Probabilistic Aspects of Computer Science: DTMC

Serge Haddad

LMF, ENS Paris-Saclay & CNRS & INRIA

MPRI M1

Plan

1. Discrete Event Systems

Renewal Processes with Arithmetic Distribution

Discrete Time Markov Chains (DTMC)

Finite DTMC

A path of a discrete-event system

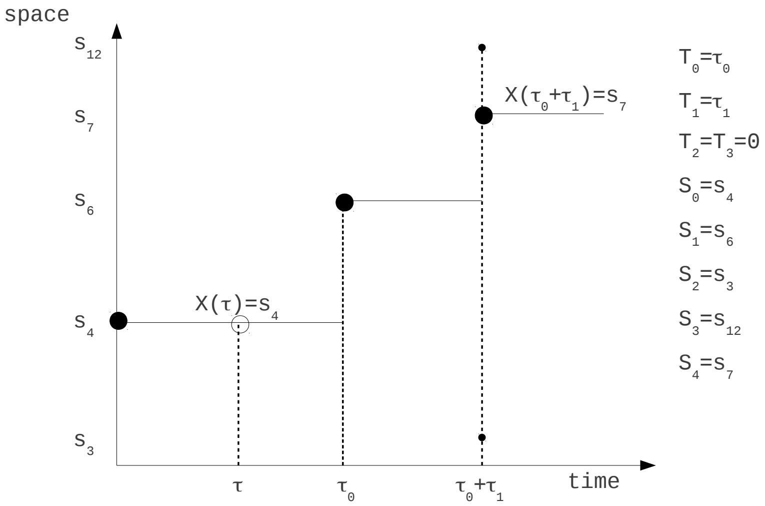

Observe that s_3 and s_{12} are hidden to the observations $\{X(\tau)\}$.

Discrete-event systems (DES)

Informally

An execution of a discrete-event system is an infinite sequence of events: $e_1, e_2, \ldots$ occurring after some (possibly null) delay.

More formally

A discrete-event system is defined by two families of random variables:
(do you need some explanations?)

- $S_0, S_1, S_2, \ldots$ such that S_0 is the initial state and S_i is the state of the system after the occurrence of e_i.
- $T_0, T_1, T_2, \ldots$ such that T_0 is the elapsed time before the occurrence of e_0 and T_i is the elapsed time between the occurrences of e_i and e_{i+1}.

DES: basic properties and notations

A DES is non Zeno *almost surely* iff $\mathbf{Pr}(\sum_{i \in \mathbb{N}} T_i = \infty) = 1$.
(cf Achilles and the tortoise)

When a DES is non Zeno:

- $N(\tau) \stackrel{\text{def}}{=} \min(\{n \mid \sum_{k \leq n} T_k > \tau\})$, the number of events up to time τ, is defined *for almost every sample*.
- $X(\tau) \stackrel{\text{def}}{=} S_{N(\tau)}$ is the observable state at time τ.

When $\mathbf{Pr}(S_0 = s) = 1$, one says that the process *starts* in s.

Analysis of DES

Two kinds of analysis

- **Transient analysis.** Computation of measures depending on the elapsed time since the initial state:

  for instance π_τ, the distribution of X_τ.

- **Steady-state analysis.** Computation of measures depending on the long-run behaviour of the system:

  for instance $\pi_\infty \stackrel{\text{def}}{=} \lim_{\tau \to \infty} \pi_\tau$.
  (requires to establish its existence)

Analysis of a web server

- What is the probability that a connection is established within 10s?
- What is the mean number of clients on the long run?

Performance indices

Definition

- A *performance index* is a function from states to numerical values.
- The measure of an index f w.r.t. to a state distribution π is given by:
  $$\sum_{s \in S} \pi(s) \cdot f(s)$$
- When range of f is $\{0, 1\}$ it is an *atomic property* and its measure can be rewritten:
  $$\sum_{s \models f} \pi(s)$$

Analysis of a web server

- Let S' be the subset of states such that the server is available.
  Then the probability that the server is available at time 10 is:
  $$\sum_{s \in S'} \pi_{10}(s)$$
- Let s be a state and $cl(s)$ be the number of clients in s.
  Then the mean number of clients on the long run is:
  $$\sum_{s \in S} \pi_{\infty}(s) \cdot cl(s)$$

Renewal process

A renewal process is a very simple case of DES.

- It has a single state.
- The time intervals between events are integers obtained by sampling i.i.d. (independent and identically distributed) random variables.
- *Renewal instants* are the instants corresponding to the occurrence of events.

Using a lamp

- A bulb is used by some lamp.
- The bulb seller provides some information about the quality of the bulb.

for $n > 0$, f_n is the probability that the bulb duration is n days

Discrete time Markov chain (DTMC)

A DTMC is a stochastic process which fulfills:

- For all n, T_n is the constant 1.
- The process is *memoryless* (X_n denotes $X(n)$)

$$\mathbf{Pr}(X_{n+1} = s_j \mid X_0 = s_{i_0}, ..., X_{n-1} = s_{i_{n-1}}, X_n = s_i)$$

$$= \mathbf{Pr}(X_{n+1} = s_j \mid X_n = s_i) = \mathtt{P}[i,j] \stackrel{\text{def}}{=} p_{ij} \quad \text{independent of } n$$

The behavior of a DTMC is defined by X_0 and $\mathtt{P}$

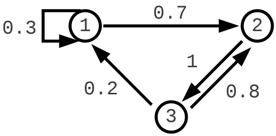

P

$$\begin{pmatrix} 0.3 & 0.7 & 0.0 \\ 0.0 & 0.0 & 1.0 \\ 0.2 & 0.8 & 0.0 \end{pmatrix}$$

Two infinite DTMCs

A random walk

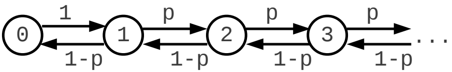

A simulation of renewal process

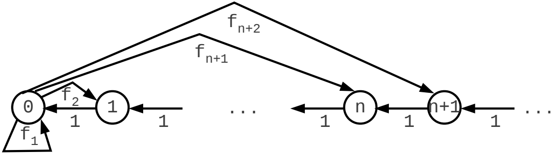

Transient analysis of a DTMC

The transient analysis is easy *(and effective in the finite case)* :

$$\pi_n = \pi_0 \cdot \mathtt{P}^n \text{ with } \pi_n \text{ the distribution of } X_n$$

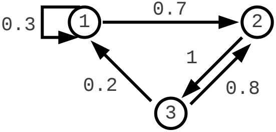

$$\mathbf{P} = \begin{pmatrix} 0.3 & 0.7 & 0.0 \\ 0.0 & 0.0 & 1.0 \\ 0.2 & 0.8 & 0.0 \end{pmatrix}$$

$$\pi_0 = \begin{pmatrix} 1.0 & 0.0 & 0.0 \end{pmatrix} \qquad \pi_1 = \begin{pmatrix} 0.3 & 0.7 & 0.0 \end{pmatrix}$$

$$\pi_2 = \begin{pmatrix} 0.09 & 0.21 & 0.70 \end{pmatrix} \qquad \pi_3 = \begin{pmatrix} 0.167 & 0.623 & 0.21 \end{pmatrix}$$

Plan

Discrete Event Systems

2 Renewal Processes with Arithmetic Distribution

Discrete Time Markov Chains (DTMC)

Finite DTMC

Analysis of renewal process

Let u_n be the probability that n is a renewal instant.

What can be expected about u_n when n goes to ∞?

Using a lamp

- Let $f_2 = 0.4$ and $f_3 = 0.6$
- Then $u_0 = 1$ (initial bulb), $u_1 = 0$
- $u_2 = f_2 = 0.4$, $u_3 = f_3 = 0.6$ (one change)
- $u_4 = u_2 f_2 = 0.16$, $u_5 = u_3 f_2 + u_2 f_3 = 0.48$ (two changes)

The renewal equation

$$u_n = u_0 f_n + \cdots + u_{n-1} f_1 \text{ when } n > 0 \qquad (1)$$

Is there a limit?

$u_{10} = 0.3696, u_{20} = 0.38867558, u_{30} = 0.38423714, u_{40} = 0.38463453, u_{50} = 0.38461546$

What is its value?

The average renewal time

The average renewal time is: $\mu \stackrel{\text{def}}{=} \sum_{i \in \mathbb{N}} i f_i$
(μ may be infinite)

Intuitively, during μ time units, there is on average one renewal instant.

So one could expect: $\lim_{n \to \infty} u_n = \mu^{-1} \stackrel{\text{def}}{=} \eta$.
(with the convention $\infty^{-1} \stackrel{\text{def}}{=} 0$)

Example (continued)

- $\mu = 0.4 \cdot 2 + 0.6 \cdot 3 = 2.6$
- $\eta = 0.38461538$ very close to $u_{50} = 0.38461546$

Is there always a limit?

Let $f_2 = 1$, then $\lim_{n\to\infty} u_{2n} = 1$ and $\lim_{n\to\infty} u_{2n+1} = 0$.

Periodicity

- The *periodicity* of the renewal process is: $\gcd(\{i \mid f_i > 0\})$.
- When the periodicity is 1 the renewal process is said *aperiodic*.

The renewal theorem

When the process is aperiodic, $\lim_{n\to\infty} u_n$ exists and is equal to η.

Another renewal equation

Let $\rho_k \overset{\text{def}}{=} \sum_{i>k} f_i$ be the probability that the duration before a new renewal instant is strictly greater than k.

$$\mu = \sum_{i \in \mathbb{N}} i f_i = \sum_{i \in \mathbb{N}} \sum_{0 \le k < i} f_i = \sum_{k \in \mathbb{N}} \sum_{i > k} f_i = \sum_{k \in \mathbb{N}} \rho_k$$

Let L_{nk} (with $k \le n$) be the event "The last renewal instant in $[0, n]$ is k"

Observe that: $\mathbf{Pr}(L_{nk}) = u_k \rho_{n-k}$.

Letting n fixed and summing over k, one obtains:

$$\rho_0 u_n + \rho_1 u_{n-1} + \cdots + \rho_n u_0 = 1 \qquad (2)$$

A first sufficient condition

If $\limsup_{n\to\infty} u_n \leq \eta$ then $\lim_{n\to\infty} u_n$ exists and is equal to η.

Sketch of proof (for μ finite)

Let $u_{n_1}, u_{n_2}, \ldots$ be an arbitrary subsequence converging toward η' ($\leq \eta$).

Let us select some integer $r > 0$ and some $\varepsilon > 0$.

There exists m such that:

$$\forall n_i \geq m \; |u_{n_i} - \eta'| \leq \varepsilon \wedge \forall 1 \leq r' \leq r \; u_{n_i - r'} - \eta \leq \varepsilon$$

Using (2) for $n_i \geq m$:

$$\rho_0(\eta' + \varepsilon) + (\eta + \varepsilon) \sum_{r'=1}^{r} \rho_{r'} + \sum_{r' > r} \rho_{r'} \geq 1$$

Letting ε goes to 0, $\rho_0 \eta' + \eta \sum_{r'=1}^{r} \rho_{r'} + \sum_{r'>r} \rho_{r'} \geq 1$

Letting r goes to ∞ (recall that $\rho_0 = 1$), $\eta' + \eta(\mu - 1) \geq 1$

Which can be rewritten as: $\eta' - \eta \geq 0$ implying $\eta' = \eta$

Standard results

Let $a_1, \dots, a_k$ be natural integers whose $\gcd$ is 1.

Then there exists n_0 such that:

$$\forall n \geq n_0 \; \exists \alpha_1, \dots, \alpha_k \in \mathbb{N} \; n = a_1\alpha_1 + \cdots + a_k\alpha_k$$

Sketch of proof

- Using Euclid algorithm, there exists $y_1, \dots, y_k \in \mathbb{Z}$ such that:
  $$1 = a_1y_1 + \cdots + a_ky_k$$
- Let us note $s = a_1 + \cdots + a_k$ and $x = \sup_i |y_i|(s-1)$.
- Let $n \geq xs$ and perform the Euclidian division of n by s.

Let $(x_{n,m})_{n,m\in\mathbb{N}}$ be a bounded family of reals.

Then there exists an infinite sequence of indices $m_1 < m_2 < \cdots$ such that:

For all $n \in \mathbb{N}$ the subsequence $(x_{n,m_k})_{k\in\mathbb{N}}$ is convergent.

Sketch of proof

- Build nested subsequences of indices $(m_k^n)_{k\in\mathbb{N}}$ such that: $(x_{n,m_k^n})_{k\in\mathbb{N}}$ is convergent.
- Pick the "diagonal" subsequence of indices $m_k = m_k^k$.

The key lemma

Let f be an aperiodic distribution and $(w_n)_{n\in\mathbb{N}}$ such that $\forall n$, $w_n \leq w_0$ and:

$$w_n = \sum_{k=1}^{\infty} f_k w_{n+k} \qquad (3)$$

Then for all n, $w_n = w_0$.

Sketch of proof

Let $A = \{k \mid f_k > 0\}$, $B = \{k \mid \exists k_1, \ldots, k_n \in A\ \exists \alpha_1, \ldots, \alpha_n \in \mathbb{N}\ k = \sum \alpha_i k_i\}$.
There exists n_0 such that $[n_0, \infty[\subseteq B$.

$w_0 = \sum_{k=1}^{\infty} f_k w_k \leq w_0 \sum_{k=1}^{\infty} f_k = w_0$ implying $w_k = w_0$ for all $k \in A$.

Let $k \in A$,
$w_k = \sum_{k'=1}^{\infty} f_{k'} w_{k+k'} \leq w_0 \sum_{k'=1}^{\infty} f_{k'} = w_0$ implying $w_{k+k'} = w_0$
for all $k, k' \in A$.
Iterating the process, one gets $w_k = w_0$ for all $k \in B$.

$w_{n_0-1} = \sum_{k=1}^{\infty} f_k w_{n_0-1+k} = w_0 \sum_{k=1}^{\infty} f_k = w_0$
Iterating this process, one concludes.

The dominated convergence theorem

Let $(a_m)_{m\in\mathbb{N}}$, $(v_m)_{m\in\mathbb{N}}$ and $(u_{m,n})_{m,n\in\mathbb{N}}$ be sequences of non negative reals.

Assume that:

- $\sum_{m\in\mathbb{N}} a_m v_m < \infty$;
- $\forall m, n\ u_{m,n} \leq v_m$;
- $\forall m\ \lim_{n\to\infty} u_{m,n} = \ell_m$.

Interpretation.

$(a_m)_{m\in\mathbb{N}}$ is a measure over $\mathbb{N}$, $(v_m)_{m\in\mathbb{N}}$ maps $\mathbb{N}$ to $\mathbb{R}_+$ integrable w.r.t. $(a_m)_{m\in\mathbb{N}}$.

For all n, $(u_{m,n})_{m,n\in\mathbb{N}}$ maps $\mathbb{N}$ to $\mathbb{R}_+$ and is bounded by $(v_m)_{m\in\mathbb{N}}$.

These mappings converge to ℓ_m.

Then:

$$\lim_{n\to\infty} \sum_{m\in\mathbb{N}} a_m u_{m,n} = \sum_{m\in\mathbb{N}} a_m \ell_m$$

Proof of the renewal theorem (μ finite)

Let $\nu \overset{\text{def}}{=} \limsup_{n\to\infty} u_n$ and $(r_m)_{m\in\mathbb{N}}$ such that $\nu = \lim_{m\to\infty} u_{r_m}$.
Define $u_{n,m} = u_{r_m - n}$ if $n \leq r_m$ and $u_{n,m} = 0$ otherwise.

There exist $m_1, m_2, \dots$ such that for all n,
$(u_{n,m_k})_{k\in\mathbb{N}}$ converges to a limit $w_n \leq w_0 = \nu$.

Equation (1) can be rewritten as: $u_{n,m_k} = \sum_{i=1}^{\infty} f_i u_{n+i,m_k}$
Letting k go ∞ yields: $w_n = \sum_{i=1}^{\infty} f_i w_{n+i}$ *(by dominated convergence theorem)*
Hence for all n, $w_n = \nu$.

Rewritting (2) for $n = m_k$ gives: $\rho_0 u_{0,m_k} + \rho_1 u_{1,m_k} + \cdots + \rho_{r_{m_k}} u_{r_{m_k},m_k} = 1$

For all fixed r, $\rho_0 u_{0,m_k} + \rho_1 u_{1,m_k} + \cdots + \rho_r u_{r,m_k} \leq 1$
Letting m_k go to ∞, $\nu \sum_{r'=0}^{r} \rho_{r'} \leq 1$ and letting r go to ∞, $\nu\mu \leq 1 \Leftrightarrow \nu \leq \eta$.

Generalizations (1)

Periodicity

Nothing more than a change of scale

> Let f be a distribution of period p with (non necessarily finite) mean μ.
> Then $\lim_{n\to\infty} u_{np} = p\mu^{-1}$ and for all n such that $n \mod p \neq 0$, $u_n = 0$.

Defective renewal process: $\sum_{n\in\mathbb{N}} f_n \leq 1$

$1 - \sum_{n\in\mathbb{N}} f_n$ is the probability that there will be no next renewal instant.
(e.g. a perfect bulb)

> The mean number of renewal instants, $\sum_{n\in\mathbb{N}} u_n$, is finite iff $\sum_{n\in\mathbb{N}} f_n < 1$.
> In this case,
> $$\sum_{n\in\mathbb{N}} u_n = \frac{1}{1 - \sum_{n\in\mathbb{N}} f_n} \text{ (so } \lim_{n\to\infty} u_n = 0\text{)}$$

Generalizations (2)

Delayed renewal process: The first renewal instant is no more 0 but follows a possibly defective distribution $\{b_n\}$ with $B \stackrel{\text{def}}{=} \sum_{n \in \mathbb{N}} b_n \leq 1$.

(e.g. initially there is a possibly perfect bulb in the lamp)

Let v_n be the probability of a renewal instant at time n in the delayed process.

- If $\lim_{n\to\infty} u_n = \omega$ then $\lim_{n\to\infty} v_n = B\omega$
- If $U \stackrel{\text{def}}{=} \sum_{n\to\infty} u_n$ is finite then $\sum_{n\to\infty} v_n = BU$

Instantaneous renewal process: $0 < f_0 < 1$

(e.g. f_0 is the probability that a new bulb is initially faulty)

u_n is now the mean number of renewals at time n.

Let f'_n be a standard renewal process with $f'_n = \frac{f_n}{1-f_0}$ for $n > 0$. Then:

$$u_n = \frac{u'_n}{1 - f_0}$$

Plan

Discrete Event Systems

Renewal Processes with Arithmetic Distribution

3 Discrete Time Markov Chains (DTMC)

Finite DTMC

Notations for a DTMC

$p^n_{i,j}$, the probability to reach in n steps state j from state i.

$f^n_{i,j}$, the probability to reach in n steps state j from state i *for the first time*

$f_{i,j} \stackrel{\text{def}}{=} \sum_{n\in\mathbb{N}} f^n_{i,j}$, the probability to reach j from i

$\mu_i \stackrel{\text{def}}{=} \sum_{n\in\mathbb{N}} n f^n_{i,i}$, the mean return time in i (only relevant if $f_{i,i} = 1$).

A renewal-like equation

$$\forall n > 0 \qquad p^n_{i,j} = \sum_{m=0}^{n} f^m_{i,j} p^{n-m}_{j,j}$$

DTMC and renewal processes

Some renewal processes

- The visits of a state i starting from i, is a renewal process with renewal distribution $\{f_{i,i}^n\}_n$. $p_{i,i}^n$ is the probability of a renewal instant at time n.
- The visits of a state i starting from j, is a delayed renewal process with delay distribution $\{f_{j,i}^n\}_n$. $p_{j,i}^n$ is the probability of a renewal instant at time n.

Classification of states w.r.t. the associated renewal process

- A state i is *transient* if $f_{i,i} < 1$,
  the probability of a return after a visit is strictly less than one.
- A state is *null recurrent* if $f_{i,i} = 1$ and $\mu_{i,i} = \infty$,
  the probability of a return after a visit is 1 and the mean return time is ∞.
- A state is *positive recurrent* if $f_{i,i} = 1$ and $\mu_{i,i} < \infty$,
  the probability of a return after a visit is 1 and the mean return time is finite.
- A state is *aperiodic* if its associated process is aperiodic.
  A state is *ergodic* if it is aperiodic and positive recurrent.

Example of classification

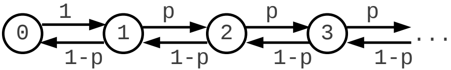

The value of p is critical.

- All states have period 2.
- If $p > \frac{1}{2}$ then all states are transient.
- If $p = \frac{1}{2}$ then all states are null recurrent.
- If $p < \frac{1}{2}$ then all states are positive recurrent.

Structure of a DTMC

A closed subset

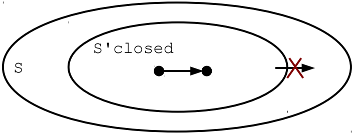

The closure of a state i

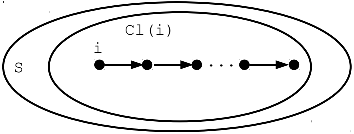

An irreducible DTMC: for all i, j

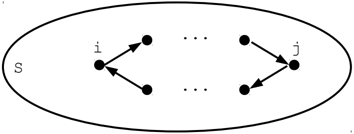

Irreducibility

All states of an irreducible DTMC are of the same kind.

Sketch of proof

Let i, j be two states.
There exist r and s such that $p^r_{i,j} > 0$ and $p^s_{j,i} > 0$. Observe that:

$$p^{n+r+s}_{i,i} \geq p^r_{i,j} p^n_{j,j} p^s_{j,i} \tag{4}$$

Transient vs recurrent. If $\sum_{n \in \mathbb{N}} p^n_{i,i}$ is finite then $\sum_{n \in \mathbb{N}} p^n_{j,j}$ is finite.

Null recurrence. If $\lim_{n \to \infty} p^n_{i,i} = 0$ then $\lim_{n \to \infty} p^n_{j,j} = 0$.

Periodicity. Assume i has periodicity $t \geq 1$.
Using (4), $r + s$ is a multiple of t.
So if n is not a multiple of t, then $p^n_{j,j} = 0$.
Thus the periodicity of j is a multiple of t.

Periodicity

Let $\mathcal{C}$ be irreducible with periodicity p. Then $S = S_0 \uplus S_1 \uplus \ldots \uplus S_{p-1}$ with:

$$\forall k \ \forall i \in S_k \ \forall j \in S \ p_{i,j} > 0 \Rightarrow j \in S_{(k+1) \bmod p}$$

Furthermore p is the greatest integer fulfilling this property.

Sketch of proof

Let i be a state. For all $0 \leq k < p$,
$S_k \stackrel{\text{def}}{=} \{j \mid \text{ there is a path from } i \text{ to } j \text{ with length equal to } lp + k \text{ for some } l\}$
Irreducibility implies $S = S_0 \cup S_1 \cup \ldots \cup S_{p-1}$
By definition, $\forall j \in S_k \ \forall j' \in S \ p_{j,j'} > 0 \Rightarrow j' \in S_{(k+1) \bmod p}$

Assume there exists $j \in S_k \cap S_{k'}$ with $k \neq k'$.
Since there is a path from j to i,
there exists at least a path from i to i whose length is not a multiple of p,
leading to a contradiction.

Assume there is some $p' > p$ fulfilling this property.
Then the periodicity of the renewal process is a multiple of p',
leading to another contradiction.

Characterization of ergodicity

Definitions.

- A distribution π is *invariant* if $\pi\mathbf{P} = \pi$.
- There is a *steady-state* distribution π for π_0 if $\pi = \lim_{n\to\infty} \pi_0\mathbf{P}^n$.

Let $\mathcal{C}$ be an irreducible DTMC whose states are aperiodic.

Existence of an invariant steady-state distribution.

Assume that the states of $\mathcal{C}$ are positive recurrent.
Then the limits $\lim_{n\to\infty} p^n_{j,i}$ exist and are equal to μ_i^{-1}.
Furthermore:

$$\sum_{i\in S} \mu_i^{-1} = 1 \text{ and } \forall i \ \mu_i^{-1} = \sum_{j\in S} \mu_j^{-1} p_{j,i}$$

Unicity of an invariant distribution.

Conversely, assume there exists $\mathbf{u}$ such that:
For all i, $\mathbf{u}_i \geq 0$, $\sum_{i\in S} \mathbf{u}_i = 1$ and $\mathbf{u}_i = \sum_{j\in S} \mathbf{u}_j p_{j,i}$.
Then for all i, $u_i = \mu_i^{-1}$.

Proof of existence

Applying renewal theory on the delayed renewal process:
$\lim_{n\to\infty} p_{j,i}^n = f_{j,i}\mu_i^{-1} = \mu_i^{-1}$ *(recurrence of states implies $f_{j,i} = 1$)*

Since $\mathbf{P}^n$ is a transition matrix, $\forall n\ \forall i\ \ \sum_{j\in S} p_{i,j}^n = 1$
Letting n go to infinity, yields: $\sum_{j\in S} \mu_j^{-1} \leq 1$

One has: $\sum_{j\in S} p_{k,j}^n p_{j,i} = p_{k,i}^{n+1}$
Letting n go to infinity, yields: $\sum_{j\in S} \mu_j^{-1} p_{j,i} \leq \mu_i^{-1}$

Summing $\sum_{j\in S} \mu_j^{-1} p_{j,i} \leq \mu_i^{-1}$ over i, one obtains:

$$\sum_{i\in S} \mu_i^{-1} \geq \sum_{i\in S}\sum_{j\in S} \mu_j^{-1} p_{j,i} = \sum_{j\in S} \mu_j^{-1} \sum_{i\in S} p_{j,i} = \sum_{j\in S} \mu_j^{-1}$$

One has equality of sums, so also equality of terms: $\mu_i^{-1} = \sum_{j\in S} \mu_j^{-1} p_{j,i}$

By iteration: $\mu_i^{-1} = \sum_{j\in S} \mu_j^{-1} p_{j,i}^n$
Letting n go to infinity, yields: *(by dominated convergence theorem)*
$\mu_i^{-1} = \sum_{j\in S} \mu_j^{-1}\mu_i^{-1}$ implying $\sum_{j\in S} \mu_j^{-1} = 1$

Proof of unicity

Assume there exists $\mathbf{u}$ such that:

$$\sum_{i \in S} \mathbf{u}_i = 1 \text{ and for all } i, \mathbf{u}_i \geq 0, \ \mathbf{u}_i = \sum_{j \in S} \mathbf{u}_j p_{j,i}$$

Let us pick some $\mathbf{u}_i > 0$, by iteration of the last equation: $\mathbf{u}_i = \sum_{j \in S} \mathbf{u}_j p^n_{j,i}$

Since the states of the chain are aperiodic, $\lim_{n \to \infty} p^n_{j,i}$ exists for all j.
Letting n go to ∞, $\mathbf{u}_i = \sum_{j \in S} \mathbf{u}_j \lim_{n \to \infty} p^n_{j,i}$ *(by dominated convergence theorem)*

There exists j such that $\lim_{n \to \infty} p^n_{j,i} > 0$.
So i is positive recurrent and then ergodic (since aperiodic).

So all states are ergodic implying that the limits of $p^n_{j,i}$ are μ_i^{-1}.
The previous equation can be rewritten as:

$$\mathbf{u}_i = \sum_{j \in S} \mathbf{u}_j \mu_i^{-1} = \mu_i^{-1}$$

Characterization of positive recurrence

Let $\mathcal{C}$ be irreducible. Then $\mathcal{C}$ is positive recurrent iff there exists $\mathbf{u}$ such that:
For all i, $\mathbf{u}[i] > 0$, $\mathbf{u} = \mathbf{u} \cdot \mathbf{P}$ and $\sum_{i \in S} \mathbf{u}[i] < \infty$

Sketch of necessity proof

Let p be the periodicity of $\mathcal{C}$ and $S = S_0 \uplus \ldots \uplus S_{p-1}$.

$\mathbf{P}^p$ defines an ergodic chain over S_0.

So there exists a positive vector $\mathbf{u}_0$ indexed by S_0 with:

$$\sum_{s \in S_0} \mathbf{u}_0[s] = 1 \text{ and } \mathbf{u}_0 \cdot \mathbf{P}^p = \mathbf{u}_0$$

Define for $0 < k < p$ and $s \in S_k$:

$$\mathbf{u}_k[s] = \sum_{s' \in S_{k-1}} \mathbf{u}_{k-1}[s']\mathbf{P}[s', s]$$

Then $\mathbf{u} = (\mathbf{u}_0, \ldots, \mathbf{u}_{p-1})$ is the required vector.

Sufficiency proof

Let p be the periodicity of $\mathcal{C}$ and $S = S_0 \uplus \ldots \uplus S_{p-1}$.

Assume there exists $\mathbf{u}$ such that:

$$\text{For all } i, \mathbf{u}[i] > 0, \mathbf{u} = \mathbf{u} \cdot \mathbf{P} \text{ and } \sum_{i \in S} \mathbf{u}[i] < \infty$$

Decompose $\mathbf{u}$ as $(\mathbf{u}_0, \ldots, \mathbf{u}_{p-1})$.
Then $\mathbf{P}^p$ defines an aperiodic irreducible chain $\mathcal{C}'$ over S_0 and $\mathbf{u}_0 \cdot \mathbf{P}^p = \mathbf{u}_0$.
So $\mathcal{C}'$ is ergodic thus positive recurrent.

Let $s \in S_0$

- the probability of a return to s is the same in $\mathcal{C}$ and $\mathcal{C}'$.
  So s is recurrent in $\mathcal{C}$.
- the mean return time s in $\mathcal{C}$ is p times the mean return time to s in $\mathcal{C}'$.
  So s is positive recurrent in $\mathcal{C}$.

Characterization of recurrence

Let $\mathcal{C}$ be irreducible (S=$\mathbb{N}$). Then it is recurrent iff $\mathbf{0}$ is the single solution of:

$$\forall i > 0 \; x[i] = \sum_{j>0} \mathbf{P}[i,j]x[j] \wedge 0 \leq x[i] \leq 1 \quad (5)$$

Sketch of proof

Let $pin^n[i]$ be the probability to stay in $\mathbb{N}^*$ during n steps starting from i.

$pin[i] \stackrel{\text{def}}{=} \lim_{n\to\infty} pin^n[i]$ is the probability to never meet 0 starting from i.

By one-step reasoning : $pin^{n+1}[i] = \sum_{j>0} \mathbf{P}[i,j]pin^n[j]$

By dct: $pin[i] = \sum_{j>0} \mathbf{P}[i,j]pin[j]$

Let $\mathbf{x}$ be a solution of (5)

One has: $\forall i \; \mathbf{x}[i] \leq 1 = pin^0[i]$

By induction, the equality holds for all n:

$\mathbf{x}[i] = \sum_{j>0} \mathbf{P}[i,j]\mathbf{x}[j] \leq \sum_{j>0} \mathbf{P}[i,j]pin^n[j] = pin^{n+1}[i]$

Letting n goes to infinity, $\forall i \; \mathbf{x}[i] \leq pin[i]$

Solving $\mathbf{u} = \mathbf{u} \cdot \mathbf{P}$ in a recurrent chain

Let ${}_r p_{ij}^{(n)}$ be the probability that starting from i one reaches j after n transitions without ever visiting r except initially (so ${}_r p_{ij}^{(0)} = \mathbf{1}_{i=j}$).

Let ${}_r \pi_{ij} \stackrel{\text{def}}{=} \sum_{n \in \mathbb{N}} {}_r p_{ij}^{(n)}$ be the mean number of visits of j without visiting r.
(finite in an irreducible chain)

Let $\mathcal{C}$ be irreducible and recurrent.
Then, up to a scalar factor, $({}_r \pi_{ri})_{i \in S}$ is the single solution of:

$$\mathbf{u} = \mathbf{u} \cdot \mathbf{P} \qquad (6)$$

Proof

Existence

Due to the initial visit, ${}_r\pi_{rr} = 1$

For $i \neq r$, ${}_rp_{ri}^{(n+1)} = \sum_{j \in S} {}_rp_{rj}^{(n)} p_{ji}$

Summing over n: $\sum_{n \geq 1} {}_rp_{ri}^{(n)} = \sum_{n \geq 0} \sum_{j \in S} {}_rp_{rj}^{(n)} p_{ji}$

Since ${}_rp_{ri}^{(0)} = 0$: ${}_r\pi_{ri} = \sum_{j \in S} {}_r\pi_{rj} p_{ji}$

$\sum_{j \in S} {}_rp_{rj}^{(n)} p_{jr}$ is the probability of a first return to r at the $n + 1^{th}$ transition.
So $\sum_{j \in S} {}_r\pi_{rj} p_{jr}$ is the probability of a return to r.
Since r is recurrent: $\sum_{j \in S} {}_r\pi_{rj} p_{jr} = 1 = {}_r\pi_{rr}$

Unicity

Let $\mathbf{u}$ fulfill $\mathbf{u} = \mathbf{u} \cdot \mathbf{P}$ and for all i, $\mathbf{u}_i \geq 0$.
So $\mathbf{u}_i = 0$ implies $\mathbf{u}_j = 0$ for all j such that $p_{ji} > 0$.
Using irreducibility, either $\mathbf{u}$ is null or all its components are strictly positive.

Let $\mathbf{u}$ be strictly positive and apply a scalar factor so that $\mathbf{u}_r = 1$. For $i \neq r$:
$\mathbf{u}_i = p_{ri} + \sum_{j \neq r} \mathbf{u}_j p_{ji} = p_{ri} + \sum_{j \neq r} \left(p_{rj} + \sum_{k \neq r} \mathbf{u}_k p_{kj} \right) p_{ji} = p_{ri} + {}_rp_{ri}^{(2)} + \sum_{k \neq r} \mathbf{u}_k \cdot {}_rp_{ki}^{(2)}$
By induction: $\mathbf{u}_i = p_{ri} + {}_rp_{ri}^{(2)} + \cdots + {}_rp_{ri}^{(n)} + \sum_{j \neq r} \mathbf{u}_j \cdot {}_rp_{ji}^{(n)}$
So $\mathbf{u}_i \geq {}_r\pi_{ri}$ and $(\mathbf{u}_i - {}_r\pi_{ri})_{i \in S}$ is another solution.
Since $\mathbf{u}_r - {}_r\pi_{rr} = 0$, for all i, $\mathbf{u}_i = {}_r\pi_{ri}$.

Summary of characterizations

Status	Characterization
Recurrent	$\mathbf{0}$ is the single solution of: $\forall i \in S \setminus s_0 \; \mathbf{u}_i = \sum_{j \in S \setminus \{s_0\}} p_{i,j} \mathbf{u}_j$ and $0 \leq \mathbf{u}_i \leq 1$ In this case given any $r \in S$, $\forall \mathbf{u} \geq 0 \; \mathbf{u} \cdot \mathbf{P} = \mathbf{u} \Leftrightarrow \exists \alpha \geq 0 \; \forall s \; \mathbf{u}[s] = \alpha \cdot {}_r\pi_{rs}$
Positive recurrent	$\exists! \mathbf{u} > 0 \; \mathbf{u} \cdot \mathbf{P} = \mathbf{u} \wedge \sum_{i \in S} \mathbf{u}_i = 1$ *($\mathbf{u}$ is the steady-state distribution when the* DTMC *is aperiodic)*
Period is p	$S = S_0 \uplus S_1 \uplus \ldots \uplus S_{p-1}$ with $\forall r < p \; \forall i \in S_r \; \forall j \in S \; p_{i,j} > 0 \Rightarrow j \in S_{r+1 \mod p}$ and p is the greatest integer fulfilling this property.

Analysis of a random walk (1)

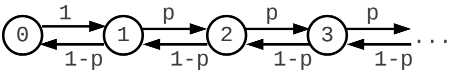

The chain has period 2.

Recurrence versus transience

$$x_1 = px_2 \text{ and } \forall i \geq 2 \; x_i = px_{i+1} + (1-p)x_{i-1}$$

It can be rewritten as:

$$x_1 = px_2 \text{ and } \forall i \geq 2 \; x_{i+1} - x_i = \frac{1-p}{p}(x_i - x_{i-1})$$

By induction:

$$x_i = x_1 + (x_2 - x_1)\sum_{j=0}^{i-2}\left(\frac{1-p}{p}\right)^j = x_1\left(1 + \frac{1-p}{p}\sum_{j=0}^{i-2}\left(\frac{1-p}{p}\right)^j\right) = x_1\left(\sum_{j=0}^{i-1}\left(\frac{1-p}{p}\right)^j\right)$$

Thus if $p \leq \frac{1}{2}$ then the x_i's are unbounded.

Otherwise the x_i's are bounded by $x_1 \frac{p}{2p-1}$.

Analysis of a random walk (2)

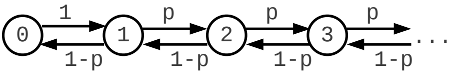

Positive versus null recurrence

$$x_0 = (1-p)x_1 \text{ and } x_1 = (1-p)x_2 + x_0 \text{ and } \forall i \geq 2\ x_i = (1-p)x_{i+1} + px_{i-1}$$

So $x_2 = \frac{px_1}{1-p}$ and by induction:

$$x_{i+1} = \frac{x_i - px_{i-1}}{1-p} = \frac{x_i - (1-p)x_i}{1-p} = \frac{px_i}{1-p}$$

Assume $x_0 > 0$. (otherwise $\mathbf{x} = 0$ is not a distribution)

- When $p = \frac{1}{2}$, $x_2 = x_1$ and by induction, for $i \geq 1$, $x_i = x_1$. So $\sum_{i \in \mathbb{N}} x_i = \infty$.
- When $p < \frac{1}{2}$ for $i \geq 1$, $x_i = x_1 \left(\frac{p}{1-p}\right)^{i-1}$

  Thus $\sum_{i \in \mathbb{N}} x_i = x_1(1 - p + \frac{1-p}{1-2p})$ is finite.

Plan

Strongly connected components (scc)

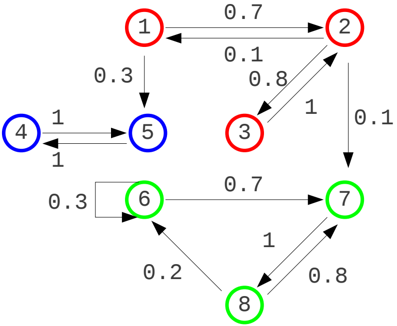

A scc S' is a maximal subset of vertices such that:
for all $i, j \in S'$ there is a path from i to j.

A scc S' is *terminal* if there is no path from S' to $S \setminus S'$.

Computation of scc's in linear time by the algorithm of Tarjan.

Classification of states

Every terminal scc is an irreducible DTMC.
The transient states are the states of the non terminal scc's.
The states of terminal scc's are positive recurrent.

Sketch of proof

Let i belonging to a non terminal scc.
There is a path from i to j outside the scc.
There is no path from j to i.
So i is transient.

Let $\mathbf{P}$ be the transition matrix of a terminal scc S'.
Pick some state $i \in S'$.
For all n, $\sum_{j \in S'} p^n_{i,j} = 1$.
Since this is a finite sum, there exists j such that:
$p^n_{i,j}$ does not converge to 0 when n goes to infinity.
So j is positive recurrent.

Computing the periodicity: an example

Adaptation of a tree covering construction

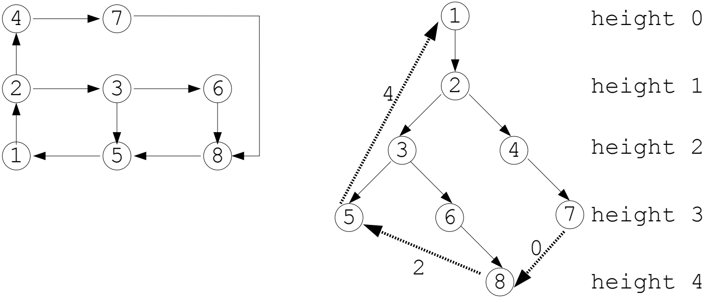

periodicity=gcd(0,2,4)=2

Computing the periodicity: the algorithm

Input G, a strongly connected graph whose set of vertices is $\{1, \ldots, n\}$
Output p, the periodicity of G
Data i, j integers, $Height$ an array of size n, Q a queue

For i **from** 1 **to** n **do** $Height[i] \leftarrow \infty$
$p \leftarrow 0$
$Height[0] \leftarrow 0$
$\textbf{InsertQueue}(Q, 0)$
While $\textbf{not EmptyQueue}(Q)$ **do**
  $i \leftarrow \textbf{ExtractQueue}(Q)$
  For $(i, j) \in G$ **do**
    If $Height[j] = \infty$ **then**
      $Height[j] \leftarrow Height[i] + 1$
      $\textbf{InsertQueue}(Q, j)$
    Else
      $p \leftarrow \gcd(p, Height[i] - Height[j] + 1)$
Return(p)

Proof of the algorithm

Let p be the periodicity, p' the $\gcd$ of the edge labels and r the root.

Given two paths with same source and destination,
the difference between the lengths of these paths must be a multiple of p.

Let (u, v) be an edge with non null label.
Let σ_u (resp. σ_v), the path from r to u (resp. v) along the tree.
$\sigma_u(u, v)$ is a path from r to v.
The difference between the lengths of the two paths is: $Height[u] - Height[v] + 1$
Thus $p|Height[u] - Height[v] + 1$ and so $p|p'$.

Let S'_i for $0 \leq i < p'$ be defined by:

$$s \in S'_i \text{ iff } Height[s] \bmod p' = i$$

An edge of the tree joins a vertex of S'_i to a vertex of $S'_{i+1 \bmod p'}$
An edge (u, v) out of the tree joins $u \in S'_{Height[u] \bmod p'}$ to $v \in S'_{Height[v] \bmod p'}$
$Height[u] - Height[v] + 1 \mod p' = 0 \Rightarrow Height[v] \bmod p' = Height[u] + 1 \bmod p'$
Using the characterization of periodicity, $p' \leq p$.

Matrices and vectors of a DTMC

Let $\{\mathcal{C}_1, \ldots, \mathcal{C}_k\}$ be the subchains associated with terminal scc.

Let π_i be the steady-state distribution of $\mathcal{C}_i$ supposed to be aperiodic.

Let T be the set of transient states.

Let $\mathbf{P}_{T,T}$ (resp. $\mathbf{P}_{T,i}$) be the transition matrix from T to T (resp. S_i).

$$\mathbf{P}_{T,T} = \begin{pmatrix} 0.0 & 0.7 & 0.0 \\ 0.1 & 0.0 & 0.8 \\ 0.0 & 0.2 & 0.0 \end{pmatrix}$$

$$\mathbf{P}_{T,1}.\mathbf{1}^T = \begin{pmatrix} 0.0 & 0.3 \\ 0.0 & 0.0 \\ 0.0 & 0.4 \end{pmatrix} \begin{pmatrix} 1.0 \\ 1.0 \end{pmatrix} = \begin{pmatrix} 0.3 \\ 0.0 \\ 0.4 \end{pmatrix}$$

$$\mathbf{P}_{T,2}.\mathbf{1}^T = \begin{pmatrix} 0.0 & 0.0 & 0.0 \\ 0.0 & 0.1 & 0.0 \\ 0.4 & 0.0 & 0.0 \end{pmatrix} \begin{pmatrix} 1.0 \\ 1.0 \\ 1.0 \end{pmatrix} = \begin{pmatrix} 0.0 \\ 0.1 \\ 0.4 \end{pmatrix}$$

$T=\{1,2,3\}, C_1=\{4,5\}, C_2=\{6,7,8\}$

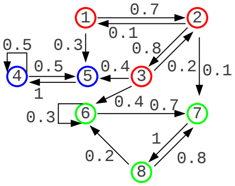

$\pi_1=(2/3,1/3)$

$\pi_2=(1/8,7/16,7/16)$

Steady-state distribution

Let $\mathcal{C}$ be a finite DTMC with initial distribution π_0 whose terminal scc's are aperiodic. Then there exists a steady-state distribution:

$$\pi_\infty \stackrel{\text{def}}{=} \sum_{i=1}^{k} \left(\left(\pi_{0,i} + \pi_{0,T} \left(\mathbf{Id} - \mathbf{P}_{T,T} \right)^{-1} \cdot \mathbf{P}_{T,i} \right) \cdot \mathbf{1}^T \right) \pi_i$$

where $\pi_{0,i}$ (resp. $\pi_{0,T}$) is π_0 restricted to states of $\mathcal{C}_i$ (resp. T) and π_i is the steady-state distribution of $\mathcal{C}_i$.

Sketch of proof

$\pi_\infty \stackrel{\text{def}}{=} \sum_{i=1}^{k} \Pr(\text{to reach } \mathcal{C}_i) \cdot \pi_i$
$\Pr(\text{to reach } \mathcal{C}_i) = \sum_{s \in S} \pi_0(s) \cdot \pi'_{\mathcal{C}_i}(s)$ where $\pi'_{\mathcal{C}_i}(s) = \Pr(\text{to reach } \mathcal{C}_i \mid S_0 = s)$

- When state $s \in \mathcal{C}_i$, then $\pi'_{\mathcal{C}_i}(s) = 1$ and $\pi'_{\mathcal{C}_j}(s) = 0$ for $j \neq i$
- The probability of paths from a transient state s along T to $\mathcal{C}_i$ of length $n+1$ is:
$$\left((\mathbf{P}_{T,T})^n \cdot \mathbf{P}_{T,i} \cdot \mathbf{1}^T \right)[s]$$

$(\sum_{n \geq 0} (\mathbf{P}_{T,T})^n)[i,j]$ is the (finite) mean number of visits of j starting from i.
For every n_0: $(\sum_{n \leq n_0} (\mathbf{P}_{T,T})^n)(\mathbf{Id} - \mathbf{P}_{T,T}) = \mathbf{Id} - (\mathbf{P}_{T,T})^{n_0+1}$
Since $\lim_{n\to\infty} (\mathbf{P}_{T,T})^n = 0$, letting n_0 go to infinity establishes the result.

Regular matrices

A matrix $\mathbf{M}$ is *positive* if for all i, j, $\mathbf{M}[i,j] > 0$.

A matrix $\mathbf{M}$ is *non negative* if for all i, j, $\mathbf{M}[i,j] \geq 0$.

A non negative square matrix $\mathbf{M}$ is *regular* if for some k, $\mathbf{M}^k$ is positive.

The transition matrix of an ergodic DTMC is regular.

Sketch of proof

Let $s \stackrel{\text{def}}{=} |S|$ and $i \in S$.

There is a n_0 such that for all $n \geq n_0$, $p_{ii}^n > 0$.

Furthermore for all j, j', there are $m, m' \leq s - 1$ such that $p_{ji}^m > 0$ and $p_{ij'}^{m'} > 0$.

So for all j, j' and $n \geq n_0 + 2(s-1)$, one has $p_{jj'}^n > 0$.

Convergence rate

Let $\mathcal{C}$ be a finite ergodic DTMC, π_n its distribution at time n and π_∞ its steady-state distribution. Then there exists some $0 < \lambda < 1$ such that:

$$\|\pi_\infty - \pi_n\| = O(\lambda^n)$$

Sketch of proof

Let Π_∞ be the square matrix where every row is a copy of π_∞.

$\Pi_\infty \mathbf{P} = \Pi_\infty$ since $\pi_\infty \mathbf{P} = \pi_\infty$ and for every transition matrix $\mathbf{P}'$, $\mathbf{P}'\Pi_\infty = \Pi_\infty$

Since $\mathbf{P}^k$ is positive, there is some $0 < \delta < 1$ such that $\forall i, j\ \mathbf{P}^k[i,j] \geq \delta \Pi_\infty[i,j]$

Let $\theta \stackrel{\text{def}}{=} 1 - \delta$ and $\mathbf{Q} \stackrel{\text{def}}{=} \frac{1}{\theta}\mathbf{P}^k - \frac{1-\theta}{\theta}\Pi_\infty$

$\mathbf{Q}$ is a transition matrix and fulfills: $\mathbf{P}^k = \theta \mathbf{Q} + (1-\theta)\Pi_\infty$

Let us prove that: $\forall n\ \mathbf{P}^{kn} = \theta^n \mathbf{Q}^n + (1-\theta^n)\Pi_\infty$ ($\Leftrightarrow \mathbf{P}^{kn} - \Pi_\infty = \theta^n(\mathbf{Q}^n - \Pi_\infty)$)

$$\Pi_\infty \mathbf{Q} = \frac{1}{\theta}\Pi_\infty \mathbf{P}^k - \frac{1-\theta}{\theta}\Pi_\infty \Pi_\infty = \frac{1}{\theta}\Pi_\infty - \frac{1-\theta}{\theta}\Pi_\infty = \Pi_\infty$$

$$\mathbf{P}^{kn+k} = (\theta^n \mathbf{Q}^n + (1-\theta^n)\Pi_\infty)(\theta \mathbf{Q} + (1-\theta)\Pi_\infty)$$

$$= \theta^{n+1}\mathbf{Q}^{n+1} + ((1-\theta^n)\theta + (1-\theta)\theta^n + (1-\theta^n)(1-\theta))\Pi_\infty$$

Multiplying by $\mathbf{P}^j$ with $0 \leq j < k$: $\forall n\ \forall j < k\ \mathbf{P}^{kn+j} - \Pi_\infty = \theta^n(\mathbf{Q}^n \mathbf{P}^j - \Pi_\infty)$

Multiplying by π_0: $\forall n\ \forall j < k\ \pi_{kn+j} - \pi_\infty = \theta^n(\pi_0 \mathbf{Q}^n \mathbf{P}^j - \pi_\infty)$