HAL open science

Markov Chains Competing for Transitions: Application to Large-Scale Distributed Systems

Emmanuelle Anceaume, François Castella, Romaric Ludinard, Bruno Sericola

To cite this version:

Emmanuelle Anceaume, François Castella, Romaric Ludinard, Bruno Sericola. Markov Chains Competing for Transitions: Application to Large-Scale Distributed Systems. [Research Report] 2010, pp.29. inria-00485667

HAL Id: inria-00485667

https://inria.hal.science/inria-00485667v1

Submitted on 25 May 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Markov Chains Competing for Transitions: Application to Large-Scale Distributed Systems

Emmanuelle Anceaume — François Castella — Romaric Ludinard — Bruno Sericola

N° 1953

Mai 2010

Rapport de recherche

ISSN 0249-6399 ISRN INRIA/RR--1953--FR+ENG

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE
INRIA
centre de recherche RENNES - BRETAGNE ATLANTIQUE

Markov Chains Competing for Transitions: Application to Large-Scale Distributed Systems

Emmanuelle Anceaume *, François Castella †, Romaric Ludinard ‡, Bruno Sericola ‡

Thème : Réseaux et télécommunications - Systèmes et services distribués - Modélisation, simulation et analyse numérique
Équipes-Projets Adept, Dionysos et Ipso

Rapport de recherche n° 1953 — Mai 2010 — 29 pages

Abstract: We consider the behavior of a stochastic system composed of several identically distributed, but non independent, discrete-time absorbing Markov chains competing at each instant for a transition. The competition consists in determining at each instant, using a given probability distribution, the only Markov chain allowed to make a transition. We analyze the first time at which one of the Markov chains reaches its absorbing state. We obtain its distribution and its expectation and we propose an algorithm to compute these quantities. We also exhibit the asymptotic behavior of the system when the number of Markov chains goes to infinity. Actually, this problem comes from the analysis of large-scale distributed systems and we show how our results apply to this domain.

Key-words: Markov Chains, Competing Markov Chains, Asymptotic Analysis, Large-Scale Distributed Systems

* IRISA - CNRS, Rennes, France
† Université de Rennes 1, Rennes, France
‡ INRIA Rennes - Bretagne Atlantique, Rennes, France

Centre de recherche INRIA Rennes – Bretagne Atlantique
IRISA, Campus universitaire de Beaulieu, 35042 Rennes Cedex
Téléphone : +33 2 99 84 71 00 — Télécopie : +33 2 99 84 71 71

Chaînes de Markov en compétition pour les transitions: application aux systèmes distribués grande échelle

Résumé : Nous considérons un système stochastique composé de plusieurs chaînes de Markov à temps discret, absorbantes, identiques mais non indépendantes, et en compétition à chaque instant pour une transition. La compétition consiste à déterminer à chaque instant la seule chaîne de Markov autorisée à faire une transition. On analyse le premier instant auquel une des chaînes de Markov est absorbée. On obtient sa distribution et sa moyenne et l'on propose un algorithme pour calculer ces quantités. On exhibe de plus le comportement asympotique du système quand le nombre de chaînes de Markov tend vers l'infini. En fait, ce problème vient de l'analyse des systèmes distribués grande échelle et l'on montre comment nos résultats s'appliquent à ce domaine.

Mots-clés : Chaînes de Markov, chaînes de Markov en compétition, analyse asymptotique, systèmes distribués grande échelle

1 Introduction

Competing Markov chains generally compete over a set of resources, see for instance [5] and the references therein. The resulting process is then a multi-dimensional Markov chain based on the Cartesian product of the states spaces and on competition rules over resources. When a resource is owned by a chain it affects the values of transition probabilities of the other components of the process and the measure of interest is the stationary probability distribution. In this paper, the Markov chains do not compete for resources but for a transition. More precisely, we consider a stochastic system composed of n identically distributed, but non independent, discrete-time absorbing Markov chains competing at each instant for a transition. The competition consists in determining at each instant, using a given probability mass function of dimension n, the only Markov chain allowed to make a transition.

For this system, we analyze the first time Θ_n at which one of the n Markov chains reaches its absorbing state. We study the distribution and the expectation of this random variable and we propose an algorithm to compute these quantities. We exhibit the asymptotic behavior when the number n of Markov chains goes to infinity and we use this behavior to approximate the distribution of Θ_n for large values of n. Actually, this problem comes from the analysis of large-scale distributed systems and we show how our results apply to this domain.

The remainder part of the paper is organized as follows. In the next section, we describe the model, the notation, we give the transition probability matrix of the global Markov chain composed of the n joined identically distributed local Markov chains and we obtain its transient state distribution. We then get an expression for the distribution of the first instant Θ_n at which one of the n Markov chains reaches its absorbing state. In Section 3, we show how this expression leads to a recurrence relation which allows us to easily compute the distribution of Θ_n and its expectation. In Section 4, we use a transform based on generating functions to analyze the asymptotic behavior of the distribution and the expectation of Θ_n as n goes to infinity. We then show how this analysis leads to an accurate approximation of the distribution of Θ_n for large values of n. Section 5 is devoted to an application from which the problem originally stems: we study the behavior of a cluster-based large-scale distributed system.

2 Transient State Analysis

We consider a homogeneous discrete-time Markov chain $X = \{X_k, k \geq 0\}$ with finite state space S composed of a set of transient states denoted by B and an absorbing state denoted by a. The transition probability matrix P of can thus be decomposed as

$$P = \begin{pmatrix} Q & v \\ 0 & 1 \end{pmatrix},$$

where Q is the submatrix of dimension $|B| \times |B|$ containing the transitions between states of B. In the same way, v is the column vector with $|B|$ entries representing the transitions from the transient states to the absorbing state. We suppose that the initial state is in B, i.e. $\mathbb{P}\{X_0 \in B\} = 1$, and we denote by α the row vector of dimension $|B|$ containing the initial probability distribution,

i.e. for every $i \in B$,

$$\alpha_i = \mathbb{P}\{X_0 = i\}.$$

We denote by Θ_1 the total time spent in B before reaching the absorbing state or equivalently the first instant at which the absorbing state a is reached. We have

$$\Theta_1 = \inf\{k \geq 0 \mid X_k = a\}.$$

The complementary cumulative distribution function of Θ_1 is easily derived as, see for instance [12] or [7],

$$\mathbb{P}\{\Theta_1 > k\} = \mathbb{P}\{X_k \in B\} = \alpha Q^k \mathbb{1}, \tag{1}$$

where $\mathbb{1}$ is the column vector of dimension $|B|$ with all components equal to 1 and I is the identity matrix of the right dimension. Since all the states of B are transient the matrix $I - Q$ is invertible and the expectation of Θ_1 is given by

$$E(\Theta_1) = \alpha(I - Q)^{-1}\mathbb{1}. \tag{2}$$

Let us now consider, for $n \geq 1$, n Markov chains denoted by $X^{(1)}, \ldots, X^{(n)}$ identical to X, i.e. with the same state space S, the same transition probability matrix P and the same initial probability distribution α. These n Markov chains compete at each instant to make a transition using the probability mass function $\pi(n) = (p_{1,n}, \ldots, p_{n,n})$.

From these n Markov chains, we construct a new Markov chain denoted by $Y = \{Y_k, k \geq 0\}$ as follows. The state space of Y is equal to S^n and $Y_k = (X_k^{(1)}, \ldots, X_k^{(n)})$. A transition in the Markov chain Y corresponds to a transition in only one of the Markov chains $X^{(1)}, \ldots, X^{(n)}$, all the others staying in the same state. The Markov chain that makes the transition is chosen with the probability mass function $\pi(n)$, which means that Markov chain $X^{(\ell)}$ makes the transition with probability $p_{\ell,n}$. We suppose without any loss of generality that, for every $\ell = 1, \ldots, n$, we have $0 < p_{\ell,n} < 1$.

The transition probability matrix of Y, which we denote by R, is thus given, for every $(i_1, \ldots, i_n)$ and $(j_1, \ldots, j_n) \in S^n$, by

$$R((i_1, \ldots, i_n), (j_1, \ldots, j_n)) = \begin{cases} \displaystyle\sum_{\ell=1}^{n} p_{\ell,n} P_{i_\ell, i_\ell} & \text{if } (i_1, \ldots, i_n) = (j_1, \ldots, j_n) \\ p_{\ell,n} P_{i_\ell, j_\ell} & \text{if } \exists!\ \ell \text{ such that } i_\ell \neq j_\ell \\ 0 & \text{otherwise.} \end{cases}$$

Indeed, in the Markov chain Y, the states that can be reached in one step from state $(i_1, \ldots, i_n)$ are either state $(i_1, \ldots, i_n)$ if the chosen Markov chain $X^{(\ell)}$ makes a loop over state i_ℓ (the others being static) or state $(i_1, \ldots, i_\ell - 1, j_\ell, i_\ell + 1, \ldots, i_n)$ if the chosen Markov $X^{(\ell)}$ makes a transition from state i_ℓ to state j_ℓ (the others being static). This shows that the states of Markov chain Y are all transient except state $(a, \ldots, a)$ which is absorbing.

We denote by β the initial probability distribution of Y and we assume that the $X_0^{(\ell)}$'s are independent, namely

$$\beta(j_1, \ldots, j_n) = \prod_{\ell=1}^{n} \mathbb{P}\{X_0^{(\ell)} = j_\ell\}.$$

This provides the value

$$\beta(j_1,\dots,j_n)=\prod_{\ell=1}^{n}\mathbb{P}\{X_0=j_\ell\}=\prod_{\ell=1}^{n}\alpha_{j_\ell}. \tag{3}$$

The following theorem gives an expression for the distribution of Y_k. For every $k\geq 0$ and $\ell\geq 1$, we introduce the set $S_{k,\ell}$ defined by

$$S_{k,\ell}=\{\underline{k}=(k_1,\dots,k_\ell)\in\mathbb{N}^\ell \mid k_1+\cdots+k_\ell=k\}.$$

Theorem 1 *For every $k\geq 0$, $n\geq 1$ and $(j_1,\dots,j_n)\in S^n$, we have*

$$\mathbb{P}\{Y_k=(j_1,\dots,j_n)\}=\sum_{\underline{k}\in S_{k,n}}\frac{k!}{k_1!\cdots k_n!}\prod_{r=1}^{n}p_{r,n}^{k_r}\mathbb{P}\{X_{k_r}=j_r\}. \tag{4}$$

Proof. The proof is a recurrence over k. The result is immediate for $k=0$ from relation (3). Suppose relation (4) is true for integer $k-1$. We have

$$\begin{aligned}
&\mathbb{P}\{Y_k=(j_1,\dots,j_n)\}\\
&=\sum_{(i_1,\dots,i_n)\in S^n}R((i_1,\dots,i_n),(j_1,\dots,j_n))\mathbb{P}\{Y_{k-1}=(i_1,\dots,i_n)\}\\
&=\sum_{h=1}^{n}p_{h,n}P_{j_h,j_h}\mathbb{P}\{Y_{k-1}=(j_1,\dots,j_n)\}\\
&\quad+\sum_{h=1}^{n}p_{h,n}\sum_{i_h\in S\setminus\{j_h\}}P_{i_h,j_h}\mathbb{P}\{Y_{k-1}=(j_1,\dots,j_{h-1},i_h,j_{h+1},\dots,j_n)\}\\
&=\sum_{h=1}^{n}p_{h,n}\sum_{i_h\in S}P_{i_h,j_h}\mathbb{P}\{Y_{k-1}=(j_1,\dots,j_{h-1},i_h,j_{h+1},\dots,j_n)\}.
\end{aligned}$$

Using the recurrence hypothesis, we get

$$\begin{aligned}
&\mathbb{P}\{Y_k=(j_1,\dots,j_n)\}\\
&=\sum_{h=1}^{n}p_{h,n}\sum_{i_h\in S}P_{i_h,j_h}\sum_{\underline{k}\in S_{k-1,n}}\frac{(k-1)!}{k_1!\cdots k_n!}p_{h,n}^{k_h}\mathbb{P}\{X_{k_h}=i_h\}\\
&\quad\times\prod_{r=1,r\neq h}^{n}p_{r,n}^{k_r}\mathbb{P}\{X_{k_r}=j_r\}\\
&=\sum_{h=1}^{n}\sum_{\underline{k}\in S_{k-1,n}}\frac{(k-1)!}{k_1!\cdots k_n!}p_{h,n}^{k_h+1}\mathbb{P}\{X_{k_h+1}=j_h\}\prod_{r=1,r\neq h}^{n}p_{r,n}^{k_r}\mathbb{P}\{X_{k_r}=j_r\}.
\end{aligned}$$

For every $h=1,\dots,n$, we introduce the subset $U_{h,k}$ of $S_{k,n}$ defined by

$$U_{h,k}=\{\underline{k}\in S_{k,n}\mid k_m\leq k-1 \text{ whenever } m\neq h\}.$$

The variable change $k_h:=k_h-1$ leads to

$$\mathbb{P}\{Y_k=(j_1,\dots,j_n)\}=\sum_{h=1}^{n}\sum_{\underline{k}\in U_{h,k}}\frac{k_h(k-1)!}{k_1!\cdots k_n!}\prod_{r=1}^{n}p_{r,n}^{k_r}\mathbb{P}\{X_{k_r}=j_r\}.$$

Introducing the subset V_k of $S_{k,n}$ defined by

$$V_k = \{\underline{k} \in S_{k,n} \mid k_m \leq k-1 \text{ for } m = 1, \ldots, n\},$$

we can write

$$U_{h,k} = V_k \cup \{u_h\},$$

where $u_h = (0, \ldots, 0, k, 0, \ldots, 0)$ is the row vector of dimension n with the h-th entry equal to k and the others equal to 0. We thus get

$$\begin{aligned}\mathbb{P}\{Y_k = (j_1, \ldots, j_n)\} &= \sum_{\underline{k} \in V_k} \sum_{h=1}^{n} \frac{k_h (k-1)!}{k_1! \cdots k_n!} \prod_{r=1}^{n} p_{r,n}^{k_r} \mathbb{P}\{X_{k_r} = j_r\} \\ &\quad + \sum_{h=1}^{n} p_{h,n}^{k} \mathbb{P}\{X_k = j_h\} \prod_{r=1, r \neq h}^{n} \mathbb{P}\{X_0 = j_r\}.\end{aligned}$$

For every $\underline{k} \in V_k$, we have

$$\sum_{h=1}^{n} \frac{k_h (k-1)!}{k_1! \cdots k_n!} = \frac{(k-1)!}{k_1! \cdots k_n!} \sum_{h=1}^{n} k_h = \frac{k!}{k_1! \cdots k_n!},$$

so we obtain

$$\begin{aligned}\mathbb{P}\{Y_k = (j_1, \ldots, j_n)\} &= \sum_{\underline{k} \in V_k} \frac{k!}{k_1! \cdots k_n!} \prod_{r=1}^{n} p_{r,n}^{k_r} \mathbb{P}\{X_{k_r} = j_r\} \\ &\quad + \sum_{h=1}^{n} p_{h,n}^{k} \mathbb{P}\{X_k = j_h\} \prod_{r=1, r \neq h}^{n} \mathbb{P}\{X_0 = j_r\} \\ &= \sum_{\underline{k} \in S_{k,n}} \frac{k!}{k_1! \cdots k_n!} \prod_{r=1}^{n} p_{r,n}^{k_r} \mathbb{P}\{X_{k_r} = j_r\}.\end{aligned}$$

This is the desired relation. ■

Using this result, we obtain the following corollary which gives the transient state probabilities of each Markov chain $X^{(h)}$ at the k-th transition of the global process Y.

Corollary 2 *For every $h = 1, \ldots, n$, $k \geq 0$ and $j \in S$, we have*

$$\mathbb{P}\{X_k^{(h)} = j\} = \sum_{\ell=0}^{k} \binom{k}{\ell} p_{h,n}^{\ell} (1 - p_{h,n})^{k-\ell} \mathbb{P}\{X_\ell = j\}.$$

Proof. Clearly the distribution on $X_k^{(h)}$ depends on h only through the parameter $p_{h,n}$, so to simplify the writing, we make the proof for $h = n$. Using Theorem 1, we have

$$
\begin{aligned}
&\mathbb{P}\{X_k^{(n)} = j\} \\
&= \mathbb{P}\{Y_k \in S \times \cdots \times S \times \{j\}\} \\
&= \sum_{j_1 \in S} \cdots \sum_{j_{n-1} \in S} \sum_{\underline{k} \in S_{k,n}} \frac{k!}{k_1! \cdots k_n!} \prod_{r=1}^{n-1} p_{r,n}^{k_r} \mathbb{P}\{X_{k_r} = j_r\} p_{n,n}^{k_n} \mathbb{P}\{X_{k_n} = j\} \\
&= \sum_{\underline{k} \in S_{k,n}} \frac{k!}{k_1! \cdots k_n!} p_{n,n}^{k_n} \mathbb{P}\{X_{k_n} = j\} \prod_{r=1}^{n-1} p_{r,n}^{k_r} \\
&= \sum_{k_n=0}^{k} \binom{k}{k_n} p_{n,n}^{k_n} (1 - p_{n,n})^{k-k_n} \mathbb{P}\{X_{k_n} = j\} \\
&\quad \times \sum_{\underline{k} \in S_{k-k_n,n-1}} \frac{(k-k_n)!}{k_1! \cdots k_{n-1}!} \prod_{r=1}^{n-1} \left(\frac{p_{r,n}}{1 - p_{n,n}} \right)^{k_r} \\
&= \sum_{k_n=0}^{k} \binom{k}{k_n} p_{n,n}^{k_n} (1 - p_{n,n})^{k-k_n} \mathbb{P}\{X_{k_n} = j\},
\end{aligned}
$$

which completes the proof. ■

The following corollary provides the distribution of the first instant Θ_n at which one of the n Markov chains $X^{(1)}, \ldots, X^{(n)}$ gets absorbed. More formally, Θ_n is defined as

$$\Theta_n = \inf\{k \geq 0 \mid \exists r \text{ such that } X_k^{(r)} = a\}.$$

Corollary 3 *For every $k \geq 0$ and $n \geq 1$, we have*

$$\mathbb{P}\{\Theta_n > k\} = \sum_{\underline{k} \in S_{k,n}} \frac{k!}{k_1! \cdots k_n!} \prod_{r=1}^{n} p_{r,n}^{k_r} \alpha Q^{k_r} \mathbb{1}. \tag{5}$$

Proof. For every $k \geq 0$ and $n \geq 1$, we have

$$\Theta_n > k \Longleftrightarrow X_k^{(1)} \in B, \ldots, X_k^{(n)} \in B.$$

We thus obtain from Theorem 1 and Relation (1).

$$
\begin{aligned}
\mathbb{P}\{\Theta_n > k\} &= \mathbb{P}\{X_k^{(1)} \in B, \ldots, X_k^{(n)} \in B\} \\
&= \sum_{(j_1,\ldots,j_n) \in B^n} \mathbb{P}\{Y_k = (j_1, \ldots, j_n)\} \\
&= \sum_{\underline{k} \in S_{k,n}} \frac{k!}{k_1! \cdots k_n!} \sum_{(j_1,\ldots,j_n) \in B^n} \prod_{r=1}^{n} p_{r,n}^{k_r} \mathbb{P}\{X_{k_r} = j_r\} \\
&= \sum_{\underline{k} \in S_{k,n}} \frac{k!}{k_1! \cdots k_n!} \prod_{r=1}^{n} p_{r,n}^{k_r} \mathbb{P}\{X_{k_r} \in B\} \\
&= \sum_{\underline{k} \in S_{k,n}} \frac{k!}{k_1! \cdots k_n!} \prod_{r=1}^{n} p_{r,n}^{k_r} \alpha Q^{k_r} \mathbb{1},
\end{aligned}
$$

which completes the proof. ■

Clearly the complexity for the computation of $\mathbb{P}\{\Theta_n > k\}$ using relation (5) is exponential. A solution to this problem is obtained in the following section.

3 Computational Aspects

We consider the computation of the distribution and of the expectation of Θ_n. We have seen in Section 1 that the states of the global Markov chain Y are all transient except state $(a, \ldots, a)$ which is absorbing. This means that the expectation $E(\Theta_n)$ is finite and given, for every $n \geq 1$, by

$$E(\Theta_n) = \sum_{k=0}^{\infty} \mathbb{P}\{\Theta_n > k\}. \tag{6}$$

The following theorem gives a recursive expression for the distribution Θ_n.

Theorem 4 *For every $k \geq 0$, $n \geq 2$, we have*

$$\mathbb{P}\{\Theta_n > k\} = \sum_{\ell=0}^{k} \binom{k}{\ell} p_{n,n}^{\ell} \left(1 - p_{n,n}\right)^{k-\ell} \alpha Q^{\ell} \mathbb{1} \mathbb{P}\{\Theta_{n-1} > k - \ell\}, \tag{7}$$

where the probability mass distribution $\pi(n-1) = (p_{1,n-1}, \ldots, p_{n-1,n-1})$ associated with Θ_{n-1} is defined, for $r = 1, \ldots, n-1$, by

$$p_{r,n-1} = \frac{p_{r,n}}{1 - p_{n,n}}.$$

Proof. For every $k \geq 0$ and $n \geq 2$, we have from Corollary 3

$$\begin{aligned}
&\mathbb{P}\{\Theta_n > k\} \\
&= \sum_{\underline{k} \in S_{k,n}} \frac{k!}{k_1! \cdots k_n!} \prod_{r=1}^{n} \left(p_{r,n}\right)^{k_r} \alpha Q^{k_r} \mathbb{1} \\
&= \sum_{k_n=0}^{k} \left(p_{n,n}\right)^{k_n} \frac{\alpha Q^{k_n} \mathbb{1}}{k_n!} \sum_{\underline{k} \in S_{k-k_n,n-1}} \frac{k!}{k_1! \cdots k_{n-1}!} \prod_{r=1}^{n-1} \left(p_{r,n}\right)^{k_r} \alpha Q^{k_r} \mathbb{1} \\
&= \sum_{k_n=0}^{k} \binom{k}{k_n} \left(p_{n,n}\right)^{k_n} \alpha Q^{k_n} \mathbb{1} \sum_{\underline{k} \in S_{k-k_n,n-1}} \frac{(k-k_n)!}{k_1! \cdots k_{n-1}!} \prod_{r=1}^{n-1} \left(p_{r,n}\right)^{k_r} \alpha Q^{k_r} \mathbb{1} \\
&= \sum_{k_n=0}^{k} \binom{k}{k_n} \left(p_{n,n}\right)^{k_n} \left(1 - p_{n,n}\right)^{k-k_n} \alpha Q^{k_n} \mathbb{1} \\
&\quad \times \sum_{\underline{k} \in S_{k-k_n,n-1}} \frac{(k-k_n)!}{k_1! \cdots k_{n-1}!} \prod_{r=1}^{n-1} \left(\frac{p_{r,n}}{1 - p_{n,n}}\right)^{k_r} \alpha Q^{k_r} \mathbb{1} \\
&= \sum_{k_n=0}^{k} \binom{k}{k_n} \left(p_{n,n}\right)^{k_n} \left(1 - p_{n,n}\right)^{k-k_n} \alpha Q^{k_n} \mathbb{1} \\
&\quad \times \sum_{\underline{k} \in S_{k-k_n,n-1}} \frac{(k-k_n)!}{k_1! \cdots k_{n-1}!} \prod_{r=1}^{n-1} \left(p_{r,n-1}\right)^{k_r} \alpha Q^{k_r} \mathbb{1} \\
&= \sum_{k_n=0}^{k} \binom{k}{k_n} \left(p_{n,n}\right)^{k_n} \left(1 - p_{n,n}\right)^{k-k_n} \alpha Q^{k_n} \mathbb{1} \mathbb{P}\{\Theta_{n-1} > k - k_n\},
\end{aligned}$$

which completes the proof. ■

This result shows that the computation of $\mathbb{P}\{\Theta_n > k\}$ can be done using a simple recurrence with a *polynomial* complexity.

We show in the following how to stop computation at a step K, to be determined, in order to make sure that for every $i = 1, \ldots, n$, $\mathbb{P}\{\Theta_i > k\} \leq \varepsilon$ for every $k \geq K$ where ε is a predetermined error tolerance.

Without any loss of generality, we suppose that the probabilities $p_{r,n}$ are ordered such that

$$p_{1,n} \geq p_{2,n} \geq \cdots \geq p_{n,n}.$$

The following result holds true.

Lemma 5 *For every $i = 1, \ldots, n-1$ and $k \geq 0$, we have*

$$(p_{i,i}Q + (1 - p_{i,i})I)^k \, \mathbb{1} \leq (p_{i+1,i+1}Q + (1 - p_{i+1,i+1})I)^k \, \mathbb{1},$$

where the inequality between vectors is meant entrywise.

Proof. The result is clearly true for $k = 0$. From Theorem 4, we have

$$p_{i,i} - p_{i+1,i+1} = \frac{p_{i,i+1}}{1 - p_{i+1,i+1}} - p_{i+1,i+1} = \frac{p_{i,i+1} - p_{i+1,i+1} + p_{i+1,i+1}^2}{1 - p_{i+1,i+1}} \geq 0.$$

The matrix Q being substochastic, we have $Q\mathbb{1} - \mathbb{1} \leq 0$ with the strict inequality for at least one entry, Here 0 is the null column vector having relevant dimension. Introducing the notation $Q_i = p_{i,i}Q + (1 - p_{i,i})I$, we get

$$Q_i\mathbb{1} - Q_{i+1}\mathbb{1} = (p_{i,i} - p_{i+1,i+1})(Q\mathbb{1} - \mathbb{1}) \leq 0,$$

which means that the result is true for $k = 1$. Suppose now that the result is true for integer $k - 1$, with $k \geq 2$. Since matrices Q_i and Q_{i+1} commute, we have

$$(Q_i)^k\mathbb{1} = Q_i(Q_i)^{k-1}\mathbb{1} \leq Q_i(Q_{i+1})^{k-1}\mathbb{1} = (Q_{i+1})^{k-1}Q_i\mathbb{1} \leq (Q_{i+1})^k\mathbb{1}.$$

This completes the proof. ■

Theorem 6 *For every $n \geq 1$, for every $\varepsilon \in (0, 1)$, we have*

$$\max_{i=1,\ldots,n} \mathbb{P}\{\Theta_i > k\} \leq \varepsilon \text{ for every } k \geq K_1,$$

where

$$K_1 = \inf\left\{k \geq 0 \;\middle|\; \sum_{\ell=0}^{k} \binom{k}{\ell} p_{n,n}^{\ell}(1 - p_{n,n})^{k-\ell}\alpha Q^{\ell}\mathbb{1} \leq \varepsilon\right\}.$$

Proof. For every $i = 1, \ldots, n$, we have

$$\begin{aligned}
\mathbb{P}\{\Theta_i > k\} &= \sum_{\ell=0}^{k} \binom{k}{\ell} p_{i,i}^{\ell} (1-p_{i,i})^{k-\ell} \alpha Q^{\ell} \mathbb{1} \mathbb{P}\{\Theta_{i-1} > k-\ell\} \\
&\leq \sum_{\ell=0}^{k} \binom{k}{\ell} p_{i,i}^{\ell} (1-p_{i,i})^{k-\ell} \alpha Q^{\ell} \mathbb{1} \\
&= \alpha \left(p_{i,i} Q + (1-p_{i,i}) I\right)^k \mathbb{1} \\
&\leq \alpha \left(p_{n,n} Q + (1-p_{n,n}) I\right)^k \mathbb{1} \qquad \text{(from Lemma 5)} \\
&= \sum_{\ell=0}^{k} \binom{k}{\ell} p_{n,n}^{\ell} (1-p_{n,n})^{k-\ell} \alpha Q^{\ell} \mathbb{1}.
\end{aligned}$$

Note that matrix $p_{n,n}Q + (1-p_{n,n})I$ is substochastic, i. e. $(p_{n,n}Q + (1-p_{n,n})I)\mathbb{1} \leq \mathbb{1}$ with the strict inequality for at least one entry. This means in particular that $\alpha \left(p_{n,n}Q + (1-p_{n,n})I\right)^k \mathbb{1}$ is decreasing with k and

$$\lim_{k \longrightarrow \infty} \alpha \left(p_{n,n}Q + (1-p_{n,n})I\right)^k \mathbb{1} = 0,$$

So, for a fixed $\varepsilon \in (0,1)$ and by definition of integer K_1 we have that for every $i = 1, \ldots, n$,

$$\mathbb{P}\{\Theta_i > k\} \leq \varepsilon, \text{ for every } k \geq K_1,$$

which completes the proof. ■

The same kind of results holds for the evaluation of the expected values $E(\Theta_i)$ for which the truncation of the series (6) is needed.

Theorem 7 *For every $n \geq 1$, for every $\varepsilon \in (0,1)$,*

$$0 \leq \max_{i=1,\ldots,n} \left(E(\Theta_i) - \sum_{k=0}^{K_2-1} \mathbb{P}\{\Theta_i > k\} \right) \leq \varepsilon,$$

where

$$K_2 = \inf \left\{ k \geq 0 \;\middle|\; \frac{1}{p_{n,n}} \sum_{\ell=0}^{k} \binom{k}{\ell} p_{n,n}^{\ell} (1-p_{n,n})^{k-\ell} \alpha (I-Q)^{-1} Q^{\ell} \mathbb{1} \leq \varepsilon \right\}.$$

Proof. We introduce the notation

$$e_i = E(\Theta_i) - \sum_{k=0}^{K_2-1} \mathbb{P}\{\Theta_i > k\}.$$

We then have, for every $i = 1, \ldots, n$,

$$\begin{aligned}
e_i &= \sum_{k=K_2}^{\infty} \mathbb{P}\{\Theta_i > k\} \\
&= \sum_{k=K_2}^{\infty} \sum_{\ell=0}^{k} \binom{k}{\ell} p_{i,i}^{\ell} (1-p_{i,i})^{k-\ell} \alpha Q^{\ell} \mathbb{1} \mathbb{P}\{\Theta_{i-1} > k-\ell\} \\
&\leq \sum_{k=K_2}^{\infty} \sum_{\ell=0}^{k} \binom{k}{\ell} p_{i,i}^{\ell} (1-p_{i,i})^{k-\ell} \alpha Q^{\ell} \mathbb{1} \\
&= \sum_{k=K_2}^{\infty} \alpha \left(p_{i,i} Q + (1-p_{i,i}) I\right)^k \mathbb{1} \\
&\leq \sum_{k=K_2}^{\infty} \alpha \left(p_{n,n} Q + (1-p_{n,n}) I\right)^k \mathbb{1} \qquad \text{(from Lemma 5)} \\
&= \alpha \left(I - (p_{n,n} Q + (1-p_{n,n}) I)\right)^{-1} \left(p_{n,n} Q + (1-p_{n,n}) I\right)^{K_2} \mathbb{1} \\
&= \frac{1}{p_{n,n}} \alpha (I-Q)^{-1} \left(p_{n,n} Q + (1-p_{n,n}) I\right)^{K_2} \mathbb{1} \\
&= \frac{1}{p_{n,n}} \sum_{\ell=0}^{K_2} \binom{K_2}{\ell} p_{n,n}^{\ell} (1-p_{n,n})^{K_2-\ell} \alpha (I-Q)^{-1} Q^{\ell} \mathbb{1} \\
&\leq \varepsilon \text{ by definition of integer } K_2.
\end{aligned}$$

which means that $\max_{i=1,\ldots,n} e_i \leq \varepsilon$. ■

The computation of the distribution of Θ_n can thus be done recursively using relations (1) and (7). If the probability mass function $\pi(n)$ is uniform, i.e. if $p_{r,n} = 1/n$ for every $r = 1, \ldots, n$, then it is easily checked that the probability mass function $\pi(i)$ is also uniform for every $i \leq n$. In this case, we obtain at the same cost not only the distribution of Θ_n but also the distribution of all Θ_i's, for $i \leq n$, using the stopping criterion given in Theorem 6. If the expectation of Θ_n is also needed then we may use the stopping criterion given in Theorem 7.

4 Asymptotic Analysis

This section is devoted to the analysis of the distribution on Θ_n when n is large. This is generally the case in practice for large-scale distributed systems which are studied in the next section.

First of all, the following result can be deduced easily from Theorem 4.

Theorem 8 *For every $n \geq 1$, we have*

$$E(\Theta_n) \leq \frac{E(\Theta_{n-1})}{1-p_{n,n}},$$

and

$$E(\Theta_n) \leq \frac{E(\Theta_1)}{p_{1,n}}.$$

Proof. Using Theorem 4, we get

$$\begin{aligned}
E(\Theta_n) &= \sum_{k=0}^{\infty} \mathbb{P}\{\Theta_n > k\} \\
&= \sum_{k=0}^{\infty} \sum_{\ell=0}^{k} \binom{k}{\ell} p_{n,n}^{\ell} \left(1 - p_{n,n}\right)^{k-\ell} \alpha Q^{\ell} \mathbb{1} \mathbb{P}\{\Theta_{n-1} > k - \ell\} \\
&= \sum_{k=0}^{\infty} \sum_{\ell=0}^{k} \binom{k}{\ell} p_{n,n}^{k-\ell} \left(1 - p_{n,n}\right)^{\ell} \alpha Q^{k-\ell} \mathbb{1} \mathbb{P}\{\Theta_{n-1} > \ell\} \\
&= \sum_{\ell=0}^{\infty} \left(1 - p_{n,n}\right)^{\ell} \mathbb{P}\{\Theta_{n-1} > \ell\} \sum_{k=\ell}^{\infty} \binom{k}{\ell} p_{n,n}^{k-\ell} \alpha Q^{k-\ell} \mathbb{1} \\
&= \sum_{\ell=0}^{\infty} \left(1 - p_{n,n}\right)^{\ell} \mathbb{P}\{\Theta_{n-1} > \ell\} \sum_{k=0}^{\infty} \binom{k+\ell}{\ell} p_{n,n}^{k} \alpha Q^{k} \mathbb{1} \\
&\leq \sum_{\ell=0}^{\infty} \left(1 - p_{n,n}\right)^{\ell} \mathbb{P}\{\Theta_{n-1} > \ell\} \sum_{k=0}^{\infty} \binom{k+\ell}{\ell} p_{n,n}^{k} \\
&= \sum_{\ell=0}^{\infty} \left(1 - p_{n,n}\right)^{\ell} \mathbb{P}\{\Theta_{n-1} > \ell\} \left(1 - p_{n,n}\right)^{-(\ell+1)} \\
&= \frac{E(\Theta_{n-1})}{1 - p_{n,n}}.
\end{aligned}$$

Recursively, we get

$$E(\Theta_n) \leq \frac{E(\Theta_{n-1})}{1 - p_{n,n}} \leq \frac{E(\Theta_{n-2})}{(1 - p_{n,n})(1 - p_{n-1,n-1})} = \frac{E(\Theta_{n-2})}{1 - p_{n,n} - p_{n-1,n}}$$

and so,

$$E(\Theta_n) \leq \frac{E(\Theta_1)}{p_{1,n}}.$$

which completes the proof. ■

When the distribution $\pi(n)$ is uniform, i.e. when $p_{i,n} = 1/n$, we obtain

$$E(\Theta_n) \leq \frac{nE(\Theta_{n-1})}{n-1} \leq nE(\Theta_1).$$

We consider the following transform. For every $n \geq 1$ and $x \in \mathbb{R}$, we introduce the function $F_n(x)$ defined by

$$F_n(x) = \sum_{k=0}^{\infty} \frac{x^k}{k!} \mathbb{P}\{\Theta_n > k\}.$$

The function F_n is defined for every $x \in \mathbb{R}$ and an explicit expression is given in the following theorem.

Theorem 9 *For every $n \geq 1$ and $x \in \mathbb{R}$, we have*

$$F_n(x) = \prod_{\ell=1}^{n} \alpha e^{Qxp_{\ell,n}} \mathbb{1}, \tag{8}$$

and, for every $k \in \mathbb{N}$,

$$\mathbb{P}\{\Theta_n > k\} = F_n^{(k)}(0), \tag{9}$$

where $F_n^{(k)}$ is the k-th derivative of function F_n with respect to x.

Proof. From Relation (7), we get, using the fact that $\mathbb{P}\{\Theta_1 > \ell\} = \alpha Q^\ell \mathbb{1}$,

$$\begin{aligned}
F_n(x) &= \sum_{k=0}^{\infty} \frac{x^k}{k!} \sum_{\ell=0}^{k} \binom{k}{\ell} p_{n,n}^\ell \left(1 - p_{n,n}\right)^{k-\ell} \mathbb{P}\{\Theta_1 > \ell\} \mathbb{P}\{\Theta_{n-1} > k - \ell\} \\
&= \sum_{\ell=0}^{\infty} \frac{x^\ell}{\ell!} p_{n,n}^\ell \mathbb{P}\{\Theta_1 > \ell\} \sum_{k=\ell}^{\infty} \frac{x^{k-\ell}}{(k-\ell)!} \left(1 - p_{n,n}\right)^{k-\ell} \mathbb{P}\{\Theta_{n-1} > k - \ell\} \\
&= \sum_{\ell=0}^{\infty} \frac{x^\ell}{\ell!} p_{n,n}^\ell \mathbb{P}\{\Theta_1 > \ell\} \sum_{k=0}^{\infty} \frac{x^k}{k!} \left(1 - p_{n,n}\right)^{k} \mathbb{P}\{\Theta_{n-1} > k\} \\
&= F_1(x p_{n,n}) F_{n-1}(x(1 - p_{n,n})).
\end{aligned}$$

This leads by induction, for $i = 1, \ldots, n$, to

$$\begin{aligned}
F_n(x) &= F_1(x p_{n,n}) F_1(x p_{n-1,n}) F_{n-2}(x(1 - p_{n,n} - p_{n-1,n})) \\
&= \cdots \\
&= \left[\prod_{\ell=n-i+1}^{n} F_1(x p_{\ell,n}) \right] F_{n-i}(x(1 - p_{n,n} - \ldots - p_{n-i+1,n})),
\end{aligned}$$

where we used the relations

$$p_{r,n-1} = \frac{p_{r,n}}{1 - p_{n,n}}.$$

We deduce, for $i = n$,

$$F_n(x) = \prod_{\ell=1}^{n} F_1(x p_{\ell,n}).$$

Now, since

$$F_1(x) = \sum_{k=0}^{\infty} \frac{x^k}{k!} \mathbb{P}\{\Theta_1 > k\} = \sum_{k=0}^{\infty} \frac{x^k}{k!} \alpha Q^k \mathbb{1} = \alpha e^{Qx} \mathbb{1},$$

we obtain

$$F_n(x) = \prod_{\ell=1}^{n} \alpha e^{Q x p_{\ell,n}} \mathbb{1}.$$

For the second part of the theorem, we simply write, by definition of the function F_n

$$F_n^{(h)}(x) = \sum_{k=h}^{\infty} \frac{x^{k-h}}{(k-h)!} \mathbb{P}\{\Theta_n > k\} = \sum_{k=0}^{\infty} \frac{x^k}{k!} \mathbb{P}\{\Theta_n > k + h\},$$

which gives the result by taking $x = 0$. ■

This result not only shows that $\mathbb{P}\{\Theta_n > 0\} = 1$ as expected, but also that, for every $n \geq 1$, we have

$$\mathbb{P}\{\Theta_n > 1\} = \alpha Q \mathbb{1}.$$

4.1 Uniform distribution

We consider in this section the case where the probability mass function $\pi(n)$ is uniform. The following theorem gives, in that case, the asymptotic distribution of Θ_n.

Theorem 10 *If $p_{\ell,n} = 1/n$ for every $\ell = 1, \ldots, n$, then for every $x \in \mathbb{R}$, we have*

$$\lim_{n \longrightarrow \infty} F_n(x) = e^{\alpha Q \mathbb{1} x}.$$

Proof. If $p_{\ell,n} = 1/n$ for every $\ell = 1, \ldots, n$, we have from relation (8),

$$F_n(x) = \left(\alpha e^{Qx/n} \mathbb{1}\right)^n,$$

which can also be written as

$$F_n(x) = e^{n \ln\left(\alpha e^{Qx/n} \mathbb{1}\right)} = e^{n(\alpha Q \mathbb{1} x/n + \varepsilon(1/n)/n))} = e^{\alpha Q \mathbb{1} x + \varepsilon(1/n)},$$

where ε is a function satisfying $\lim_{n \longrightarrow \infty} \varepsilon(1/n) = 0$.

This completes the proof. ■

This result suggests that the limit of $\mathbb{P}\{\Theta_n > k\}$ when n goes to infinity should be equal to $(\alpha Q \mathbb{1})^k$. Actually, if $p_{\ell,n} = 1/n$ for every $\ell = 1, \ldots, n$, we easily get from relations (8) and (9)

$$\mathbb{P}\{\Theta_n > k\} = F_n^{(k)}(0) = \left. \frac{d^k \left(\alpha e^{Qx/n} \mathbb{1}\right)^n}{dx^k} \right|_{x=0}.$$

This gives for instance when $n \geq 4$,

$$\begin{aligned}
F_n^{(1)}(x) &= \left(\alpha e^{Qx/n} \mathbb{1}\right)^{n-1} \alpha Q e^{Qx/n} \mathbb{1}, \\
F_n^{(2)}(x) &= \frac{n-1}{n} \left(\alpha e^{Qx/n} \mathbb{1}\right)^{n-2} \left(\alpha Q e^{Qx/n} \mathbb{1}\right)^2 \\
&\quad + \frac{1}{n} \left(\alpha e^{Qx/n} \mathbb{1}\right)^{n-1} \alpha Q^2 e^{Qx/n} \mathbb{1}, \\
F_n^{(3)}(x) &= \frac{(n-1)(n-2)}{n^2} \left(\alpha e^{Qx/n} \mathbb{1}\right)^{n-3} \left(\alpha Q e^{Qx/n} \mathbb{1}\right)^3 \\
&\quad + \frac{3(n-1)}{n^2} \left(\alpha e^{Qx/n} \mathbb{1}\right)^{n-2} \alpha Q e^{Qx/n} \mathbb{1} \alpha Q^2 e^{Qx/n} \mathbb{1} \\
&\quad + \frac{1}{n^2} \left(\alpha e^{Qx/n} \mathbb{1}\right)^{n-1} \alpha Q^3 e^{Qx/n} \mathbb{1}, \\
F_n^{(4)}(x) &= \frac{(n-1)(n-2)(n-3)}{n^3} \left(\alpha e^{Qx/n} \mathbb{1}\right)^{n-4} \left(\alpha Q e^{Qx/n} \mathbb{1}\right)^4 \\
&\quad + \frac{6(n-1)(n-2)}{n^3} \left(\alpha e^{Qx/n} \mathbb{1}\right)^{n-3} \left(\alpha Q e^{Qx/n} \mathbb{1}\right)^2 \alpha Q^2 e^{Qx/n} \mathbb{1} \\
&\quad + \frac{3(n-1)}{n^3} \left(\alpha e^{Qx/n} \mathbb{1}\right)^{n-2} \left(\alpha Q^2 e^{Qx/n} \mathbb{1}\right)^2 \\
&\quad + \frac{4(n-1)}{n^3} \left(\alpha e^{Qx/n} \mathbb{1}\right)^{n-2} \alpha Q e^{Qx/n} \mathbb{1} \alpha Q^3 e^{Qx/n} \mathbb{1} \\
&\quad + \frac{1}{n^3} \left(\alpha e^{Qx/n} \mathbb{1}\right)^{n-1} \alpha Q^4 e^{Qx/n} \mathbb{1},
\end{aligned}$$

and so, for $x = 0$,

$$\begin{aligned}
\mathbb{P}\{\Theta_n > 1\} &= \alpha Q \mathbb{1}, \\
\mathbb{P}\{\Theta_n > 2\} &= \frac{n-1}{n}(\alpha Q \mathbb{1})^2 + \frac{1}{n}\alpha Q^2 \mathbb{1}, \\
\mathbb{P}\{\Theta_n > 3\} &= \frac{(n-1)(n-2)}{n^2}(\alpha Q \mathbb{1})^3 + \frac{3(n-1)}{n^2}\alpha Q \mathbb{1} \alpha Q^2 \mathbb{1} + \frac{1}{n^2}\alpha Q^3 \mathbb{1}, \\
\mathbb{P}\{\Theta_n > 4\} &= \frac{(n-1)(n-2)(n-3)}{n^3}(\alpha Q \mathbb{1})^4 + \frac{6(n-1)(n-2)}{n^3}(\alpha Q \mathbb{1})^2 \alpha Q^2 \mathbb{1} \\
&\quad + \frac{3(n-1)}{n^3}\left(\alpha Q^2 \mathbb{1}\right)^2 + \frac{4(n-1)}{n^3}\alpha Q \mathbb{1} \alpha Q^3 \mathbb{1} + \frac{1}{n^3}\alpha Q^4 \mathbb{1}.
\end{aligned}$$

An explicit expression of $F_n^{(k)}(x)$ can be obtained using the Faà Di Bruno formula, see for instance [3] or [14].

Theorem 11 (Faà Di Bruno formula) *When functions f and g are smooth enough, we have*

$$\frac{d^k}{dx^k} f(g(x)) = k! \sum_{j=1}^{k} \frac{f^{(j)}(g(x))}{j!} \sum_{\underline{m} \in T_{j,k}} \frac{j!}{m_1! \cdots m_k!} \prod_{\ell=1}^{k} \left(\frac{g^{(\ell)}(x)}{\ell!} \right)^{m_\ell},$$

where $T_{j,k}$ is defined by

$$T_{j,k} = \left\{ \underline{m} = (m_1, \ldots, m_k) \,\middle|\, \sum_{\ell=1}^{k} m_\ell = j \text{ and } \sum_{\ell=1}^{k} \ell m_\ell = k \right\}. \tag{10}$$

Taking $g(x) = \alpha e^{Qx/n} \mathbb{1}$ and $f(x) = x^n$, we get

$$f^{(j)}(x) = \frac{n!}{(n-j)!} x^{n-j} 1_{\{j \le n\}} \text{ and } g^{(\ell)}(x) = \frac{1}{n^\ell} \alpha Q^\ell e^{Qx/n} \mathbb{1},$$

which leads, for every $n \ge 1$ and $k \ge 1$, to $F_n(x) = f(g(x))$ and to

$$F_n^{(k)}(x) = \frac{k!}{n^k} \sum_{j=1}^{k \wedge n} \binom{n}{j} \left(\alpha e^{Q \frac{x}{n}} \mathbb{1} \right)^{n-j} \sum_{\underline{m} \in T_{j,k}} \frac{j!}{m_1! \cdots m_k!} \prod_{\ell=1}^{k} \left(\frac{\alpha Q^\ell e^{Qx/n} \mathbb{1}}{\ell!} \right)^{m_\ell}.$$

Thus, taking $x = 0$, we recover

$$\mathbb{P}\{\Theta_n > k\} = \frac{k!}{n^k} \sum_{j=1}^{k \wedge n} \binom{n}{j} \sum_{\underline{m} \in T_{j,k}} \frac{j!}{m_1! \cdots m_k!} \prod_{\ell=1}^{k} \left(\frac{\alpha Q^\ell \mathbb{1}}{\ell!} \right)^{m_\ell}. \tag{11}$$

Note that by taking $g(x) = e^{x/n}$ and $f(x) = x^n$, we obtain $f(g(x)) = e^x$ and so, for every $n \ge 1$, we also recover

$$\frac{k!}{n^k} \sum_{j=1}^{k \wedge n} \binom{n}{j} \sum_{\underline{m} \in T_{j,k}} \frac{j!}{m_1!(1!)^{m_1} m_2!(2!)^{m_2} \cdots m_k!(k!)^{m_k}} = 1. \tag{12}$$

The limiting behavior of $\mathbb{P}\{\Theta_n > k\}$ is given by the following theorem.

Theorem 12 *If $p_{\ell,n} = 1/n$ for every $\ell = 1, \ldots, n$, then for every $k \geq 0$, we have*

$$\lim_{n \longrightarrow \infty} \mathbb{P}\{\Theta_n > k\} = (\alpha Q \mathbb{1})^k$$

and,

$$\lim_{n \longrightarrow \infty} E(\Theta_n) = \begin{cases} \dfrac{1}{1 - \alpha Q \mathbb{1}} & \text{if} \quad \alpha Q \mathbb{1} < 1 \\ \infty & \text{if} \quad \alpha Q \mathbb{1} = 1. \end{cases}$$

Proof. We extend the definition of the function F_n to the whole complex plane $\mathbb{C}$. For $z \in \mathbb{C}$, we have

$$F_n(z) = \sum_{k=0}^{\infty} \frac{z^k}{k!} \mathbb{P}\{\Theta_n > k\},$$

and we consider the unit circle $C = \{z \in \mathbb{C} \mid |z| = 1\}$. The function $F_n(z)$ being analytic on $\mathbb{C}$, we have, for every $p \geq 0$,

$$\begin{aligned} \oint_C \frac{F_n(z)}{z^{p+1}} dz &= i \int_0^{2\pi} F_n(e^{i\eta}) e^{-ip\eta} d\eta \\ &= \sum_{k=0}^{\infty} \frac{\mathbb{P}\{\Theta_n > k\}}{k!} \int_0^{2\pi} i e^{i(k-p)\eta} d\eta \\ &= 2i\pi \frac{\mathbb{P}\{\Theta_n > p\}}{p!}, \end{aligned}$$

since

$$\int_0^{2\pi} i e^{i(k-p)\eta} d\eta = \begin{cases} 0 & \text{if} \quad k \neq p \\ 2i\pi & \text{if} \quad k = p. \end{cases}$$

We thus have, for every $p \geq 0$,

$$\mathbb{P}\{\Theta_n > p\} = \frac{p!}{2i\pi} \oint_C \frac{F_n(z)}{z^{p+1}} dz.$$

We have shown in Theorem 10 that $F_n(z)$ converges pointwise in z to $e^{\alpha Q \mathbb{1} z}$ when n goes to infinity. Moreover, we have

$$|F_n(z)| \leq e^{|z|} \text{ and } \oint_C \frac{e^{|z|}}{|z|^{p+1}} |dz| = e \oint_C |dz| = e \int_0^{2\pi} d\eta = 2e\pi.$$

So, from the dominated convergence theorem, we get

$$\begin{aligned} \lim_{n \longrightarrow \infty} \mathbb{P}\{\Theta_n > p\} &= \frac{p!}{2i\pi} \oint_C \frac{e^{\alpha Q \mathbb{1} z}}{z^{p+1}} dz \\ &= \frac{p!}{2i\pi} \sum_{k=0}^{\infty} \frac{(\alpha Q \mathbb{1})^k}{k!} \int_0^{2\pi} i e^{i(k-p)\eta} d\eta \\ &= (\alpha Q \mathbb{1})^p, \end{aligned}$$

which completes the first part of the proof.

For the second part, we recall that

$$E(\Theta_n) = \sum_{k=0}^{\infty} \mathbb{P}\{\Theta_n > k\}.$$

We first consider the case where $\alpha Q \mathbb{1} = 1$. From the previous part of the proof, we have, for every $k \geq 0$,

$$\lim_{n \longrightarrow \infty} \mathbb{P}\{\Theta_n > k\} = 1.$$

Using Fatou's lemma for series, we get

$$\liminf_{n \longrightarrow \infty} \sum_{k=0}^{\infty} \mathbb{P}\{\Theta_n > k\} \geq \sum_{k=0}^{\infty} \lim_{n \longrightarrow \infty} \mathbb{P}\{\Theta_n > k\} = \infty,$$

so $\lim_{n \longrightarrow \infty} E(\Theta_n) = \infty$, which completes the second part of the proof when $\alpha Q \mathbb{1} = 1$.

We now consider the case where $\alpha Q \mathbb{1} < 1$. Using relations (11) and (12), we have

$$\begin{aligned}
&\mathbb{P}\{\Theta_n > k\} \\
&= \frac{k!}{n^k} \sum_{j=1}^{k \wedge n} \binom{n}{j} \sum_{\underline{m} \in T_{j,k}} \frac{j!}{m_1!(1!)^{m_1} m_2!(2!)^{m_2} \cdots m_k!(k!)^{m_k}} \prod_{\ell=1}^{k} \left((\alpha Q^\ell \mathbb{1})^{1/\ell} \right)^{\ell m_\ell} \\
&\leq \left(\max_{\ell=1,\ldots,k} (\alpha Q^\ell \mathbb{1})^{1/\ell} \right)^k \\
&\leq \left(\sup_{\ell \geq 1} (\alpha Q^\ell \mathbb{1})^{1/\ell} \right)^k .
\end{aligned}$$

If we show that $\sup_{\ell \geq 1} (\alpha Q^\ell \mathbb{1})^{1/\ell} < 1$ then the result follows using the dominated convergence theorem for series. Let us show that $\sup_{\ell \geq 1} (\alpha Q^\ell \mathbb{1})^{1/\ell} < 1$. First of all, since $\alpha Q \mathbb{1} < 1$, we have $\alpha Q^\ell \mathbb{1} < 1$ for any $\ell \geq 1$, hence

$$\left(\alpha Q^\ell \mathbb{1} \right)^{1/\ell} < 1 \text{ for every } \ell \geq 1. \tag{13}$$

We introduce the infinite norm for matrices, defined for every matrix M, not necessarily square, by

$$||M|| = \sup_i \sum_j |M_{i,j}|.$$

We then have $||\alpha|| = 1$, $||\mathbb{1}|| = 1$ and for every $\ell \geq 1$,

$$(\alpha Q^\ell \mathbb{1})^{1/\ell} \leq \left(||\alpha|| ||\mathbb{1}|| ||Q^\ell|| \right)^{1/\ell} = ||Q^\ell||^{1/\ell}. \tag{14}$$

It is well-know, see for instance [11], that the spectral radius $\rho(Q)$ of the matrix Q is satisfies

$$\rho(Q) = \lim_{\ell \longrightarrow \infty} ||Q^\ell||^{1/\ell} < 1. \tag{15}$$

From relations (14) and (15) we obtain

$$\limsup_{\ell \longrightarrow \infty} (\alpha Q^\ell \mathbb{1})^{1/\ell} < 1. \tag{16}$$

From relations (13) and (16), we easily get that $\sup_{\ell\geq 1}(\alpha Q^\ell \mathbb{1})^{1/\ell} < 1$. By the dominated convergence theorem for series, we obtain

$$\lim_{n\longrightarrow\infty} E(\Theta_n) = \sum_{k=0}^{\infty}(\alpha Q\mathbb{1})^k = \frac{1}{1-\alpha Q\mathbb{1}}.$$

This completes the proof. ■

In the sequel, we consider the last terms of relation (11). They can be easily computed. They also lead to a new way of computing the distribution of Θ_n for large values of n, small values of k and for any Markov chain X.

For every $1 \leq k \leq n$, relation (11) can be written as

$$\mathbb{P}\{\Theta_n > k\} = \frac{k!}{n^k}\sum_{j=1}^{k}\binom{n}{j}u_{j,k}(Q), \tag{17}$$

where

$$u_{j,k}(Q) = \sum_{\underline{m}\in T_{j,k}} \frac{j!}{m_1!(1!)^{m_1}m_2!(2!)^{m_2}\cdots m_k!(k!)^{m_k}}\prod_{\ell=1}^{k}\left(\alpha Q^\ell\mathbb{1}\right)^{m_\ell}.$$

It is easily checked that the j-th term in the sum is asymptotically equal to $O(1/n^{k-j})$, i.e.

$$\frac{k!}{n^k}\binom{n}{j}u_{j,k}(Q) = O(1/n^{k-j}).$$

Nevertheless, the constant involved in this expansion depends on k and is very large even for relatively small values of k. To see this, we explicitly compute below the last 5 terms of relation (17). By definition of the set $T_{j,k}$, we have for every $k \geq 1$,

$$\begin{aligned}
T_{k,k} &= \{(k,0,\ldots,0)\},\\
T_{k-1,k} &= \{(k-2,1,0,\ldots,0)\},\\
T_{k-2,k} &= \{(k-4,2,0,\ldots,0),(k-3,0,1,\ldots,0)\},\\
T_{k-3,k} &= \{(k-6,3,0,\ldots,0),(k-5,1,1,0,\ldots,0),(k-4,0,0,1,0\ldots,0)\},\\
T_{k-4,k} &= \{(k-8,4,0,\ldots,0),(k-7,2,1,0,\ldots,0),(k-6,1,0,1,0,\ldots,0),\\
&\quad (k-6,0,2,0,\ldots,0),(k-5,0,0,0,1,0\ldots,0)\},
\end{aligned}$$

where each vector having a negative entry is withdrawn from the corresponding set. We then have, for every $k \geq 1$,

$$\begin{aligned}
u_{k,k}(Q) &= (\alpha Q \mathbb{1})^k, \\
u_{k-1,k}(Q) &= \frac{k-1}{2} (\alpha Q \mathbb{1})^{k-2} \alpha Q^2 \mathbb{1} 1_{\{k\geq 2\}}, \\
u_{k-2,k}(Q) &= \frac{(k-2)(k-3)}{8} (\alpha Q \mathbb{1})^{k-4} \left(\alpha Q^2 \mathbb{1}\right)^2 1_{\{k\geq 4\}} \\
&\quad + \frac{k-2}{6} (\alpha Q \mathbb{1})^{k-3} \alpha Q^3 \mathbb{1} 1_{\{k\geq 3\}}, \\
u_{k-3,k}(Q) &= \frac{(k-3)(k-4)(k-5)}{48} (\alpha Q \mathbb{1})^{k-6} \left(\alpha Q^2 \mathbb{1}\right)^3 1_{\{k\geq 6\}} \\
&\quad + \frac{(k-3)(k-4)}{12} (\alpha Q \mathbb{1})^{k-5} \alpha Q^2 \mathbb{1} \alpha Q^3 \mathbb{1} 1_{\{k\geq 5\}} \\
&\quad + \frac{k-3}{24} (\alpha Q \mathbb{1})^{k-4} \alpha Q^4 \mathbb{1} 1_{\{k\geq 4\}}, \\
u_{k-4,k}(Q) &= \frac{(k-4)(k-5)(k-6)(k-7)}{384} (\alpha Q \mathbb{1})^{k-8} \left(\alpha Q^2 \mathbb{1}\right)^4 1_{\{k\geq 8\}} \\
&\quad + \frac{(k-4)(k-5)(k-6)}{48} (\alpha Q \mathbb{1})^{k-7} \left(\alpha Q^2 \mathbb{1}\right)^2 \alpha Q^3 \mathbb{1} 1_{\{k\geq 7\}} \\
&\quad + \frac{(k-4)(k-5)}{48} (\alpha Q \mathbb{1})^{k-6} \alpha Q^2 \mathbb{1} \alpha Q^4 \mathbb{1} 1_{\{k\geq 6\}} \\
&\quad + \frac{(k-4)(k-5)}{72} (\alpha Q \mathbb{1})^{k-6} \left(\alpha Q^3 \mathbb{1}\right)^2 1_{\{k\geq 6\}} \\
&\quad + \frac{k-4}{120} (\alpha Q \mathbb{1})^{k-5} \alpha Q^5 \mathbb{1} 1_{\{k\geq 5\}}.
\end{aligned}$$

Note that specifying the dependence on Q allows us to extract the sum of the coefficients of the $u_{j,k}(Q)$ by simply writing $u_{j,k}(I)$ since we have $\alpha Q^\ell \mathbb{1} = 1$ when $Q = I$. This leads to

$$\mathbb{P}\{\Theta_n > k\} = \frac{k!}{n^k} \sum_{j=k-4}^{k} \binom{n}{j} u_{j,k}(Q) + e_{n,k},$$

where $e_{n,k}$ satisfies using relation (12),

$$e_{n,k} = \frac{k!}{n^k} \sum_{j=1}^{k-5} \binom{n}{j} u_{j,k}(Q) \leq \frac{k!}{n^k} \sum_{j=1}^{k-5} \binom{n}{j} u_{j,k}(I) = 1 - \frac{k!}{n^k} \sum_{j=k-4}^{k} \binom{n}{j} u_{j,k}(I).$$

We denote by $b_{n,k}$ this last bound, i.e.

$$b_{n,k} = 1 - \frac{k!}{n^k} \sum_{j=k-4}^{k} \binom{n}{j} u_{j,k}(I).$$

We then have

$$0 \leq \mathbb{P}\{\Theta_n > k\} - \frac{k!}{n^k} \sum_{j=k-4}^{k} \binom{n}{j} u_{j,k}(Q) \leq b_{n,k}.$$

It can be easily shown that $b_{n,k}$ is increasing with k, so for a fixed error tolerance ε and for every n, we can easily compute the maximum value k^* of k such that the error is less than or equal to ε, i.e.

$$k^* = \sup\{k \in \mathbb{N} \mid b_{n,k} \leq \varepsilon\}.$$

Figure 1 gives some values of k^* for different values of ε and n. It shows for instance that, for every $k \leq 53$, we have

$$0 \leq \mathbb{P}\{\Theta_{1000} > k\} - \frac{k!}{1000^k} \sum_{j=k-4}^{k} \binom{1000}{j} u_{j,k}(Q) \leq 10^{-2}.$$

ε	10^{-2}	10^{-3}	10^{-4}
$n = 10^2$	18	15	12
$n = 10^3$	53	41	32
$n = 10^4$	162	124	97
$n = 10^5$	508	387	301

Fig. 1: Values of k^* for different values of ε and n.

The above upper bound might seem unnecessarily pessimistic, since all terms $\alpha Q^\ell \mathbb{1}$ have been upper bounded by 1 here, a crude upper-bound. For that reason, one may argue a bit further, and investigate how the above estimate might be improved.

The sequence $\alpha Q^\ell \mathbb{1}$ being decreasing with ℓ from 1 to 0, we denote by u the greatest integer such that $\alpha Q^u \mathbb{1} = 1$, i.e

$$u = \sup\{\ell \geq 0 \mid \alpha Q^\ell \mathbb{1} = 1\}.$$

Besides, due to the fact that $\rho(Q)$, the spectral radius of Q, satisfies $\rho(Q) < 1$, taking an ε such that $\rho(Q) < \rho(Q) + \epsilon < 1$, it is clear that there exists a constant $C_\varepsilon > 0$ such that for any $\ell \geq 0$, we have $\alpha Q^\ell \mathbb{1} \leq C_\varepsilon(\rho(Q) + \varepsilon)^\ell$. Therefore, gathering the results, it becomes now clear that, choosing a q such that $\rho(Q) < q < 1$, there is an integer ℓ_0 such that

$$\alpha Q^\ell \mathbb{1} \leq 1 \text{ whenever } 0 \leq \ell \leq \ell_0,$$
$$\alpha Q^\ell \mathbb{1} \leq q^\ell \text{ whenever } \ell \geq \ell_0 + 1.$$

Let us now estimate $\mathbb{P}\{\Theta_n > k\}$ using these (close to optimal) bounds. Using Corollary 3, we have for every $k \geq 0$ and $n \geq 1$,

$$\mathbb{P}\{\Theta_n > k\} = \frac{1}{n^k} \sum_{\underline{k} \in S_{k,n}} \frac{k!}{k_1! \cdots k_n!} \prod_{r=1}^{n} \alpha Q^{k_r} \mathbb{1} \leq A_{n,k},$$

where

$$\begin{aligned} A_{n,k} &= \frac{1}{n^k} \sum_{\underline{k} \in S_{k,n}} \frac{k!}{k_1! \cdots k_n!} \prod_{r=1}^{n} \left(1_{\{k_r \leq \ell_0\}} + q^{k_r} 1_{\{k_r \geq \ell_0+1\}}\right) \\ &= \frac{q^k}{n^k} \sum_{\underline{k} \in S_{k,n}} \frac{k!}{k_1! \cdots k_n!} \prod_{r=1}^{n} \left(\frac{1}{q^{k_r}} 1_{\{k_r \leq \ell_0\}} + 1_{\{k_r \geq \ell_0+1\}}\right). \end{aligned}$$

Let us now estimate $A_{n,k}$.

On the one hand, we have the following lower bound

$$A_{n,k} \geq \frac{q^k}{n^k} \sum_{\underline{k} \in S_{k,n}} \frac{k!}{k_1! \cdots k_n!} = q^k.$$

On the other hand, and in the similar spirit, we may as well write the following upper bound

$$\begin{aligned} A_{n,k} &\leq \frac{q^k}{n^k} \sum_{\underline{k} \in S_{k,n}} \frac{k!}{k_1! \cdots k_n!} \prod_{r=1}^{n} \frac{1}{q^{\ell_0}} \\ &= \frac{q^{k-n\ell_0}}{n^k} \sum_{\underline{k} \in S_{k,n}} \frac{k!}{k_1! \cdots k_n!} \\ &= q^{k-n\ell_0}. \end{aligned}$$

Thus, we have proved that $A_{n,k}$ anyhow satisfies the two bounds

$$q^k \leq A_{n,k} \leq q^{k-n\ell_0}.$$

For large values of n, i.e. when $n \longrightarrow \infty$, the upper bound is compatible with the information we otherwise have at hand, namely

$$\mathbb{P}\{\Theta_n > k\} \longrightarrow \alpha Q^k \mathbb{1}.$$

Needless to say, the obtained bound on $A_{n,k}$ in this case does not improve on our previous crude estimate. For large values of k at variance, we have proved that

$$\mathbb{P}\{\Theta_n > k\} \leq q^{k-n\ell_0},$$

which establishes that $\mathbb{P}\{\Theta_n > k\}$ goes exponentially fast to zero as $k \longrightarrow \infty$. This constitutes a new piece of information.

4.2 Speeding-up the computation

The computation of the distribution and of the expectation of Θ_n using Theorem 4 is useful if these quantities are needed for all the values of i from 1 to n. This leads to an algorithm with a complexity in $O(n)$ which makes the computation prohibitive for very large values of n. In the following theorem, we generalize the result of Theorem 4 to deal with very large values of n.

Theorem 13 *For every $k \geq 0$, $n \geq 2$, we have*

$$\mathbb{P}\{\Theta_n > k\} = \sum_{\ell=0}^{k} \binom{k}{\ell} s_{u,n}^{\ell} (1 - s_{u,n})^{k-\ell} \mathbb{P}\{\Theta_u > \ell\} \mathbb{P}\{\Theta'_{n-u} > k - \ell\} \qquad (18)$$

where the random variable Θ'_{n-u} is defined exactly as Θ_{n-u} but with the probability mass distribution

$$\left(\frac{p_{u+1,n}}{1 - s_{u,n}}, \ldots, \frac{p_{n,n}}{1 - s_{u,n}} \right) \text{ and } s_{u,n} = \sum_{r=1}^{u} p_{r,n}.$$

Proof. Let $n \geq 2$ and u such that $1 \leq u \leq n-1$. For $r = 1, \ldots, u$, we have $p_{r,u} = p_{r,n}/s_{u,n}$. For every $k \geq 0$, we have from Corollary 3

$$
\begin{aligned}
&\mathbb{P}\{\Theta_n > k\} \\
&= \sum_{\underline{k} \in S_{k,n}} \frac{k!}{k_1! \cdots k_n!} \prod_{r=1}^{n} (p_{r,n})^{k_r} \, \alpha Q^{k_r} \mathbb{1} \\
&= \sum_{\ell=0}^{k} \sum_{\underline{k} \in S_{\ell,u}} \sum_{(k_{u+1},\ldots,k_n) \in S_{k-\ell,n-u}} \frac{k!}{k_1! \cdots k_n!} \prod_{r=1}^{n} (p_{r,n})^{k_r} \, \alpha Q^{k_r} \mathbb{1} \\
&= \sum_{\ell=0}^{k} \binom{k}{\ell} \sum_{\underline{k} \in S_{\ell,u}} \frac{\ell!}{k_1! \cdots k_u!} \prod_{r=1}^{u} (p_{r,n})^{k_r} \, \alpha Q^{k_r} \mathbb{1} \\
&\quad \times \sum_{(k_{u+1},\ldots,k_n) \in S_{k-\ell,n-u}} \frac{(k-\ell)!}{k_{u+1}! \cdots k_n!} \prod_{r=u+1}^{n} (p_{r,n})^{k_r} \, \alpha Q^{k_r} \mathbb{1} \\
&= \sum_{\ell=0}^{k} \binom{k}{\ell} s_{u,n}^{\ell} (1 - s_{u,n})^{k-\ell} \sum_{\underline{k} \in S_{\ell,u}} \frac{\ell!}{k_1! \cdots k_u!} \prod_{r=1}^{u} (p_{r,u})^{k_r} \, \alpha Q^{k_r} \mathbb{1} \\
&\quad \times \sum_{(k_{u+1},\ldots,k_n) \in S_{k-\ell,n-u}} \frac{(k-\ell)!}{k_{u+1}! \cdots k_n!} \prod_{r=u+1}^{n} \left(\frac{p_{r,n}}{1 - s_{u,n}} \right)^{k_r} \alpha Q^{k_r} \mathbb{1} \\
&= \sum_{\ell=0}^{k} \binom{k}{\ell} s_{u,n}^{\ell} (1 - s_{u,n})^{k-\ell} \mathbb{P}\{\Theta_u > \ell\} \mathbb{P}\{\Theta'_{n-u} > k - \ell\}
\end{aligned}
$$

which completes the proof. ■

Corollary 14 *If $p_{r,n} = 1/n$ for every $r = 1, \ldots, n$, then for every $k \geq 0$ and u such that $1 \leq u \leq n-1$, we have*

$$
\mathbb{P}\{\Theta_n > k\} = \sum_{\ell=0}^{k} \binom{k}{\ell} \left(\frac{u}{n}\right)^{\ell} \left(1 - \frac{u}{n}\right)^{k-\ell} \mathbb{P}\{\Theta_u > \ell\} \mathbb{P}\{\Theta_{n-u} > k - \ell\}
$$

and, for every $m \geq 1$, we have

$$
\mathbb{P}\{\Theta_{2^m} > k\} = \frac{1}{2^k} \sum_{\ell=0}^{k} \binom{k}{\ell} \mathbb{P}\{\Theta_{2^{m-1}} > \ell\} \mathbb{P}\{\Theta_{2^{m-1}} > k - \ell\}. \tag{19}
$$

Proof. If the probability mass distribution is uniform then $s_{u,n} = u/n$ and $\Theta'_{n-u} = \Theta_{n-u}$. The second relation is simply obtained by taking $n = 2^m$ and $u = 2^{m-1}$. ■

Relation (19) is interesting for very large values of n. Indeed, the complexity for the computation of the distribution and of the expectation of Θ_n is now $O(\log_2 n)$.

Note also that this relation can be split into 2 identical sums plus the central term corresponding to $\ell = k/2$ when k is even.

5 Application to Large-Scale Distributed Systems

Large-scale distributed systems are systems populated by thousand or millions of heterogeneous nodes. These systems keep evolving according to nodes failures, departures, or the introduction of new nodes. To handle the scale and dynamics of these systems, peer-to-peer overlays, commonly abbreviated as P2P overlays, have been proposed. P2P overlays are virtual networks built on top of a physical network. Nodes of the overlay, usually called peers, communicate among each other along the edges of the overlay by using the communication primitives provided by the underlying network (e.g., IP network service). The algorithms that peers use to choose their neighbors and to route messages define the overlay topology. The topology of unstructured overlays conforms random graphs, i.e., relationships among peers are mostly set according to a random process, and routing is not constrained. Most popular P2P systems (e.g., Gnutella, kaZaA) rely on unstructured overlays. On the other hand, structured overlays (also called Distributed Hash Tables (DHTs)) are such that peers self-organize in structured graphs (e.g., hypercube, torus). Typically, peers self-organize according to a distance function based on their identifier, allowing to partition the identifier space among all the peers of the system. Most proposed structured overlays (e.g. [6, 9, 10, 13, 16, 15]) are highly satisfactory in terms of efficiency and scalability when evolving in weakly dynamic environments (i.e., their key-based routing interface guarantees operations whose complexity in messages and latency usually scale logarithmically with system size). However, in the presence of very frequent peers connections and/or disconnections, a very large number of join and leave operations are locally triggered generating accordingly multiple and concurrent maintenance traffic. Ensuring routing tables consistency quickly becomes unbearable, leading to misrouting, and to possible partitioning of the system as argued by Locher et al. [8].

A different approach to build a churn-resilient overlay network has been proposed with cluster-based structured overlays (e.g. [1, 2, 4, 8]). In these overlays, peers which are close to each other according to a given proximity metric group together into clusters. These clusters form the vertices of the structured topology. Clusters size is lower and upper bounded. The lower bound is usually tuned to satisfy some constraint based on the assumed failure model (for instance, a lower bound greater than or equal to $3f + 1$ allows Byzantine tolerant agreement protocols to be run despite up to f malicious peers). The upper bound is typically in $\mathcal{O}(logN)$ to meet scalability requirements, where N is the current number of peers in the system. Most of the maintenance traffic is captured by clusters which severely limits both topology changes and communication overhead imposed by the above solutions. On the other hand, to keep clusters size bounded, these clusters have to `split` into smaller clusters or `merge` with their closest clusters each time they become over populated or under populated. Handling these operations requires high synchronization among involved peers. Clearly this may tend to overwhelm the benefit of cluster-based DHTs if both operations are too frequent with respect to join and leave events. The remaining part of this section is devoted to this analysis.

The Markov chain X modelling the behavior of one cluster is depicted in Figure 1 in which $q = 1 - p$ and $p \in (0, 1)$.

The transition probability p means that a new peer joins the cluster while the transition probability q means that a peer leaves the cluster. The transition

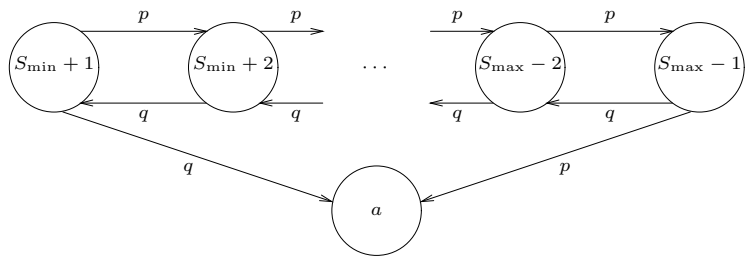

Figure 1: Markov chain model of one cluster.

from state $S_{\min}+1$ to the absorbing state expresses that the cluster has reached its minimal size $S_{\min}$ and that it has to merge with another cluster. In the same way the transition from state $S_{\max}-1$ to the absorbing state means that the cluster has reached its maximal size $S_{\max}$ and that it must split into two smaller clusters. The initial distributions α that we consider are the unit row vectors e_j for $j = S_{\min}+1, \ldots, S_{\max}-1$. So, the initial distribution $\alpha = e_j$ means that $X_0 = j$ with probability 1.

The matrix Q which gives the transitions between the transient states of X is thus a tri-diagonal matrix where non-zero entries are $Q_{i,i+1} = p$ and $Q_{i,i-1} = q = 1-p$. We suppose that the probability mass function $\pi(n)$ is uniform, i.e. that $p_{i,n} = 1/n$, for every $i = 1, \ldots, n$. With these values, the limiting behavior of respectively the distribution and the expectation of Θ_n are given from Theorem 12 for every $k \geq 0$, by

$$\lim_{n\longrightarrow\infty} \mathbb{P}\{\Theta_n > k\} = \begin{cases} p^k & \text{if} \quad X_0 = S_{\min}+1 \\ 1 & \text{if} \quad X_0 = j, \text{ for } S_{\min}+2 \leq j \leq S_{\max}-2 \\ (1-p)^k & \text{if} \quad X_0 = S_{\max}-1 \end{cases}$$

and

$$\lim_{n\longrightarrow\infty} E(\Theta_n) = \begin{cases} \dfrac{1}{1-p} & \text{if} \quad X_0 = S_{\min}+1 \\ \infty & \text{if} \quad X_0 = j, \text{ for } S_{\min}+2 \leq j \leq S_{\max}-2 \\ \dfrac{1}{p} & \text{if} \quad X_0 = S_{\max}-1. \end{cases}$$

For the numerical evaluations, we have chosen $p = 1/2$. With this value, we easily get, when $\alpha = e_j$,

$$E(\Theta_1) = (j - S_{\min})(S_{\max} - j).$$

We have also chosen $S_{\min} = 4$ and $S_{\max} = 16$ which implies that the number of transient states is equal to 11.

Figure 2 shows the distribution of Θ_n for different values of n when the initial state X_0 of each cluster is state $S_{\min}+1 = 5$. It also shows the limiting distribution given by

$$\lim_{n\longrightarrow\infty} \mathbb{P}\{\Theta_n > k\} = \frac{1}{2^k}.$$

It is worthwhile noting that the limiting distribution is reached very quickly for small values of n.

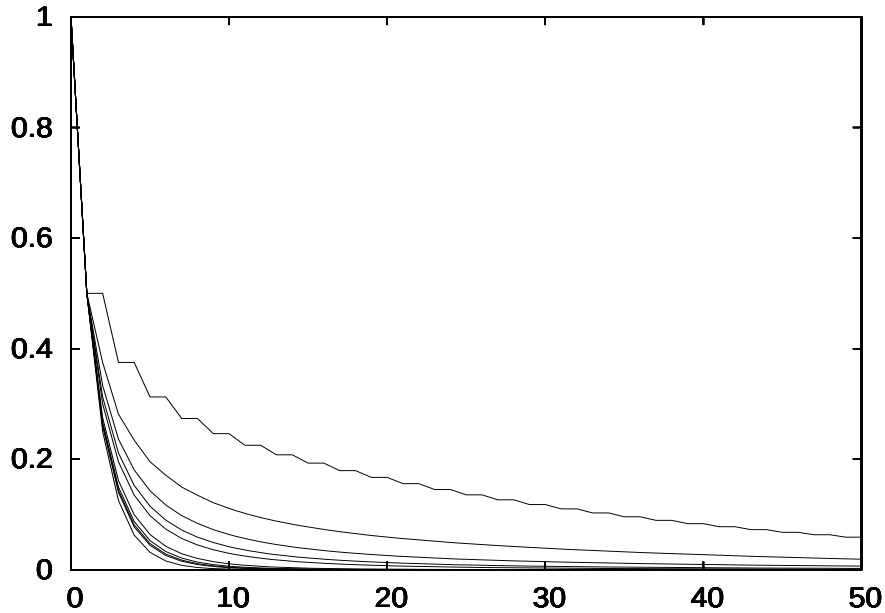

Figure 2: From top to the bottom: $\mathbb{P}\{\Theta_n > k\}$ when $X_0 = S_{\min} + 1 = 5$, for $n = 1, 2, 3, 4, 5, 10, 15, 20, 25, \infty$, as functions of k.

In the same way, Figure 3 shows the distribution of Θ_n for different values of n when the initial state X_0 is state $(S_{\min} + S_{\max})/2 = 10$. In that case the limiting distribution is given by

$$\lim_{n \longrightarrow \infty} \mathbb{P}\{\Theta_n > k\} = 1.$$

The figure shows that it is reached rather slowly. Note the scaling difference between the k-axis in Figures 2 and 3.

Figure 4 shows the expected value of Θ_n for different values of the initial state and for n from 1 to 25. We can observe the limiting behavior of $E(\Theta_n)$ which is equal to 2 when $X_0 = 5$ and equal to ∞ when $X_0 = 6, 7, 8, 9, 10$. Here again it is reached very quickly when $X_0 = 5$.

Figure 5 shows the distribution of Θ_n for different large values of n when the initial state X_0 of each cluster is state $(S_{\min} + S_{\max})/2 = 10$. This figure has been obtained using the recurrence described in Corollary 14.

Figure 6 shows the expected value of Θ_n for different values of the initial state and for $n = 1, 2, 4, 8, \ldots, 2^{15} = 32768$. Here again, we can observe the limiting behavior of $E(\Theta_n)$ which is equal to ∞ when $X_0 = 5, 6, 7, 8, 9, 10$. The value of $E(\Theta_n)$ when $X_0 = 5$ is practically invisible because it is too close to the x-axis. Indeed, in that case limiting behavior of $E(\Theta_n)$ is equal to 2.

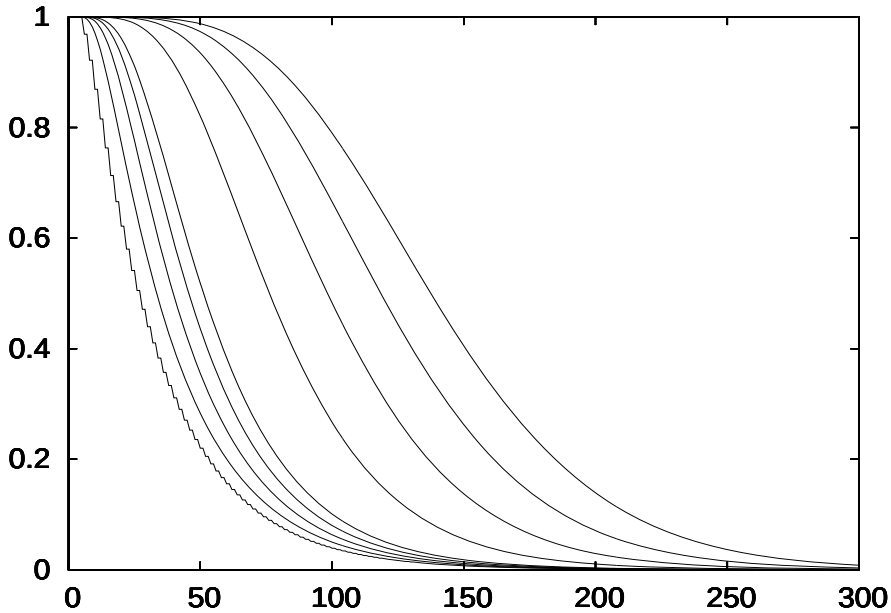

Figure 3: From bottom to the top: $\mathbb{P}\{\Theta_n > k\}$ when $X_0 = (S_{\max} + S_{\min})/2 = 10$, for $n = 1, 2, 3, 4, 5, 10, 15, 20, 25$, as functions of k.

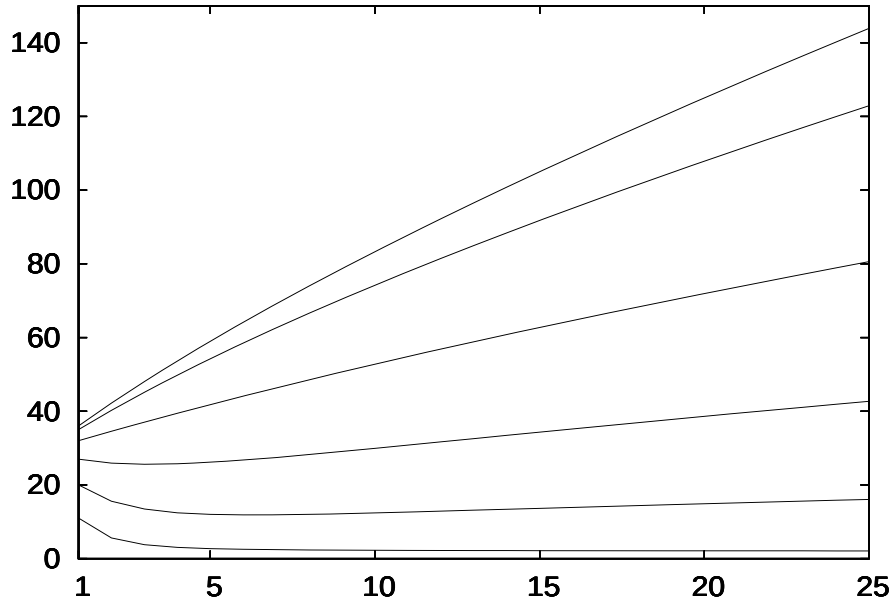

Figure 4: From bottom to the top: $E(\Theta_n)$ for $X_0 = 5, 6, 7, 8, 9, 10$, as functions of n.

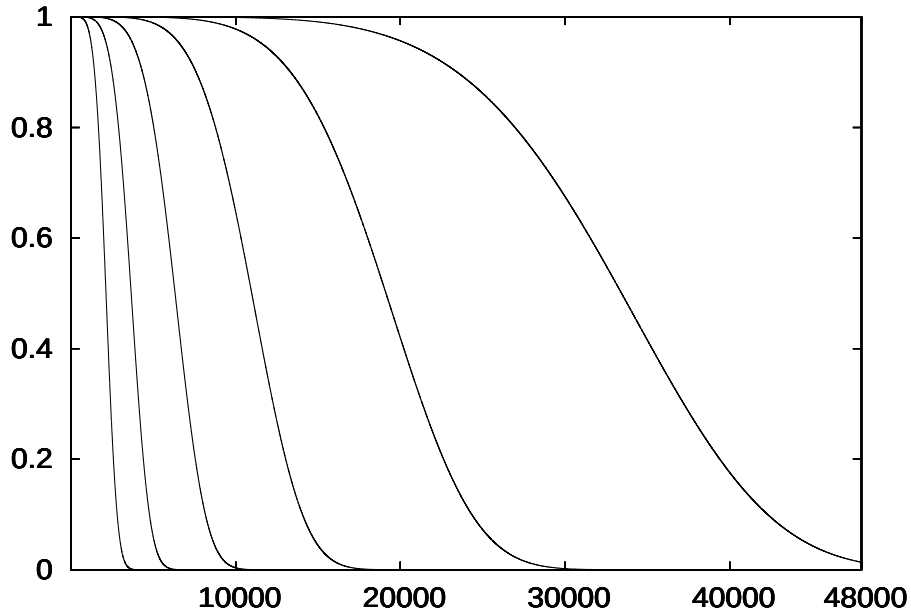

Figure 5: From bottom to the top: $\mathbb{P}\{\Theta_n > k\}$ when $X_0 = (S_{\max} + S_{\min})/2 = 10$, for $n = 2^{10}, 2^{11}, 2^{12}, 2^{13}, 2^{14}, 2^{15} = 32768$, as functions of k.

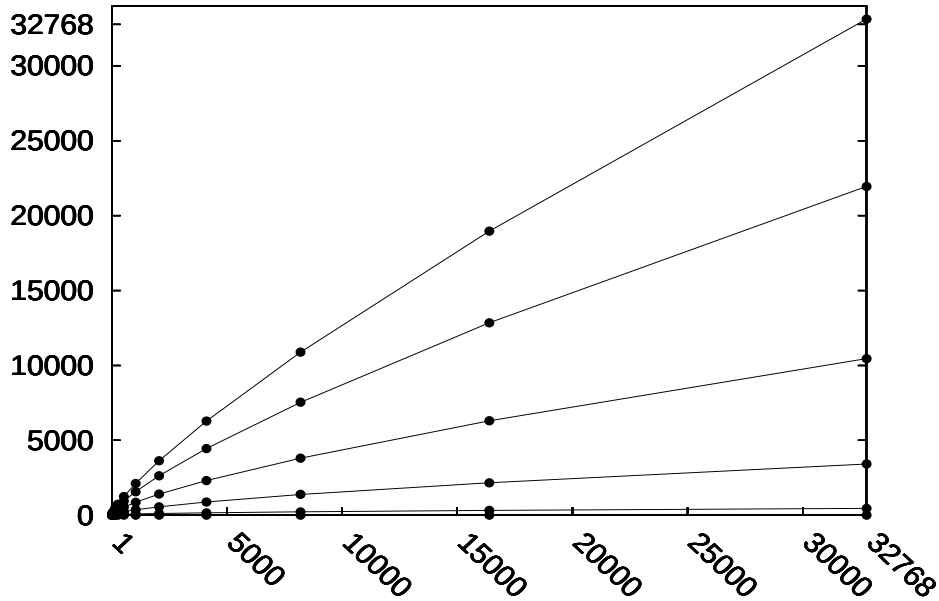

Figure 6: From bottom to the top: $E(\Theta_n)$ for $X_0 = 5, 6, 7, 8, 9, 10$, as functions of n, for $n = 1, 2, 4, 8, \ldots, 2^{15}$.

References

[1] Anceaume, E., Brasiliero, F., Ludinard, R. and Ravoaja, A.(2008). Peercube: an hypercube-based p2p overlay robust against collusion and churn. In *IEEE Proceedings of the International Conference on Self Autonomous and Self Organizing Systems (SASO).*

[2] Awerbuch, B. and Scheideler, C.(2004). Group spreading: A protocol for provably secure distributed name service. In *Procs of the International Colloquium on Automata, Languages and Programming (ICALP).*

[3] Beaugendre, P. (2005). Un opérateur d'extension linéaire explicite. Revue de la Filière Mathématique 116(1).

[4] Fiat, A., Saia, J. and Young, M.(2005). Making chord robust to byzantine attacks. In *Proceedings of the Annual European Symposium on Algorithms (ESA).*

[5] Fourneau, J.-M. (2008). Discrete time Markov chains competing over resources: product form steady-state distribution. In *Proceedings of the 5th International Conference on the Quantitative Evaluation of Systems (QEST'08)*, Saint-Malo, France, 147-156.

[6] Harvey, N, Jones, M. B., Saroui, S., Theimer, M. and Wolman, A.(2003). Skipnet: A scalable overlay network with practical locality properties. In *Proceedings of the 4th USENIX Symposium on Internet Technologies and Systems (USITS).*

[7] Karlin, S. and Taylor, H. M. (1981). A Second Course in Stochastic Processes. Academic Press.

[8] Locher, T., Schmid, S. and Wattenhofer, R.(2006). eQuus: A provably robust and locality-aware peer-to-peer system. In *Proceedings of the International Conference on Peer-to-Peer Computing (P2P).*

[9] Malkhi, D., Naor, M. and Ratajzcak, D.(2003). Viceroy: Scalable emulation of butterfly networks for distributed hash tables. In *Proceedings of the Annual Symposium on Principles of distributed computing (PODC).*

[10] Manku, G. S., Bawa, M. and Raghavan, P.(2003). Symphony: Distributed hashing in a small world. In *Proceedings of the 4th USENIX Symposium on Internet Technologies and Systems (USITS).*

[11] Meyer, C. D. (2000). Matrix Analysis and Applied Linear Algebra. SIAM.

[12] Neuts, M. F. (1981). Matrix-Geometric Solutions in Stochastic Models: An Algorithmic Approach. Johns Hopkins University Press.

[13] Ratnasamy, S., Francis, P., Handley, M., Karp, R. and Shenker, S.(2001). A scalable content-addressable network. In *Proceedings of the ACM SIGCOMM.*

[14] Riordan, J. (1967). An Introduction to Combinatorial Analysis. Wiley.

[15] Rowstron, A. and Druschel, P.(2001). Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In *Proceedings of the International Conference on Distributed Systems Platforms (Middleware).*

[16] Stoica, I., Liben-Nowell, D., Morris, R., Karger, D., Dabek, F., Kaashoek, M. F. and Balakrishnan, H.(2001). Chord: A scalable peer-to-peer lookup service for internet applications. In *Proceedings of the ACM SIGCOMM.*

Centre de recherche INRIA Rennes – Bretagne Atlantique
IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Grenoble – Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Nancy – Grand Est : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex
Centre de recherche INRIA Paris – Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex
Centre de recherche INRIA Saclay – Île-de-France : Parc Orsay Université - ZAC des Vignes : 4, rue Jacques Monod - 91893 Orsay Cedex
Centre de recherche INRIA Sophia Antipolis – Méditerranée : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
http://www.inria.fr
ISSN 0249-6399