The Swedish Leader Election Protocol: Analysis and Variations

Guy Louchard
U. Libre Bruxelles
Belgium

Conrado Martínez
U. Politècnica Catalunya
Spain

Helmut Prodinger
U. Stellenbosch
South Africa

ANALCO, San Francisco, January 2011

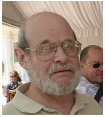

G. Louchard

H. Prodinger

Introduction

- A common task in distributed computing is to choose a leader among n agents in a decentralized manner
- A typical protocol requires each agent flipping a biased coin: if the outcome is heads (with probability q), proceed to next round, if the outcome is tails (with probability $p = 1 - q$) the agent gets out of the process
- If a single agent “survives” after a certain number of rounds, it is declared the leader

Introduction

Introduction

Introduction

Introduction

Introduction

Introduction

Introduction

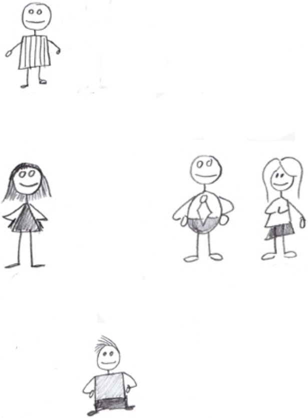

Introduction

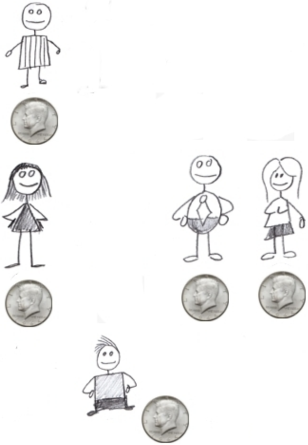

Introduction

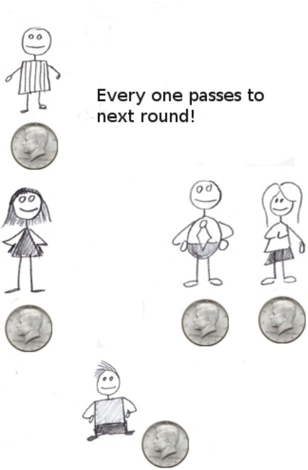

Introduction

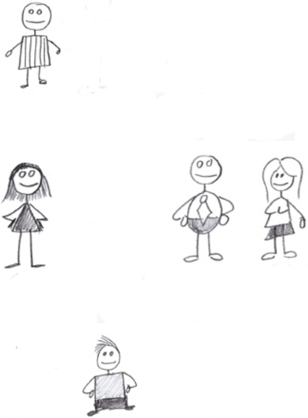

Introduction

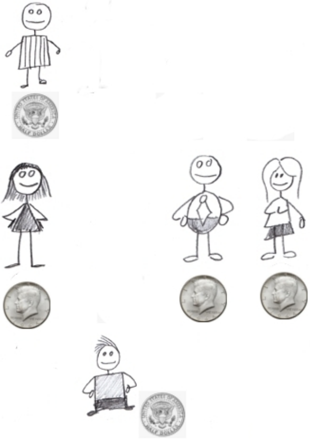

Introduction

Introduction

Introduction

Introduction

Introduction

Introduction

Introduction

Introduction

Example

- $n = 14$
- Number of rounds = 6 (R_n)
- Number of coin flippings = 14+8+5+5+3+3=28 (F_n)
- Number of stalled rounds = 1
- Number of null rounds = 1 (I_n)

Introduction

- The Swedish election protocol introduces a new parameter τ, the maximum number of consecutive null rounds
- If more than τ null rounds occur in a row the protocol fails to declare a leader
- Other variants might restrict the maximum number of consecutive stalled rounds, the total number of null rounds (consecutive or not), etc.
- In a practical setting, bounding the number of stalled rounds corresponds to setting time-outs

Introduction

- The model is inspired in the k-silent elimination protocol by the Swedish researchers L. Bondesson, T. Nilsson and G. Wikstrand
- The case $\tau \to \infty$ is the classical Leader Election Protocol
- R. Kalpathy and H. Mahmoud have investigated a similar problem with $\tau = 1$

Introduction

- We use standard techniques (analytic Poissonization-depoissonization, Mellin transforms, etc.) to analyze the protocol
- The asymtotic analysis of the quantities of interest involves these unknown quantities! E.g., the probability of success $S_n := S_n(\tau)$ is

$$S_n = C(q, \tau) + \delta(\log_Q n) + O(1/n), \qquad \text{as } n \to \infty,$$

where $Q = 1/q$, $L = \log Q$,

$$C(q, \tau) = \frac{1}{L}\left(qp^{\tau} + \sum_{k>0} \frac{S_k}{k}\left(p^k - \frac{q^k p^{\tau k}}{(1-p^{\tau+1})^k}\right)\right),$$

and $\delta(x)$ is a periodic function of "small" amplitude (depending on q and τ) and period 1, also involving the unknown S_k's

Introduction

- But only a few exact values of the unknowns (of S_k in the example) are actually needed to get useful asymptotic estimates, as the error term that we incur if we discard all but the first few terms in the summations is very small
- For all practical purposes, it suffices to compute S_n exactly for n up to, say, $N = 20$, using the exact recurrence and use the approximation given by the first N terms of the summation

Introduction

Dashed lines: Exact value of S_n
Solid lines: Approximation $C + \delta$ with $N = 20$ terms

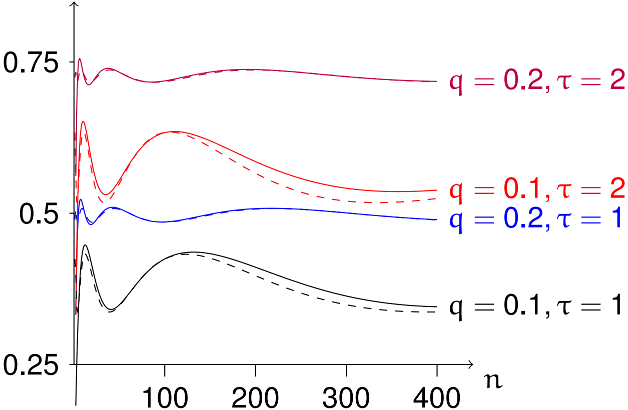

What's next

1. A sketch of the methods: computing the probability of success of the protocol
2. Other results
3. Final remarks

Probability of success

- τ: maximum number of consecutive null rounds; if there are $\tau + 1$ consecutive null rounds, the protocol fails
- $S_n(t)$: probability of success of the protocol when $t - 1$ additional consecutive null rounds is allowed; $S_n := S_n(\tau)$
- $S_n(0) = 0$ if $n \geqslant 2$; $S_1(t) = 1$ if $t > 0$

Probability of success

The recurrence for $S_n(t)$ when $n \geqslant 2$ and $t > 0$:

$$S_n(t) = \sum_{1 \leqslant j \leqslant n} \binom{n}{j} p^{n-j} q^j S_j(\tau) + p^n S_n(t-1), \qquad t > 0, n \geqslant 2,$$

where q is the probability of heads (agents passes to next round)

Probability of success

- Define $K_n(\tau) = \sum_{1 \leqslant j \leqslant n} \binom{n}{j} p^{n-j} q^j S_j(\tau)$, hence

$$S_n(t) = K_n(\tau) + p^n S_n(t-1) = K_n(\tau) + p^n K_n(\tau) + p^{2n} S_n(t-2)$$
$$= \cdots = K_n(\tau)\big(1 + p^n + p^{2n} + \cdots + p^{(k-1)n}\big) + p^{kn} S_n(t-k)$$

- Therefore, for $S_n := S_n(\tau)$

$$S_n = \frac{1 - p^{\tau n}}{1 - p^n} \sum_{j=1}^{n} \binom{n}{j} p^{n-j} q^j S_j, \qquad n \geqslant 2,$$

and $S_1 = 1$.

Probability of success: The “pipeline”

$$\begin{tikzcd}
S_n = \cdots \arrow[d, "\boxed{S(z) = \sum_{n \geqslant 0} S_n \frac{z^n}{n!}}"'] \\
S(z) = \cdots
\end{tikzcd}$$

Probability of success: The “pipeline”

$$\begin{tikzcd}
S_n = \cdots \arrow[d, "S(z) = \sum_{n \geqslant 0} S_n \frac{z^n}{n!}"'] \\
S(z) = \cdots \arrow[d, "\text{Poissonize: } \hat{S}(z) = e^{-z} S(z)"'] \\
\hat{S}(z) = \cdots
\end{tikzcd}$$

Probability of success: The “pipeline”

$$\begin{tikzcd}
S_n = \cdots \arrow[d, "S(z) = \sum_{n \geqslant 0} S_n \frac{z^n}{n!}"'] & & \\
S(z) = \cdots \arrow[d, "\text{Poissonize: } \hat{S}(z) = e^{-z} S(z)"'] & & \\
\hat{S}(z) = \cdots \arrow[rr, bend right, "\text{Mellin: } S^*(s) = \int_0^\infty \hat{S}(z) z^{s-1} \, dz"'] & & S^*(s) = \cdots
\end{tikzcd}$$

Probability of success: The “pipeline”

$$\begin{tikzcd}
S_n = \cdots \arrow[d, "S(z) = \sum_{n \geqslant 0} S_n \frac{z^n}{n!}"'] & \\
S(z) = \cdots \arrow[d, "\text{Poissonize: } \hat{S}(z) = e^{-z} S(z)"'] & \hat{S}(z) \sim \cdots \ (z \to \infty) \\
\hat{S}(z) = \cdots \arrow[r, bend right, "\text{Mellin: } S^*(s) = \int_0^\infty \hat{S}(z) z^{s-1} \, dz"'] & S^*(s) = \cdots \arrow[u, "\text{Invert Mellin: } \hat{S}(z) = \frac{1}{2\pi \mathbf{i}} \int_{-c-\mathbf{i}\infty}^{-c+\mathbf{i}\infty} S^*(s) z^{-s} \, ds"]
\end{tikzcd}$$

Probability of success: The “pipeline”

$$\begin{tikzcd}
S_n = \cdots \arrow[d, "S(z) = \sum_{n \geqslant 0} S_n \frac{z^n}{n!}"'] & S_n \sim \hat{S}(n) + \text{l.o.t} \\
S(z) = \cdots \arrow[d, "\text{Poissonize: } \hat{S}(z) = e^{-z} S(z)"'] & \hat{S}(z) \sim \cdots \ (z \to \infty) \arrow[u, "\text{Depoissonize}"] \\
\hat{S}(z) = \cdots \arrow[r, bend right, "\text{Mellin: } S^*(s) = \int_0^\infty \hat{S}(z) z^{s-1}\, dz"'] & S^*(s) = \cdots \arrow[u, "\text{Invert Mellin: } \hat{S}(z) = \frac{1}{2\pi i} \int_{-c-i\infty}^{-c+i\infty} S^*(s) z^{-s}\, ds"]
\end{tikzcd}$$

Probability of success

Step #1: Translate the recurrence into a functional equation over the EGF

$$S_n = \frac{1 - p^{\tau n}}{1 - p^n} \sum_{j=1}^{n} \binom{n}{j} p^{n-j} q^j S_j, \qquad n \geqslant 2$$

$$\Downarrow$$

$$S(z) - S(pz) = e^{pz} S(qz) - e^{p^{\tau+1} z} S(qp^\tau z) + qp^\tau z$$

with $S(z) = \sum_{n \geqslant 0} S_n \, z^n / n!$

Probability of success

Step #2: Poissonize

$$S(z) - S(pz) = e^{pz} S(qz) - e^{p^{\tau+1} z} S(qp^{\tau} z) + qp^{\tau} z$$

$$\Downarrow$$

$$\hat{S}(z) - e^{-z} S(pz) = \hat{S}(qz) - e^{-z(1-p^{\tau+1})} S(qp^{\tau} z) + qp^{\tau} z e^{-z}$$

with $\hat{S}(z) = e^{-z} S(z)$

Probability of success

Step #3: “Mellinize”

$$\hat{S}(z) - \hat{S}(qz) = e^{-z}S(pz) - e^{-z(1-p^{\tau+1})}S(qp^{\tau}z) + qp^{\tau}ze^{-z}$$

$$\Downarrow$$

$$S^*(s) = \frac{1}{1-q^{-s}}\Big(qp^{\tau}\Gamma(s+1) \\ + \mathcal{M}\left\{e^{-z}S(pz) - e^{-z(1-p^{\tau+1})}S(qp^{\tau}z); s\right\}\Big)$$

with $S^*(s) = \mathcal{M}\left\{\hat{S}(z); s\right\} = \int_0^{\infty} \hat{S}(z)z^{s-1}\,dz$

Probability of success

Step #4: "Demellinize" (via residue computations)

$$S^*(s) = \frac{1}{1-q^{-s}}\Big(qp^{\tau}\Gamma(s+1) \\ + \mathcal{M}\left\{e^{-z}S(pz) - e^{-z(1-p^{\tau+1})}S(qp^{\tau}z); s\right\}\Big)$$

$$\Downarrow$$

$$\hat{S}(z) = \frac{1}{2\pi\mathbf{i}}\int_{-\frac{1}{2}-\mathbf{i}\infty}^{-\frac{1}{2}+\mathbf{i}\infty} S^*(s)z^{-s}\,ds \\ = \sum_{\text{poles } \sigma} \text{Res}(S^*(s)z^{-s}; s=\sigma) + \text{error terms}$$

Probability of success

Step #5: Depoissonize (check that conditions on growth rate of $\hat{S}(z)$ are met to apply analytic depoissonization)

$$
\begin{aligned}
S_n &\sim \hat{S}(n) + O(1/n) \\
&= \frac{1}{L}\left(qp^{\tau} + \sum_{k>0} \frac{S_k}{k}\left(p^k - \frac{q^k p^{\tau k}}{(1-p^{\tau+1})^k} \right) \right) + \delta(\log_Q n) \\
&\quad + O(1/n)
\end{aligned}
$$

$$
\begin{aligned}
\delta(x) = \frac{1}{L}\sum_{j\neq 0} e^{-2x\pi \mathbf{i} j}\Big(& qp^{\tau}\Gamma(\chi_j+1) \\
& + \sum_{k>0} \frac{S_k}{k!}\Gamma(\chi_j+k)\big(p^k - \frac{q^k p^{\tau k}}{(1-p^{\tau+1})^{\chi_j+k}}\big)\Big),
\end{aligned}
$$

where $Q = 1/q$ and $L = \log Q$.

Probability of success

The asymptotic estimate for S_n involves the sequence S_n itself. The trivial bound $S_n \geqslant 1$ can be used to easily show

1. The sum
$$\sum_{k>0} \frac{S_k}{k} \left(p^k - \frac{q^k p^{\tau k}}{(1-p^{\tau+1})^k} \right)$$
that appears in the constant term of S_n converges
2. The error term when we take only the first N terms of this sum in numerical computations is $O(p^N)$; it suffices thus to use a few terms to get very precise asymptotic estimations

Probability of success

Similar arguments apply with regard the fluctuation $\delta(x)$, which has “small amplitude”, mean 0 and period 1.
For other quantities we need similar considerations about the convergence of infinite sums appearing in the asymptotic estimates, and about error terms when we use a few terms of those infinite sums to do the numerical computations

Other results

We have applied the same methodology to investigate the expectation of several other parameters, which in all cases satisfy the recurrence:

$$X_n(t) = \sum_{j=1}^{n} \binom{n}{j} p^{n-j} q^j X_j(\tau) \\ + p^n X_n(t-1) + T_n, \qquad t > 0, n \geqslant 2,$$

for some suitably chosen toll sequence T_n

The toll sequence, together with the initial values $X_1 = X_1(\tau)$, T_1 and $X_n(0)$, characterize many different parameters of the protocol

Other results

1. Number of rounds R_n: $T_n = 1$ if $n \geqslant 2$,
   $R_1 = R_n(0) = T_1 = 0$
2. Number of null rounds I_n: $T_n = p^n$ if $n \geqslant 2$,
   $T_1 = I_1 = I_n(0) = 0$
3. Number of flipped coins F_n: $T_n = n$ if $n \geqslant 2$,
   $T_1 = F - 1 = F_n(0) = 0$
4. Leftovers L_n: $T_n = 0$, $L_1 = 0$, $L_n(0) = n$ if $n \geqslant 2$

Leftovers are the players still active in the last non-null round when the protocol stops (either with success or with a failure)

Other results

- Many steps of the “pipeline” can be applied in a generic way
- The fundamental strip of definition of the Mellin transform will differ from one problem to the other
- And so the poles for residue computation, error terms, etc.

Other results

- Number of rounds

$$R_n = \log_Q n + \frac{\gamma}{L} + \frac{1}{2} - \frac{1}{L} + \frac{1}{L}\big(p^{\tau} + \log(1 - p^{\tau})\big) \\ + C(R; q, \tau) + \delta_R(\log_Q n) + O(n^{-1} \log n).$$

- Number of null rounds

$$I_n = 1 - \frac{p}{L} + \frac{1}{L}\big(p^{\tau+1} + \log(1 - p^{\tau+1})\big) \\ + C(I; q, \tau) + \delta_I(\log_Q n) + O(1/n).$$

Other results

- Number of coin flips

$$F_n = \frac{n}{p} + O(1).$$

- Number of leftovers

$$L_n = \frac{1}{L}\left(\frac{1}{1-p^\tau} - \frac{1}{1-p^{\tau+1}} - p^\tau + p^{\tau+1}\right) + C(L; q, \tau) + \delta_L(\log_Q n) + O(1/n).$$

Other results

- Three of the parameters involve the constant terms of the form

$$C(A;q,\tau) := \frac{1}{L}\sum_{k\geqslant 1}\frac{A_k}{k}\left(p^k - \frac{q^k p^{\tau k}}{(1-p^{\tau+1})^k}\right)$$

- They also involve fluctuations $\delta_R(x)$, $\delta_I(x)$ and $\delta_L(x)$ whose Fourier coefficients can be explicitly computed

Final remarks

- The standard analytic Poissonization-Depoissonization works well to analyze this extension of the classical leader election protocol; this new protocol is of independent interest because of potential practical applications
- Other extensions, like restrictions on the number of consecutive stalled rounds, total number of null rounds (consecutive or not), etc. can also be analyzed using the same methodology
- The asymptotic analyis can be carried out even without an explicit solution for the Mellin transform

Final remarks

- We are now working a longer journal version with our new results about the probability distributions of several of the parameters discussed here
- Moreover, we derive results for the probability distributions conditioned on success and on failure of the protocol

Thank you for your attention!