Probabilistic Model Checking

Stefan Pranger

19. 05. 2022

1

Communication Protocol

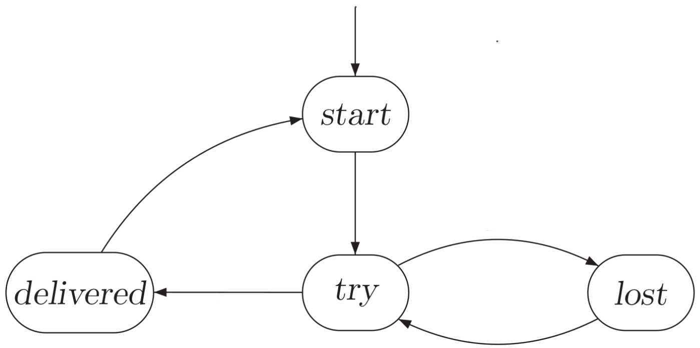

But $\mathcal{M}, start \models \exists \mathbf{G}\ \neg delivered$?

or $\mathcal{M}, start \models \forall \mathbf{F}\ delivered$?

Communication Protocol

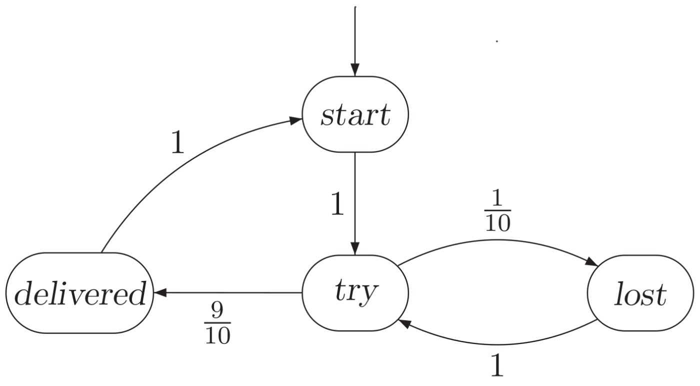

But $\mathcal{M}, \mathit{start} \models \exists\mathbf{G}\ \neg \mathit{delivered}$?

or $\mathcal{M}, \mathit{start} \models \forall\mathbf{F}\ \mathit{delivered}$?

Does not make sense with probabilities! $\rightarrow$ We *need* new descriptions for properties.

We *have* different models.

Markov Chains

Markov Chain $\mathcal{M} = (S, \mathbb{P}, s_0, AP, L)$

- S a set of states and initial state s_0,
- $\mathbb{P} : S \times S \to [0, 1]$, s.t.

$$\sum_{s' \in S} \mathbb{P}(s, s') = 1 \; \forall s \in S$$

- AP set of atomic states and $L : S \to 2^{AP}$ a labelling function.

Representation

How to represent a MC in code?

Representation

How to represent a MC in code?

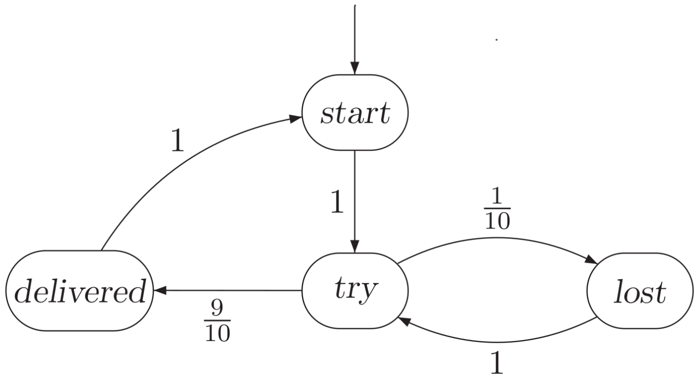

$$\longmapsto$$

$$\mathbf{A} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & \frac{1}{10} & \frac{9}{10} \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix}$$

Representation

How to represent a MC in code?

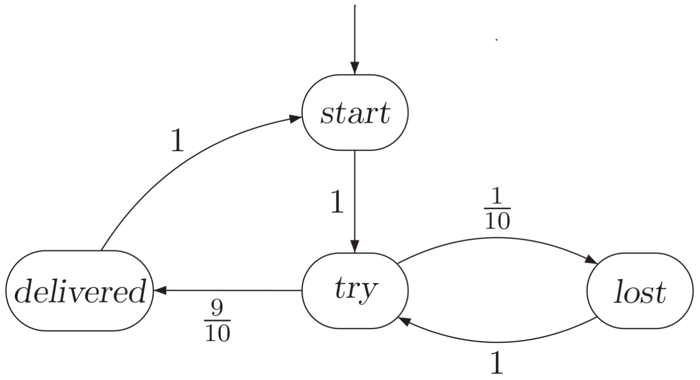

$$\longmapsto \quad \mathbf{A} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & \frac{1}{10} & \frac{9}{10} \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix}$$

- A *path* $\pi = s_0 s_1 s_2 \ldots \in S^\omega$, s.t. $\mathbb{P}(s_i, s_{i+1}) > 0, \forall i \geq 0$
- $Paths(\mathcal{M})$ is the set of all paths in $\mathcal{M}$ and
- $Paths_{fin}(\mathcal{M})$ is the set of all finite path fragments in $\mathcal{M}$.

Model Checking via $\mathcal{M}$

- Explicit CTL model checking allows *qualitative* model checking.
- We want to do *quantitative* model checking.
    - How *likely* is the system to fail?

$$Pr(\mathcal{M}, s \models \mathbf{F}\ s_{error})$$

    - Whats the *probability* of my message to arrive after infinitely many tries?

$$Pr(\mathcal{M}, s \models \mathbf{F}\ \text{delivered})$$

Events and Paths

In order to talk about probabilities of certain paths we need to talk about probability spaces.

- Outcomes = $\{HH, HT, TH, TT\}$
- Events = $\{HH\}, \{HT\}, \{TH\}, \{TT\}$

We could, for example, be interested in the events where H is thrown first = $\{HH\}, \{HT\}$.

What is a possible outcome in a specific Markov Chain $\mathcal{M}$?

Events and Paths

In order to talk about probabilities of certain paths we need to talk about probability spaces.

- Outcomes = $\{HH, HT, TH, TT\}$
- Events = $\{HH\}, \{HT\}, \{TH\}, \{TT\}$

We could, for example, be interested in the events where H is thrown first = $\{HH\}, \{HT\}$.

What is a possible outcome in a specific Markov Chain $\mathcal{M}$?

$\rightarrow$ an infinite path $\pi \in Paths(\mathcal{M})$!

- Outcomes = $Paths(\mathcal{M})$
- Events *of interest* are $\hat{\pi}_1, \hat{\pi}_2, \ldots \in Paths_{fin}(\mathcal{M})$ *that satisfy our property*
- Formally we introduce the *cylinder set* of a prefix:

$$Cyl(\hat{\pi}_i) = \{\pi \in Paths(\mathcal{M}) \mid \hat{\pi}_i \in \text{pref}(\pi)\}$$

Events and Paths

What is a possible outcome in a specific Markov Chain $\mathcal{M}$?

$\rightarrow$ an infinite path $\pi \in Paths(\mathcal{M})$!

- Outcomes = $Paths(\mathcal{M})$
- Events *of interest* are $\hat{\pi}_1, \hat{\pi}_2, \ldots \in Paths_{fin}(\mathcal{M})$ *that satisfy our property*
- Formally we introduce the *cylinder set* of a prefix:

$$Cyl(\hat{\pi}_i) = \{\pi \in Paths(\mathcal{M}) \mid \hat{\pi}_i \in \mathrm{pref}(\pi)\}$$

- The probability of one event of interest is then:

$$Pr(Cyl(\hat{\pi}_i)) = Pr(Cyl(s_0 s_1 \ldots s_n)) = \prod_{0 \leq i < n} \mathbb{P}(s_i, s_{i+1})$$

Example from the Blackboard

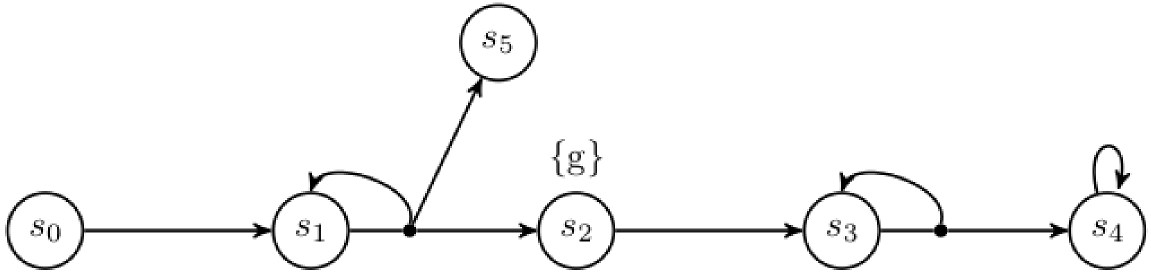

We are interested in all the finite path fragments $\hat{\pi}$ that satisfy '$\mathbf{F}g$':

They can be characterized by $\Pi_{\mathbf{F}g} = \{\hat{\pi} = s_0(s_1)^n s_2 \mid n \in \mathbb{N}\}$

Via a similar analysis we can see that there is no finite path fragment satisfying '$\mathbf{FG}g$', i.e. $\Pi_{\mathbf{GF}g} = \emptyset$

Modelling $\mathcal{M}$ with Code

We need to *bake* our models into *"code"* for a model checker.

A well-established language for that is the **PRISM**-language:

We need to describe the states and transitions of $\mathcal{M}$:

- In order to describe states we need variables:

```
x : [0..2] init 0;
b : bool init false;
```

- Transitions are modelled via so-called *commands*:

```
[] x=0 -> 0.8:(x'=0) + 0.2:(x'=1);
[] x1=0 & x2>0 & x2<10 -> 0.5:(x1'=1)&(x2'=x2+1) + 0.5:(x1'=2)&(x2'=x2-1);
```

Modelling $\mathcal{M}$ with Code

- Transitions are modelled via so-called *commands*:

```
[] x=0 -> 0.8:(x'=0) + 0.2:(x'=1);
```

A command consists of:

- The *guard* `x=0` describes the behaviour of $\mathcal{M}$ when the `x` equals 0.
- It is followed by a list of (*state*-) updates associated with their probabilities.

Note that the updates are indicated via a tick: `(x'=0)`.

Modelling $\mathcal{M}$ with Code

```
dtmc

label "success" = delivered=1;
label "lost" = lost=1;

module msg_delivery
    start: [0..1] init 1;
    try: [0..1] init 0;
    lost: [0..1] init 0;
    delivered: [0..1] init 0;


    [] start=1      -> 1: (start'=0) & (try'=1);
    [] try=1        -> 0.1: (try'=0) & (lost'=1) +
                       0.9: (try'=0) & (delivered'=1);
    [] lost=1       -> 1: (lost'=0) & (try'=1);
    [] delivered=1  -> 1: (delivered'=0) & (start'=1);

endmodule
```

Reachability Probabilities

Let $B \subseteq S$ be a set of states. We are interested in

$$Pr(\mathcal{M}, s_0 \models \mathbf{F} B).$$

Reachability Probabilities

Let $B \subseteq S$ be a set of states. We are interested in

$$Pr(\mathcal{M}, s_0 \models \mathbf{F}B).$$

We can characterize all path fragments π that satisfy $\mathbf{F}B$ with the set

$$\Pi_{\mathbf{F}B} = Paths_{fin}(\mathcal{M}) \cap (S \setminus B)^* B$$

All $\hat{\pi} \in \Pi_{\mathbf{F}B}$ are pairwise disjoint, hence:

$$Pr(\mathcal{M}, s_0 \models \mathbf{F}B) = \sum_{\hat{\pi} \in \Pi_{\mathbf{F}B}} Pr(Cyl(\hat{\pi}))$$

Computing $Pr(\mathcal{M}, s_0 \models \mathbf{F} B)$

We want an algorithmic way to compute the reachability probability.

Let x_s be the probability to reach B from s and $\widetilde{S} \subseteq S \setminus B$ be the set of states from which B is reachable.

We compute the probability $\mathcal{M}, s \models \mathbf{F} B$ via:

- The probability to reach B in one step: $\sum_{u \in B} \mathbb{P}(s, u)$
- and the probability to reach B via a path fragment $s\ t\ \ldots\ u$: $\sum_{t \in \widetilde{S}} \mathbb{P}(s, t) \cdot x_t$
- Together

$$x_s = \sum_{u \in B} \mathbb{P}(s, u) + \sum_{t \in \widetilde{S}} \mathbb{P}(s, t) \cdot x_t$$

Computing $Pr(\mathcal{M}, s_0 \models \mathbf{F} B)$

For $\mathbf{x} = (x_s)_{s \in \widetilde{S}}$ we want to compute

$$\mathbf{x} = \mathbf{A}\mathbf{x} + \mathbf{b},$$

where:

- $\mathbf{A}$ is the matrix of $\mathcal{M}_{\widetilde{S}}$ and
- $\mathbf{b} = (b_s)_{s \in \widetilde{S}}$ contains the probabilities to reach B in one step.

We rewrite this problem into:

$$(\mathbf{Id} - \mathbf{A})\mathbf{x} = \mathbf{b}.$$

Back to the Communication Protocol

Back to the Communication Protocol

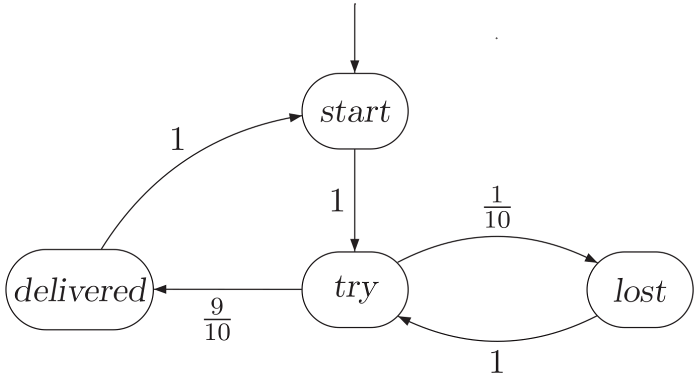

$$
\begin{aligned}
x_{start} &= x_{try} \\
x_{try} &= \frac{1}{10} x_{lost} + \frac{9}{10} \\
x_{lost} &= x_{try}
\end{aligned}
\qquad
\begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & -\frac{1}{10} \\ 0 & -1 & 1 \end{bmatrix} \cdot \mathbf{x} = \begin{pmatrix} 0 \\ \frac{9}{10} \\ 0 \end{pmatrix}
$$

Back to the Communication Protocol

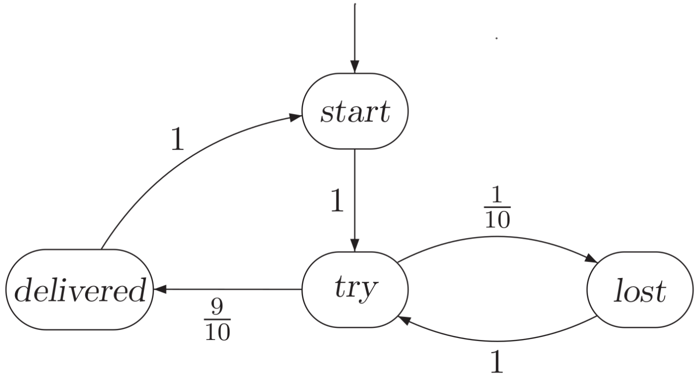

$$\begin{aligned} x_{start} &= x_{try} \\ x_{try} &= \frac{1}{10} x_{lost} + \frac{9}{10} \\ x_{lost} &= x_{try} \end{aligned}$$

$$\begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & -\frac{1}{10} \\ 0 & -1 & 1 \end{bmatrix} \cdot \mathbf{x} = \begin{pmatrix} 0 \\ \frac{9}{10} \\ 0 \end{pmatrix}$$

Complexity? Improvements? Unique Solution?

Computing $Pr(\mathcal{M}, s_0 \models \mathbf{F} B)$

$$(\mathbf{Id} - \mathbf{A})\mathbf{x} = \mathbf{b}$$

might have more than one solution.

We want to find the *least solution* in $[0, 1]^{\widetilde{S}}$. For that we consider *constrained reachability*:

$$\mathcal{M}, s \models C \ \mathbf{U}^{\leq n} B$$

where $C \ \mathbf{U}^{\leq n} B$ means that B should be reached within n steps while only passing through states in C.

Computing $Pr(\mathcal{M}, s_0 \models \mathbf{F} B)$

$$(\mathbf{Id} - \mathbf{A})\mathbf{x} = \mathbf{b}$$

might have more than one solution.

We want to find the *least solution* in $[0,1]^{\widetilde{S}}$. For that we consider *constrained reachability*:

$$\mathcal{M}, s \models C \ \mathbf{U}^{\leq n} B$$

where $C \ \mathbf{U}^{\leq n} B$ means that B should be reached within n steps while only passing through states in C.

First, some analysis of the problem:

- $B \subseteq S_{=1} \subseteq \{s \in S \mid Pr(s \models C \ \mathbf{U} \ B) = 1\}$,
- $S \setminus (C \cup B) \subseteq S_{=0} \subseteq \{s \in S \mid Pr(s \models C \ \mathbf{U} \ B) = 0\}$ and
- $S_? = S \setminus (S_{=1} \cup S_{=0})$

Computing $Pr(\mathcal{M}, s_0 \models \mathbf{F} B)$

We still need to handle $S_?$-states for which we compute the least solution.

$$\mathbf{x}^{(n+1)} = \mathbf{A}\mathbf{x}^{(n)} + \mathbf{b}, \text{with } \mathbf{x}^{(0)} = \mathbf{0}$$

where

$$\mathbf{x}^{(n)} = (x_s)_{s \in S_?} \text{ and } x_s^{(n)} = \mathcal{M}, s \models C \; \mathbf{U}^{\leq n} S_{=1}$$

Computing $Pr(\mathcal{M}, s_0 \models \mathbf{F} B)$

We still need to handle $S_?$-states for which we compute the least solution.

$$\mathbf{x}^{(n+1)} = \mathbf{A}\mathbf{x}^{(n)} + \mathbf{b}, \text{ with } \mathbf{x}^{(0)} = \mathbf{0}$$

where

$$\mathbf{x}^{(n)} = (x_s)_{s \in S_?} \text{ and } x_s^{(n)} = \mathcal{M}, s \models C \ \mathbf{U}^{\leq n} S_{=1}$$

This gives us a recipe to compute $Pr(\mathcal{M}, s_0 \models \mathbf{F} B)$:

- Run a graph-based algorithm to determine $S_{=0}$, $S_{=1}$ and $S_?$.
- Compute the probabilities to reach $S_{=1}$ from $S_?$.

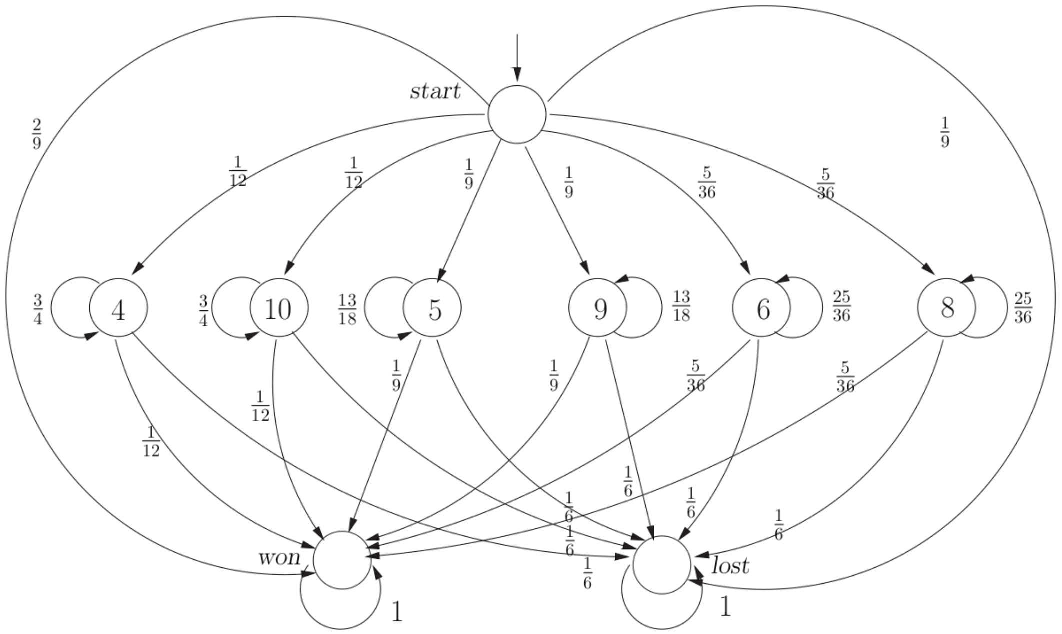

start
4
10
5
9
6
8
won
lost
1
1
$\frac{2}{9}$
$\frac{1}{9}$
$\frac{1}{12}$
$\frac{5}{36}$
$\frac{3}{4}$
$\frac{13}{18}$
$\frac{25}{36}$
$\frac{1}{6}$

Transient State Probabilities

We will consider a slightly different algorithm:

$$\mathbf{A}^n = \mathbf{A} \cdot \mathbf{A} \cdot \mathbf{A} \cdot \mathbf{A} \cdot \dots \cdot \mathbf{A}$$

contains the probability to be in state t after n steps in entry $\mathbf{A}^n(s, t)$.

We call

$$\Theta_n^{\mathcal{M}}(t) = \sum_{s \in S} \mathbf{A}^n(s, t)$$

the *transient state probability* for state t.

Transient State Probabilities

Let's consider $(\Theta_n^{\mathcal{M}}(t))_{s \in S}$, the vector of transient state probabilities for the nth step.

We can compute $Pr(\mathcal{M}, s_0 \models \mathbf{F}^{\leq n} B)$ in a modified Markov chain:

$$\mathcal{M}_B = (S, s_0, \mathbb{P}_B, AP, L)$$

where:

- $\mathbb{P}_B(s,t) = \mathbb{P}(s,t)$ if $s \notin B$
- $\mathbb{P}_B(s,s) = 1$ if $s \in B$
- $\mathbb{P}_B(s,t) = 0$ if $s \in B$ and $t \notin B$

i.e. all $s \in B$ become sinks and B cannot be left anymore.

Transient State Probabilities

- $\mathbb{P}_B(s,t) = \mathbb{P}(s,t)$ if $s \notin B$
- $\mathbb{P}_B(s,s) = 1$ if $s \in B$
- $\mathbb{P}_B(s,t) = 0$ if $s \in B$ and $t \notin B$

i.e. all $s \in B$ become sinks and B cannot be left anymore.

We then have

$$Pr(\mathcal{M}, s \models \mathbf{F}^{\leq n} B) = Pr(\mathcal{M}_B, s \models \mathbf{F}^{=n} B)$$

and therefore

$$Pr(\mathcal{M}, s \models \mathbf{F}^{\leq n} B) = \sum_{t \in B} \Theta_n^{\mathcal{M}_B}(t)$$

Computing $Pr(\mathcal{M}, s \models \mathbf{F}^{\leq n} B)$ via Transient State Probabilities

We have the following algorithm to compute $Pr(\mathcal{M}, s \models \mathbf{F}^{\leq n} B)$:

- $\Theta_0^{\mathcal{M}}(t) = \mathbf{e}_i$, i.e. the unit vector with 1 at the ith position and 0 else.
- For $k = 0$ up to $n - 1$: $\Theta_{k+1}^{\mathcal{M}}(t) = \mathbf{A} \cdot \Theta_k^{\mathcal{M}}(t)$
- $Pr(\mathcal{M}, s \models \mathbf{F}^{\leq n} B) = \sum_{t \in B} \Theta_n^{\mathcal{M}_B}(t)$

Extra

Let $\mathcal{M} = (S = \{s_0, s_1, s_2, \ldots s_9\}, s_0 = s_0, \mathbb{P}, \{0, 1, 2, \ldots, 9\}, L)$ be a MC with

$$\mathbb{P} = \frac{1}{10} \begin{bmatrix} 1 & \ldots & 1 \\ \vdots & \ddots & \vdots \\ 1 & \ldots & 1 \end{bmatrix}$$

Further let $f : S^\omega \to [0, 1)$ s.t.

$$f(\pi) = f(s_0 s_1 s_2 \ldots) = 0.L(s_1)L(s_2)\ldots$$

where $L(s_i) = i$

$f^{-1} : [0, 1) \to S^\omega$ can be defined similarly. Hence we have a bijection between S^ω and $[0, 1)$ and therefore there must be uncountably infinite many $\pi \in S^\omega$.