

Brief Announcement: Leader Election in SINR Model with Arbitrary Power Control

Magnús M. Halldórsson*
ICE-TCS, Reykjavik University
101 Reykjavik
Iceland
mmh@ru.is

Stephan Holzer†
TDS Group, MIT
77 Massachusetts Ave
Cambridge, MA
holzer@mit.edu

Evangelia Anna Markatou‡
TDS Group, MIT
77 Massachusetts Ave
Cambridge, MA
markatou@mit.edu

ABSTRACT

In this article, we study the leader election problem in the Signal-to-Interference-plus-Noise-Ratio (SINR) model where nodes can adjust their transmission power. We show that in this setting it is possible to solve the leader election problem in two communication rounds, with high probability. Previously, it was known that $\Omega(\log n)$ rounds were sufficient and necessary when using uniform power, where n is the number of nodes in the network.

We then examine how much power control is needed to achieve fast leader election. We show that any 2-round leader election algorithm in the SINR model running correctly w.h.p. requires a power range $2^{\Omega(n)}$ even when n is known. We match this with an algorithm that uses power range $2^{\tilde{\Theta}(n)}$, when n is known and $2^{\tilde{O}(n^{1.5})}$ when n is not known. We also explore tradeoffs between time and power used, and show that to elect a leader in t rounds, a power range $\exp(n^{1/\Theta(t)})$ is sufficient and necessary.

CCS CONCEPTS

•Networks → Network algorithms; •Theory of computation → Distributed algorithms;

KEYWORDS

SINR; leader election; power control; capture effect

ACM Reference format:

Magnús M. Halldórsson, Stephan Holzer, and Evangelia Anna Markatou. 2017. Brief Announcement: Leader Election in SINR Model with Arbitrary Power Control. In *Proceedings of PODC'17, July 25-27, 2017, Washington, DC, USA*, 3 pages.

DOI: <http://dx.doi.org/10.1145/3087801.3087851>

1 INTRODUCTION

In this article we explore what can be accomplished in an SINR network by utilizing *power control*, the ability of nodes to transmit with variable transmission power, and the *capture effect*, a property

of SINR networks, where a transmission can be successful while other transmissions within the communication range occur in the same round.

We study the leader election problem as our vehicle to explore this frontier. Leader election, the problem of electing a *unique leader* node among the nodes in a network, is one of the oldest and most studied problems in distributed computing. It is the ultimate way of breaking symmetry within radio networks in an initially unknown system, and is frequently used as a preliminary step in more complex communication tasks.

The best solution known for this problem in an SINR network has $O(\log n)$ runtime with high probability (w.h.p.) using uniform transmission power [1]. In the classical radio network model, the leader election problem requires $\Theta(\log^2 n)$ rounds w.h.p. [2]. Fine-man et al. [1] prove a $\Omega(\log n)$ lower bound for the leader election problem in the SINR model using uniform power and suggest that it may be possible to develop algorithms with faster runtimes using power control.

Indeed, by using power control we are able to improve on the $\Omega(\log n)$ previous bound and achieve an $O(1)$ algorithm for leader election.

2 MODEL AND PROBLEM STATEMENT

Let V be a set of n nodes, deployed in a single-hop network, that represent wireless devices. Every node can communicate with any other node using transmission power P , in absence of interference from other nodes. Time is divided into synchronous rounds. In each round, a node v can either transmit a message of size $O(\log n)$ with some power P_v , or listen. Node $v \in V$ can receive a message transmitted by node $u \in V$, iff v is listening and

$$\text{SINR}(u, v, I) = \frac{\frac{P_u}{d(u, v)^\alpha}}{N + \sum_{w \in I} \frac{P_w}{d(w, v)^\alpha}} \geq \beta, \quad (1)$$

where I is the set of other nodes $\notin \{u, v\}$ that transmit simultaneously.

Here, $P_u \in \mathbb{R}_{\geq 1}$ is the transmission power of node u , $d(u, v)$ is the distance between nodes v and u , α is the path-loss exponent, N is the non-zero ambient noise, and β is a hardware-dependent minimum SINR threshold required for a successful message reception. Our algorithms work for any $\beta > 0$, and the lower bounds require $\beta \geq 2$.

We denote by R the ratio of the longest to shortest distance between any two nodes in the network. We assume that the shortest distance between two nodes in the network is 1. Similar to [1], we

*Supported by grants 152679-05 and 174484-05 from the Icelandic Research Fund.

†Supported by grant AFOSR FA9550-13-1-0042.

‡Supported by grants NSF CCF-1461559 and NSF CCF-0939370

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

PODC'17, July 25-27, 2017, Washington, DC, USA

© 2017 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4992-5/17/07.

DOI: <http://dx.doi.org/10.1145/3087801.3087851>

assume that R is bounded by a polynomial in n , $R \leq n^c$, for some $c \in \mathbb{N}$. Let γ be a constant such that $\gamma \geq (c\alpha + 2) \log \beta$. We assume that the nodes know or can infer (an upper bound on) γ .

The $\tilde{O}$ -notation omits logarithmic factors. All logs are base 2. We consider that an event happens with high probability (w.h.p.) if it happens with probability greater than $1 - 1/n$.

PROBLEM 1 (LEADER ELECTION PROBLEM). *Given n nodes in a network, elect exactly one node (called the leader), with all nodes knowing whether or not they were elected to be the leader.*

3 2-ROUND LEADER ELECTION

3.1 The Essence of Our Algorithm

Below we present a high level description of the key ideas behind our algorithm.

- (1) **Geometric random variable:** The nodes use a geometric random variable k to count the tails flipped in a sequence of coin flips before the first heads is flipped. This geometric random variable allows some nodes to approximate n with no prior knowledge of the instance. More specifically, at least one and at most $16 \log n$ nodes flip a coin more than $\log n - \log \log n - 3$ times.
- (2) **Random IDs:** Each node chooses an *ID* (identification number) randomly using k . The geometric random variable k ensures that exactly one node v holds the maximum *ID*, which allows node v to break the symmetry of the network and stand out as the leader.
- (3) **Feedback:** In order to inform all nodes of the leader node v , we split the set of nodes V into listeners and competitors. The competitors compete for the leader position during the first round. The listeners inform the competitors of the winner during the second round. Both rounds use the same protocol with different message contents.
- (4) **The loudest node wins:** Each broadcasting node v determines its transmission power by evaluating power function $f(ID_v) = P \cdot ID_v^{\gamma ID_v}$ using its identification number, ID_v . Transmission power function f ensures that all listening nodes receive a message exactly from the node with the largest *ID*, as long as that *ID* is unique.

In summary, a *geometric random variable* allows the nodes to approximate n with no prior knowledge of the instance, *random IDs* ensure that the node v with the maximum *ID* stands out, *arbitrary transmission power* allows node v to inform the other nodes it is the leader, and *feedback* makes sure that all nodes know who the leader node is.

3.2 Leader Election Algorithm

The algorithm proceeds as follows. Initially, each node v flips a coin (a Bernoulli random variable) to determine its *role*: a *competitor* if heads are flipped, and a *listener* if tails. It then computes a geometric random variable (r.v.) k_v , which counts the tails flipped in a sequence of coin flips before the first heads is flipped. The ID of the node, ID_v , is an integer selected uniformly at random from the range $[J, 2 \cdot J]$, where $J = g(k_v) := 2^{k_v} k_v^4$. Finally, the power P_v that v uses for broadcast is given by $f(ID_v) := P \cdot ID_v^{\gamma ID_v}$, where

P is the minimum power needed to reach all nodes in the network (overcoming the ambient noise).

During round 1, competitors transmit their ID using the assigned power P_v , which is to be received by the listeners. In round 2, the roles are reversed, as the listeners report back the ID of the purported leader that they received.

3.3 Analysis

We proceed by showing that the highest power used by a competitor is sufficient to overpower all the other competitors, ensuring that this competitor is heard by all the listeners. Identical arguments hold for the reporting back in round 2.

To this end, we first show that there is a competitor whose geometric r.v. is nearly $\log n$, and at most a polylogarithmic number of competitors have that large value.

LEMMA 3.1. *Let $k_1 := \log n - \log \log n - 3$. For at least one and at most $O(\log n)$ competitors v it holds that $k_v \geq k_1$ with probability greater than $1 - \frac{1}{6n}$.*

We then show that all the $O(\log n)$ IDs at the high end of the spectrum are unique, i.e., selected by a single node.

LEMMA 3.2. *A single competitor receives the highest ID with probability greater than $1 - \frac{1}{20n}$, given that at least one node calculated $k_v \geq k_1$.*

The difference in power used by nodes with different ID ensures that the competitor with highest ID will overpower all the other competitors and be heard by all the listeners.

LEMMA 3.3. *The competitor v with the highest ID_v is received by all the listeners, if $ID_v \geq g(k_1)$.*

Using the above lemmas, we can prove Theorem 3.4.

THEOREM 3.4. *The 2-round leader election algorithm terminates with all nodes agreeing on a common leader, with probability at least $1 - \frac{1}{2n}$.*

Remark Radios operate in *full-duplex* when they can transmit and receive simultaneously (effectively by subtracting their transmitted signal from the received one). This is a powerful assumption since full-duplex is hard to implement but the technology has been progressing in recent years. With full-duplex, a single round suffices for the leader election algorithm. Namely, our arguments apply unchanged to the success of reception by the other competitors.

4 RANGE OF POWER NEEDED FOR A 2-ROUND LEADER ELECTION

Power control is the essential property that allows our algorithms to work. That begs the question *how much* power control is needed?

We say that an algorithm uses a *power range* X if the powers assigned fall in the range $[P, \dots, X \cdot P]$. The basic question is then how does the power range need to grow as a function of n for leader election to work correctly.

4.1 Upper Bound

If the power range is bounded, we may assume that the nodes know the upper bound of the range, P_{max} . Thus, the algorithm would automatically truncate the power assigned to be at most P_{max} . Using this observation, we prove the following theorem.

THEOREM 4.1. *Our 2-round leader election algorithm can work correctly with a power range of $2^{\tilde{O}(n^{1.5})}$, with probability greater than $1 - \frac{1}{n}$.*

If nodes know n , we can work with a smaller power range as follows: We can first sample the nodes with probability $\Theta(\log n/n)$, and have each selected node select ID uniformly at random from the range $[J, 2J]$, where $J = n \log^2 n$. The power used is $f(ID_v)$ as before, and the arguments are otherwise the same. This results in a power range of $2^{\tilde{\Theta}(n)}$.

PROPOSITION 4.2. *When nodes know n , a power range of $2^{\tilde{O}(n)}$ suffices.*

4.2 Lower Bound

We show that an exponential-size power range is actually necessary for any leader election protocol running in (at most) two rounds. We prove the following theorem assuming that the nodes are in a uniform metric space, and dividing the available range of power into subranges, each within factor 2. We then calculate the minimum power range needed such that the highest subrange used is used by exactly one node.

THEOREM 4.3. *Any 2-round leader election algorithm in the SINR model running correctly w.h.p. requires a power range $2^{\Omega(n)}$. This holds even if the nodes know n , the number of nodes in the network, and the nodes are located in a uniform metric (where all distances are equal).*

Observe that for the case of known n , we obtain an essentially tight bound of $2^{\tilde{\Theta}(n)}$ on the needed power range.

5 TRADING TIME FOR POWER

In this section, we explore how much the power range can be reduced by increasing the round complexity. We present a multi-round protocol that requires limited power range and derive a lower bound on the power range required by any t -round leader election algorithm, for $t \geq 2$.

5.1 Multi-Round Protocol

When a smaller power range is available, we can give a protocol that uses a larger number of rounds.

Our multi-round algorithm simply repeats the 2-round algorithm t times, for a given number $t \geq 1$, but using a slower-growing power function. Namely, we change the ID-selection function to $g_t(k) = 2^k k^{3t+1}$, and the power function to $f_k(ID_v) = P \cdot ID_v^{Y(ID_v)^{1/t}}$. After each repetition, each competitor v updates its $leader_v$ value to the largest among those heard so far.

To argue correctness, we first observe that it suffices to succeed in one of the round-pairs.

OBSERVATION 1. *If, in some round-pair, all receivers hear from a particular node v , and the senders all get informed of v as a leader, then the algorithm successfully terminates with v as leader.*

We then prove a counterpart of Lemma 3.2.

LEMMA 5.1. *With probability at least $1/(2n^{1/t})$, a single node received a higher ID in a given round than all the others.*

Correctness now follows from the above observations.

THEOREM 5.2. *For each $t \geq 1$, there is a $2t$ -round algorithm that elects a leader with probability at least $1 - 1/n$.*

5.2 Lower Bound

Using Theorem 4.3 and assuming that the nodes are in a uniform metric space, we prove a lower bound for multi-round protocols.

THEOREM 5.3. *Any t -round leader election algorithm in the SINR model running correctly w.h.p. requires a power range $2^{\Omega(t^{-1}\sqrt{n})}$, $t \geq 2$. This holds even if the nodes know n , the number of nodes in the network, and the nodes are located in a uniform metric (where all distances are equal).*

6 CONCLUSIONS AND ACKNOWLEDGMENTS

We have shown that power control can yield the ultimate speedup for leader election in the SINR model. This is thanks to the capture effect, which is the crucial property in which SINR differs from graphs-based models.

It would be exciting to see these techniques applied more widely. Multi-hop settings and more restricted power ranges are natural directions to examine, as well as problems beyond leader election. In general, the value of power control and the capture effect is still not fully understood.

We would like to thank Hsin-Hao Su, Nancy Lynch, Evangelos Markatos and Catherine Chronaki for the helpful comments and discussions.

REFERENCES

- [1] Jeremy T Fineman, Seth Gilbert, Fabian Kuhn, and Calvin Newport. Contention resolution on a fading channel. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*, pages 155–164. ACM, 2016.
- [2] Calvin C. Newport. Radio network lower bounds made easy. *CoRR*, abs/1405.7300, 2014.