Analysis of an on-line algorithm for page importance computation

Nelly Litvak
University of Twente
n.litvak@ewi.utwente.nl

Joint work with Philippe Robert (INRIA, France)

Outline

- PageRank algorithm
- On-line methods
- Convergence of the algorithm
- The Cat and Mouse process

Ranking schemes

- A user types a query to find relevant pages.
- Problem: Normally, there are hundreds of relevant pages.
  In which order should we list the pages for the user??

Ranking schemes

- A user types a query to find relevant pages.
- Problem: Normally, there are hundreds of relevant pages.
  In which order should we list the pages for the user??
- The ‘best’-text-match-techniques quickly became insufficient (besides, they are not user-friendly)

Ranking schemes

- A user types a query to find relevant pages.
- Problem: Normally, there are hundreds of relevant pages.
  In which order should we list the pages for the user??
- The 'best'-text-match-techniques quickly became insufficient (besides, they are not user-friendly)
- S. Brin and L. Page (1998), J.M. Kleinberg (1999)

  Idea: List most important and popular pages first. Define the importance through the hyperlink structure

- A user types a query to find relevant pages.
- Problem: Normally, there are hundreds of relevant pages.
  In which order should we list the pages for the user??
- The ‘best’-text-match-techniques quickly became insufficient (besides, they are not user-friendly)
- S. Brin and L. Page (1998), J.M. Kleinberg (1999)

  Idea: List most important and popular pages first. Define the importance through the hyperlink structure

  S. Brin, L. Page, R. Motwami and T. Winograd (1998) The PageRank citation ranking: bringing order to the web.

- PageRank π_i of page i is the long run fraction of time that a random surfer spends on page i.
- 'Easily bored surfer' model. With probability c (=0.85), a surfer follows a randomly chosen outgoing link. Otherwise, he/she jumps to a random page.

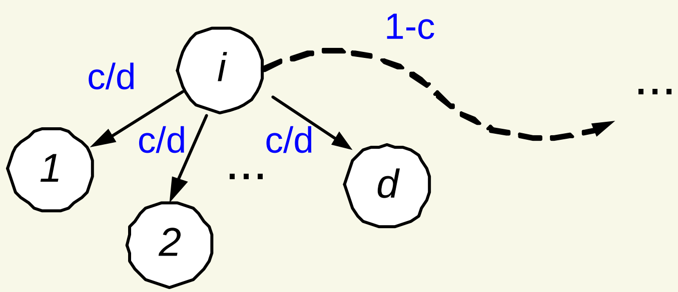

$$\pi_i = \sum_{j \to i} \frac{c}{d_j}\pi_j + \frac{1-c}{N}$$

N – # pages pages
d_j – out-degree of j
$c < 1 \Rightarrow$ solution exists

- PageRank π_i of page i is the long run fraction of time that a random surfer spends on page i.
- 'Easily bored surfer' model. With probability c (=0.85), a surfer follows a randomly chosen outgoing link. Otherwise, he/she jumps to a random page.

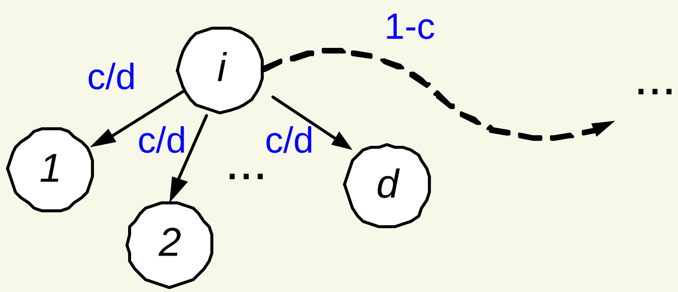

$$\pi_i = \sum_{j \to i} \frac{c}{d_j}\pi_j + \frac{1-c}{N}$$

N – # pages pages
d_j – out-degree of j
$c < 1 \Rightarrow$ solution exists

Page is important if many important pages link to it!

- n is the total number of pages
- $P = (p_{ij})$ - hyperlink matrix

$$p_{ij} = \begin{cases} 1/d_i & \text{if } j \text{ is one of the } d_i \text{ outgoing links of } i \\ 1/n & \text{if } d_i = 0 \\ 0 & \text{otherwise} \end{cases}$$

- Modified transition matrix: $\tilde{P} = cP + (1-c)(1/N)E$
  E is an $n \times n$ matrix consisting of one's, $c = 0.85$
- PageRank vector: $\pi\tilde{P} = \pi, \quad \pi\underline{1} = 1$

- n is the total number of pages
- $P = (p_{ij})$ - hyperlink matrix

$$p_{ij} = \begin{cases} 1/d_i & \text{if } j \text{ is one of the } d_i \text{ outgoing links of } i \\ 1/n & \text{if } d_i = 0 \\ 0 & \text{otherwise} \end{cases}$$

- Modified transition matrix: $\tilde{P} = cP + (1-c)(1/N)E$
  E is an $n \times n$ matrix consisting of one's, $c = 0.85$
- PageRank vector: $\pi\tilde{P} = \pi, \quad \pi\underline{1} = 1$

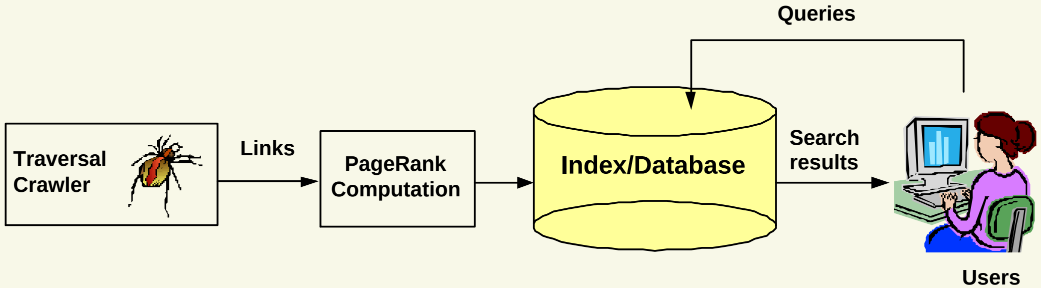

- Power Iterations: $\pi^{(0)} = (1/N, \dots, 1/N); \quad \pi^{(t)} = \pi^{(t-1)}\tilde{P}, \ t > 0$

  Accuracy of the order c^t (50–100 iterations with $c = 0.85$)

PageRank evaluation problem

- PageRank is a stationary distribution of a huge Markov chain: state space is the set of all Web pages.

PageRank evaluation problem

- PageRank is a stationary distribution of a huge Markov chain: state space is the set of all Web pages.
- The task of PageRank computation trigged enormous developments in the well established area of numerical solution of large MC.

PageRank evaluation problem

- PageRank is a stationary distribution of a huge Markov chain: state space is the set of all Web pages.
- The task of PageRank computation trigged enormous developments in the well established area of numerical solution of large MC.
  - Advanced linear algebra methods to speed up power iterations (off-line)

PageRank evaluation problem

- PageRank is a stationary distribution of a huge Markov chain: state space is the set of all Web pages.
- The task of PageRank computation trigged enormous developments in the well established area of numerical solution of large MC.
    - Advanced linear algebra methods to speed up power iterations (off-line)
    - Monte Carlo methods (off-line or on-line)

PageRank evaluation problem

- PageRank is a stationary distribution of a huge Markov chain: state space is the set of all Web pages.
- The task of PageRank computation trigged enormous developments in the well established area of numerical solution of large MC.
    - Advanced linear algebra methods to speed up power iterations (off-line)
    - Monte Carlo methods (off-line or on-line)
    - Other non-trivial on-line methods. One such method by Abiteboul, Preda and Cobena (1999) we will discuss today.

On-line algorithm by Abiteboul, Preda, Cobena (1999)

- Consider a crawler that performs a Markov random walk $\{C(t),\ t = 0, 1, \ldots\}$ on directed graph of N nodes.

On-line algorithm by Abiteboul, Preda, Cobena (1999)

- Consider a crawler that performs a Markov random walk $\{C(t),\ t=0,1,\ldots\}$ on directed graph of N nodes.
- Transition matrix P, where $p_{ij}>0$ iff there is a link (edge) from i to j. The goal is to define the stationary distribution π.

On-line algorithm by Abiteboul, Preda, Cobena (1999)

- Consider a crawler that performs a Markov random walk $\{C(t),\ t = 0, 1, \ldots\}$ on directed graph of N nodes.
- Transition matrix P, where $p_{ij} > 0$ iff there is a link (edge) from i to j. The goal is to define the stationary distribution π.
- Algorithm:
  - At $t = 0$, each node receives an equal amount $1/N$ of cash.

On-line algorithm by Abiteboul, Preda, Cobena (1999)

- Consider a crawler that performs a Markov random walk $\{C(t),\ t=0,1,\ldots\}$ on directed graph of N nodes.
- Transition matrix P, where $p_{ij}>0$ iff there is a link (edge) from i to j. The goal is to define the stationary distribution π.
- Algorithm:
  - At $t=0$, each node receives an equal amount $1/N$ of cash.
  - Each time the crawler visits a node i, the node distributes all its cash among its outgoing links proportional to p_{ij}.

On-line algorithm by Abiteboul, Preda, Cobena (1999)

- Consider a crawler that performs a Markov random walk $\{C(t),\ t = 0, 1, \ldots\}$ on directed graph of N nodes.
- Transition matrix P, where $p_{ij} > 0$ iff there is a link (edge) from i to j. The goal is to define the stationary distribution π.
- Algorithm:
  - At $t = 0$, each node receives an equal amount $1/N$ of cash.
  - Each time the crawler visits a node i, the node distributes all its cash among its outgoing links proportional to p_{ij}.
  - $X_i(t)$: the amount of cash at node i at time t

On-line algorithm by Abiteboul, Preda, Cobena (1999)

- Consider a crawler that performs a Markov random walk $\{C(t),\ t = 0, 1, \ldots\}$ on directed graph of N nodes.
- Transition matrix P, where $p_{ij} > 0$ iff there is a link (edge) from i to j. The goal is to define the stationary distribution π.
- Algorithm:
  - At $t = 0$, each node receives an equal amount $1/N$ of cash.
  - Each time the crawler visits a node i, the node distributes all its cash among its outgoing links proportional to p_{ij}.
  - $X_i(t)$: the amount of cash at node i at time t
  - History of a node: $H_i(t) = \sum_{s=0}^{t} X_i(s)\mathbf{1}_{[C(s)=i]}$, amount of cash distributed by node i on $[0, t]$
    Total history: $H(t) = \sum_{i=1}^{n} H_i(t)$

On-line algorithm by Abiteboul, Preda, Cobena (1999)

University of Twente
Enschede - The Netherlands

- Consider a crawler that performs a Markov random walk $\{C(t),\ t = 0, 1, \ldots\}$ on directed graph of N nodes.
- Transition matrix P, where $p_{ij} > 0$ iff there is a link (edge) from i to j. The goal is to define the stationary distribution π.
- Algorithm:
  - At $t = 0$, each node receives an equal amount $1/N$ of cash.
  - Each time the crawler visits a node i, the node distributes all its cash among its outgoing links proportional to p_{ij}.
  - $X_i(t)$: the amount of cash at node i at time t
  - History of a node: $H_i(t) = \sum_{s=0}^{t} X_i(s)\mathbf{1}_{[C(s)=i]}$, amount of cash distributed by node i on $[0, t]$
    Total history: $H(t) = \sum_{i=1}^{n} H_i(t)$
  - The estimator of π_i at time t is

$$\pi_i(t) = \frac{H_i(t) + X_i(t)}{H(t) + 1}.$$

Convergence

- $X_i(t)$: the amount of cash at node i at time t
- History of a node: $H_i(t) = \sum_{k=0}^{t} X_i(k)\mathbf{1}_{[C(k)=i]}$, amount of cash distributed by node i on $[0,t]$. Total history: $H(t) = \sum_{i=1}^{n} H(t)$
- The estimator of π_i at time t is $\pi_i(t) = (H_i(t) + X_i(t))/(H(t)+1)$.

Convergence

- $X_i(t)$: the amount of cash at node i at time t
- History of a node: $H_i(t) = \sum_{k=0}^{t} X_i(k)\mathbf{1}_{[C(k)=i]}$, amount of cash distributed by node i on $[0,t]$. Total history: $H(t) = \sum_{i=1}^{n} H(t)$
- The estimator of π_i at time t is $\pi_i(t) = (H_i(t) + X_i(t))/(H(t)+1)$.

Lemma. $H(t) \to \infty$ w.p. 1.

Convergence

- $X_i(t)$: the amount of cash at node i at time t
- History of a node: $H_i(t) = \sum_{k=0}^{t} X_i(k)\mathbf{1}_{[C(k)=i]}$, amount of cash distributed by node i on $[0,t]$. Total history: $H(t) = \sum_{i=1}^{n} H(t)$
- The estimator of π_i at time t is $\pi_i(t) = (H_i(t) + X_i(t))/(H(t)+1)$.

Lemma. $H(t) \to \infty$ w.p. 1.
Pf. Cover time: a time needed to visit all states and come back to the original state. The cover time is finite w.p. 1. During one cover time cycle, $H(t)$ grows by at least 1 cash unit.

Convergence

- $X_i(t)$: the amount of cash at node i at time t
- History of a node: $H_i(t) = \sum_{k=0}^{t} X_i(k)\mathbf{1}_{[C(k)=i]}$, amount of cash distributed by node i on $[0,t]$. Total history: $H(t) = \sum_{i=1}^{n} H(t)$
- The estimator of π_i at time t is $\pi_i(t) = (H_i(t) + X_i(t))/(H(t)+1)$.

Lemma. $H(t) \to \infty$ w.p. 1.

Pf. Cover time: a time needed to visit all states and come back to the original state. The cover time is finite w.p. 1. During one cover time cycle, $H(t)$ grows by at least 1 cash unit.

Theorem. $\pi_i(t) \to \pi$ a.s. as $t \to \infty$.

Convergence

- $X_i(t)$: the amount of cash at node i at time t
- History of a node: $H_i(t) = \sum_{k=0}^{t} X_i(k)\mathbf{1}_{[C(k)=i]}$, amount of cash distributed by node i on $[0,t]$. Total history: $H(t) = \sum_{i=1}^{n} H(t)$
- The estimator of π_i at time t is $\pi_i(t) = (H_i(t) + X_i(t))/(H(t)+1)$.

Lemma. $H(t) \to \infty$ w.p. 1.
Pf. Cover time: a time needed to visit all states and come back to the original state. The cover time is finite w.p. 1. During one cover time cycle, $H(t)$ grows by at least 1 cash unit.

Theorem. $\pi_i(t) \to \pi$ a.s. as $t \to \infty$.
Pf. From the APC paper

$$H_i(t) + C_i(t) = 1/n + \sum_j p_{ji} H_j(t), \quad i = 1, \ldots, n$$

Divide by $H(t)$: $\pi_i(t) = [1/n - C_i(t)]/H(t) + \sum_j p_{ji}\pi_j(t), \quad i = 1, \ldots, n$
Solution: $\pi_i(t) = \pi_i + \left[[(\mathbf{1}/n - \mathbf{C}(t))/H(t)]\left(\sum_{k=0}^{\infty} P^k - \mathbf{1}^T\pi\right)\right]_i$
(here $\mathbf{1}$ is a row-vector of ones)

Speed of convergence

University of Twente
Enschede - The Netherlands

- The speed of convergence is determined by the term $[(\mathbf{1}/n - \mathbf{C}(t))/H(t)]$, with $H(t)$ in denominator. Thus, the algorithm converges as $1/H(t)$ when $t \to \infty$.

Speed of convergence

- The speed of convergence is determined by the term $[(\mathbf{1}/n - \mathbf{C}(t))/H(t)]$, with $H(t)$ in denominator. Thus, the algorithm converges as $1/H(t)$ when $t \to \infty$.
- $H(t) = \sum_{s=0}^{t} \sum_i X_i(s) \mathbf{1}_{[C(s)=i]}$. The cash process $\{X_i(t)\}$ is determining for the speed of convergence of the algorithm.

Speed of convergence

- The speed of convergence is determined by the term $[(\mathbf{1}/n - \mathbf{C}(t))/H(t)]$, with $H(t)$ in denominator. Thus, the algorithm converges as $1/H(t)$ when $t \to \infty$.
- $H(t) = \sum_{s=0}^{t} \sum_i X_i(s) \mathbf{1}_{[C(s)=i]}$. The cash process $\{X_i(t)\}$ is determining for the speed of convergence of the algorithm.
- But how can we characterize the cash process? For instance, can we prove that $X_i(t)$ converges, or, at least, $\mathbb{E}[X_i(t)]$ converges?

Speed of convergence

- The speed of convergence is determined by the term $[(\mathbf{1}/n-\mathbf{C}(t))/H(t)]$, with $H(t)$ in denominator. Thus, the algorithm converges as $1/H(t)$ when $t\to\infty$.
- $H(t)=\sum_{s=0}^{t}\sum_i X_i(s)\mathbf{1}_{[C(s)=i]}$. The cash process $\{X_i(t)\}$ is determining for the speed of convergence of the algorithm.
- But how can we characterize the cash process? For instance, can we prove that $X_i(t)$ converges, or, at least, $\mathbb{E}[X_i(t)]$ converges?
- The cash process is inconvenient for analysis. Idea: to translate it into an easier process.

Cat and mouse game

Cat and mouse game

- Cat performs the Markov random walk $C(t)$, transition matrix P.

- Cat performs the Markov random walk $C(t)$, transition matrix P.
- At time t, the mouse is at the node $M(t)$.

Cat and mouse game

- Cat performs the Markov random walk $C(t)$, transition matrix P.
- At time t, the mouse is at the node $M(t)$.
- The mouse makes a move only when found by the cat, the event $[C(t) = M(t)]$. In this case, the mouse makes one step, using the same transition matrix P.

Cat and mouse game

- Cat performs the Markov random walk $C(t)$, transition matrix P.
- At time t, the mouse is at the node $M(t)$.
- The mouse makes a move only when found by the cat, the event $[C(t) = M(t)]$. In this case, the mouse makes one step, using the same transition matrix P.

Relation to the cash process

- Cat performs the Markov random walk $C(t)$, transition matrix P.
- At time t, the mouse is at the node $M(t)$.
- The mouse makes a move only when found by the cat, the event $[C(t) = M(t)]$. In this case, the mouse makes one step, using the same transition matrix P.

$\mathcal{F}_t$: σ-field generated by $C(0)$, $C(1)$, $\ldots$, $C(t)$.

Relation to the cash process

- Cat performs the Markov random walk $C(t)$, transition matrix P.
- At time t, the mouse is at the node $M(t)$.
- The mouse makes a move only when found by the cat, the event $[C(t) = M(t)]$. In this case, the mouse makes one step, using the same transition matrix P.

$\mathcal{F}_t$: σ-field generated by $C(0)$, $C(1)$, ..., $C(t)$.

Theorem. For $t \geq 0$,

$$(X_i(t), 1 \leq i \leq N) \stackrel{\text{dist.}}{=} (\mathbb{P}[M(t) = i \mid \mathcal{F}_{t-1}], 1 \leq i \leq N)\,.$$

In particular, for $1 \leq i \leq N$,

$$\mathbb{E}(X_i(t)) = \mathbb{P}(M(t) = i).$$

$C(t)$ – cat position; $M(t)$ – mouse position; $X_i(t)$ – cash on page i.
We need to prove that: $\mathbb{E}(X_i(t)) = \mathbb{P}(M(t) = i)$.

$$\mathbb{P}(M(t+1) = i \mid \mathcal{F}_t) = \sum_{j \neq i} 1_{\{C(t)=j\}} p_{j,i} \times \mathbb{P}(M(t) = j \mid \mathcal{F}_{t-1}) + 1_{\{C(t) \neq i\}} \mathbb{P}(M(t) = i \mid \mathcal{F}_{t-1}).$$

On the other hand, for the cash

$$\mathbb{E}[X_i(t+1)] = \sum_{j \neq i} 1_{\{C(t)=j\}} p_{j,i} \times \mathbb{E}[X_j(t)] + 1_{\{C(t) \neq i\}} \mathbb{E}[X_i(t)].$$

Same equation! Statement follows by induction in t.

The MC $(C(t), M(t))$ has a transition matrix $q(\cdot,\cdot)$ given by

$$q_{(i,j),(k,j)} = p_{i,k} \qquad \text{if } i \neq j;$$

$$q_{(j,j),(k,k')} = p_{j,k}p_{j,k'}.$$

Stationary distribution of the Cat and Mouse chain-1

The MC $(C(t), M(t))$ has a transition matrix $q(\cdot,\cdot)$ given by

$$q_{(i,j),(k,j)} = p_{i,k} \qquad \text{if } i \neq j;$$
$$q_{(j,j),(k,k')} = p_{j,k}p_{j,k'}.$$

$\nu = \mathbb{P}(C_\infty = \cdot\,, M_\infty = \cdot)$ is the invariant distribution. Then

$$\nu(i,j) = \sum_{k \neq j} \nu(k,j)p_{k,i} + \sum_{k} \nu(k,k)p_{k,i}p_{k,j}$$

Stationary distribution of the Cat and Mouse chain-1

University of Twente
Enschede - The Netherlands

The MC $(C(t), M(t))$ has a transition matrix $q(\cdot,\cdot)$ given by

$$
\begin{aligned}
q_{(i,j),(k,j)} &= p_{i,k} && \text{if } i \neq j; \\
q_{(j,j),(k,k')} &= p_{j,k}p_{j,k'}.
\end{aligned}
$$

$\nu = \mathbb{P}(C_\infty = \cdot\,, M_\infty = \cdot)$ is the invariant distribution. Then

$$\nu(i,j) = \sum_{k\neq j} \nu(k,j)p_{k,i} + \sum_k \nu(k,k)p_{k,i}p_{k,j}$$

Note: $\sum_j \nu(i,j) = \pi_i$. Summing over i we get $\sum\limits_k \nu(k,k)p_{k,j} = \nu(j,j)$, and therefore that there exists some constant c such that,

$$\nu(j,j) = c\pi_j, \quad j = 1,\ldots,N.$$

Stationary distribution of the Cat and Mouse chain-1

The MC $(C(t), M(t))$ has a transition matrix $q(\cdot,\cdot)$ given by

$$q_{(i,j),(k,j)} = p_{i,k} \qquad \text{if } i \neq j;$$
$$q_{(j,j),(k,k')} = p_{j,k}p_{j,k'}.$$

$\nu = \mathbb{P}(C_\infty = \cdot\,, M_\infty = \cdot)$ is the invariant distribution. Then

$$\nu(i,j) = \sum_{k\neq j} \nu(k,j)p_{k,i} + \sum_k \nu(k,k)p_{k,i}p_{k,j}$$

Note: $\sum_j \nu(i,j) = \pi_i$. Summing over i we get $\sum\limits_k \nu(k,k)p_{k,j} = \nu(j,j)$, and therefore that there exists some constant c such that,

$$\nu(j,j) = c\pi_j, \quad j = 1,\ldots,N.$$

We have $\mathbb{E}(X_j(t)|C_t = j) = \mathbb{P}(M_t = j|C_t = j) \to c$ as $t \to \infty$. On average, the ‘transaction’ is c at each step. Thus, $\mathbb{E}[H(t)] \sim ct$.

Denote by

$$p^*_{i,j} = \frac{\pi_j}{\pi_i} p_{j,i}$$

the transition matrix of the reversed Markov chain $(C^*(t))$ associated to $(C(t))$. Let

$$T_j = \inf\{t > 0 : C(t) = j\}.$$

Stationary distribution of the Cat and Mouse chain-2

Denote by

$$p^*_{i,j} = \frac{\pi_j}{\pi_i} p_{j,i}$$

the transition matrix of the reversed Markov chain $(C^*(t))$ associated to $(C(t))$. Let

$$T_j = \inf\{t > 0 : C(t) = j\}.$$

Proposition. The stationary distribution of the mouse is

$$\mathbb{P}(M(\infty) = j) = c\mathbb{E}_\pi \left[p_{C(0),j}\, T_j\right]$$

and

$$c = \left[\sum_k \mathbb{E}_\pi \left[p_{C(0),k}\, T_k\right]\right]^{-1}$$

The case of reversible Markov chain

It is assumed for the moment that $(C(t))$ is a reversible Markov chain, i.e. $\pi_i p_{ij} = \pi_j p_{ji}$

$$\begin{aligned}\mathbb{E}_\pi\left[p_{C(0),j}\,T_j\right] &= \sum_i \pi_i p_{ij}\mathbb{E}_i(T_j) = \sum_i \pi_j p_{ji}\mathbb{E}_i(T_j)\\ &= \pi_j\mathbb{E}_j(T_j-1) = 1-\pi(y),\end{aligned}$$

Consequently,

$$c = \frac{1}{N-1} \text{ and } \mathbb{P}(S_\infty = y) = \frac{1-\pi(y)}{N-1}.$$

Tetali (1994) showed that if $(C(t))$ is a general recurrent Markov chain, then

$$\sum_k \mathbb{E}_\pi\left[p_{C(0),k}\,T_k\right] \le N-1. \tag{1}$$

It follows that the value $c = 1/(N-1)$ obtained for reversible chains, is the minimal possible value of c.

Cat and Mouse on the M/M/1 Markov chain - 1

- $C(t)$ is a simple random walk on $\mathbb{Z}_+$ reflected at zero. If the current state is $x > 0$, then the next state is $x+1$ or $x-1$ w.p. $p = \lambda/(\lambda+\mu)$ or $q = \lambda/(\lambda+\mu)$. If the current state is 0, then the next state is 1 with probability one.

Cat and Mouse on the M/M/1 Markov chain - 1

- $C(t)$ is a simple random walk on $\mathbb{Z}_+$ reflected at zero. If the current state is $x > 0$, then the next state is $x+1$ or $x-1$ w.p. $p = \lambda/(\lambda+\mu)$ or $q = \lambda/(\lambda+\mu)$. If the current state is 0, then the next state is 1 with probability one.
- This is a slotted version of the $M/M/1$ queue with arrival rate λ and service rate μ. We assume that the system is stable, that is, $\rho = \lambda/\mu = p/q < 1$.

Cat and Mouse on the M/M/1 Markov chain - 1

- $C(t)$ is a simple random walk on $\mathbb{Z}_+$ reflected at zero. If the current state is $x > 0$, then the next state is $x+1$ or $x-1$ w.p. $p = \lambda/(\lambda+\mu)$ or $q = \lambda/(\lambda+\mu)$. If the current state is 0, then the next state is 1 with probability one.
- This is a slotted version of the $M/M/1$ queue with arrival rate λ and service rate μ. We assume that the system is stable, that is, $\rho = \lambda/\mu = p/q < 1$.
- Assume that initially, the mouse is at some remote position $x \to \infty$, while the cat just left the neighborhood of x and went back to zero. The time needed for the cat to reach zero is approximately linear in x. The time needed to come back to x, multiplied by ρ^x, converges to an exponential random variable with parameter $(\mu-\lambda)^2/\mu$.

Cat and Mouse on the M/M/1 Markov chain - 1

- $C(t)$ is a simple random walk on $\mathbb{Z}_+$ reflected at zero. If the current state is $x > 0$, then the next state is $x+1$ or $x-1$ w.p. $p = \lambda/(\lambda+\mu)$ or $q = \lambda/(\lambda+\mu)$. If the current state is 0, then the next state is 1 with probability one.
- This is a slotted version of the $M/M/1$ queue with arrival rate λ and service rate μ. We assume that the system is stable, that is, $\rho = \lambda/\mu = p/q < 1$.
- Assume that initially, the mouse is at some remote position $x \to \infty$, while the cat just left the neighborhood of x and went back to zero. The time needed for the cat to reach zero is approximately linear in x. The time needed to come back to x, multiplied by ρ^x, converges to an exponential random variable with parameter $(\mu-\lambda)^2/\mu$.
- Note that for a finite reversible chain, $c = 1/(N-1)$ where N is the number of states. For $N = \infty$ we obtain that the chain $(C(t), M(t))$ is null-recurrent. The question is: where does the mouse stay?

- The cat and the mouse meet at some remote state x. Possibilities:

Mouse movements

- The cat and the mouse meet at some remote state x. Possibilities:
  - mouse up, cat up (p^2)

Mouse movements

- The cat and the mouse meet at some remote state x. Possibilities:
  - mouse up, cat up (p^2)
  - mouse down, cat down (q^2)

Mouse movements

- The cat and the mouse meet at some remote state x. Possibilities:
  - mouse up, cat up (p^2)
  - mouse down, cat down (q^2)
  - mouse down, cat up (pq), they will meet again soon

Mouse movements

- The cat and the mouse meet at some remote state x. Possibilities:
  - mouse up, cat up (p^2)
  - mouse down, cat down (q^2)
  - mouse down, cat up (pq), they will meet again soon
  - mouse up cat down (pq), the cat leaves 'forever' w.p. $(1-\rho^2)$.

Mouse movements

- The cat and the mouse meet at some remote state x. Possibilities:
  - mouse up, cat up (p^2)
  - mouse down, cat down (q^2)
  - mouse down, cat up (pq), they will meet again soon
  - mouse up cat down (pq), the cat leaves 'forever' w.p. $(1-\rho^2)$.
- Mouse makes a geometric number G of steps, where

$$\mathbb{P}(G=k)=[1-pq(1-\rho^2)]^k pq(1-\rho^2), \quad k=0,1,\ldots$$

and then the cat leaves for a long time.

Mouse movements

- The cat and the mouse meet at some remote state x. Possibilities:
  - mouse up, cat up (p^2)
  - mouse down, cat down (q^2)
  - mouse down, cat up (pq), they will meet again soon
  - mouse up cat down (pq), the cat leaves 'forever' w.p. $(1-\rho^2)$.
- Mouse makes a geometric number G of steps, where

$$\mathbb{P}(G=k) = [1-pq(1-\rho^2)]^k pq(1-\rho^2), \quad k=0,1,\dots$$

and then the cat leaves for a long time. A step X distributed as follows:

$$X = \begin{cases} +1 & \text{with probability } \frac{p^2+pq\rho^2}{1-pq(1-\rho^2)} = \frac{p^2(1+\rho)}{1-pq(1-\rho^2)}, \\ -1 & \text{with probability } \frac{q^2+pq}{1-pq(1-\rho^2)} = \frac{q}{1-pq(1-\rho^2)}. \end{cases}$$

The last step of the mouse is always of size $+1$. After that the cat and the mouse will meet again after $\approx$ exponentially distributed time with parameter $\rho^y(\mu-\lambda)^2/\mu$, at the mouse's current position y.

‘Free’ process

- ‘Free’ continuous-time Markov process $(\tilde{M}(t))$: At state x, the transition rate of $(\tilde{M}(t))$ is $\rho^x(\mu-\lambda)^2/\mu$. At each transition, the process $(\tilde{M}(t))$ makes a jump of a random size

$$\Delta = 1 + \sum_{i=1}^{G} X_i$$

- ‘Free’ continuous-time Markov process $(\tilde{M}(t))$: At state x, the transition rate of $(\tilde{M}(t))$ is $\rho^x(\mu-\lambda)^2/\mu$. At each transition, the process $(\tilde{M}(t))$ makes a jump of a random size

$$\Delta = 1 + \sum_{i=1}^{G} X_i$$

- $\mathbb{E}[\Delta] = \mathbb{E}[G]\mathbb{E}[X] + 1 = -\rho^{-1} < 0$, and $\mathbb{E}[u^\Delta]$ is well-defined if

$$u_1 = \frac{(1-\sqrt{1-4p^2})q}{2p^2} < u < \frac{(1+\sqrt{1-4p^2})q}{2p^2} = u_2.$$

$\mathbb{E}[u^\Delta]$ is well defined on the interval $[1, 1/\rho]$.

‘Free’ process

- ‘Free’ continuous-time Markov process $(\tilde{M}(t))$: At state x, the transition rate of $(\tilde{M}(t))$ is $\rho^x(\mu-\lambda)^2/\mu$. At each transition, the process $(\tilde{M}(t))$ makes a jump of a random size

$$\Delta = 1 + \sum_{i=1}^{G} X_i$$

- $\mathbb{E}[\Delta] = \mathbb{E}[G]\mathbb{E}[X] + 1 = -\rho^{-1} < 0$, and $\mathbb{E}[u^\Delta]$ is well-defined if

$$u_1 = \frac{(1-\sqrt{1-4p^2})q}{2p^2} < u < \frac{(1+\sqrt{1-4p^2})q}{2p^2} = u_2.$$

  $\mathbb{E}[u^\Delta]$ is well defined on the interval $[1, 1/\rho]$.
- After the first transition of size Δ, the transition rate becomes $\rho^{x+\Delta}(\mu-\lambda)^2/\mu$. The expected time until the next transition is

$$\mathbb{E}[\rho^{-(x+\Delta)}\mu/(\mu-\lambda)^2] = \rho^{-x}\mu\mathbb{E}[\rho^{-\Delta}]/(\mu-\lambda)^2 = \rho^{-x}\mu/(\mu-\lambda)^2 \ (!)$$

Asymptotic behavior of the scaled process-1

- Consider following scaled process:

$$\bar{S}_x(t) = \frac{\tilde{S}([\rho^{-x}\mu/(\mu-\lambda)^2]t)}{x}, \quad t \geq 0.$$

Asymptotic behavior of the scaled process-1

- Consider following scaled process:

$$\bar{S}_x(t) = \frac{\tilde{S}([\rho^{-x}\mu/(\mu-\lambda)^2]t)}{x}, \quad t \geq 0.$$

Lemma. For any $t \geq 0$, $\lim_{x\to\infty} \bar{M}_x(t) = \begin{cases} 1, & t < W; \\ -\infty, & t \geq W. \end{cases}$

where $W = \sum_{n=0}^{\infty}[\rho^{-(\Delta_1+\ldots+\Delta_n)}]E_n$ with where $E_1, E_2, \ldots$ i.i.d. exponential(1) random variables, independent of the Δ_i's

Asymptotic behavior of the scaled process-1

- Consider following scaled process:

$$\bar{S}_x(t) = \frac{\tilde{S}([\rho^{-x}\mu/(\mu-\lambda)^2]t)}{x}, \quad t \geq 0.$$

Lemma. For any $t \geq 0$, $\lim_{x\to\infty} \bar{M}_x(t) = \begin{cases} 1, & t < W; \\ -\infty, & t \geq W. \end{cases}$

where $W = \sum_{n=0}^{\infty}[\rho^{-(\Delta_1+\ldots+\Delta_n)}]E_n$ with where $E_1, E_2, \ldots$ i.i.d. exponential(1) random variables, independent of the Δ_i's

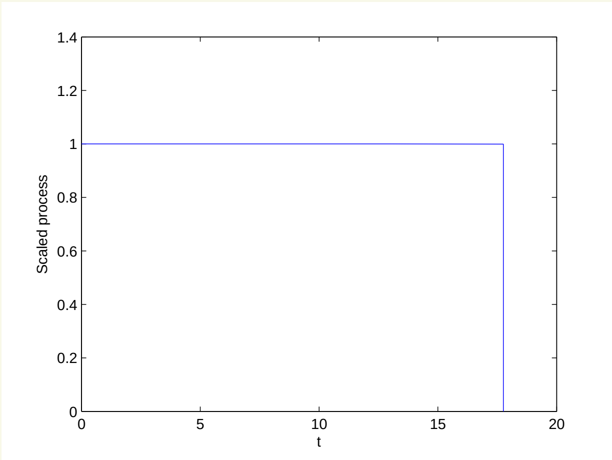

Asymptotic behavior of the scaled process-1

- Consider following scaled process:

$$\bar{S}_x(t) = \frac{\tilde{S}([\rho^{-x}\mu/(\mu-\lambda)^2]t)}{x}, \quad t \geq 0.$$

Lemma. For any $t \geq 0$, $\lim_{x\to\infty} \bar{M}_x(t) = \begin{cases} 1, & t < W; \\ -\infty, & t \geq W. \end{cases}$

where $W = \sum_{n=0}^{\infty}[\rho^{-(\Delta_1+\ldots+\Delta_n)}]E_n$ with where $E_1, E_2, \ldots$ i.i.d. exponential(1) random variables, independent of the Δ_i's

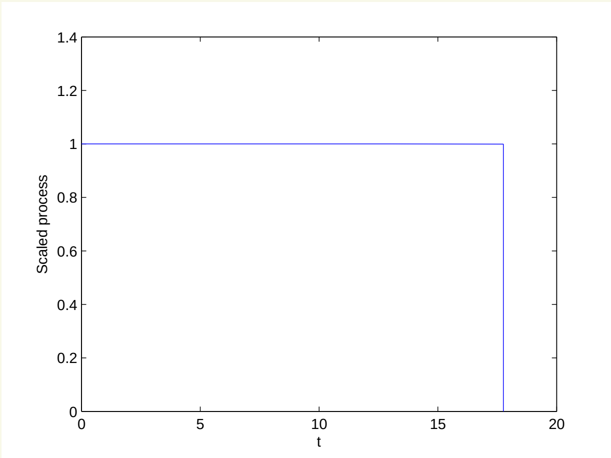

Intuition: It takes time to get on some distance from x but then the drop happens in 'no time'

Asymptotic behavior of the scaled process-2

- For the proof, we need to show that: 1) the scaled process converges to 1 before time W and to 0 after time W; 2) the time W is finite.

Asymptotic behavior of the scaled process-2

- For the proof, we need to show that: 1) the scaled process converges to 1 before time W and to 0 after time W; 2) the time W is finite.
- The random variable W satisfies

$$W \stackrel{d}{=} \rho^{-\Delta} W + E$$

This type of random variables is well-known in literature.

Asymptotic behavior of the scaled process-2

- For the proof, we need to show that: 1) the scaled process converges to 1 before time W and to 0 after time W; 2) the time W is finite.
- The random variable W satisfies

$$W \stackrel{d}{=} \rho^{-\Delta} W + E$$

This type of random variables is well-known in literature.
- W is extremely heavy-tailed: $\mathbb{E}[W^s] < \infty$ for

$$-\log(u_1)/\log(\rho) - 1 < s < 1.$$

Asymptotic behavior of the scaled process-2

- For the proof, we need to show that: 1) the scaled process converges to 1 before time W and to 0 after time W; 2) the time W is finite.
- The random variable W satisfies

$$W \stackrel{d}{=} \rho^{-\Delta} W + E$$

  This type of random variables is well-known in literature.
- W is extremely heavy-tailed: $\mathbb{E}[W^s] < \infty$ for

$$-\log(u_1)/\log(\rho) - 1 < s < 1.$$

- With other random walk of the cat ($M/M/\infty$, symmetric r.w., etc.) the mouse behavior is entirely different

Asymptotic behavior of the scaled process-2

- For the proof, we need to show that: 1) the scaled process converges to 1 before time W and to 0 after time W; 2) the time W is finite.
- The random variable W satisfies

$$W \stackrel{d}{=} \rho^{-\Delta} W + E$$

This type of random variables is well-known in literature.

- W is extremely heavy-tailed: $\mathbb{E}[W^s] < \infty$ for

$$-\log(u_1)/\log(\rho) - 1 < s < 1.$$

- With other random walk of the cat ($M/M/\infty$, symmetric r.w., etc.) the mouse behavior is entirely different
- Work in progress...

Asymptotic behavior of the scaled process-2

- For the proof, we need to show that: 1) the scaled process converges to 1 before time W and to 0 after time W; 2) the time W is finite.
- The random variable W satisfies

$$W \stackrel{d}{=} \rho^{-\Delta} W + E$$

  This type of random variables is well-known in literature.
- W is extremely heavy-tailed: $\mathbb{E}[W^s] < \infty$ for

$$-\log(u_1)/\log(\rho) - 1 < s < 1.$$

- With other random walk of the cat ($M/M/\infty$, symmetric r.w., etc.) the mouse behavior is entirely different
- Work in progress...
- I hope you liked the cat and mouse...

Asymptotic behavior of the scaled process-2

- For the proof, we need to show that: 1) the scaled process converges to 1 before time W and to 0 after time W; 2) the time W is finite.
- The random variable W satisfies

$$W \stackrel{d}{=} \rho^{-\Delta} W + E$$

This type of random variables is well-known in literature.

- W is extremely heavy-tailed: $\mathbb{E}[W^s] < \infty$ for

$$-\log(u_1)/\log(\rho) - 1 < s < 1.$$

- With other random walk of the cat ($M/M/\infty$, symmetric r.w., etc.) the mouse behavior is entirely different
- Work in progress...
- I hope you liked the cat and mouse...
- Thank you for your attention!