Time and Communication Complexity of Leader Election in Anonymous Networks

Dariusz R. Kowalski
Augusta Univ.

Miguel A. Mosteiro
Pace Univ.

ICDCS 2021

Leader Election

Some variants:

- Implicit vs. Explicit (who knows the leader)
- Irrevocable vs. Revocable (whether decision is final)
- Known vs. Unknown n

Ad-hoc Network Model

- **Static connected network :**
  - fixed set of n nodes and m links
  - there is a path between every pair of nodes
- **Network knowledge:**
  - no identifiers or labels, only port numbers
  - we consider known and unknown n
- **Synchronous communication :** in each round
  - a node broadcasts a message to its neighbors
  - receives the messages of its neighbors
  - executes some local computation
- **CONGEST communication :** in each round
  - O(log n) bits through each link

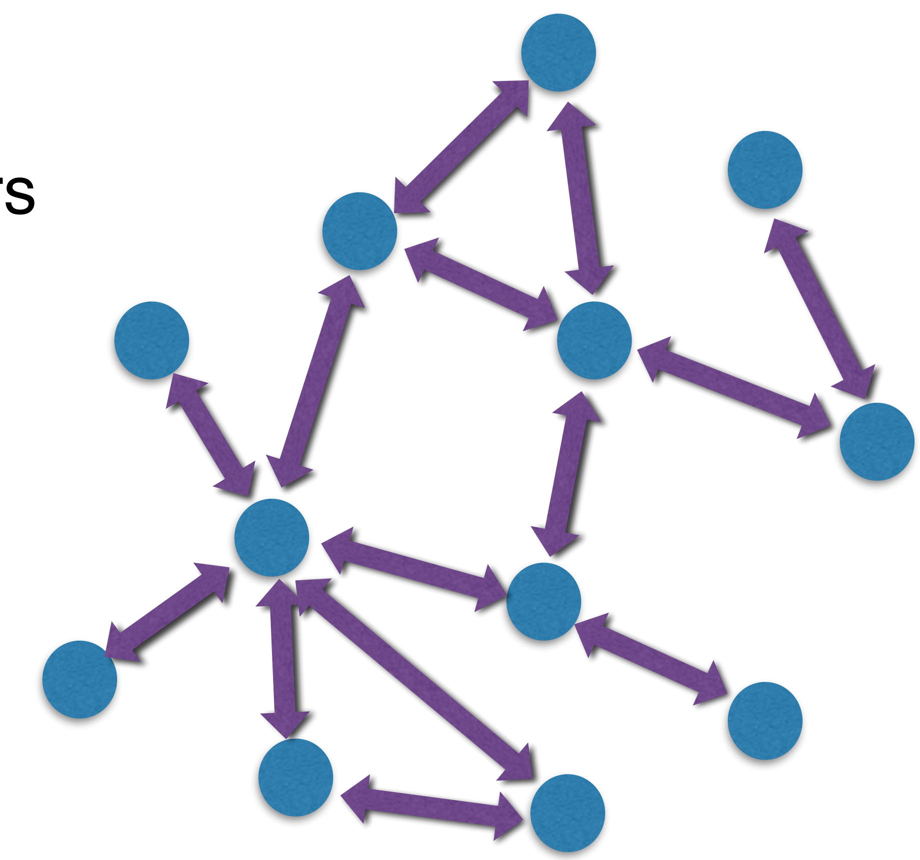

Algorithms

- Performance metric? time and message (energy) complexity.
- Deterministic LE? not possible in anonymous network [Angluin, STOC'80]
- Randomized LE? two scenarios:
  - Known n? ✔
  - Unknown n? ✗ we show that no algorithm stops
    - » So, how about Revocable LE with unknown n? ✔

Theorem 1. *For* $x \in \widetilde{\Theta}\left(\min\{\sqrt{\frac{n\log n}{\Phi t_{mix}}}, \frac{n}{t_{mix}\Phi\log n}\}\right)$, *the leader election algorithm elects a unique leader and uses* $\widetilde{O}(\min\{\sqrt{\frac{nt_{mix}}{\Phi}}, \frac{n}{\Phi\log n}\})$ *point to point messages/bits of communication in the CONGEST model with known (a linear upper bound on)* n, *whp. It works in time* $O(t_{mix}\log n + \Phi^{-1}\log^3 n)$.

Theorem 2. *For any non-decreasing positive integer function* $T(n)$ *and any constant* $0 < c < 1$, *there is no algorithm solving Leader Election problem in time* $T(n)$ *with probability* c, *in the setting without known number of nodes* n.

Theorem 3. *For* $0 < \epsilon \leq 1$ *and* $0 < \xi < 1$, *after running Blind Leader Election with Certificates via Diffusion with Thresholds on a network with* $n > 1$ *nodes* $r(k) = \frac{8k^{2(1+\epsilon)}}{i(G)^2}\log(k^{2(1+\epsilon)}) + k^{1+\epsilon}\log(2k)$, $p(k) = \frac{\ln 2}{k^{1+\epsilon}}$, $\tau(k) = 1 - \frac{1}{k^{1+\epsilon}-1}$, $f(k) = \frac{4\sqrt{2}\ln(k^{1+\epsilon}/\xi)}{(\sqrt{2}-1)^2}$, *the explicit Revocable Leader Election problem is solved with probability at least* $1 - 1/n^{\log(8/5)} - 2\xi$, *with* $O(\frac{n^{4(1+\epsilon)}}{i(G)^2}\log^5 n)$ *time and* $O(\frac{n^{4(1+\epsilon)}}{i(G)^2}m\log^5 n)$ *messages, where* m *is the number of links.*

Randomized Leader Election

known	succes wp 1	succes whp	succes wp $1-o(1)$	succes wp constant
n, D	$O(m)$ exp. msgs, $O(D)$ exp time [14]			
n, Φ, t_{mix}		$\widetilde{O}(\min\{\sqrt{nt_{mix}/\Phi}, n/(\Phi \log n)\})$ **msgs,** $O(t_{mix} \log n + \Phi^{-1} \log^3 n)$ **time** **[this work]**		
n		$O(t_{mix}\sqrt{n} \log^{7/2} n)$ msgs, $O(t_{mix} \log^2 n)$ time [4]	$\exists G : \Omega(\sqrt{n}/\phi^{3/4})$ exp. msgs [4]	$\exists G : \Omega(m)$ exp. msgs [14]
		$O(m \min(\log\log n, D))$ exp. msgs, $O(D)$ time [14]		$\exists G : \Omega(D)$ time [14]
		$O(m + n \log n)$ msgs, $O(D \log n)$ time [14]		$O(m)$ exp. msgs, $O(D)$ time [14]
-		$O(n^{4(2+\epsilon)} m \log^5 n)$ **msgs,** $O(n^{4(2+\epsilon)} \log^5 n)$ **time** **[this work] (*)**		$\forall$ 2-connected G : $\exists$ labeling : $\Omega(m)$ exp. msgs [4] $\forall T(n)$: $\nexists$ LE alg in time $T(n)$ **[this work]**
$i(G)$		$O(\frac{n^{4(1+\epsilon)}}{i(G)^2} m \log^5 n)$ **msgs,** $O(\frac{n^{4(1+\epsilon)}}{i(G)^2} \log^5 n)$ **time** **[this work] (*)**		

Better time and message complexity for $\mathbf{\Phi}^{-1} = o(t_{mix} / \log n)$. Otherwise, up to log and polylog factor larger.

Refs:

(*) Revocable, Explicit

$\epsilon > 0$: any arbitrarily small constant

n : number of nodes

m : number of links

D : diameter

$\mathbf{\Phi}$: graph conductance

t_{mix} : random-walk mixing time

$i(G)$: isoperimetric number

[4] S. Gilbert, P. Robinson, and S. Sourav, "Leader election in well-connected graphs," in *PODC* 2018.

[14] S. Kutten, G. Pandurangan, D. Peleg, P. Robinson, and A. Trehan, "On the complexity of universal leader election," *J. ACM*, 2015.

Randomized Leader Election

known	succes wp 1	succes whp	succes wp $1-o(1)$	succes wp constant
n, D	$O(m)$ exp. msgs, $O(D)$ exp time [14]			
n, Φ, t_{mix}		$\widetilde{O}(\min\{\sqrt{n t_{mix}/\Phi}, n/(\Phi \log n)\})$ **msgs,** $O(t_{mix} \log n + \Phi^{-1} \log^3 n)$ **time** **[this work]**		
n		$O(t_{mix}\sqrt{n} \log^{7/2} n)$ msgs, $O(t_{mix} \log^2 n)$ time [4]	$\exists G : \Omega(\sqrt{n}/\phi^{3/4})$ exp. msgs [4]	$\exists G : \Omega(m)$ exp. msgs [14]
		$O(m \min(\log\log n, D))$ exp. msgs, $O(D)$ time [14]		$\exists G : \Omega(D)$ time [14]
		$O(m + n \log n)$ msgs, $O(D \log n)$ time [14]		$O(m)$ exp. msgs, $O(D)$ time [14]
-		$O(n^{4(2+\epsilon)} m \log^5 n)$ **msgs,** $O(n^{4(2+\epsilon)} \log^5 n)$ **time** **[this work] (*)**		$\forall$ 2-connected G : $\exists$ labeling : $\Omega(m)$ exp. msgs [4] $\forall T(n)$: $\nexists$ LE alg in time $T(n)$ **[this work]**
$i(G)$		$O(\frac{n^{4(1+\epsilon)}}{i(G)^2} m \log^5 n)$ **msgs,** $O(\frac{n^{4(1+\epsilon)}}{i(G)^2} \log^5 n)$ **time** **[this work] (*)**		

Nearly optimal for reasonable expansion $t_{mix} = \widetilde{\Theta}(\Phi^{-1}) \leq \widetilde{O}(D)$, yielding $O(n^{1/2}/\Phi)$ msgs.

Refs:

(*) Revocable, Explicit

$\epsilon > 0$: any arbitrarily small constant

n : number of nodes

m : number of links

D : diameter

Φ : graph conductance

t_{mix} : random-walk mixing time

$i(G)$: isoperimetric number

[4] S. Gilbert, P. Robinson, and S. Sourav, "Leader election in well-connected graphs," in *PODC* 2018.

[14] S. Kutten, G. Pandurangan, D. Peleg, P. Robinson, and A. Trehan, "On the complexity of universal leader election," *J. ACM*, 2015.

Randomized Leader Election

known	succes wp 1	succes whp	succes wp $1-o(1)$	succes wp constant
n, D	$O(m)$ exp. msgs, $O(D)$ exp time [14]			
n, Φ, t_{mix}		$\widetilde{O}(\min\{\sqrt{nt_{mix}/\Phi}, n/(\Phi \log n)\})$ **msgs,** $O(t_{mix} \log n + \Phi^{-1} \log^3 n)$ **time** **[this work]**		
n		$O(t_{mix}\sqrt{n} \log^{7/2} n)$ msgs, $O(t_{mix} \log^2 n)$ time [4]	$\exists G : \Omega(\sqrt{n}/\phi^{3/4})$ exp. msgs [4]	$\exists G : \Omega(m)$ exp. msgs [14]
		$O(m \min(\log\log n, D))$ exp. msgs, $O(D)$ time [14]		$\exists G : \Omega(D)$ time [14]
		$O(m + n \log n)$ msgs, $O(D \log n)$ time [14]		$O(m)$ exp. msgs, $O(D)$ time [14]
-		$O(n^{4(2+\epsilon)} m \log^5 n)$ **msgs,** $O(n^{4(2+\epsilon)} \log^5 n)$ **time** **[this work] (*)**		$\forall$ 2-connected G : $\exists$ labeling : $\Omega(m)$ exp. msgs [4] $\forall T(n)$: $\nexists$ LE alg in time $T(n)$ **[this work]**
$i(G)$		$O(\frac{n^{4(1+\epsilon)}}{i(G)^2} m \log^5 n)$ **msgs,** $O(\frac{n^{4(1+\epsilon)}}{i(G)^2} \log^5 n)$ **time** **[this work] (*)**		

For unknown n, no algorithm stops even with constant probability.

Refs:

(*) Revocable, Explicit

$\epsilon > 0$: any arbitrarily small constant

n : number of nodes

m : number of links

D : diameter

$\mathbf{\Phi}$: graph conductance

t_{mix} : random-walk mixing time

$i(G)$: isoperimetric number

[4] S. Gilbert, P. Robinson, and S. Sourav, "Leader election in well-connected graphs," in *PODC* 2018.

[14] S. Kutten, G. Pandurangan, D. Peleg, P. Robinson, and A. Trehan, "On the complexity of universal leader election," *J. ACM*, 2015.

Randomized Leader Election

known	succes wp 1	succes whp	succes wp $1-o(1)$	succes wp constant
n, D	$O(m)$ exp. msgs, $O(D)$ exp time [14]			
$n, \Phi,$ t_{mix}		$\widetilde{O}(\min\{\sqrt{nt_{mix}/\Phi}, n/(\Phi \log n)\})$ **msgs,** $O(t_{mix} \log n + \Phi^{-1} \log^3 n)$ **time** **[this work]**		
n		$O(t_{mix}\sqrt{n} \log^{7/2} n)$ msgs, $O(t_{mix} \log^2 n)$ time [4]	$\exists G : \Omega(\sqrt{n}/\phi^{3/4})$ exp. msgs [4]	$\exists G : \Omega(m)$ exp. msgs [14]
		$O(m \min(\log \log n, D))$ exp. msgs, $O(D)$ time [14]		$\exists G : \Omega(D)$ time [14]
		$O(m + n \log n)$ msgs, $O(D \log n)$ time [14]		$O(m)$ exp. msgs, $O(D)$ time [14]
-		$O(n^{4(2+\epsilon)} m \log^5 n)$ **msgs,** $O(n^{4(2+\epsilon)} \log^5 n)$ **time** **[this work] (*)**		$\forall$ 2-connected G : $\exists$ labeling : $\Omega(m)$ exp. msgs [4] $\forall T(n) : \nexists$ LE alg in time $T(n)$ **[this work]**
$i(G)$		$O(\frac{n^{4(1+\epsilon)}}{i(G)^2} m \log^5 n)$ **msgs,** $O(\frac{n^{4(1+\epsilon)}}{i(G)^2} \log^5 n)$ **time** **[this work] (*)**		

Explicit Revocable Leader Election with unknown n

Refs:

(*) Revocable, Explicit

$\epsilon > 0$: any arbitrarily small constant

n : number of nodes

m : number of links

D : diameter

$\boldsymbol{\Phi}$: graph conductance

t_{mix} : random-walk mixing time

$i(G)$: isoperimetric number

[4] S. Gilbert, P. Robinson, and S. Sourav, "Leader election in well-connected graphs," in *PODC* 2018.

[14] S. Kutten, G. Pandurangan, D. Peleg, P. Robinson, and A. Trehan, "On the complexity of universal leader election," *J. ACM*, 2015.

Leader Election Algorithm

Known *n* :

1. Each node chooses ID at random.
2. Each node chooses to be candidate at random.
3. Each candidate spans a “territory” tree by Cautious Broadcast, only nodes in less populated branches extend the tree.
4. Each candidate probes its territory by multiple Lazy Random Walks, larger ID wins.
5. Convergecast along the spanning tree of each territory of winning candidate ID (largest ID that hit the territory).

Our analysis shows:

candidates learn IDs of other candidates,

candidate that does not receive any larger ID becomes the leader,

all with high probability.

Leader Election Algorithm

Known n :

Cautious Broadcast key ingredients:

- The candidate spans a tree broadcasting its ID.
- Tree nodes choose new neighbors at random to expand the tree, but only in sparse branches.
- Tree nodes maintain: parent and children port numbers, number of nodes in subtree, and spanning status (active or passive).
- The candidate “controls” the size of the tree by activating or de-activating the expansion as needed.

Revocable Leader Election Algorithm

Unknown *n* :

Blind Leader Election with Certificates
via Diffusion with Thresholds key ingredients:

For each size estimate $k=1,2,4,8,\ldots$

1. Certification: to check if k is still too low to choose ID.
   A. Diffusion: nodes share some potential values for a number of rounds to decide if k is low if some thresholds are reached.
   B. Dissemination: of status for a number of rounds, if k is large enough all nodes receive.
2. Decision:
   - each node that did not choose ID and did not detect k as low chooses ID, storing k as "certificate".
   - each node updates leader ID and certificate (initially self) if a larger certificate or same with smaller ID is received.

Our analysis shows:

some node will not choose ID until the estimate is large enough
(we use the estimate as a "certificate" of uniqueness).

Open Questions

- Remove knowledge of Φ and/or t_{mix}
- Improve upper and/or lower bounds.
- Asynchronous protocol.