad \text{the quotient of } \textit{LTS} \text{ under } \sim$$

where

- $S' = S/\!\sim \; = \; \{ \, [s]_\sim \mid s \in S \, \}$ with $[s]_\sim \; = \; \{ \, s' \in S \mid s \sim s' \, \}$

  $\Rightarrow$ *states are equivalence classes under* $\sim$

- $\longrightarrow'$ is defined by:
  $$\frac{s \xrightarrow{\alpha} s'}{[s]_\sim \xrightarrow{\alpha}{}' [s']_\sim}$$

- $s_0' = [s_0]_\sim$, the equivalence class of the initial state in *LTS*

note that *LTS* $\sim$ *LTS*$/\!\sim$ Why?

Strong bisimulation revisited

Let $P : S \times \mathit{Act} \times 2^S \to \{0, 1\}$ be a predicate such that for $S' \subseteq S$:

$$P(s, \alpha, S') = \begin{cases} 1 & \text{if } \exists s' \in S'.\ s \xrightarrow{\alpha} s' \\ 0 & \text{otherwise} \end{cases}$$

Let $\mathit{LTS} = (S, \mathit{Act}, \longrightarrow, s_0)$ and R an *equivalence relation* on S.

Then: R is a *strong bisimulation* on S if for $(s, s') \in R$:

$$P(s, \alpha, C) \;=\; P(s', \alpha, C) \quad \text{for all} \quad C \text{ in } S/R \text{ and } \alpha \in \mathit{Act}$$

this definition is equivalent to the previous one (exercise)

Probabilistic bisimulation: intuition

- Strong bisimulation is used to compare labeled transition systems
- Strongly bisimilar states exhibit the same step-wise behaviour
- We like to adapt bisimulation to discrete-time Markov chains
- This yields a probabilistic variant of strong bisimulation
- When do two DTMC states exhibit the same step-wise behaviour?
- Key: if their transition probability for each equivalence class coincides

for technical reasons, consider a slight generalization of DTMCs

Fully probabilistic system

A *fully probabilistic system* (FPS) is a pair $\mathcal{D} = (S, \mathbf{P})$ where:

- S is a countable set of states
- $\mathbf{P} : S \times S \to [0, 1]$ is a *probability matrix* satisfying

$$\sum_{s' \in S} \mathbf{P}(s, s') \in [0, 1] \quad \text{for all} \quad s \in S$$

If $\sum_{s' \in S} \mathbf{P}(s, s') = 1$, state s is called *stochastic*; otherwise, s is *sub-stochastic*

Deadlocks

- The probability to move from s to (a state in) $C \subseteq S$:

$$\mathbf{P}(s, C) \ = \ \sum_{s' \in C} \mathbf{P}(s, s')$$

- Let $\mathbf{P}(s, \bot) = 1 - \ \mathbf{P}(s, S)$
  - the probability to stay forever in s without performing any transition
  - although $\bot$ is not a "real" state (i.e., $\bot \notin S$), it may be regarded as a *deadlock*
  - $\bot$ is treated in the next lecture as an auxiliary state

$$s \text{ is stochastic} \quad \text{iff} \quad \mathbf{P}(s, \bot) = 0 \quad \text{iff} \quad \mathbf{P}(s, S) = 1$$

Discrete-time Markov chain

A DTMC is an FPS where *no* state is sub-stochastic:

$$\mathbf{P}(s, S) \;=\; 1 \quad \text{for all} \quad s \in S$$

Probabilistic bisimulation

- Let $\mathcal{D} = (S, \mathbf{P})$ be a FPS and R an equivalence relation on S
- R is a *probabilistic bisimulation* on S if for any $(s, s') \in R$:

$$\mathbf{P}(s, C) = \mathbf{P}(s', C) \quad \text{for all} \quad C \text{ in } S/R$$

- s and s' are *probabilistic bisimilar* (or: lumping equivalent), $s \sim_p s'$, if:

  there exists a probabilistic bisimulation R on S with $(s, s') \in R$

it follows that: $s \sim_p s' \quad \Rightarrow \quad \mathbf{P}(s, \perp) = \mathbf{P}(s', \perp)$

Example

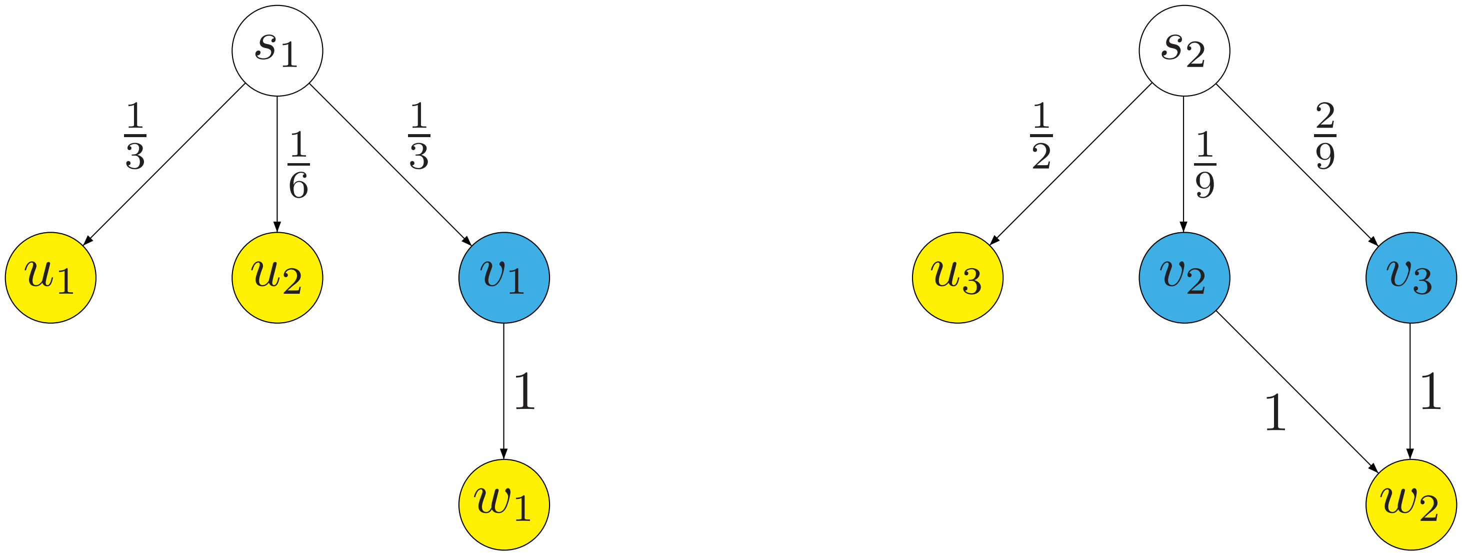
s_1
$\frac{1}{3}$
$\frac{1}{6}$
$\frac{1}{3}$
u_1
u_2
v_1
1
w_1
s_2
$\frac{1}{2}$
$\frac{1}{9}$
$\frac{2}{9}$
u_3
v_2
v_3
1
1
w_2

Another example

Quotient FPS under $\sim_p$

For $\mathcal{D} = (S, \mathbf{P})$ and probabilistic bisimulation $\sim_p \subseteq S \times S$ let

$$\mathcal{D}/\sim_p = (S', \mathbf{P}'), \qquad \text{the quotient of } \mathcal{D} \text{ under } \sim_p$$

where

- $S' = S/\sim_p = \{ [s]_{\sim_p} \mid s \in S \}$ with $[s]_{\sim_p} = \{ s' \in S \mid s \sim_p s' \}$
- $\mathbf{P}' : S' \times S' \rightarrow [0, 1]$ is defined by:

$$\mathbf{P}'([s]_{\sim_p}, [s']_{\sim_p}) = \mathbf{P}(s, [s']_{\sim_p})$$

if an initial distribution is given on $\mathcal{D}$, then: $\underline{p}'_C(0) = \sum_{s \in C} \underline{p}_s(0)$ for each $C \in S/\sim_p$

Craps

- Roll two dice and bet on outcome

- Come-out roll ("pass line" wager):
  - outcome 7 or 11: win
  - outcome 2, 3, or 12: loss ("craps")
  - any other outcome: roll again (outcome is "point")

- Repeat until 7 or the "point" is thrown:
  - outcome 7: loss ("seven-out")
  - outcome the point: win
  - any other outcome: roll again

A DTMC model of Craps

- Come-out roll:
  - 7 or 11: win
  - 2, 3, or 12: loss
  - else: roll again
- Next roll(s):
  - 7: loss
  - point: win
  - else: roll again

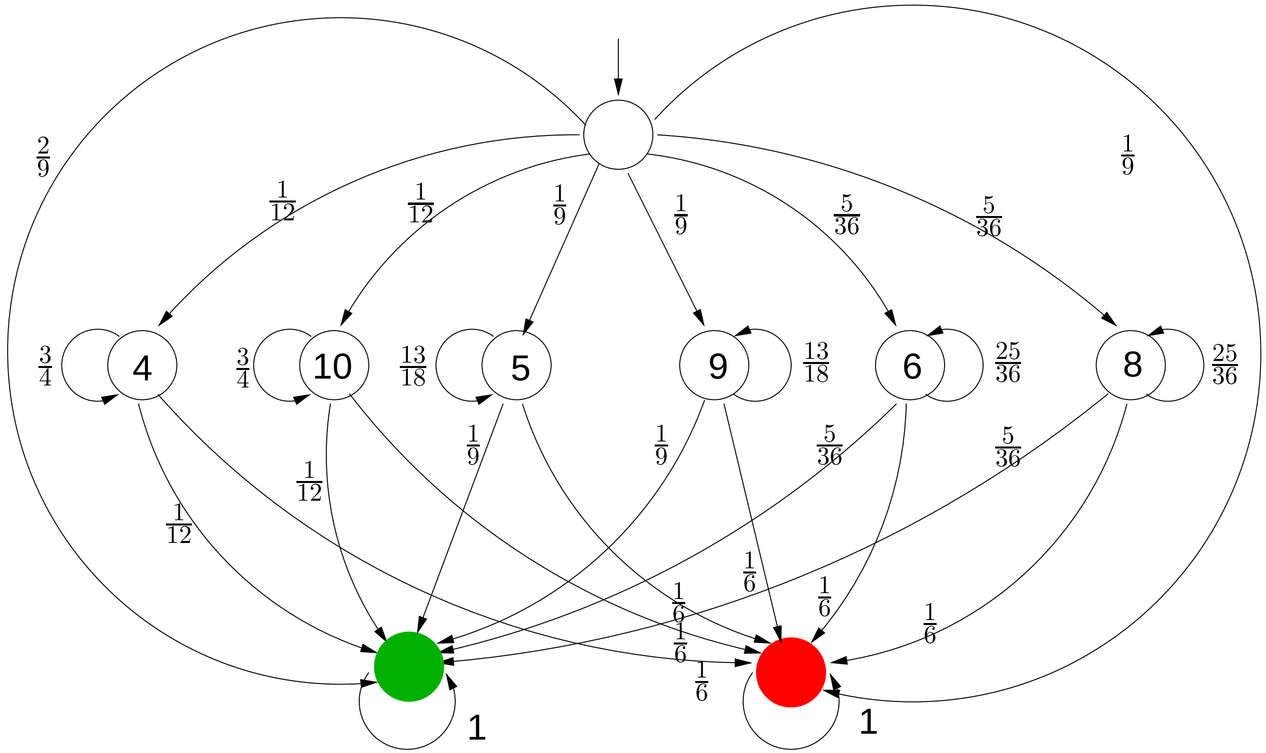

Quotient DTMC

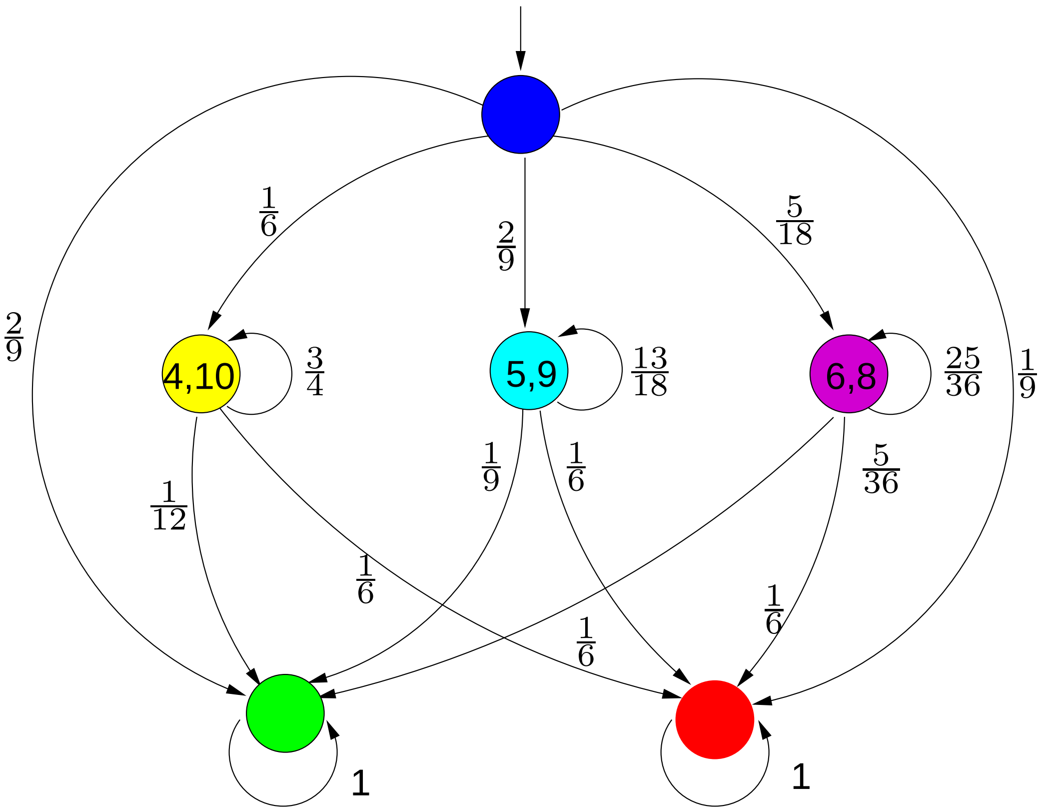

Crowds protocol (Reiter & Rubin, 1998)

- A protocol for anonymous web browsing (variants: mCrowds, BT-Crowds)
- Hide user's communication by random routing within a crowd
  - sender selects a crowd member randomly using a uniform distribution
  - selected router flips a biased coin:
    - with probability $1 - p$: direct delivery to final destination
    - otherwise: select a next router randomly (uniformly)
  - once a routing path has been established, use it until crowd changes
- Rebuild routing paths on crowd changes (R times)
- Probable innocence:
  - probability real sender is discovered $< \frac{1}{2}$ if $N \geqslant \frac{p}{p-\frac{1}{2}} \cdot (c+1)$
  - where N is crowd's size and c is number of corrupt crowd members

Crowds protocol

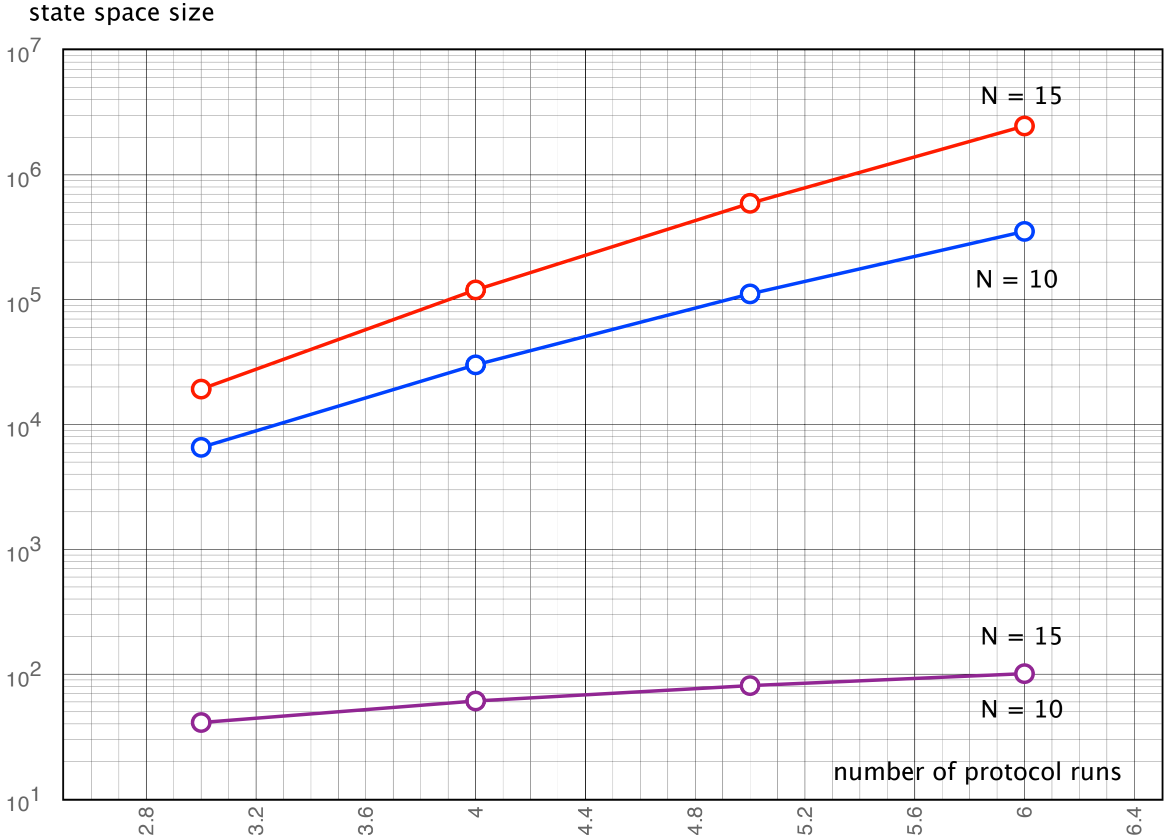

state space reductions for bisimulation quotient

Initial distribution

- Let $\mathcal{D} = (S, \mathbf{P})$ be an FPS with initial distribution $\underline{p}(0) : S \to [0,1]$
- Let $\mathcal{D}_0 = (S_0, \mathbf{P}_0)$ be obtained from $\mathcal{D}$ by adding a new initial state:
  - $S_0 = S \cup \{ s_0 \}$ with $s_0 \notin S$
  - $\mathbf{P}_0(s, s') = \begin{cases} \mathbf{P}(s, s') & \text{if } s \neq s_0, s' \neq s_0 \\ \underline{p}_{s'}(0) & \text{if } s = s_0, s' \neq s_0 \\ 0 & \text{otherwise} \end{cases}$
- Two FPSs with initial distribution are bisimilar, $(\mathcal{D}, \underline{p}(0)) \sim_p (\mathcal{D}', \underline{p}'(0))$
  - if there exists a probabilistic bisimulation R on $S_0 \uplus S_0'$ with $(s_0, s_0') \in R$

Preservation of state probabilities

- Let $\mathcal{D} = (S, \mathbf{P})$ be an FPS with initial distribution $\underline{p}(0)$ and $\mathcal{D}_0/\sim_p$ the quotient under $\sim_p$
- For any $C \in S_0/\sim_p$ we have:

$$\underline{p}'_C(n) = \sum_{s \in C} \underline{p}_s(n) \qquad \text{for any } n \geqslant 0$$

- If the limiting distribution exists, then it follows:

$$\underline{p}'_C = \lim_{n\to\infty} \underline{p}'_C(n) = \lim_{n\to\infty} \sum_{s \in C} \underline{p}_s(n) = \sum_{s \in C} \underline{p}_s$$

Preservation of reachability probabilities

For any equivalence class $C \in S_{\sim_p}$:

$$s \sim_p s' \;\Rightarrow\; \underbrace{\Pr\left\{s \overset{\leqslant n}{\rightsquigarrow} C\right\}}_{p(s,n,C)} = \Pr\left\{s' \overset{\leqslant n}{\rightsquigarrow} C\right\} \quad \text{for any } n \geqslant 0$$

where the probability to reach C within at most n steps is:

$$p(s,n,C) = \begin{cases} 1 & \text{if } s \in C \\ \sum_{s' \in S} \mathbf{P}(s,s') \cdot p(s', n{-}1, C) & \text{if } s \notin C \text{ and } n > 0 \\ 0 & \text{otherwise} \end{cases}$$

this can be generalized by forbidding paths that visit states in $B \in S_{\sim_p}$ prior to reaching C

Crowds protocol

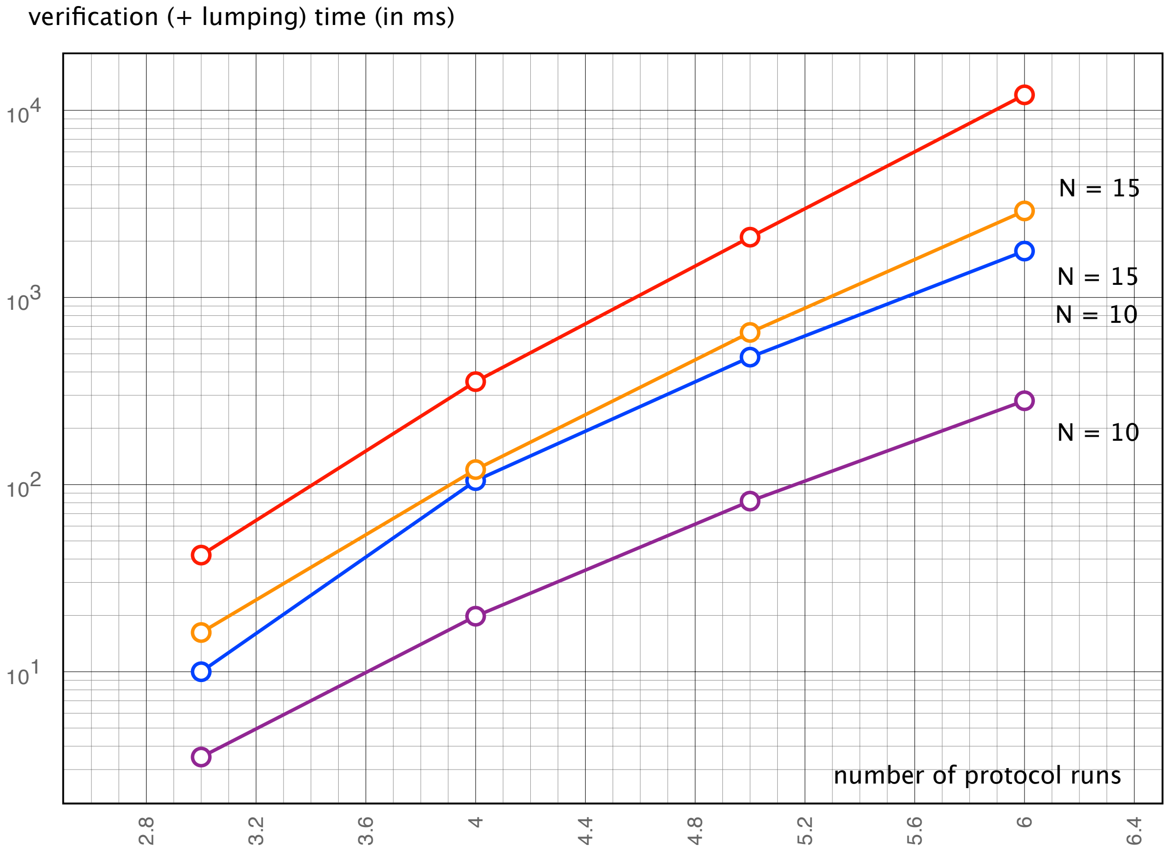

run times for eventually observer the real sender more than once