Modeling and Verification of Probabilistic Systems

Joost-Pieter Katoen

Lehrstuhl für Informatik 2
Software Modeling and Verification Group

http://moves.rwth-aachen.de/teaching/ws-1819/movep18/

October 15, 2018

Overview

Summary of previous lecture

What are Markov chains?

- A discrete-time Markov chain (DTMC) is a time-homogeneous Markov process with discrete parameter T and discrete state space S.
- State residence times are geometrically distributed.
- Alternative: a DTMC $\mathcal{D}$ is a tuple $(S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ with:
  - state space S
  - transition probability function $\mathbf{P}$
  - initial distribution ι_{init}

What are transient probabilities?

- $\Theta_n^{\mathcal{D}}(s)$ is the probability to be in state s after n steps.
- These transient probabilities satisfy: $\Theta_n^{\mathcal{D}} = \iota_{\text{init}} \cdot \mathbf{P}^n$.

Aim of this lecture

How to determine reachability probabilities?

Three major steps

1. What are reachability probabilities? I mean, precisely.
   This requires a bit of measure theory. Sorry for that.
2. Reachability probabilities = unique solution of linear equation system.
3. Bounded reachability probabilities = transient probabilities[1].

[1] in a slightly modified DTMC.

Recall Knuth's die

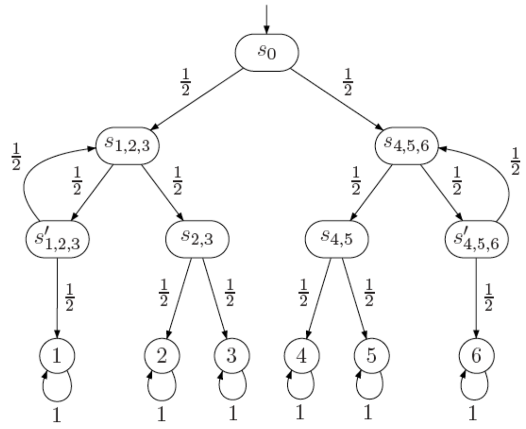

Heads = "go left"; tails = "go right". Does this DTMC model a six-sided die?

Overview

Paths

State graph

The *state graph* of DTMC $\mathcal{D}$ is a digraph $G = (V, E)$ with V the states of $\mathcal{D}$, and $(s, s') \in E$ iff $\mathbf{P}(s, s') > 0$.

Let $Pre(s)$ be the *predecessors* of s, $Pre^*(s)$ its reflexive and transitive closure.

Paths

Paths in $\mathcal{D}$ are infinite paths in its state graph.

$Paths(\mathcal{D})$ denotes the set of paths in $\mathcal{D}$, and $Paths^*(\mathcal{D})$ its finite prefixes.

Some events of interest

Let DTMC $\mathcal{D}$ with (possibly infinite) state space S.

(Simple) reachability

Eventually reach a state in $G \subseteq S$. Formally:

$$\Diamond G \;=\; \{\, \pi \in Paths(\mathcal{D}) \mid \exists i \in \mathbb{N}.\, \pi[i] \in G \,\}$$

Invariance, i.e., always stay in state in G:

$$\Box G \;=\; \{\, \pi \in Paths(\mathcal{D}) \mid \forall i \in \mathbb{N}.\, \pi[i] \in G \,\} \;=\; \overline{\Diamond \overline{G}}.$$

Constrained reachability

Or "reach-avoid" properties where states in $F \subseteq S$ are forbidden:

$$\overline{F} \,\mathsf{U}\, G \;=\; \{\, \pi \in Paths(\mathcal{D}) \mid \exists i \in \mathbb{N}.\, \pi[i] \in G \wedge \forall j < i.\, \pi[j] \notin F \,\}$$

More events of interest

Repeated reachability

Repeatedly visit a state in G; formally:

$$\Box\Diamond G \;=\; \{\, \pi \in Paths(\mathcal{D}) \mid \forall i \in \mathbb{N}.\, \exists j \geqslant i.\, \pi[j] \in G \,\}$$

Persistence

Eventually reach in a state in G and always stay there; formally:

$$\Diamond\Box G \;=\; \{\, \pi \in Paths(\mathcal{D}) \mid \exists i \in \mathbb{N}.\, \forall j \geqslant i.\, \pi[j] \in G \,\}$$

Overview

Recall: Measurable space

Sample space

A *sample space* Ω of a chance experiment is a set of elements that have a 1-to-1 relationship to the possible outcomes of the experiment.

σ-algebra

A *σ-algebra* is a pair $(\Omega, \mathcal{F})$ with $\Omega \neq \varnothing$ and $\mathcal{F} \subseteq 2^{\Omega}$ a collection of subsets of sample space Ω such that:

1. $\Omega \in \mathcal{F}$
2. $A \in \mathcal{F} \;\Rightarrow\; \Omega - A \in \mathcal{F}$ — complement
3. $(\forall i \geqslant 0.\ A_i \in \mathcal{F}) \;\Rightarrow\; \bigcup_{i \geqslant 0} A_i \in \mathcal{F}$ — countable union

The elements in $\mathcal{F}$ of a σ-algebra $(\Omega, \mathcal{F})$ are called *events*.
The pair $(\Omega, \mathcal{F})$ is called a *measurable space*.

Let Ω be a set. $\mathcal{F} = \{\varnothing, \Omega\}$ yields the smallest σ-algebra; $\mathcal{F} = 2^{\Omega}$ yields the largest one.

What's the probability of infinite paths?

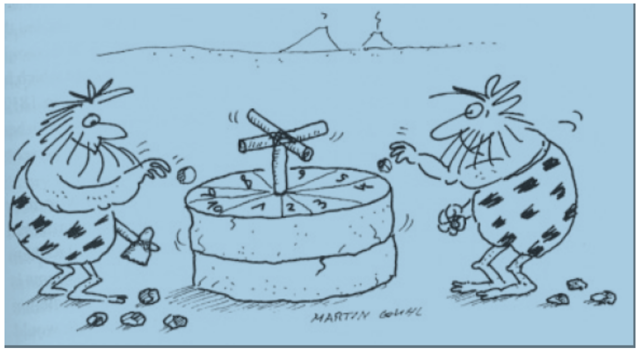

Probability space

Probability space

A *probability space* $\mathcal{P}$ is a structure $(\Omega, \mathcal{F}, Pr)$ with:

- $(\Omega, \mathcal{F})$ is a σ-algebra, and
- $Pr : \mathcal{F} \to [0, 1]$ is a *probability measure*, i.e.:
  1. $Pr(\Omega) = 1$, i.e., Ω is the certain event
  2. $Pr\left(\bigcup_{i \in I} A_i\right) = \sum_{i \in I} Pr(A_i)$ for any $A_i \in \mathcal{F}$ with $A_i \cap A_j = \varnothing$ for $i \neq j$

The events in $\mathcal{F}$ of a probability space $(\Omega, \mathcal{F}, Pr)$ are called *measurable*.

Paths and probabilities

To reason quantitatively about the behavior of a DTMC, we need to define a probability space over its paths.

Intuition

For a given state s in DTMC $\mathcal{D}$:

- Outcomes := set of all infinite paths starting in s.
- Events := subsets of these outcomes.
- These events are defined using cylinder sets.
- Cylinder set of a finite path := set of all its infinite continuations.

Probability measure on DTMCs

Cylinder set

The *cylinder set* of finite path $\hat{\pi} = s_0\, s_1 \ldots s_n \in Paths^*(\mathcal{D})$ is defined by:

$$Cyl(\hat{\pi}) \;=\; \{\, \pi \in Paths(\mathcal{D}) \mid \hat{\pi} \text{ is a prefix of } \pi \,\}$$

The cylinder set spanned by finite path $\hat{\pi}$ thus consists of all infinite paths that have prefix $\hat{\pi}$.

Probability space of a DTMC

The set of events of the probability space DTMC $\mathcal{D}$ contains all cylinder sets $Cyl(\hat{\pi})$ where $\hat{\pi}$ ranges over all finite paths in $\mathcal{D}$.

Probability measure on DTMCs

Cylinder set

The cylinder set of finite path $\hat{\pi} = s_0\, s_1 \ldots s_n \in Paths^*(\mathcal{D})$ is defined by:

$$Cyl(\hat{\pi}) \;=\; \{\, \pi \in Paths(\mathcal{D}) \mid \hat{\pi} \text{ is a prefix of } \pi \,\}$$

Probability measure

Pr is the unique *probability measure* defined by:

$$Pr(Cyl(s_0 \ldots s_n)) \;=\; \iota_{\text{init}}(s_0) \cdot \mathbf{P}(s_0\, s_1 \ldots s_n)$$

where $\mathbf{P}(s_0\, s_1 \ldots s_n) \;=\; \prod_{0 \leqslant i < n} \mathbf{P}(s_i, s_{i+1})$ for $n > 0$ and $\mathbf{P}(s_0) = 1$.

Measurability

Measurability theorem

Events $\lozenge G$, $\square G$, $\overline{F} \, \mathsf{U} \, G$, $\square\lozenge G$ and $\lozenge\square G$ are measurable on any DTMC.

Proof:

To show this, every event has to be expressed as allowed operations (complement and/or countable unions) of the events — our cylinder sets!— of a DTMC.

Note that $\square G = \overline{\lozenge \overline{G}}$ and $\lozenge\square G = \overline{\square\lozenge\overline{G}}$.

It remains to prove the measurability for the remaining three cases.

Proof for $\lozenge G$

Which event does $\lozenge G$ exactly mean?

the union of all cylinders $Cyl(s_0 \ldots s_n)$ where

$s_0 \ldots s_n$ is a finite path in $\mathcal{D}$ with $s_0, \ldots, s_{n-1} \notin G$ and $s_n \in G$, i.e.,

$$\lozenge G \;=\; \bigcup_{s_0 \ldots s_n \in Paths^*(\mathcal{D}) \cap (S \setminus G)^* G} Cyl(s_0 \ldots s_n)$$

Thus $\lozenge G$ is measurable.

As all cylinder sets are pairwise disjoint, its probability is defined by:

$$\begin{aligned} Pr(\lozenge G) &= \sum_{s_0 \ldots s_n \in Paths^*(\mathcal{D}) \cap (S \setminus G)^* G} Pr\big(Cyl(s_0 \ldots s_n)\big) \\ &= \sum_{s_0 \ldots s_n \in Paths^*(\mathcal{D}) \cap (S \setminus G)^* G} \iota_{\text{init}}(s_0) \cdot \mathbf{P}(s_0 \ldots s_n) \end{aligned}$$

A similar proof strategy applies to the case $\overline{F} \, \mathsf{U} \, G$.

Proof for $\square\lozenge G$

Reachability probabilities: Knuth's die

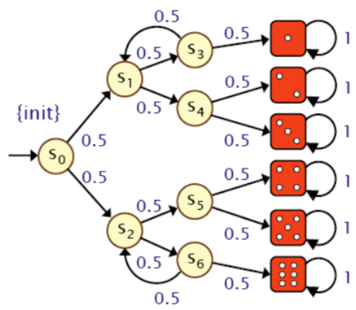

- Consider the event $\lozenge 4$
- Using the previous theorem we obtain:

  $Pr(\lozenge 4) = \sum_{s_0 \ldots s_n \in (S \setminus 4^*) 4} \mathbf{P}(s_0 \ldots s_n)$
- This yields:
  $\mathbf{P}(s_0 s_2 s_5 4) + \mathbf{P}(s_0 s_2 s_6 s_2 s_5 4) + \ldots\ldots$
- Or: $\sum_{k=0}^{\infty} \mathbf{P}(s_0 s_2 (s_6 s_2)^k s_5 4)$
- Or: $\frac{1}{8} \cdot \sum_{k=0}^{\infty} \left(\frac{1}{4}\right)^k$
- Geometric series: $\frac{1}{8} \cdot \frac{1}{1 - \frac{1}{4}} \;=\; \frac{1}{8} \cdot \frac{4}{3} \;=\; \frac{1}{6}$

There is however an simpler way to obtain reachability probabilities!

Overview

Reachability probabilities in finite DTMCs

Problem statement

Let $\mathcal{D}$ be a DTMC with finite state space S, $s \in S$ and $G \subseteq S$.

Aim: determine $Pr(s \models \Diamond G) = Pr_s(\Diamond G) = Pr_s\{\pi \in Paths(s) \mid \pi \in \Diamond G\}$

where Pr_s is the probability measure in $\mathcal{D}$ with single initial state s.

Characterisation of reachability probabilities

- Let variable $x_s = Pr(s \models \Diamond G)$ for any state s
  - if G is not reachable from s, then $x_s = 0$
  - if $s \in G$ then $x_s = 1$
- For any state $s \in Pre^*(G) \setminus G$:

$$x_s = \underbrace{\sum_{t \in S \setminus G} \mathbf{P}(s,t) \cdot x_t}_{\text{reach } G \text{ via } t \in S \setminus G} + \underbrace{\sum_{u \in G} \mathbf{P}(s,u)}_{\text{reach } G \text{ in one step}}$$

Reachability probabilities: Knuth's die

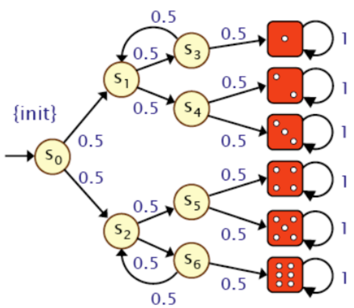

- Consider the event $\Diamond 4$
- Using the previous characterisation we obtain:

$x_1 = x_2 = x_3 = x_5 = x_6 = 0$ and $x_4 = 1$

$x_{s_1} = x_{s_3} = x_{s_4} = 0$

$x_{s_0} = \frac{1}{2} x_{s_1} + \frac{1}{2} x_{s_2}$

$x_{s_2} = \frac{1}{2} x_{s_5} + \frac{1}{2} x_{s_6}$

$x_{s_5} = \frac{1}{2} x_5 + \frac{1}{2} x_4$

$x_{s_6} = \frac{1}{2} x_{s_2} + \frac{1}{2} x_6$

- Gaussian elimination yields:

$x_{s_5} = \frac{1}{2}$, $x_{s_2} = \frac{1}{3}$, $x_{s_6} = \frac{1}{6}$, and $\boxed{x_{s_0} = \frac{1}{6}}$

Linear equation system

Reachability probabilities as linear equation system

- Let $S_? = Pre^*(G) \setminus G$, the states that can reach G by > 0 steps
- $\mathbf{A} = \big(\mathbf{P}(s,t)\big)_{s,t \in S_?}$, the transition probabilities in $S_?$
- $\mathbf{b} = (b_s)_{s \in S_?}$, the probs to reach G in 1 step, i.e., $b_s = \sum_{u \in G} \mathbf{P}(s,u)$

Then: $\mathbf{x} = (x_s)_{s \in S_?}$ with $x_s = Pr(s \models \Diamond G)$ is the unique solution of:

$$\mathbf{x} = \mathbf{A} \cdot \mathbf{x} + \mathbf{b} \quad \text{or} \quad (\mathbf{I} - \mathbf{A}) \cdot \mathbf{x} = \mathbf{b}$$

where $\mathbf{I}$ is the identity matrix of cardinality $|S_?| \times |S_?|$.

Reachability probabilities: Knuth's die

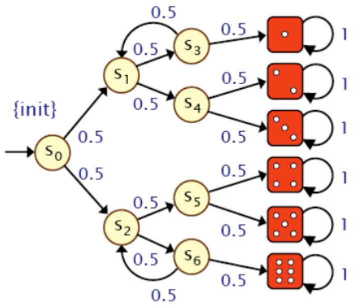

- Consider the event $\lozenge 4$
- $S_? = \{ s_0, s_2, s_5, s_6 \}$

$$\begin{pmatrix} 1 & -\frac{1}{2} & 0 & 0 \\ 0 & 1 & -\frac{1}{2} & -\frac{1}{2} \\ 0 & 0 & 1 & 0 \\ 0 & -\frac{1}{2} & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} x_{s_0} \\ x_{s_2} \\ x_{s_5} \\ x_{s_6} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ \frac{1}{2} \\ 0 \end{pmatrix}$$

- Gaussian elimination yields:

$x_{s_5} = \frac{1}{2}$, $x_{s_2} = \frac{1}{3}$, $x_{s_6} = \frac{1}{6}$, and $\boxed{x_{s_0} = \frac{1}{6}}$

Constrained reachability probabilities

Problem statement

Let $\mathcal{D}$ be a DTMC with finite state space S, $s \in S$ and $\overline{F}, G \subseteq S$.

Aim: $Pr(s \models \overline{F} \,\mathsf{U}\, G) \;=\; Pr_s(\overline{F} \,\mathsf{U}\, G) \;=\; Pr_s\{ \pi \in Paths(s) \mid \pi \models \overline{F} \,\mathsf{U}\, G \}$

where Pr_s is the probability measure in $\mathcal{D}$ with single initial state s.

Characterisation of constrained reachability probabilities

- Let variable $x_s = Pr(s \models \overline{F} \,\mathsf{U}\, G)$ for any state s
  - if G is not reachable from s via $\overline{F}$, then $x_s = 0$
  - if $s \in G$ then $x_s = 1$
- For any state $s \in (Pre^*(G) \cap \overline{F}) \setminus G$:

$$x_s \;=\; \sum_{t \in S \setminus G} \mathbf{P}(s, t) \cdot x_t \;+\; \sum_{u \in G} \mathbf{P}(s, u)$$

Proof

In the previous characterisation we basically set:

- $S_{=1} \;=\; G$
- $S_{=0} \;=\; \{ s \in S \mid Pr(\overline{F} \,\mathsf{U}\, G) = 0 \}$
- $S_? \;=\; S \setminus (S_{=0} \cup S_{=1})$

In fact any partition of S satisfying the following constraints will do:

- $G \subseteq S_{=1} \subseteq \{ s \in S \mid Pr(\overline{F} \,\mathsf{U}\, G) = 1 \}$
- $F \setminus G \subseteq S_{=0} \subseteq \{ s \in S \mid Pr(\overline{F} \,\mathsf{U}\, G) = 0 \}$
- $S_? \;=\; S \setminus (S_{=0} \cup S_{=1})$

In practice, $S_{=0}$ and $S_{=1}$ should be chosen as large as possible, as then $S_?$ is of minimal size, and the smallest linear equation system needs to be solved.

Thus $S_{=0} \;=\; \{ s \in S \mid Pr(\overline{F} \,\mathsf{U}\, G) = 0 \}$ and $S_{=1} \;=\; \{ s \in S \mid Pr(\overline{F} \,\mathsf{U}\, G) = 1 \}$.

These sets can easily be determined in linear time by a graph analysis.

Iteratively computing reachability probabilities

Theorem

The vector $\mathbf{x} = \Big(Pr(s \models \overline{F} \, \mathrm{U} \, G) \Big)_{s \in S_?}$ is the *unique* solution of:

$$\mathbf{y} \;=\; \mathbf{A}{\cdot}\mathbf{y} + \mathbf{b}$$

with $\mathbf{A}$ and $\mathbf{b}$ as defined before.

Furthermore, let:

$$\mathbf{x}^{(0)} = \mathbf{0} \quad \text{and} \quad \mathbf{x}^{(i+1)} \;=\; \mathbf{A}{\cdot}\mathbf{x}^{(i)} + \mathbf{b} \text{ for } 0 \leqslant i.$$

Then:

1. $\mathbf{x}^{(n)}(s) \;=\; Pr(s \models \overline{F} \, \mathrm{U}^{\leqslant n} \, G)$ for $s \in S_?$
2. $\mathbf{x}^{(0)} \leqslant \mathbf{x}^{(1)} \leqslant \mathbf{x}^{(2)} \leqslant \ldots \leqslant \mathbf{x}$
3. $\mathbf{x} = \lim_{n \to \infty} \mathbf{x}^{(n)}$

where $\overline{F} \, \mathrm{U}^{\leqslant n} G$ contains those paths that reach G via $\overline{F}$ within n steps.

Remark

Iterative algorithms to compute x

There are various algorithms to compute $\mathbf{x} = \lim_{n \to \infty} \mathbf{x}^{(n)}$ where:

$$\mathbf{x}^{(0)} = \mathbf{0} \quad \text{and} \quad \mathbf{x}^{(i+1)} \;=\; \mathbf{A}{\cdot}\mathbf{x}^{(i)} + \mathbf{b} \text{ for } 0 \leqslant i.$$

Then:

1. $\mathbf{x}^{(n)}(s) \;=\; Pr(s \models \lozenge^{\leqslant n} \, G)$ for $s \in S_?$
2. $\mathbf{x}^{(0)} \leqslant \mathbf{x}^{(1)} \leqslant \mathbf{x}^{(2)} \leqslant \ldots \leqslant \mathbf{x}$ and $\mathbf{x} = \lim_{n \to \infty} \mathbf{x}^{(n)}$

The Power method computes vectors $\mathbf{x}^{(0)}, \mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \ldots$ and aborts if:

$$\max_{s \in S_?} \left| \, x_s^{(n+1)} - x_s^{(n)} \, \right| \;<\; \varepsilon \quad \text{for some small tolerance } \varepsilon$$

This technique guarantees convergence.

Alternatives: e.g., Jacobi or Gauss-Seidel, successive overrelaxation (SOR).

Example: Knuth's die

- Let $G = \{ 1, 2, 3, 4, 5, 6 \}$
- Then $Pr(s_0 \models \lozenge G) = 1$
- And $Pr(s_0 \models \lozenge^{\leqslant k} G)$ for $k \in \mathbb{N}$ is given by:

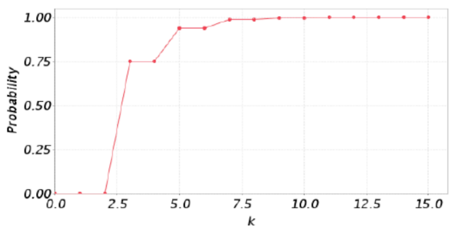

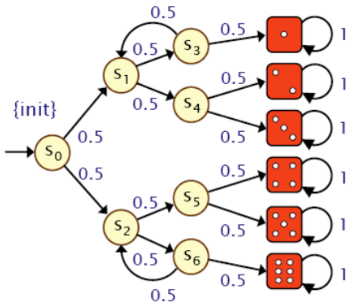

Overview

Recall: transient probability distribution

Transient distribution

$\mathbf{P}^n(s,t)$ equals the probability of being in state t after n steps given that the computation starts in s.

The probability of DTMC $\mathcal{D}$ being in state t after exactly n transitions is:

$$\Theta_n^{\mathcal{D}}(t) \;=\; \sum_{s \in S} \iota_{\text{init}}(s) \cdot \mathbf{P}^n(s,t) \;=$$

The function $\Theta_n^{\mathcal{D}}$ is the *transient state distribution* at epoch n of $\mathcal{D}$.

When considering $\Theta_n^{\mathcal{D}}$ as vector $(\Theta_n^{\mathcal{D}})_{t \in S}$ we have:

$$\Theta_n^{\mathcal{D}} \;=\; \iota_{\text{init}} \cdot \underbrace{\mathbf{P} \cdot \mathbf{P} \cdot \ldots \cdot \mathbf{P}}_{n \text{ times}} \;=\; \iota_{\text{init}} \cdot \mathbf{P}^n.$$

Computation: $\Theta_0^{\mathcal{D}} = \iota_{\text{init}}$ and $\Theta_{n+1}^{\mathcal{D}} = \Theta_n^{\mathcal{D}} \cdot \mathbf{P}$ for $n \geqslant 0$.

Reachability probabilities vs. transient probabilities

Aim

Compute $Pr(\Diamond^{\leqslant n} G)$ in DTMC $\mathcal{D}$. Observe that once a path π reaches G, then the remaining behaviour along π is not important. This suggests to make all states in G absorbing.

Let DTMC $\mathcal{D} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ and $G \subseteq S$. The DTMC $\mathcal{D}[G] = (S, \mathbf{P}_G, \iota_{\text{init}}, AP, L)$ with $\mathbf{P}_G(s,t) = \mathbf{P}(s,t)$ if $s \notin G$ and $\mathbf{P}_G(s,s) = 1$ if $s \in G$.

All outgoing transitions of $s \in G$ are replaced by a single self-loop at s.

Lemma

$$\underbrace{Pr(\Diamond^{\leqslant n} G)}_{\text{reachability in } \mathcal{D}} \;=\; \underbrace{Pr(\Diamond^{= n} G)}_{\text{reachability in } \mathcal{D}[G]} \;=\; \underbrace{\iota_{\text{init}} \cdot \mathbf{P}_G^n}_{\text{in } \mathcal{D}[G]} \;=\; \Theta_n^{\mathcal{D}[G]}$$

Constrained reachabilities vs. transient probabilities

Aim

Compute $Pr(\overline{F} \, \mathsf{U}^{\leqslant n} \, G)$ in DTMC $\mathcal{D}$. Observe (as before) that once a path π reaches G via $\overline{F}$, then the remaining behaviour along π is not important. Now also observe that once $s \in F \setminus G$ is reached, then the remaining behaviour along π is not important. This suggests to make all states in G <u>and</u> $F \setminus G$ absorbing.

Lemma

$$\underbrace{Pr(\overline{F} \, \mathsf{U}^{\leqslant n} \, G)}_{\text{reachability in } \mathcal{D}} \;=\; \underbrace{Pr(\Diamond^{= n} G)}_{\text{reachability in } \mathcal{D}[F \cup G]} \;=\; \underbrace{\iota_{\text{init}} \cdot \mathbf{P}_{F \cup G}^n}_{\text{in } \mathcal{D}[F \cup G]} \;=\; \Theta_n^{\mathcal{D}[F \cup G]}$$

Spare time tonight? Play Craps!

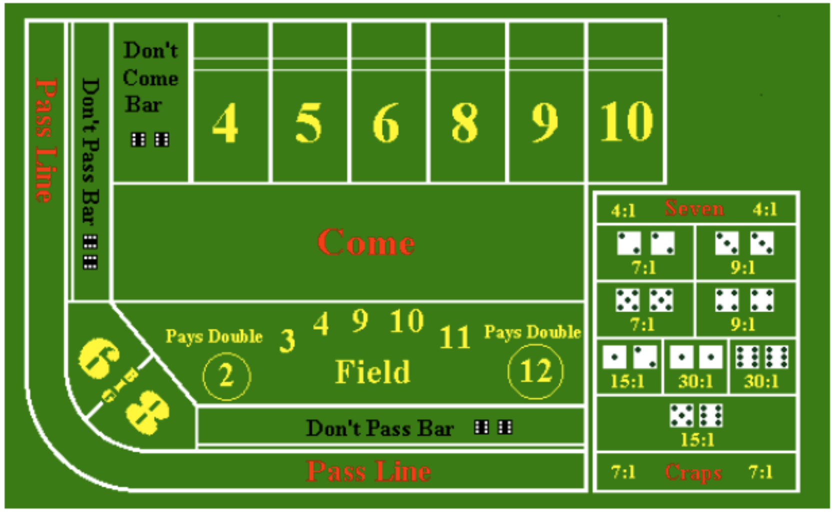

Craps

- Roll two dice and bet
- Come-out roll ("pass line" wager):
  - outcome 7 or 11: win
  - outcome 2, 3, or 12: lose ("craps")
  - any other outcome: roll again (outcome is "point")
- Repeat until 7 or the "point" is thrown:
  - outcome 7: lose ("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: lose
  - else: roll again
- Next roll(s):
  - 7: lose
  - point: win
  - else: roll again

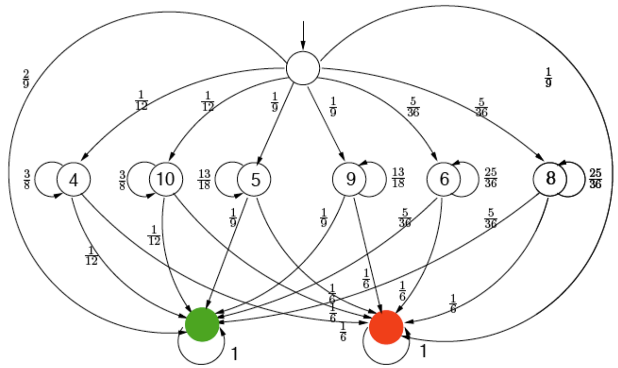

What is the probability to win the Craps game?

Summary

How to determine reachability probabilities?

1. Probabilities of sets of infinite paths defined using cylinders.
2. Events $\Diamond\, G$, $\Box\Diamond\, G$ and $\overline{F}\, \mathrm{U}\, G$ are measurable.
3. Reachability probabilities = unique solution of linear equation system.
4. Bounded reachabilities = transient probabilities in a modified DTMC.