



HAL
open science

The CENTAURS race. Performance evaluation of an election algorithm on a ring

Alain Jean-Marie, François Baccelli

► **To cite this version:**

Alain Jean-Marie, François Baccelli. The CENTAURS race. Performance evaluation of an election algorithm on a ring. [Research Report] RR-0926, INRIA. 1988. inria-00075630

HAL Id: inria-00075630

<https://inria.hal.science/inria-00075630>

Submitted on 24 May 2006

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.

INRIA

UNITÉ DE RECHERCHE
INRIA-SOPHIA ANTIPOLIS

Institut National
de Recherche
en Informatique
et en Automatique

Domaine de Voluceau
Rocquencourt
BP 105
78153 Le Chesnay Cedex
France
Tél. (1) 39 63 55 11

Rapports de Recherche

N° 926

Programme 3

THE CENTAURS RACE **Performance Evaluation of an Election** **Algorithm on a Ring**

Alain JEAN-MARIE
François BACCELLI

Novembre 1988



★ R R . 8 9 2 6 ★

THE CENTAURS RACE

Performance Evaluation of an Election Algorithm on a Ring

Alain JEAN-MARIE
François BACCELLI

INRIA-Sophia Antipolis
2004 Route des Lucioles
06565 VALBONNE Cedex, FRANCE

Draft, August 31, 1988

ABSTRACT

We evaluate several variations of a simple election algorithm on a ring structured network, both in discrete and continuous time. In each case, both the one- and the bi-directional cases are studied and compared. Using a recursive decomposition approach, we obtain various analytical representations for the distributions of the number of messages that are exchanged and derive their asymptotic behavior when the ring size grows large. In particular, it is shown that the expected number of messages is asymptotically quadratic. The evaluation of the completion time of the algorithm is also addressed. This time decomposes into a *startup time*, an *exploration time*. The *startup time* is also analyzed in detail and it is shown that its mean is of order $O(\log n)$. The analysis of the *exploration time* remains an open problem.

KEYWORDS Virtual Ring, Election Algorithms, Performance Evaluation, Complex Analysis Methods.

LA COURSE DES CENTAURES

Evaluation des performances d'un algorithme d'élection sur un anneau

Alain JEAN-MARIE
François BACCELLI

INRIA-Sophia Antipolis
2004 Route des Lucioles
06565 VALBONNE Cedex, FRANCE

RESUME

Nous évaluons les performances de plusieurs variantes d'un algorithme distribué d'élection sur un réseau en forme d'anneau, connu sous le nom d'"Algorithme des Centaures". Deux modèles probabilistes sont proposés, l'un en temps discret, l'autre en temps continu. Nous montrons que les performances des variantes uni- et bi-directionnelles de l'algorithme peuvent s'analyser au moyen d'équations de partition récursive. Nous obtenons en particulier des expressions mathématiques pour les distributions du nombre des messages échangés et du temps de complétion, ce dont nous déduisons le comportement asymptotique de l'algorithme lorsque la taille du réseau devient grande. En particulier, nous établissons que la valeur moyenne du nombre des messages est asymptotiquement quadratique. Pour le temps de complétion, nous proposons une décomposition en deux phases, la phase de démarrage et la phase d'exploration. La phase de démarrage est analysée en détail et nous montrons que sa durée moyenne est de l'ordre de $O(\log(n))$. L'analyse du temps d'exploration reste un problème ouvert.

Mots Clés Anneau virtuel, algorithmes d'élection distribuée, évaluation des performances, analyse complexe.

0/FOREWORD

Somewhere on the highlands of the Mount Olympus stands a Ring of high Towers connected by narrow paths. There live the Centaurs. At the beginning of Time, there was a Centaur egg in each Tower. It is written that when one of these eggs will hatch, the new born Centaur will restlessly wander from Tower to Tower, for the Centaurs are doomed to run and not to remember the paths they already followed. And it is said that no Centaur will ever wander counter-clockwise, lest he might be bound to Hades by the offended Gods. Nor will two Centaurs ever run abreast along the same path, for these are too narrow. And each time a Centaur will find a Centaur egg, he will destroy it without mercy, and this will be the fate of many eggs. At last, every Centaur will return back home and die in the very Tower where he was born. And it is said that the Day when the Last of the Centaurs will die out will be the day of the Complete Knowledge of the World, and therefore Its End.

1/INTRODUCTION

This paper deals with the performance analysis of an election algorithm on a ring. Elections in distributed systems occur typically when there is a need for a coordinator, for instance for a failure recovery, a result gathering or a commitment in a distributed database. Many algorithms have been proposed for general networks (for instance [LEL]). The particular topology of the ring allows dramatic simplifications and improvements of distributed algorithms, so that this topology is often used (virtual rings). Several election algorithms have been described in the literature, presenting successive improvements for the message-complexity ([CHA],[KAI]).

The election algorithm we shall study has been proposed by C. Kaiser ([KAI]), and is described in a somewhat more technical way in Section 2. It is suited for a ring structured asynchronous network with communication by messages under the assumption that messages are neither lost nor altered. All the sites and the links are failure free and their identities are distinct. Sites may decide independently to initiate an election. The algorithm is described for one and bi-directional rings.

There are two performance measures of interest: the number of messages exchanged between two consecutive sites, which will be called "letters" in the sequel, and which should not to be mistaken with the notion of "election-messages" or "completion-messages", to be defined in section 2/; and the completion time. The first performance measure is important when the election algorithm is superimposed with another distributed computation that it may disturb through the traffic it generates. The second one is more important for real-time applications. Also, the ratio of these two quantities expresses the *throughput* of messages during the election, which is an important factor, for instance if the communication medium happens to be a bus. Observe that the performances of distributed algorithms are more often analyzed in terms of traffic (in number or size of letters), than in terms of time.

We introduce two stochastic models for this algorithm, one based upon a continuous time representation and the second on a discrete time one. Following C. Kaiser, we formulate these models in terms of a "Centaurs Race". It turns out that both the traffic and the time performance measures only depend on two state random variables (RV's) associated with these models: The number of Centaurs born during the race, and the date when the "last Centaur egg is destroyed".

These two models are respectively studied in section 3/ and 4/. For each of them we evaluate the traffic and the completion time measures. The cases of a one and a bi-directional ring are successively presented, and their relative merits are compared.

For the continuous time model, we give an expression for the distribution of the number of Centaurs ever born on the ring, and we show that, when properly normalized, this distribution converges to a Dirac mass as the size of the problem grows large. From this, it is established that the number of Centaurs born on the ring during one execution of the algorithm is asymptotically linear in the size of the ring, which entails that the number of letters exchanged is asymptotically quadratic.

This asymptotic result still holds in the discrete time case, as far as mean values are concerned. In addition to that, we obtain an expression for the distribution function of the date where the last egg is destroyed (which is closely related to the completion time). We also discuss its asymptotics, and show that it is of order $O(\log(n))$.

The study of these two models involves various methods based on complex variable analysis, including the resolution of both differential and functional equations and the characterisation of the asymptotic behavior of their solutions via classical singularity analysis.

2/THE ALGORITHM AND ITS MODELS

Description of the algorithm in the one direction case

The election algorithm is designed for a ring shaped asynchronous network with say n sites, $n = 1, 2, \dots$. Each site has a distinct identity, which is supposed to be constant during the whole process. The aim of the algorithm is to elect a leader site among the n sites. For instance, assuming without loss of generality that the identities can be ordered, this leader will be the site with the largest or the smallest identity. Before the execution of the algorithm, the sites do neither know the identity of the leader, nor the size of the ring. Any site may decide to launch an election at any moment. If so, this site becomes an "initiator" and sends an *election-message* to its neighbor. These will be represented by *Centaurs*, racing around a Ring of Towers. Initially, all sites are potentially initiators so that each Tower contains a "Centaur.egg".

When such an election-message reaches a new site, the identity of this new site together with the information carried by the message are used to compute the identity of a "partial leader" (for instance the site with the largest or smallest identity among those already visited by the election message since its creation) and this information is then transmitted to the next site on the ring. If the new site was not a past initiator, it learns through the arrival of the election message that an election is going on, and decides not to initiate any further election. This principle translate into the fact that a Centaur destroys the egg he might find when arriving at a new Tower. When an election-message comes back to its initiator, this one knows the identity of the global leader and sends a *completion-message*. When a site receives a completion-message, it forwards it and stops listening to messages. The election is over when everybody knows the name of the leader and no letter is between two sites.

We will not present any proof of this algorithm; it is easy to convince oneself that it eventually terminates, and that the elected site is unique.

Description of the algorithm in the bi-directional case

The principle is the same, except that an initiator sends a pair of election-messages, one clockwise and one counter-clockwise. Another site (at most two) will receive these two messages and learn the identity of the leader. This site will send a pair of completion-messages that will be forwarded until they arrive in a site that already knows the leader.

A preliminary analysis of the algorithm in the one direction case shows that if the network is overtake free, all the election-messages will accomplish a complete rotation around the ring. As every site sends one unique completion-message, the total number of letters sent will be $n \times (p + 1)$, where p denotes the total number of initiators.

When the ring is bi-directional, the number of letters may go up to $p \times (n + 2) + 2 \times n$, because two letters of the same election-message may cross each other on a link. If completion messages are emitted in one direction only, this bound reduces to $p \times (n + 2) + n$; if the communication protocol does not allow letters to cross each other, the number of letters will still be $n \times (p + 1)$.

The time analysis is less immediate, as many events can occur simultaneously. It turns out that the direct analysis of the completion time of the algorithm *via* stochastic models is a very difficult problem.

However, if one assumes that all letters can be delivered within a fixed interval of time $[\delta, \Delta]$, then it is plain that the duration of the election, computed from the instant of the first initiation, will not be less than $n\delta$ (an election-message at least has to visit all sites) but will not exceed $2n\Delta$ units.

A more careful analysis shows that, still assuming some bounds on the letter transit time, the completion time is bounded above by: $(n + 1)\Delta + L$, where L denotes the time at which the *last egg*

is destroyed, either by hatching or by the arrival of a live Centaur. To show this, notice that at date L , all sites are aware that an election is in progress (but none of them knows who is the leader). This means that all the sites have seen at least one Centaur passing. After another $n\Delta$ units of time, all the Centaurs will have completed their turn and have generated a "completion Centaur" that will be located at least where the corresponding Centaur was at time L . Therefore, at time $n\Delta + L$, all the sites know the identity of the leader, and at time $(n + 1)\Delta + L$, the eventually remaining completion-Centaurs will have stopped.

Therefore, it turns out that the time and message complexity of the algorithm depends upon two key quantities, respectively the time L at which the last egg disappears, and the number p of Centaurs ever born in the ring. Their evaluation is the purpose of the paper.

To analyze the problem, two stochastic (markovian) models arise naturally, depending on whether the representation of time is chosen to be continuous or discrete:

In continuous time, the eggs hatching time are supposed to be mutually independent and exponentially distributed RV's with mean $1/\nu$. The sojourn times of the Centaurs in Towers are also assumed to be mutually independent exponentially distributed RV's with mean $1/\tau$. Furthermore, the two sequences of hatching and sojourn times are supposed to be mutually independent. Similar independence assumptions will be adopted in the discrete time case, with the exponential distributions replaced by Bernoulli distributions with parameters that will also be denoted by ν and τ for the hatching and moving probabilities respectively.

In both cases, the brute force approach would consist in taking as state variables the set of all possible distributions of eggs and Centaurs on the ring. This setting which would lead to the analysis of the *transient* evolution of a Markov chain with several absorbing states (or state subsets), is mostly unpractical, even with n small and will not be pursued here. Instead, we propose to use the non-overtaking assumption to recursively decompose a problem of size n in several problems of smaller size. This recursive partitioning approach is proved to be quite fruitful for deriving asymptotic results.

The continuous time model is more suited to capture the behavior of a totally asynchronous system, where the letters transmission times are not *a priori* bounded. It also has the advantage of providing a simple solution to the message-complexity problem. However, in this model, transition delays are not bounded, so that the above preliminary analysis does not apply directly. It is tantalizing to apply the same reasoning "in mean", with Δ being the mean transition delay. Unfortunately, this is incorrect, since this evaluation involves maxima of RV's, the expected values of which cannot be derived from the mean of the message delays alone. In this case, the time behavior of the algorithm after time L (the startup time) remains an open problem as to the writing of this paper.

On the other hand, the discrete time model has the advantage of taking in account the "unit of time assumption" made above (actually the deterministic transitions case); it also brings some light on the more difficult time-complexity problem.

3/THE CONTINUOUS TIME PROBLEM

3.1/The One-Direction Case

Initially (at time $t = 0$), all the Towers of the ring are assumed to contain an egg. The eggs hatching times are independent and identically distributed exponential RV's with mean $1/\nu$. A Centaur stays in a site for an exponential sojourn time with mean $1/\tau$ and then instantaneously moves to the next site, provided that this one contains no Centaur. Equivalently, one might consider that it takes an exponentially distributed communication time with parameter τ for a Centaur to travel between two Towers, and that only one Centaur can be on the path between two Towers at a given time. When there is another Centaur in (resp. on the way to) the next site, it is assumed that the Centaur waits until this Centaur leaves (resp. reaches) the next site, and then starts his own exponential delay. This enforces the non-overtaking property, and also provides a stronger "non-interacting" property: when a Centaur happens to catch another one up, the stochastic behavior of the latter one is not affected.

Observe that, as far as the number of births is concerned, the problem reduces to the analysis of a line of Towers with a Centaur in the first one and Centaur eggs in the others. This follows from the fact that in the continuous time model, simultaneous events occur with zero probability, so that Centaurs will hatch one at a time with probability one. The number of births on the ring is then simply equal to the number of births on the line (counting the one already born). Similarly, the time to the last egg destruction on the ring will be obtained from the corresponding time on the line by adding the hatching time of the first born Centaur.

The key point of the method is that, due to the assumptions made above, a Centaur's behavior is neither influenced by what happens *behind* him nor by what happens *before* him as long as he encounters cells containing an egg. Similarly, when a Centaur encounters either an empty cell or another Centaur, his future behavior will neither influence the number of births nor the last egg's destruction time, because he will never encounter an egg again.

Thus, the system can be viewed as a collection of independent Centaurs, each at the head of a row of cells containing an egg. If an egg hatches on this row, this simply splits the row in two sub-rows, thus decomposing the problem in problems of lesser size.

Note that this property remains true for other behaviors of the Centaur, provided that the non-overtaking and non-interacting property is respected. For instance, one may consider that Centaurs that find an empty cell start running faster. Also, as proposed in [KAI], the Centaurs may die (stop definitively) when they arrive in an empty cell with a higher identity than their own, except if they have already been traveling for a certain fixed amount of time. Similarly, when two Centaurs are in the same cell, they could fight so that only the one with the highest identity, or the oldest one, etc... would continue. These improvements would not modify the parameters p and L , but would of course influence the traffic on the ring and therefore the final completion time and the total number of letters exchanged.

3.1.1/The Number of Centaurs

For $n > 0$, let N_n be the number of births in an array of n sites with a Centaur in site 1 and an egg in each of the sites 2 through n . If D represents the event: "the Centaur moves to site 2 before any birth takes place in sites 2, 3, ..., n ", the RV's $N_n, n = 2, 3, \dots$ satisfy the recurrence relation:

$$N_n = N_{n-1} \mathbf{1}_{\{D\}} + \sum_{l=2}^n (N_{l-1} + N_{n-l+1}) \mathbf{1}_{\{\neg D \text{ and first birth in site } l\}}, \quad (3.1)$$

with $N_1 = 1$. This equation is a consequence of the memoryless property of exponential distributions, together with the non-interacting principle that allows to "forget" the Centaurs that are not in front of a non empty row of eggs, as if they had stopped at the first eggless site they found on their way.

Taking the generating function in the complex parameter z , $|z| < 1$ in (3.1) leads to the functional equation:

$$\begin{aligned} P_n(z) &\triangleq \mathbb{IE}(z^{N_n}) \\ &= \mathbb{IP}(D) \mathbb{IE}(z^{N_{n-1}}) + \mathbb{IP}(-D) \sum_{l=2}^n \mathbb{IP}(\text{birth in } l | -D) \mathbb{IE}(z^{N_{l-1}}) \mathbb{IE}(z^{N_{n-l+1}}), \\ &= \frac{\tau}{\tau + (n-1)\nu} P_{n-1}(z) + \frac{(n-1)\nu}{\tau + (n-1)\nu} \sum_{l=1}^{n-1} \frac{1}{n-1} P_l(z) P_{n-l}(z). \end{aligned}$$

Denoting as ρ the ratio τ/ν , we thus obtain that the polynomials P_n , $n = 1, 2, \dots$ satisfy the recurrence

$$\begin{aligned} P_1(z) &= z \\ \forall n > 1, \quad (\rho + n - 1)P_n(z) &= \rho P_{n-1}(z) + \sum_{l=1}^{n-1} P_l(z) P_{n-l}(z), \end{aligned} \quad (3.2)$$

which provides a first (but painful) way of computing the value of $P_n(z)$, $n = 1, 2, \dots$

We focus now on the mean number of births: let $f_n = \mathbb{IE}(N_n)$. Taking the derivative of (3.2) at $z = 1$, one gets:

$$(\rho + n - 1)f_n = \rho f_{n-1} + 2 \sum_{l=1}^{n-1} f_l, \quad (3.3)$$

which, together with $f_1 = 1$, allows a recursive computation of the real numbers f_n , $n = 1, 2, \dots$

This together with standard computations show that the generating function $\phi(t) = \sum_1^{\infty} f_n t^n$ satisfies the first order differential equation:

$$\phi'(t) = \rho + \left(\rho + \frac{1-\rho}{t} + \frac{2}{t-1}\right) \phi(t), \quad |t| < 1. \quad (3.4)$$

Observe that looking for the coefficient of t^n in (3.4) now leads to the simplified recurrence relation:

$$(\rho + n - 1)f_n = (n + 2\rho)f_{n-1} - \rho f_{n-2}, \quad (3.4')$$

that we could also have derived directly from (3.3) by subtracting rank n and rank $n-1$ equations.

We can easily integrate (3.4), and by using the boundary condition $\phi(0) = 0$, we obtain the following representation for ϕ :

$$\phi(t) = \frac{e^{\rho t} t^{1-\rho}}{(1-t)^2} \rho \int_0^t (1-u)^2 e^{-\rho u} u^{\rho-1} du. \quad (3.5)$$

Extracting the coefficient of t^n in $\phi(t)$ leads to a closed expression for f_n (we skip some painful computations):

$$f_n(\rho) = \frac{n+\rho}{1+\rho} + \sum_{k=2}^{n-1} (n-k) \frac{\rho^{k-1}}{(\rho+1) \dots (\rho+k)}, \quad (3.6)$$

from which various bounds and asymptotics are easily derived. In particular, we have the following asymptotic behavior for f_n when n grows to infinity:

$$f_n \sim n K(\rho) \quad n \rightarrow \infty, \quad (3.7)$$

where

$$K(\rho) = \rho e^\rho \int_0^1 (1-u)^2 e^{-\rho u} u^{\rho-1} du \quad (3.8a)$$

$$= \frac{1}{1+\rho} + \sum_{n=2}^{\infty} \frac{\rho^n}{\prod_{i=0}^n (\rho+i)} \quad (3.8b)$$

$$= \sum_{n=2}^{\infty} n(n-1) \frac{\rho^{n-2}}{\prod_{i=1}^n (\rho+i)}. \quad (3.8c)$$

The asymptotics (3.7) and the expression (3.8a) are obtained directly from the integral representation (3.5) of ϕ by applying a Tauberian theorem (for instance, [FEL, p. 423]) which is recalled in Appendix A-1 for easier reference. Note that the sequence $\{f_n\}_{n \geq 0}$ is obviously positive and it is increasing, owing to the fact that the RV's N_n form a *stochastically increasing* sequence (this is also proved in Appendix A-2). Equation (3.8b) is derived directly from (3.6). Finally, expression (3.8c) is obtained by expanding the integral (3.8a) in ρ .

One checks immediately that $K(0) = 1$ (if Centaurs do not move, $f_n = n$, $n = 1, 2, \dots$) and $K(+\infty) = 0$ (if eggs do not hatch, $f_n = 1$, $n = 1, 2, \dots$). Note also the aesthetic aspect of the particular value $K(1) = e - 2$.

We show in figure 1 the values of the ratio f_n/n for various values of n and ρ , with the graph of $K(\rho)$.

An asymptotic analysis, based on formula (3.8a) leads to the following expansion:

$$K(\rho) = \frac{\sqrt{\pi}}{\sqrt{2\rho}} - \frac{2}{3\rho} + O\left(\frac{1}{\rho^{3/2}}\right). \quad (3.9)$$

The outline of this derivation is given in Appendix A-3.

In order to study the complete distribution function of N_n , we introduce the generating function of the polynomials $P_n(z)$, $n = 1, 2, \dots$:

$$\Pi(t, z) \triangleq \sum_{n=1}^{\infty} P_n(z) t^n, \quad |t| < 1, |z| \leq 1.$$

Making use of (3.2), one obtains easily that Π satisfies the differential equation

$$t \frac{\partial \Pi}{\partial t}(t, z) = \rho t z + (1 + \rho(t-1))\Pi(t, z) + \Pi^2(t, z). \quad (3.10)$$

Differentiating (3.10) with respect to z and using: $\Pi(t, 1) = t/(1-t)$, yields (3.4) again. This equation is an homogeneous first order but second degree differential (Ricatti) equation, and does not admit, as such, a closed form solution in general. However, using the change of function

$$\Pi(t, z) = -t \frac{h'(t, z)}{h(t, z)}, \quad (3.11a)$$

where differentiation is understood with respect to the variable t , it is easy to check that (3.10) can be transformed into the linear second order equation

$$th''(t, z) + \rho(1-t)h'(t, z) + \rho zh(t, z) = 0. \quad (3.11b)$$

Applying standard methods ([INC] pp .158-), we look for solutions of the form

$$h(t, z) = t^\alpha \sum_{i=0}^{\infty} h_i(z) t^i,$$

where the functions $h_i(z)$, $i = 0, 1, \dots$ are coefficients to be determined. We find that the constant α and these functions must satisfy the relations:

$$\begin{cases} \alpha = 0 \text{ or } 1 - \rho \\ h_{n+1}(z) = \rho \frac{\alpha + n - z}{(n+1)(n + \rho + 2\alpha)} h_n(z), \quad \forall n \geq 1, \end{cases}$$

and $h_0(z)$ is an undefined constant.

If $\rho \neq 1$, it is plain that we have two independent solutions to (3.10), and therefore all its solutions. Making use of the condition $\Pi(t, 1) = t/(1-t)$, we check that the right solution corresponds to $\alpha = 0$, and is given by:

$$h_n(z) = \frac{\rho^n}{n!} \prod_{i=0}^{n-1} \frac{i-z}{i+\rho} \quad n \geq 0, \quad (3.12)$$

and $h_0(z) = 1$.

It can be checked that the series $h(t, z)$ converges at least for all $|z| \leq 1$ and $|t| < 1$. Observe also that if z is a positive integer, $h(t, z)$ is a polynomial in t . Note that $h(t, 1) = 1 - t$.

If $\rho = 1$, this method gives only one solution, denoted $h_0(t, z)$. The other ones have to be searched in the form $k(t, z)h_0(t, z)$, and the function k is given as an integral involving h_0 . The details are not relevant here. Actually, it is better to avoid this case, and apply the results of the case $\rho \neq 1$ "by continuity".

We consider now the asymptotics of the moments of N_n as $n \rightarrow \infty$. The successive derivatives of the generating function $P_n(z)$, taken at $z = 1$ are the *factorial moments* of N_n :

$$\frac{d^k P_n}{dz^k}(1) = \sum_{p=k}^{\infty} p(p-1) \dots (p-k+1) \text{IP}(N_n = p). \quad (3.13)$$

Factorial moments and natural moments are obviously related: developing the right-hand-side of (3.13), we have:

$$\frac{d^k P}{dz^k}(1) = \sum_{i=0}^k s_{k,i} \text{IE}(N_n^i), \quad (1.14)$$

where the $s_{k,i}$'s (which are called the Stirling numbers of first kind), are constants independent of n . Thus, asymptotics of the factorial moments will translate into asymptotics of the natural moments. Note that $s_{k,k} = 1$.

From (3.11a), it is plain that, for all k ,

$$\frac{\partial^k \Pi}{\partial z^k}(t, z) = -t \frac{R_k(t, z)}{h^{k+1}(t, z)},$$

where $R_k(t, z)$, $|t| < 1, k > 0$ is the analytic function of $z, |z| \leq 1$, given by the recurrence:

$$R_k(t, z) = \frac{\partial R_{k-1}}{\partial z}(t, z)h(t, z) - kR_{k-1}(t, z)\frac{\partial h}{\partial z}(t, z),$$

with $R_0(t, z) = h'(t, z)$.

Letting z go to 1, and using $h(t, 1) = 1 - t$, we obtain:

$$\frac{\partial^k \Pi}{\partial z^k}(t, 1) = -t \frac{R_k(t, 1)}{(1-t)^{k+1}}$$

and

$$R_k(t, 1) = \frac{\partial R_{k-1}}{\partial z}(t, 1)(1-t) - kR_{k-1}(t, 1)\frac{\partial h}{\partial z}(t, z)\Big|_{z=1}.$$

Using the stochastic increasingness of the RV's N_n , we know that the sequences $\{\mathbb{E}(N_n^k)\}_n$ are increasing for all $k \geq 1$, as well as the corresponding factorial moments. Then, it follows by the same Tauberian arguments that:

$$\begin{aligned} [t^n] \frac{\partial^k \Pi}{\partial z^k}(t, 1) &\sim -\frac{n^k}{k!} \lim_{t \rightarrow 1} R_k(t, 1) \\ &\sim -\frac{n^k}{k!} k \frac{\partial h}{\partial z}(t, z)\Big|_{z=1} \lim_{t \rightarrow 1} R_{k-1}(t, 1) \\ &\sim (nK(\rho))^k \end{aligned} \quad (3.15)$$

where we used the fact that

$$K(\rho) = -\lim_{t \rightarrow 1} \frac{\partial h}{\partial z}(t, z)\Big|_{z=1}. \quad (3.16)$$

We made use of the common notation: $[t^n]F(t)$ is the coefficient of t^n in the (formal) series expansion of the function F .

Note that, making use of (3.12), this formula provides now a fourth expression for $K(\rho)$:

$$K(\rho) = 1 - \sum_{n=2}^{\infty} \frac{\rho^n}{n(n-1)} \frac{\rho^n}{\prod_{i=1}^{n-1} i + \rho}. \quad (3.8d)$$

A simple recurrence, based on formula (3.13) will then show that the same asymptotics hold for the natural moments: this is true for $k = 1$ as the natural and the factorial moment coincide; if this is true up to order $k - 1$, then (3.13) rewrites:

$$\mathbb{E}(N_n^k) = \frac{d^k P}{dz^k}(1) + O((K(\rho)n)^{k-1}),$$

and application of (3.15) gives:

$$\forall k > 0, \quad \mathbb{E}(N_n^k) \sim (nK(\rho))^k. \quad (3.17)$$

In particular, one has:

$$\lim_{n \rightarrow \infty} \mathbb{E} \left(\left(\frac{N_n}{n} \right)^2 \right) - \left(\mathbb{E} \left(\frac{N_n}{n} \right) \right)^2 = 0,$$

that is, the limit of the variance of the normalized RV N_n/n is zero, which proves that N_n/n converges in law to a Dirac mass at point $K(\rho)$. It can be conjectured, using the analogy of this model with certain *random tries* families ([JARE]), that the distribution of $(N_n - K(\rho))/\sqrt{n}$ converges to a Gaussian distribution.

3.1.2/The Completion Time

Consider a chain of n sites with an egg in each except for the first one which contains a Centaur. Let T_n denote the time at which the last egg is destroyed (either by hatching or by the arrival of another Centaur). Let e_n be the instant of the first event (birth or move). Using the event D introduced in 3.1.1/, T_n can be shown to satisfy the recurrence:

$$T_n = e_n + T_{n-1} \mathbf{1}_{\{D\}} + \sum_{l=2}^n \max(T_{l-1}, T_{n-l+1}) \mathbf{1}_{\{-D \text{ and birth in site } l\}}, \quad (3.18)$$

with $T_1 = 0$.

Clearly, e_n is exponentially distributed with rate $\tau + (n-1)\nu$. Denoting by $F_n(t) = \mathbb{P}(T_n \leq t)$ the distribution of T_n , (3.18) implies:

$$\begin{aligned} F_n(t) &= \int_0^t \left[\frac{\tau}{\tau + (n-1)\nu} F_{n-1}(t-x) + \frac{\nu}{\tau + (n-1)\nu} \sum_{l=1}^{n-1} F_l(t-x) F_{n-l}(t-x) \right] d\mathbb{P}(e_n \leq x) \\ &= \int_0^t \left[\tau F_{n-1}(t-x) + \nu \sum_{l=1}^{n-1} F_l(t-x) F_{n-l}(t-x) \right] e^{-x(\tau+(n-1)\nu)} dx \end{aligned}$$

Introducing the generating function $\Phi(t, z) = \sum_1^\infty F_n(t) z^n$ leads to the following (and yet unsolved) functional equation:

$$\Phi(t, z) = z F_1(t) + \int_0^t [\tau z \Phi(t-x, z e^{-\nu x}) + \nu e^{\nu x} \Phi^2(t-x, z e^{-\nu x})] e^{-\tau x} dx. \quad (3.19)$$

Even in the simplest case $\tau = 0$, this equation seems untractable. However, notice that in this case, T_n writes:

$$T_n = e_n + e_{n-1} + \dots + e_2,$$

where e_n is exponentially distributed with rate $(n-1)\nu$, so that its expected value is:

$$\mathbb{E} T_n = \frac{1}{\nu} \left(1 + \frac{1}{2} + \dots + \frac{1}{n-1} \right) \sim \frac{\log n}{\nu}.$$

But it should be clear that, ν and n being fixed, $\mathbb{E} T_n$ is a decreasing function of τ . Therefore, we have:

$$\forall \tau, \nu, \forall n, \quad \mathbb{E} T_n \leq \frac{\log n}{\nu}.$$

The results of section 4.3/ indicate that for all $\tau > 0$ and $\nu > 0$, T_n is actually of order $\log n$.

3.2/The Bi-Directional Case

In the bi-directional case, each egg gives birth two twin Centaurs of opposite spin who run across the Ring in two opposite directions, and stop when they meet again. As before, the non-overtaking assumption allows us to open the ring and reduce the problem to the case of an array of sites.

Throughout this section, we shall use the same notations as in 3.1/, but distinguish the present case to the previous one by adding a "bar" to the corresponding quantities. Thus, we are interested in the distribution of \bar{N}_n , the number of births (i.e. the number of Centaur *pairs*), and \bar{T}_n the time when the last egg disappears.

For $n = 1, 2, \dots$, let \bar{N}_n be the number of births in an array of $n + 1$ sites such that there is one Centaur in site 1 and one in site $n + 1$, going in opposite directions, and an egg in each of the sites 2 through n . As before, we have:

$$\bar{N}_n = \bar{N}_{n-1} \mathbf{1}_{\{\bar{D}\}} + \sum_{l=2}^n (\bar{N}_{l-1} + \bar{N}_{n-l+1}) \mathbf{1}_{\{-\bar{D} \text{ and birth in site } l\}}, \quad (3.1')$$

where \bar{D} represents now the event: "Centaur 1 moves to site 2 or Centaur $n + 1$ moves to site n before any birth takes place".

As the probability of \bar{D} is

$$\mathbb{P}(\bar{D}) = \frac{2\tau}{2\tau + (n-1)\nu},$$

it follows that all the computations on N_n can be applied to \bar{N}_n replacing τ by 2τ , or ρ by 2ρ . Thus, we obtain

$$\bar{f}_n(\rho) = f_n(2\rho), \quad \bar{K}(\rho) = k(2\rho), \quad \bar{T}_n(\rho) = T_n(2\rho).$$

The following question comes in order: "what are the benefits of having a bi-directional ring?". Comparing the number of messages exchanged is equivalent to comparing the number of births. It is plain that the functions f_n must be decreasing, so that:

$$\forall n, \forall \rho, \quad \bar{f}_n(\rho) \leq f_n(\rho).$$

Figure 2 shows the graph of the ratio $f_n(2\rho)/f_n(\rho)$ for various values of n (including infinity) and of ρ . It turns out that this ratio goes to 1 when ρ is small (n births in both cases) but also when ρ is large (a single birth in both cases). For each fixed n , there is, by numerical evidence, a single optimal ρ that minimizes this ratio (but no known easy way to compute it). When n goes to infinity, this optimum goes to infinity extremely rapidly, and $K(2\rho)/K(\rho)$ goes to $1/\sqrt{2}$ (according to (3.9)) very slowly. In any case, we observe the bound:

$$\forall n, \forall \rho, \quad \bar{f}_n(\rho) > \frac{\sqrt{2}}{2} f_n(\rho). \quad (3.20)$$

4/THE DISCRETE TIME CASE

In a discrete time model, events can occur only at fixed instants $0, \delta t, 2\delta t, \dots, k\delta t, \dots$. The constant δt is called "time unit", or "time interval" or "time slot". Here, Centaurs may move or hatch out at each unit of time with fixed probabilities (independent of the time and the state of the system). Exponential RV's are therefore replaced by Bernoulli variables.

This time discretization may seem strange to study a system that is essentially asynchronous. Two reasons justify such a model. First, one may recover results in continuous time from results in discrete time by letting the time slot length and the parameters go to zero together in a proper way. In particular, the results that are obtained for the completion time are more complete here than in section 3/. Secondly, this model may be used as an approximation (actually an upper bound) for a continuous time system, allowing to release the exponential assumptions we made in 3/. Also, note that the assumption of deterministic transmission times is often done in the literature. The deterministic model includes this case, which allows to

The discrete time model seems more difficult to approach, mainly in that more than one event may occur at a time. Thus, the "decomposition tree" we implicitly used in section 3/ will not be unary-binary any more, and the complexity of the equations will increase. However, we were able to obtain nice results on this second case also, especially for the completion time problem, which appeared untractable in the previous section.

4.1/The One-Direction Case

4.1.1/The Number of Centaurs

Despite the possibility of multiple events, we can reduce the study to the case of a linear array of sites, as before.

Let N_n and M_n denote the number of births, respectively in a ring of size n with n eggs, and in an array of length $n - 1$ with a Centaur in site 1 and eggs everywhere else. For all $n = 1, 2, \dots$ and $k = 0, 1, \dots$, define the events:

$D_n(l_0, \dots, l_k) = \{\text{In one time slot, there are } k \text{ births in the } n\text{-site-array defining } k \text{ sub-arrays of respective lengths } l_0, \dots, l_k\}.$

As the decisions of moving and hatching are independent of each other, it may happen that, at a certain time, a Centaur decides to move to a site where an egg decides to hatch. It is necessary to decide what is the result of the event: one or two Centaurs in the site? The simplest possibilities are to decide that Centaurs always move "just before" or "just after" the eggs hatch, which results in two and one Centaur respectively. In order to take these two deterministic decisions into account in the same model, we "interpolate" them by a general Bernoulli choice, assuming that in case of such a conflict, the grownup Centaur tosses a coin, and decides to let the egg hatch with probability ϵ and to destroy it with probability $1 - \epsilon$. Therefore, we define the event:

$E = \{\text{the Centaur in site 1 destroys the egg, given that he moves to site 2 and that the egg in site 2 wants to hatch}\}.$

Then, looking for events during a time interval, we have the recurrence relation:

$$\begin{aligned}
N_n &= N_n \mathbf{1}_{\{\text{no births}\}} + \sum_{k=1}^n \sum_{l_0, \dots, l_k} (M_{l_1} + \dots + M_{l_{k-1}} + M_{l_k + l_0 - 1}) \mathbf{1}_{\{D_{n+1}(l_0, \dots, l_k)\}} \\
M_n &= \mathbf{1}_{\{\text{no move}\}} \sum_{k=0}^{n-1} \sum_{l_0, \dots, l_k} (M_{l_0} + \dots + M_{l_k}) \mathbf{1}_{\{D_n(l_0, \dots, l_k)\}} \\
&\quad + \mathbf{1}_{\{\text{move}\}} (\mathbf{1}_{\{E\}} \mathbf{1}_{\{\text{birth in 2}\}} + \sum_{k=0}^{n-2} \sum_{l_0, \dots, l_k} (M_{l_0} + \dots + M_{l_k}) \mathbf{1}_{\{D_{n-1}(l_0, \dots, l_k)\}}), \quad \forall n > 1
\end{aligned} \tag{4.1}$$

with $M_1 = 1$. The event $D_{n+1}(\dots)$ in the first equation expresses the choice of an origin on the ring.

Let τ and ν denote now the probabilities of moving and hatching during a time interval, and let $\varepsilon = \mathbb{P}(E)$. First, notice that:

$$\mathbb{P}(D_n(l_0, \dots, l_k)) = \begin{cases} \nu^k (1 - \nu)^{n-k-1} & \text{if } l_i > 0, l_0 + \dots + l_k = n \\ 0 & \text{otherwise} \end{cases}$$

If $P_n(z)$ and $Q_n(z)$ are the respective generating functions of N_n and M_n , then the following relations are easily derived from (4.1):

$$\begin{aligned}
P_n(z) &= (1 - \nu)^n P_n(z) + \sum_{k=0}^n \nu^k (1 - \nu)^{n-k} \sum_{\substack{l_0 + \dots + l_k = n+1 \\ l_i > 0}} Q_{l_1}(z) \dots Q_{l_{k-1}}(z) Q_{l_k + l_0 - 1}(z) \\
Q_n(z) &= (1 - \tau) \sum_{k=0}^{n-1} \nu^k (1 - \nu)^{n-1-k} \sum_{\substack{l_0 + \dots + l_k = n \\ l_i > 0}} Q_{l_0}(z) \dots Q_{l_k}(z) \\
&\quad + \tau(\varepsilon \nu z + 1 - \varepsilon \nu) \sum_{k=0}^{n-2} \nu^k (1 - \nu)^{n-2-k} \sum_{\substack{l_0 + \dots + l_k = n-1 \\ l_i > 0}} Q_{l_0}(z) \dots Q_{l_k}(z)
\end{aligned} \tag{4.2}$$

Let $\Pi(t, z)$ and $\Psi(t, z)$ be the generating functions of P_n and Q_n , $n > 0$ respectively. From (4.2) these functions are shown to satisfy the following functional (nonlinear and non-local) equations:

$$\begin{aligned}
\Pi(t, z) &= \Pi(\beta t, z) + \nu \beta t \frac{\frac{\partial \Psi}{\partial t}(\beta t, z)}{\beta - \nu \Psi(\beta t, z)} \\
\Psi(t, z) &= \tau t z + (1 - \tau + \tau(\varepsilon \nu z + 1 - \varepsilon \nu)) \frac{\Psi(\beta t, z)}{\beta - \nu \Psi(\beta t, z)},
\end{aligned} \tag{4.3}$$

where $\beta = 1 - \nu$. The detail of the summations are given in the Appendix A-4.

From these equations, closed form expressions can be derived for Π and Ψ , but their complexity makes them practically useless for recovering the P_n 's and the Q_n 's, and for their asymptotic analysis. Let us now look for the expected values of N_n and M_n . Introducing the functions:

$$F(t) = \sum_{n=1}^{\infty} \mathbb{E}(N_n) t^n = \frac{\partial \Pi}{\partial z}(t, 1) \quad G(t) = \sum_{n=1}^{\infty} \mathbb{E}(M_n) t^n = \frac{\partial \Psi}{\partial z}(t, 1),$$

we have, from (4.3) and the identity $\Psi(t, 1) = t/(1-t)$:

$$\begin{aligned} F(t) &= F(\beta t) + \nu t \frac{d}{dt} \left[\frac{1 - \beta t G(\beta t)}{1-t} \frac{1}{\beta} \right] \\ G(t) &= \frac{\tau t}{1-t} (1-t + \varepsilon \nu t) + (1-\tau + \tau t) \frac{G(\beta t)}{\beta} \left(\frac{1-\beta t}{1-t} \right)^2. \end{aligned} \quad (4.4)$$

Both equations of (4.4) are linear non local equations and can be solved recursively as shown in Appendix A-5. The only difficulty is to check the convergence of the iterations.

Here, we obtain:

$$F(t) = \nu t \sum_{i=0}^{\infty} \beta^i \frac{d}{du} \left[\frac{1 - \beta u G(\beta u)}{1-u} \frac{1}{\beta} \right]_{u=\beta^i t} \quad (4.5)$$

$$G(t) = \frac{\tau t}{(1-t)^2} \sum_{i=0}^{\infty} (1 - \beta^i t + \varepsilon \nu \beta^i t) (1 - \beta^i t) \prod_{j=0}^{i-1} (1 - \tau + \tau \beta^j t) \quad (4.6)$$

Following the usual convention, an empty product is equal to one.

We have not addressed the problem of recovering IEN_n and IEM_n from (4.5 - 6), but we get easily their asymptotics, as F and G turn out to have a pole of order 2 at $t = 1$:

$$\begin{aligned} \text{IEM}_n &\sim n \tau \sum_{i=1}^{\infty} (1 - \beta^i + \varepsilon \nu \beta^i) (1 - \beta^i) \prod_{j=0}^{i-1} (1 - \tau + \tau \beta^j) \\ \text{IEN}_n &\sim \text{IEM}_n \end{aligned}$$

We showed in figures 3 and 4 the values of the limit IEN_n/n , $n \rightarrow \infty$, for ν varying from 0 to 1, for values of τ ranging from 0 to 1 with step 0.1, and for $\varepsilon = 0$ and 1 respectively. Figures 5 and 6 represent the incidence of ε , τ being fixed to 0.5 and 1 respectively.

4.1.2/The Completion Time

For more convenience, we will evaluate times with the unit δt , and therefore omit it in the formulas. We shall use it again explicitly in section 4.3/.

Let U_n (resp T_n) be the time at which the last egg disappears in an array (resp a ring) of n sites in initial configuration. These two RV's satisfy the recurrences:

$$\begin{aligned} T_n &= 1 + T_n \mathbf{1}_{\{\text{no births}\}} + \sum_{k=1}^n \sum_{l_0, \dots, l_k} \max(U_{l_1}, \dots, U_{l_{k-1}}, U_{l_k + l_0 - 1}) \mathbf{1}_{\{D_{n+1}(l_0, \dots, l_k)\}} \\ U_n &= \mathbf{1}_{\{\text{no move}\}} \sum_{k=0}^{n-1} \sum_{l_0, \dots, l_k} \max(U_{l_0}, \dots, U_{l_k}) \mathbf{1}_{\{D_n(l_0, \dots, l_k)\}} \\ &+ \mathbf{1}_{\{\text{move}\}} \sum_{k=0}^{n-2} \sum_{l_0, \dots, l_k} \max(U_{l_0}, \dots, U_{l_k}) \mathbf{1}_{\{D_{n-1}(l_0, \dots, l_k)\}} \\ &+ 1, \end{aligned}$$

and $U_1 = 0$.

Introducing the generating functions: $\Theta_t(z) = \sum_1^\infty \mathbb{P}(T_n \leq t)z^n$ and $\Upsilon_t(z) = \sum_1^\infty \mathbb{P}(U_n \leq t)z^n$, we use the calculations of the previous section to get: $\forall t > 1, \forall |z| < 1$,

$$\begin{aligned}\Theta_t(z) &= \Theta_{t-1}(\beta z) + \nu \beta z \frac{\Upsilon'_{t-1}(\beta z)}{\beta - \nu \Upsilon_{t-1}(\beta z)} \\ \Upsilon_t(z) &= \tau z + (1 - \tau + \tau z) \frac{\Upsilon_{t-1}(\beta z)}{\beta - \nu \Upsilon_{t-1}(\beta z)},\end{aligned}\tag{4.7}$$

with $\Upsilon_0(z) = z$ and $\Theta_0(z) = 0$.

Here, we have a non linear and non local functional recurrence. It can be solved to obtain $\Upsilon_t(z)$ for each t , but this is not of practical interest.

Our problem is now to obtain from (4.7) information on quantities $\mathbb{P}(T_n \leq t)$, and on the asymptotic behavior of $\mathbb{E}T_n$ as n grows. It turns out that everything depends on the location of the poles of the functions Υ_t , $t \geq 0$.

We therefore concentrate on the study of the recurrence on Υ_t . To get rid of the "constant" term τz , we introduce the new function $\tilde{\Upsilon}_t(z) = (z/1-z) - \Upsilon_t(z) = \sum_1^\infty z^n \mathbb{P}(T_n > t)$. Then we use the function $U_t = (1-z)\tilde{\Upsilon}(z)$, to obtain the simplified recurrence:

$$U_t(z) = (1 - \tau + \tau z)(1 - \beta z) \frac{U_{t-1}(\beta z)}{\beta(1-z) + \nu U_{t-1}(\beta z)},\tag{4.8}$$

with initial condition $U_0(z) = z^2$. Now, notice that U_t must be a rational function. Therefore, we set $U_t = N_t/D_t$ where N_t and D_t are relatively prime polynomials. Then, with (4.8), we obtain: $\forall t \geq 1$,

$$\begin{aligned}N_t(z) &= \frac{1 - \tau + \tau z}{\beta} N_{t-1}(\beta z) \\ (1 - \beta z)D_t(z) &= (1 - z)D_{t-1}(\beta z) + \frac{\nu}{\beta} N_{t-1}(\beta z).\end{aligned}\tag{4.9}$$

If $\tau \neq 0$, the degrees of N_t and D_t are $t+2$ and t . If $\tau = 0$ (pure birth process), then $U_t(z) = \beta^t z^2 / (1 - (1 - \beta^t)z)$.

The second recurrence defines a polynomial, as its right-hand side is zero when $\beta z = 1$. This recurrences are solved to obtain: $D_1(z) = 1 - \nu z$ and for all $t > 1$,

$$\begin{aligned}N_t(z) &= z^2 \beta^t \prod_{i=0}^{t-1} (1 - \tau + \tau \beta^i z) \\ D_t(z) &= (1 - z)(1 + z\beta^t) + \frac{1}{\beta} N_{t-1}(\beta z) + \frac{\tau}{\beta} (1 - z) \sum_{i=0}^{t-2} N_i(\beta^{t-i} z).\end{aligned}\tag{4.10}$$

Note that in this last formula, $D_t(z)$ seems to have the degree of $N_{t-1}(z)$ which is $t+1$, but the term in z^{t+1} actually vanishes.

We shall now construct the functions $\Theta_t(z)$, $t \geq 0$. Rewriting the recurrence (4.7) as:

$$\Theta_t(z) = \Theta_{t-1}(\beta z) - z \frac{d}{dz} \log(\beta - \nu \Upsilon_{t-1}(\beta z)),$$

introducing $N_{t-1}(z)$ and $D_{t-1}(z)$ and making use of the recurrences (4.9), we obtain the surprising relation:

$$\Theta_t(z) - z \frac{D'_t(z)}{D_t(z)} = \Theta_{t-1}(\beta z) - \beta z \frac{D'_{t-1}(\beta z)}{D_{t-1}(\beta z)} = \dots = \Theta_0(\beta^t z) - \beta^t z \frac{D'_0(\beta^t z)}{D_0(\beta^t z)} = 0.$$

Denoting by $p_{1,t}, \dots, p_{t,t}$ the t zeroes of D_t , the distribution of T_n is simply given by:

$$\mathbb{P}(T_n \leq t) = [z^n] \Theta_t(z) = \sum_{i=1}^t \frac{1}{p_{i,t}^n}. \quad (4.11)$$

If $\tau = 0$, only one pole $p_{1,t} = 1/(1 - \beta^t)$ is finite, and this reduces to: $\mathbb{P}(T_n \leq t) = (1 - \beta^t)^n$, which direct arguments would have shown more easily indeed.

We now turn to the study of the asymptotics of $\mathbb{E}T_n$, when n grows. Note that

$$\mathbb{E}T_n = [z^n] \tilde{\Theta}(z), \quad \tilde{\Theta}(z) \triangleq \sum_{t=0}^{\infty} \left(\frac{z}{1-z} - \Theta_t(z) \right).$$

We first expose the degenerated case $\tau = 0$, which will be a guide for the general case. In this case,

$$\mathbb{E}T_n = \sum_{t=0}^{\infty} 1 - (1 - \beta^t)^n.$$

Asymptotics of this sum can be derived, for instance using Mellin transform techniques [FLA], or by probabilistic arguments, showing that $T_n/\log n$ converges weakly to a Dirac mass at point $-1/\log \beta$. One obtains:

$$\mathbb{E}T_n \sim - \frac{\log n}{\log(1 - \nu)}.$$

It is worth noting here that, as in section 3.1/, $\mathbb{E}T_n$ is a decreasing function of τ , and therefore, for all values of τ , the mean completion time is bounded by a function of order $\log(n)$.

The following analysis will prove that $\mathbb{E}T_n = O(\log(n))$ for all $0 \leq \tau < 1$ and $0 < \nu < 1$.

The first step consists in locating the zeroes of D_t . Starting from (4.10), careful boundings and the application of Rouché's theorem allow to state that for all $z \in \mathbb{C}$, $D_t(z) \rightarrow 1 - z$ as $t \rightarrow \infty$, and that for all $M > 0$, $p_{1,t}$ is the single zero of D_t in the disk of radius M , if t is sufficiently large. This zero, of course goes to 1 when t goes to infinity. Therefore, there exists some $R > 1$ such that, for t large enough:

$$[z^n] \Theta_t(z) = \frac{1}{p_{1,t}^n} + O(R^{-n}).$$

The second step consists in estimating $p_{1,t}$. The fact that $D_t(z) \rightarrow 1 - z$ suggests that $p_{1,t} \sim 1 + \phi_t$ where $\phi_t = D_t(1) = N_t(1)$ (this is Newton's method for finding zeroes or real functions). It turns out that

$$p_{1,t} = 1 + \phi_t + o(\beta^t).$$

The third step consists in investigating how ϕ_t goes to zero. We have:

$$\phi_t = \beta^t \prod_{i=0}^{t-1} (1 - \tau + \tau \beta^i) = (\beta(1 - \tau))^t \prod_{i=0}^{t-1} \left(1 + \frac{\tau}{1 - \tau} \beta^i \right) \sim \alpha (\beta(1 - \tau))^t,$$

where α is the limit of the product when t tends to infinity, which exists and is not zero, provided that $\tau < 1$ and $\beta < 1$.

The last step consists in summing the estimates of $\mathbb{P}(T_n > t)$ over t , to get:

$$\mathbb{E}T_n \sim \sum_{t=0}^{\infty} 1 - (1 - \phi_t)^n \sim \sum_{t=0}^{\infty} 1 - (1 - \alpha\beta^t(1 - \tau)^t)^n \sim -\frac{\log n}{\log((1 - \nu)(1 - \tau))}. \quad (4.12)$$

To conclude this section, we look at the limiting cases where the preceding analysis does not apply:

- The case $\nu = 1$: all the Centaurs hatch out at $t = 1$: $T_n = 1$ a.s. for all $n \geq 1$.
- The case $\nu = 0$: the Centaurs never hatch out, so that $T_n = \infty$. However, one may notice that $-1/\log((1 - \nu)(1 - \tau))$ does not go to zero when ν tends to 0 (if $\tau \neq 0$). It is the infinite product α , which would appear in a more detailed asymptotic expansion of $\mathbb{E}T_n$, which grows to infinity.
- The case $\tau = 1, \nu < 1$. In this case, $\phi_t = \beta^{t(t+1)/2}$. As above, it is easy to prove that $T_n/\sqrt{\log n}$ converges in law to a Dirac mass, so that:

$$\mathbb{E}T_n \sim \sqrt{-\frac{2 \log n}{\log(1 - \nu)}}.$$

4.2/The Bi-Directional Case

4.2.1/The Number of Centaurs

In this section, as in 3.2/, we use the previous notations with a "bar". Let \bar{N}_n be the number of births (or centaur pairs) in a ring of size n , and \bar{M}_n this number in an array of $n + 1$ sites, as in 3.2/.

The equation relating the \bar{N} 's to the \bar{M} 's is the same as in 4.1/, but the decomposition of \bar{M}_n changes. As in the one direction case, conflicts appear between moves and births in the same site. In order to avoid the complications that appear in the case $n = 2$ when two Centaurs decide to jump simultaneously in the same site, we shall set ε to zero in the remaining.

The RV's \bar{M}_n satisfy:

$$\begin{aligned} \bar{M}_n &= \mathbf{1}_{\{\text{no moves}\}} \sum_{k=0}^{n-1} (\dots) + \mathbf{1}_{\{\text{one move}\}} \sum_{k=0}^{n-2} (\dots) + \mathbf{1}_{\{\text{two moves}\}} \sum_{k=0}^{n-3} (\dots) \quad n > 2 \\ \bar{M}_2 &= \mathbf{1}_{\{\text{no moves}\}} \sum_{k=0}^1 (\dots) + \mathbf{1}_{\{\text{one move}\}} + \mathbf{1}_{\{\text{two moves}\}} \end{aligned}$$

and $\bar{M}_1 = 1$.

The corresponding generating functions then satisfy:

$$\begin{aligned} \bar{\Pi}(t, z) &= \bar{\Pi}(\beta t, z) + \nu\beta t \frac{\frac{\partial \bar{\Psi}}{\partial t}(\beta t, z)}{\beta - \nu\bar{\Psi}(\beta t, z)} \\ \bar{\Psi}(t, z) &= \tau t z(2 - \tau + \tau t) + (1 - \tau + \tau t)^2 \frac{\bar{\Psi}(\beta t, z)}{\beta - \nu\bar{\Psi}(\beta t, z)}, \end{aligned} \quad (4.13)$$

where β still denotes $1 - \nu$.

The function \bar{G} is therefore given by:

$$\bar{G}(t) = \tau t(2 - \tau + \tau t) + (1 - \tau + \tau t)^2 \frac{\bar{G}(\beta t)}{\beta} \left(\frac{1 - \beta t}{1 - t} \right)^2, \quad (4.14)$$

so that the asymptotics of \bar{N}_n are:

$$\bar{N}_n \sim \bar{M}_n \sim n\tau \sum_{i=1}^{\infty} (1 - \beta^i)^2 (2 - \tau + \tau\beta^i) \prod_{j=0}^{i-1} (1 - \tau + \tau\beta^j)^2.$$

Figure 7 compares the asymptotics of \bar{N}_n/n and N_n/n for various values of ν and τ . Their ratio is close to 1 when τ is close to 0 or ν close to 1. Note also that it is not less than its limit when ν approaches zero, a number which looks very much like $\sqrt{2}/2$ (independently of τ), as the results of the continuous case allowed to guess.

4.2.2/The Completion Time

Like in the preceding subsection, the analysis is entirely similar to that of the one direction case, with only some modifications in the polynomial coefficients of the generating functions.

Only the decomposition of \bar{U}_n changes. It is now given by:

$$\bar{U}_n = \mathbf{1}_{\{\text{no move}\}} \sum_{k=0}^{n-1} \sum_{l_0, \dots, l_k} (\dots) + \mathbf{1}_{\{\text{one move}\}} \sum_{k=0}^{n-2} \sum_{l_0, \dots, l_k} (\dots) + \mathbf{1}_{\{\text{two moves}\}} \sum_{k=0}^{n-2} \sum_{l_0, \dots, l_k} (\dots) + 1.$$

and $\bar{U}_1 = 0$.

Thus, the relation between $\bar{\Theta}_t$ and $\bar{\Upsilon}_{t-1}$ does not change, and we have:

$$\bar{\Upsilon}_t(z) = \tau z(2 - \tau + \tau z) + (1 - \tau + \tau z)^2 \frac{\bar{\Upsilon}_{t-1}(\beta z)}{\beta - \nu \bar{\Upsilon}_{t-1}(\beta z)}.$$

Now, introducing $\bar{\Upsilon}_t$, the polynomials \bar{N}_t and \bar{D}_t are:

$$\bar{N}_t(z) = \beta^t z^2 \prod_{i=1}^{t-1} (1 - \tau + \tau\beta^i z)^2$$

$$\bar{D}_t(z) = 1 - z + z(1 - z)\beta^t + \frac{\bar{N}_{t-1}(\beta z)}{\beta} + \frac{\tau}{\beta}(1 - z) \sum_{i=1}^{t-1} \bar{N}_{i-1}(\beta^{t-i+1} z)(2 - \tau + \tau\beta^{t-i} z).$$

The first zero of \bar{D}_t is located at $\bar{p}_{1,t} = 1 + \bar{\phi}_t$, with $\bar{\phi}_t \sim \phi_t^2$, and we have:

$$\mathbb{E}\bar{T}_n \sim - \frac{\log n}{\log((1 - \nu)^2(1 - \tau)^2)} \sim \frac{1}{2} \mathbb{E}T_n,$$

if $0 \leq \tau < 1$ and $0 < \nu < 1$, and

$$\mathbb{E}\bar{T}_n \sim \sqrt{- \frac{\log n}{\log(1 - \nu)}},$$

if $\tau = 1$.

4.3/From Discrete to Continuous

The purpose of this last section is to derive the asymptotics of $\mathbb{E}T_n$ in the continuous case. For this, we shall use the results of the discrete time model and adjust the parameters to let the discrete geometric laws tend to continuous exponential laws. It is understood in this subsection that we assume the existence of some “continuity property” of the system’s statistics with respect to its parameters: when the laws of the Centaur’s births and transitions converge weakly to some limit, the laws of the number of Centaurs and the completion time also converge to the corresponding laws.

Let X be a RV with a geometric distribution of parameter p and time slot δt , that is: $\mathbb{P}(X = k\delta t) = p(1-p)^{k-1}$ for all positive integer k . Then $\mathbb{P}(X > k\delta t) = (1-p)^k$, and $\mathbb{E}X = \delta t/p = 1/\lambda$ for some positive number λ . Now, replacing p by $\lambda\delta t$, then for all positive real x ,

$$\lim_{\delta t \rightarrow 0} \mathbb{P}(X > x) = \lim_{\delta t \rightarrow 0} (1 - \lambda\delta t)^{\lfloor \frac{x}{\delta t} \rfloor} = e^{-\lambda x}.$$

Thus, if the time slot goes to zero while its mean is kept constant, the geometric law converges to the exponential law with the same mean.

Therefore, if we want to obtain the expected time $\mathbb{E}T_n$ in the continuous model with exponential laws with mean $1/\lambda$ (moves) and $1/\mu$ (births), we set $\tau = \lambda\delta t$ and $\nu = \mu\delta t$, and let δt shrink to zero in (4.12), to obtain (remember that $\mathbb{E}T_n$ was evaluated as a number of time slots):

$$\lim_{\delta t \rightarrow 0} \mathbb{E}T_n = \lim_{\delta t \rightarrow 0} - \frac{\delta t \log n}{\log((1 - \mu\delta t)(1 - \lambda\delta t))} = \frac{\log n}{\lambda + \mu}.$$

Note that the same method can be applied *a priori* to recover the results for the number of centaurs in the continuous case, from the results in the discrete case. In particular, we have:

$$\lim_{\substack{\tau = \lambda\delta t \\ \nu = \mu\delta t \\ \delta t \rightarrow 0}} \left(\tau \sum_{i=1}^{\infty} (1 - \beta^i + \epsilon\nu\beta^i)(1 - \beta^i) \prod_{j=0}^{i-1} (1 - \tau + \tau\beta^j) \right) = K\left(\frac{\lambda}{\mu}\right).$$

This identity has been checked numerically with a remarkable accuracy, but proving it analytically looks hard.

5/CONCLUSIONS

We provide an analysis of the election algorithm, mainly based on a *tree-like decomposition principle*. The interesting performance measures are the total number p of sites that initiate an election, and the time T at which the last site is aware that an election is going on. These quantities are studied both in continuous and discrete time.

The formulation of the continuous time problem leads to the expression of the double generating function of the distribution of p , and n , as the solution of a second order linear differential equation. Although not solved in closed form, this equation allows to state that the law of p/n converges to a Dirac mass as n grows. The generating function of the expected value of p is obtained in closed form. On the other hand, the analysis of the distribution of T leads to an integral, nonlinear and non-local functional equation which seems untractable.

The study of the discrete time problem does not lead to differential equations, but to harmonic and non-local functional equations. The "closed form" expressions obtained for the double generating function of p and n are extremely complicated, but the asymptotics of p when n is large can still be derived, involving the solution of linear non-local equations. As in the continuous case, the expected value of p is asymptotically proportional to n . The analysis of the time T leads to an interesting non linear and non local functional recurrence, the analysis of which shows that the expected value of T is of order of $\log n$. This same logarithmic growth is then shown to hold in the continuous case.

Note that the "branching" nature of the problem indicates that the analysis may apply to some combinatorial problems, as for instance the size and the height of certain random trees which arise in searching/sorting algorithms or in conflict resolution protocols (see [FLA] again for a survey on these topics).

Future research on the subject should turn to the analysis of (in some sense) improved version of the algorithm. For instance, it is shown in [KAI], using the technique of [CHA], that killing the Centaurs who arrive in a Tower with a larger identity than their own, leads to a mean letter complexity of $n \log(p)$, p still being the number of Centaurs born on the ring. It is also mentioned that the worst time complexity is increased. One can easily conjecture that the mean time complexity also increases, and it would be interesting to know how much it does.

ACKNOWLEDGEMENTS

We thank C. Kaiser for having provided us this interesting performance evaluation problem, in the recreative form of a Centaur Game. We also thank C. Lavault for his technical remarks. The study of the asymptotics of $K(\rho)$ is largely indebted to G. Fayolle.

REFERENCES

- [CHA] CHANG E. and ROBERTS R., "An Improved Algorithm for Decentralized Extrema-Finding in Circular Configuration of Processes", *Comm. ACM*, Vol.22,5, May 1979, pp. 281-283.
- [FEL] FELLER W., "An Introduction to Probability Theory and its Applications", Vol 2, J. Wiley & sons, 1970.
- [FLA] FLAJOLET P., "Random Tree Models in the Analysis of Algorithms", Proc. PERFORMANCE'87, Brussels, Dec 7-9, 1987, North-Holland, P.-J. Courtois and G. Latouche Editors.
- [INC] INCE E.L., "Ordinary Differential Equations", *Dover Publications*, 1926.

- [JARE] JACQUET P. and REIGNER M., "*Normal Limiting Distribution of the Size of Tries*", Proc. PERFORMANCE'87, Brussels, Dec 7-9, 1987, North-Holland, P.-J. Courtois and G. Latouche Editors.
- [KAI] KAISER C., "*Election sur un Anneau par Parcours de Reconnaissance, Anneau Unidirectionnel ou Bidirectionnel*", in preparation, 1988.
- [LAV] LAVENTIEV M. and CHABAT B., "*Méthodes de la Théorie des Fonctions d'une Variable Complexe*", Editions MIR, 1977.
- [LEL] LE LANN G., "*Distributed Systems. Towards a Formal Approach*", IFIP Congress 1977, North Holland, 1977, pp. 185-187.

APPENDIX

A-1: A Tauberian Theorem

Assume that the function $F(z) = \sum_0^\infty f_n z^n$ is defined and analytic in the disc $|z| < 1$, and has an unique singularity on the unit circle, which is a pole at $z = 1$:

$$F(z) = \frac{\alpha}{(1-z)^p} + O(1) \quad z \rightarrow 1,$$

for some integer $p > 0$ and some constant α . Then, if the sequence f_n is positive and monotonic, we have:

$$f_n \sim \frac{\alpha n^{p-1}}{(p-1)!}, \quad n \rightarrow \infty,$$

for some constant α .

A-2: Stochastic Increasingness of the Sequence $\{N_n\}_n$.

We show that $\forall n \geq 1, N_n \leq_{st} N_{n+1}$, that is, $\forall k \geq 0, \mathbb{P}(N_n > k) \leq \mathbb{P}(N_{n+1} > k)$.

First, we construct a common probability space for sample paths of the Centaurs Race on lines of n and $n+1$ Towers. Then, the argument goes as follow: consider a realization ω of the Centaur race on a line of length $n+1$. Up until the moment when something happens (arrival or birth) in the $n+1$ th Tower, this realization is identical to some $\bar{\omega}$ of the Centaur race on a line of length n . After this moment, these two realization are still identical, but $N_{n+1}(\omega) = 1 + N_n(\bar{\omega})$. Thus, whether something happens in the last Tower or not, $N_{n+1}(\omega) \geq N_n(\bar{\omega})$, which implies the ordering.

A-3: Asymptotic Expansion of $K(\rho)$

The asymptotic analysis of the integral (3.8a) uses the so called *saddle point method*, a presentation of which can be found in [LAV]. We use here slightly weaker assumptions. We only expose the main steps of the method, and leave the details to the reader. First rewrite (3.8a), introducing $v = 1 - u$ and $\theta = \rho - 1$:

$$K(\theta) = (1 + \theta) \int_0^1 v^2 e^v e^{\theta[v + \log(1-v)]} dv.$$

The function $v \mapsto v + \log(1-v)$ is negative and strictly decreasing for $v \in [0, 1)$. This suggests that the principal part of the integral when $\theta \rightarrow \infty$ is located around $v = 0$. Indeed, we have:

$$\forall \delta > 0, \forall N > 0, \quad K(\theta) = (1 + \theta) \int_0^\delta v^2 e^v e^{\theta[v + \log(1-v)]} dv + o(\theta^{-N}).$$

Furthermore, this same function is equivalent to $-v^2/2$ near zero, which suggests the change of variable: $\tau = \phi(v) = \sqrt{-v - \log(1-v)}$. Then, noting $\psi = \phi^{-1}$, we have:

$$\forall \delta > 0, \forall N > 0, \quad K(\theta) = (1 + \theta) \int_0^{\phi(\delta)} e^{\psi(\tau)} \psi^2(\tau) \psi'(\tau) e^{-\tau^2 \theta} d\tau + o(\theta^{-N})$$

Now, ϕ is analytical around $v = 0$ and its derivative is different from zero, so that ψ admits a series expansion around 0, and therefore:

$$\forall N > 0, \quad e^{\psi(\tau)}\psi^2(\tau)\psi'(\tau) = \sum_{n=0}^N c_n \tau^n + o(\tau^N).$$

Therefore:

$$\forall N, \delta, \quad K(\rho) = (1 + \theta) \left[\sum_{n=0}^N c_n \int_0^{\phi(\delta)} \tau^n e^{-\tau^2 \theta} d\tau + o(\theta^{-N}) + \int_0^{\phi(\delta)} o(\tau^N) e^{-\tau^2 \theta} d\tau \right].$$

We choose now to set $\delta = \theta^{1/3}$, so that $\theta\phi^2(\delta) \sim \theta^{1/3} \rightarrow \infty$ and then:

$$\begin{aligned} K(\rho) &= (1 + \theta) \sum_{n=0}^N \frac{c_n}{2} \theta^{-\frac{n+1}{2}} \int_0^{\theta\phi(\delta)^2} e^{-\sigma} \sigma^{\frac{n-1}{2}} d\sigma + o(\theta^{-\frac{N+1}{2}}) \\ &= (1 + \theta) \sum_{n=0}^N \frac{c_n}{2} \Gamma\left(\frac{n+1}{2}\right) \theta^{-\frac{n+1}{2}} + o(\theta^{-\frac{N+1}{2}}). \end{aligned}$$

Some easy computations show that $c_0 = c_1 = 0$, $c_2 = 2\sqrt{2}$ and $c_3 = -4/3$, which leads to the desired asymptotic expansion.

A-4: Summations of 4/

All the computations of 4/ derive from the following summations. If $A(t) = \sum_{i=1}^{\infty} a_i t^i$, then:

$$\sum_{n=1}^{\infty} t^n \sum_{k=1}^n \nu^k \beta^{n-k} \sum_{\substack{l_0 + \dots + l_k = n+1 \\ l_i > 0}} a_{l_1} \dots a_{l_{k-1}} a_{l_k + l_0 - 1} = \nu t \beta \frac{A'(\beta t)}{\beta - \nu A(\beta t)} \quad (A.1)$$

$$\sum_{n=1}^{\infty} t^n \sum_{k=0}^{n-1} \nu^k \beta^{n-k-1} \sum_{\substack{l_0 + \dots + l_k = n \\ l_i > 0}} a_{l_0} \dots a_{l_k} = \frac{A(\beta t)}{\beta - \nu A(\beta t)} \quad (A.2)$$

Let us first prove (A.2):

$$\begin{aligned} \sum_{n=1}^{\infty} (\dots) &= t \sum_{k=0}^{\infty} (\nu t)^k \sum_{n=k}^{\infty} (\beta t)^{n-k} \sum_{\substack{l_0 + \dots + l_k = n-k \\ l_i \geq 0}} a_{l_0} \dots a_{l_k} \\ &= t \sum_{k=0}^{\infty} (\nu t)^k \sum_{n=0}^{\infty} (\beta t)^n \sum_{\substack{l_0 + \dots + l_k = n \\ l_i \geq 0}} a_{l_0+1} \dots a_{l_k+1} \\ &= t \sum_{k=0}^{\infty} (\nu t)^k \sum_{n=0}^{\infty} \left[\frac{A(\beta t)}{\beta t} \right]^{k+1} \\ &= \frac{1}{\nu} \sum_{k=0}^{\infty} \left[\frac{\nu}{\beta} A(\beta t) \right]^{k+1}, \end{aligned}$$

QED. Now, for (A.1):

$$\begin{aligned}
\sum_{n=1}^{\infty} (\dots) &= \sum_{k=1}^{\infty} (\nu t)^k \sum_{n=0}^{\infty} (\beta t)^n \sum_{l_0+l_k=0}^n a_{l_k+l_0+1} \sum_{\substack{l_1+\dots+l_{k-1}=n-l_0-l_{k-1} \\ l_i \geq 0}} a_{l_1+1} \dots a_{l_{k-1}+1} \\
&= \left(\sum_{l_0+l_k=0}^{\infty} (\beta z)^{l_0+l_k} a_{l_k+l_0+1} \right) \sum_{k=1}^{\infty} (\nu t)^k \sum_{n=0}^{\infty} (\beta t)^n \sum_{\substack{l_1+\dots+l_{k-1}=n \\ l_i \geq 0}} a_{l_1+1} \dots a_{l_{k-1}+1} \\
&= \left(\sum_{l=0}^{\infty} (l+1) (\beta t)^l a_{l+1} \right) \sum_{k=1}^{\infty} (\nu t)^k \left(\frac{A(t)}{\beta t} \right)^{k-1} \\
&= A'(\beta t) \frac{\nu \beta t}{\beta - \nu A(\beta t)},
\end{aligned}$$

All these computations apply to formal series, and to functions of the complex variable t in a domain where $\sum_1^{\infty} a_i t^i$ converges and where $|A(\beta t)| < \beta/\nu$.

□

A-5: Linear non-local equations

We show here how to solve the linear non local functional equations that arouse in section 4/. We do not address a general theory of these equations, as we know that our parameter and solution functions have good analyticity properties.

Let $\beta \in]0, 1[$. Let $a(z)$, $b(z)$ and $F(z)$ be analytic functions inside the unit disk, satisfying the equation:

$$F(z) = a(z) + b(z) F(\beta z).$$

Replacing recursively $F(\beta z)$ by its value, we have $\forall n \geq 0$,

$$F(z) = \sum_{i=0}^n a(\beta^i z) \prod_{j=0}^{i-1} b(\beta^j z) + F(\beta^{n+1} z) \prod_{j=0}^n b(\beta^j z),$$

where an empty product is 1 by convention.

The idea is to let n go to infinity in this expression. Due to the (strong) hypothesis we made on a , b and F (analyticity), it suffices that $|b(0)| < 1$, in order that:

- i/ the remainder goes to zero
 - ii/ the series converges absolutely for all z in the unit disk.
- Therefore: $|b(0)| < 1$ implies:

$$F(z) = \sum_{i=0}^{\infty} a(\beta^i z) \prod_{j=0}^{i-1} b(\beta^j z). \quad (\text{A.3})$$

□

Unfortunately, these condition is not satisfied in the equations we have to solve in 4/, as $b(0)$ is $(1-\tau)/\beta$ in (4.4) and $(1-\tau)^2/\beta$ in (4.14). However, we have:

If $a(0) = 0$ and $|b(0)| < 1/\beta$, then i/, ii/ and (A.3) hold.

Proof: if $a(0) = 0$, then $F(0) = 0$. Thus, the remainder rewrites:

$$\frac{F(\beta^{n+1}z)}{\beta^{n+1}} \prod_{j=0}^n \frac{b(\beta^j z)}{\beta}$$

As the first factor converges to $F'(0)$ and the second shrinks to zero, this remainder vanishes. Similarly, the series converges absolutely in the unit disk.

□




Fig 1




Fig 2




Fig 3: $\epsilon = 0$




Fig 4: $\epsilon = 1$




Fig 5: $\zeta = 0.5$






Fig 6: $\zeta = 1$



2 dir.



1 dir.

Fig 7

