# Convergence of Leader Election Algorithms

Christian LAVAULT

LIPN (CNRS UMR 7030), Un. Paris 13 France
(joint work with Svante JANSON and Guy LOUCHARD)

Quantum & Combinatorics – ZAKOPANE 2009

# *Franklin's LE Algorithm on a ring*

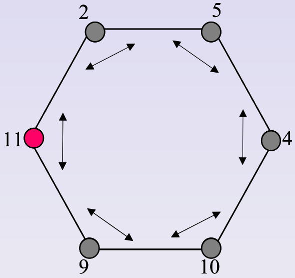


Non Oriented Ring [Franklin-82] & Oriented Ring [Dolev *et al.*-82]

# *Franklin's LE Algorithm on a ring*

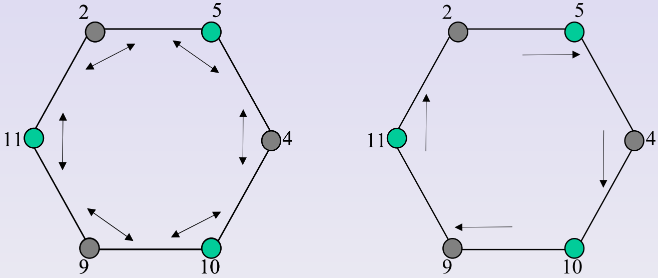


Non Oriented Ring [Franklin-82] & Oriented Ring [Dolev *et al.*-82]

# *Franklin's LE Algorithm on a ring*

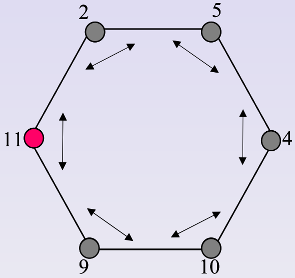


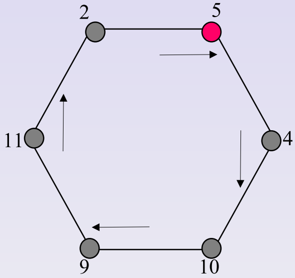


Non Oriented Ring [Franklin-82] & Oriented Ring [Dolev *et al.*-82]

# *Analysis of Franklin's Algorithm*

Worst/Average-case: $2n$ messages in each round ('virtual time') plus $n$ for termination.

- Worst-case message and 'time' complexity:
  $M_{max}(n) = 2n \times$ maximal # of rounds,
  i.e. whenever the maximal # of peaks is $\lfloor n/2 \rfloor$.

$$\text{maximal \# of rounds} = \lfloor \log_2(n) \rfloor + 1$$

$$M_{max}(n) = 2n(\lfloor \log_2(n) \rfloor + 1) = \Theta(n \log n).$$

# *Analysis of Franklin's Algorithm*

- Worst-case message and 'time' complexity:
  $M_{max}(n) = 2n \times$ maximal # of rounds,
  i.e. whenever the maximal # of peaks is $\lfloor n/2 \rfloor$.

$$\text{maximal \# of rounds} = \lfloor \log_2(n) \rfloor + 1$$

$$M_{max}(n) = 2n(\lfloor \log_2(n) \rfloor + 1) = \Theta(n \log n).$$

- Average message complexity ($n$ candidates)

Mean $\mathbb{M}$ and variance $\mathbb{V}$ of the number of peaks in a permutation $\sigma \in \mathbb{S}_n$ [Carlitz-74], [Flajolet,Sedgewick-09]:

$$\mathbb{M}(n) = (n-2)/3 \quad (n \geq 2) \quad \text{et} \quad \mathbb{V}(n) = 2(n+1)/45 \quad (n \geq 4).$$

# Analysis of Franklin's Algorithm (cont.

Distribution on a ring: $P_c(n,k) = \mathbb{P}(n \text{ candidates}, k \text{ peaks})$.

Distribution of the $k$ peaks in a permutation $\sigma \in \mathbb{S}_n$: $P(n,k)$.

$\rightsquigarrow$ Simple recurrence $P_c(n,k) = P_\ell(n-1,k-1)$ yields

$$\mathbb{M}_c(n) = n/3 \quad (n \geq 3) \quad \text{and} \quad \mathbb{V}_c(n) = 2n/45 \quad (n \geq 5).$$

# *Analysis of Franklin's Algorithm (cont.)*

Problem! The remaining peaks (candidates) compare their *Ids* from round to round, till the leader is elected. By conditioning on the number $m$ of surviving peaks in the end of round 1, the $m!$ distinct orderings of peaks within round 2 do not have the same probability.

With $n = 8$ and $m = 4$ at the end of round 1, the probability that 2 peaks are survivors in round 2 is $10/34$, whereas it is $1/3$ in the uniform case…

# *Analysis of Franklin's Algorithm (cont.)*

Problem! The remaining peaks (candidates) compare their *Ids* from round to round, till the leader is elected. By conditioning on the number $m$ of surviving peaks in the end of round 1, the $m!$ distinct orderings of peaks within round 2 do not have the same probability.

With $n = 8$ and $m = 4$ at the end of round 1, the probability that 2 peaks are survivors in round 2 is $10/34$, whereas it is $1/3$ in the uniform case...

Mean number of surviving peaks after 2 rounds is $c_2 n + o(1)$ ($c_2 \approx 0,1096868681 \lesssim 1/9$).
Franklin is not asymptotically equivalent to its variant with permutations in $\mathbb{S}_n$ ... [Janson,CL,Louchard-09]

# *Uniform Variant of Franklin's Algorithm*

Simplify the process…
In each round, each surviving peak draws a new *Id* independently at random .
The peaks' distribution remains unchanged from the 1st to last round.

$$\text{mean \# of rounds} = \log_3(n) + \phi(n) + o(1),$$

where $\phi(n)$ is periodic and $|\phi(n)| < 1$. Thus,

$$\mathbb{E}(M_n) = 2n\log_3(n) + \mathcal{O}(n) = \Theta(n\log n).$$

□

# Rules of the Game

Given some random procedure, for any set of $n \geq 2$ individuals it eliminates some (but not all) individuals.

If there is more that one survivor, repeat the procedure with the survivors until only one (the winner) remains.

What is the (random) number $X_n$ of rounds required?
($X_1 = 0$, $X_n \geq 1$ for $n \geq 2$)

# *Rules of the Game*

Given some random procedure, for any set of $n \geq 2$ individuals it eliminates some (but not all) individuals.

If there is more that one survivor, repeat the procedure with the survivors until only one (the winner) remains.

What is the (random) number $X_n$ of rounds required?
($X_1 = 0$, $X_n \geq 1$ for $n \geq 2$)

Let $N_k$ be the number of individuals remaining after round $k$:

$$X_n := \min\{k : N_k = 1\},$$

starting with $N_0 = n$ ($N_k$ is defined for all $k \geq 0$ and $N_k = 1$ for all $k \geq X_n$.)

Assume that the number $Y_n$ of survivors of a set of $n$ individuals has a distribution depending only on $n$.
We have $1 \leq Y_n \leq n$ and allow that $Y_n = n$, but we assume $\mathbb{P}(Y_n = n) < 1$ for every $n \geq 2$.

Also assume that the sequence $(N_k)_0^\infty$ is a Markov chain on $\{1, 2, \ldots\}$ and $X_n$ is the number of steps to absorption in 1.

Assume that the number $Y_n$ of survivors of a set of $n$ individuals has a distribution depending only on $n$.
We have $1 \leq Y_n \leq n$ and allow that $Y_n = n$, but we assume $\mathbb{P}(Y_n = n) < 1$ for every $n \geq 2$.

Also assume that the sequence $(N_k)_0^\infty$ is a Markov chain on $\{1, 2, \ldots\}$ and $X_n$ is the number of steps to absorption in 1.

The transition probabilities of this Markov chain are $(Y_1 = 1)$

$$P(i,j) := \mathbb{P}(Y_i = j) = \mathbb{P}(j \text{ survives of a set of } i).$$

$P(i,j) = 0$ if $j > i$ and $P(i,i) < 1$, for $i > 1$.

Assume that the number $Y_n$ of survivors of a set of $n$ individuals has a distribution depending only on $n$.
We have $1 \le Y_n \le n$ and allow that $Y_n = n$, but we assume $\mathbb{P}(Y_n = n) < 1$ for every $n \ge 2$.

Also assume that the sequence $(N_k)_0^\infty$ is a Markov chain on $\{1, 2, \dots\}$ and $X_n$ is the number of steps to absorption in 1.

The transition probabilities of this Markov chain are $(Y_1 = 1)$

$$P(i,j) := \mathbb{P}(Y_i = j) = \mathbb{P}(j \text{ survives of a set of } i).$$

$P(i,j) = 0$ if $j > i$ and $P(i,i) < 1$, for $i > 1$.

LE algorithms are such that, asymptotically, a fixed proportion is eliminated in each round ($X_n$ is expected to be $\mathcal{O}(\log n)$).

# *Conditions on* $Y_n$

For every $n \geq 1$, $Y_n$ is a random variable s.t. $1 \leq Y_n \leq n$, and $\mathbb{P}(Y_n = n) < 1$ for $n \geq 2$. Further,

- *(i)* $Y_n$ is stochastically increasing in $n$:
  $\mathbb{P}(Y_n \leq k) \geq \mathbb{P}(Y_{n+1} \leq k)$ for all $n \geq 1$ and $k \geq 1$.
  (Equivalently, we may couple $Y_n$ and $Y_{n+1}$ such that $Y_n \leq Y_{n+1}$.)

# *Conditions on* $Y_n$

For every $n \geq 1$, $Y_n$ is a random variable s.t. $1 \leq Y_n \leq n$, and $\mathbb{P}(Y_n = n) < 1$ for $n \geq 2$. Further,

- *(i)* $Y_n$ is stochastically increasing in $n$: $\mathbb{P}(Y_n \leq k) \geq \mathbb{P}(Y_{n+1} \leq k)$ for all $n \geq 1$ and $k \geq 1$. (Equivalently, we may couple $Y_n$ and $Y_{n+1}$ such that $Y_n \leq Y_{n+1}$.)
- *(ii)* For some constants $\alpha \in (0, 1)$ and $\epsilon > 0$, and a sequence $\delta_n = \mathcal{O}\left((\log n)^{-1-\epsilon}\right)$,

$$\mathbb{E}(Y_{n+1}) - \mathbb{E}(Y_n) = \alpha + \mathcal{O}(\delta_n).$$

# *Conditions on* $Y_n$

For every $n \geq 1$, $Y_n$ is a random variable s.t. $1 \leq Y_n \leq n$, and $\mathbb{P}(Y_n = n) < 1$ for $n \geq 2$. Further,

- *(i)* $Y_n$ is stochastically increasing in $n$: $\mathbb{P}(Y_n \leq k) \geq \mathbb{P}(Y_{n+1} \leq k)$ for all $n \geq 1$ and $k \geq 1$. (Equivalently, we may couple $Y_n$ and $Y_{n+1}$ such that $Y_n \leq Y_{n+1}$.)
- *(ii)* For some constants $\alpha \in (0, 1)$ and $\epsilon > 0$, and a sequence $\delta_n = \mathcal{O}\left((\log n)^{-1-\epsilon}\right)$,

$$\mathbb{E}(Y_{n+1}) - \mathbb{E}(Y_n) = \alpha + \mathcal{O}(\delta_n).$$

- *(iii)* For some $\epsilon$ and $\delta_n$ as in *(ii)*,

$$\mathbb{P}(|Y_n - \alpha n| > n\delta_n) = \mathcal{O}\left(n^{-2-\epsilon}\right).$$

# *Total Variation $d_{TV}$ and Wasserstein Distance $d_W$*

Total variation distance $d_{TV}$ between two arbitrary r.v. $X$ and $Y$:

$$d_{TV}(X,Y) := 1/2 \sum_k |\mathbb{P}(X=k) - \mathbb{P}(Y=k)|$$

and, for any two distributions $\mu$ and $\nu$,

$$d_{TV}(\mu,\nu) := \inf \{\mathbb{P}(X \neq Y) : X \sim \mu,\, Y \sim \nu\},$$

where inf is taken over all over all couplings $(X,Y)$ of $\mu$ and $\nu$.

# *Total Variation $d_{TV}$ and Wasserstein Distance $d_W$*

Total variation distance $d_{TV}$ between two arbitrary r.v. $X$ and $Y$:

$$d_{TV}(X,Y) := 1/2 \sum_k |\mathbb{P}(X=k) - \mathbb{P}(Y=k)|$$

and, for any two distributions $\mu$ and $\nu$,

$$d_{TV}(\mu,\nu) := \inf \{\mathbb{P}(X \neq Y) : X \sim \mu,\ Y \sim \nu\},$$

where inf is taken over all over all couplings $(X, Y)$ of $\mu$ and $\nu$.

$d_W$ is defined for proba. distributions with finite expectation:

$$d_W(\mu,\nu) := \inf \{\mathbb{E}|X - Y| : X \sim \mu,\ Y \sim \nu\}.$$

# *Total Variation $d_{TV}$ and Wasserstein Distance $d_W$*

Total variation distance $d_{TV}$ between two arbitrary r.v. $X$ and $Y$:

$$d_{TV}(X, Y) := 1/2 \sum_k |\mathbb{P}(X = k) - \mathbb{P}(Y = k)|$$

and, for any two distributions $\mu$ and $\nu$,

$$d_{TV}(\mu, \nu) := \inf \{\mathbb{P}(X \neq Y) : X \sim \mu,\ Y \sim \nu\},$$

where inf is taken over all over all couplings $(X, Y)$ of $\mu$ and $\nu$.

$d_W$ is defined for proba. distributions with finite expectation:

$$d_W(\mu, \nu) := \inf \{\mathbb{E}|X - Y| : X \sim \mu,\ Y \sim \nu\}.$$

For integer-valued r.v. $X$ and $Y$ (but not in general),

$$d_{TV}(X, Y) \leq d_W(X, Y).$$

# *Theorem*

Consider the LE algorithm described above. There exists a distribution function $F$ with bounded density function $f = F'$ s.t.

$$\sup_{k\in\mathbb{Z}} |\Pr(X_n \leq k) - F(k - \log_{1/\alpha}(n))| \to 0$$

or, equivalently, if $Z \sim F$,

$$d_{TV}(X_n, \lceil Z + \log_{1/\alpha}(n) \rceil) \to 0.$$

More precisely, $d_W(X_n, \lceil Z + \log_{1/\alpha}(n) \rceil) \to 0$ and

$$\sum_{k\in\mathbb{Z}} |\mathbb{P}(X_n \leq k) - F(k - \log_{1/\alpha}(n))| \to 0.$$

# *Theorem (cont.)*

As a consequence, defining $\Delta F(x) := F(x) - F(x-1)$,

$$\sup_{k \in \mathbb{Z}} \left| \mathbb{P}(X_n = k) - \Delta F(k - \log_{1/\alpha}(n)) \right| \to 0.$$

Furthermore,

$$\mathbb{E} X_n = \log_{1/\alpha}(n) + \phi(n) + o(1),$$

for a continuous function $\phi(t)$ on $(0, \infty)$ which is periodic in $\log_{1/\alpha}(t)$ ($\phi(t) = \phi(\alpha t)$), and locally Lipschitz.

# *Conclusion*

We thus do not have convergence in distribution as $n \to \infty$, but the usual type of oscillations with an asymptotic periodicity in $\log_{1/\alpha}(n)$ and convergence in distribution along subsequences such that the fractional part $\{\log_{1/\alpha}(n)\}$ converges.

This phenomenon is well-known for many other problems with integer-valued r.v. [Louchard-Prodinger08,Janson06]; it happens frequently when the variance stays bounded.