



ELSEVIER

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcsLeader election in SINR model with arbitrary power control[☆]Magnús M. Halldórsson^a, Stephan Holzer^b, Evangelia Anna Markatou^{b,*},
Nancy Lynch^b^a ICE-TCS, School of Computer Science, Reykjavik University, Iceland^b TDS Group, Massachusetts Institute of Technology, USA

ARTICLE INFO

Article history:

Received 1 November 2017

Received in revised form 6 November 2018

Accepted 17 January 2019

Available online xxxx

Keywords:

SINR

Leader election

Power control

Capture effect

ABSTRACT

We consider 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 elect a leader in two communication rounds, with high probability. Previously, it was known that $\Theta(\log n)$ 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 every 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 complement this with an algorithm that uses power range $2^{\tilde{O}(n)}$ ¹, 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 range of possible power levels of size $\exp(n^{1/\Theta(t)})$ is sufficient and necessary.

© 2019 Elsevier B.V. All rights reserved.

1. Introduction

In this article we discuss what we can accomplish in a Signal-to-Interference-plus-Noise-Ratio (SINR) network using *power control*, the ability of nodes to transmit with variable transmission power, and the *capture effect*, a property of SINR networks, where a node's transmission can be successful even if other nodes within its communication range transmit at the same time.

We study the leader election problem as a vehicle to explore this frontier. Leader election, the problem of determining a *unique leader* among the nodes in a network, is one of the oldest and most studied problems in distributed computing. It provides a strong form 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 leader election problem was originally introduced in the 1970s, with the publication of the ALOHA radio network paper [1]. In the following years, many variations of the leader election problem have been extensively studied under a variety of models and algorithmic constraints. For example, in the multiple access channel model with collision detection [12], and in the SINR model with no collision detection [8].

[☆] This work was funded by Icelandic Research Fund grants 152679-05, 174484-05, AFOSR FA9550-13-1-0042, NSF CCF-1461559 and NSF CCF-0939370.^{*} Corresponding author.E-mail addresses: mmh@ru.is (M.M. Halldórsson), holzer@mit.edu (S. Holzer), markatou@mit.edu (E.A. Markatou), lynch@csail.mit.edu (N. Lynch).¹ The \tilde{O} -notation omits logarithmic factors.<https://doi.org/10.1016/j.tcs.2019.01.024>

0304-3975/© 2019 Elsevier B.V. All rights reserved.

We consider the leader election problem in SINR networks. In the SINR model, nodes operate in synchronous rounds. In each round a node either broadcasts a message to its neighbors or listens. A node v may or may not receive a message from node u depending on the distance between u and v , the transmission power of u , and the interference generated by other broadcasting nodes, as defined in Section 3.

In the classical radio network model, the leader election problem requires $\Theta(\log^2 n)$ rounds w.h.p. [11]. Fineman et al. [8] show that $O(\log n)$ rounds are sufficient and necessary to elect a leader in SINR networks without power control. They suggest that improved bounds may be possible using power control. Indeed, we show that power control can provide the ultimate speedup.

Our contributions: We present an algorithm that solves the leader election problem in two rounds w.h.p. We also present a multi-round leader election algorithm that uses limited transmission power. Our work is complemented by nearly matching lower bounds on the transmission power range for both two-round and multi-round leader election algorithms.

2. Related work

Gallager [9] presents a good survey of early work on leader election. From the mid-eighties, there was an increased interest in the *radio network model* [5], where concurrent transmissions are lost due to collisions, and nodes do not learn if their messages are successfully received. In this classical model, the leader election problem can be solved in $\Theta(\log^2 n)$ rounds w.h.p. [7] where n is the number of nodes in the network. This bound can be improved to $\Theta(\log n)$ w.h.p. assuming that nodes can detect collisions [4], and to $O(\log n_u)$ expected rounds assuming an upper bound n_u of n [2].

In the beginning of the new millennium came a renewed interest in *fading radio networks*, especially the SINR model, which take interference into account in a more realistic way. Moscibroda and Wattenhofer [10] showed that algorithms in the fading radio networks model can achieve better runtimes than algorithms for the radio networks model on certain problems, as SINR allows for better spatial reuse.

The most efficient leader election protocol known in the SINR model is by Fineman et al. [8]. It runs in time $O(\log n + \log R)$ w.h.p. in a single-hop network using uniform transmission power, where n is the number of nodes and $R = O(\text{poly}(n))$ is the ratio between the longest and the shortest link. Fineman et al. suggest that it may be possible to achieve better performance using power control. Indeed, for problems like link scheduling and connectivity, power control has been shown to give much better performance [10]. Power control has also been used in the SINR setting to solve the link scheduling problem while conserving energy, e.g. [3], [6].

To our knowledge, there has been no published work using power control to optimize the runtime of the leader election problem, or examining the trade-offs between the required communication complexity and power range of a leader election algorithm.

3. Model and problem statement

Let V be a set of n nodes, that represent wireless devices, deployed in a single-hop network located on a metric space. 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 $v \in V$ receives a message transmitted by node $u \in V$, iff v is listening and

$$\text{SINR}(u, v, I) = \frac{P_u}{d(u, v)^\alpha} \geq \beta, \quad (1)$$

$$N + \sum_{w \in I} \frac{P_w}{d(w, v)^\alpha}$$

where I is the set of other nodes transmitting simultaneously, $d(u, v)$ is the internode distance, and α, β, N are constants. Specifically, 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 $\beta > 0$, while the lower bounds use $\beta \geq 2$.

Every node can communicate with any other node using transmission power P_{\min} , in absence of interference from other nodes, that is $\frac{P_{\min}}{Nd(u, v)^\alpha} \geq \beta, \forall u, v \in V$.

In this paper, we consider the leader election problem.

Problem 1 (Leader election problem). Given n nodes in a network, elect exactly one node (called the leader), with all nodes knowing the identity of the leader.

We denote by R the ratio of the longest to the shortest distance between nodes in the network. Similarly to [8], we assume that there exists some constant $c \in \mathbb{N}$, such that $R \leq n^c$. Let γ be a constant such that $\gamma \geq \max(1, c\alpha + 1 + \log \beta)$. We assume that the nodes know 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 $1 - 1/n$.

Definition 2 (Bernoulli(p) random variable). A Bernoulli(p) random variable X takes on the values *competitor* or *listener* such that $\Pr[X = \text{competitor}] = p$ and $\Pr[X = \text{listener}] = 1 - p$.

4. 2-round leader election algorithm

In this section, we present a 2-round leader election algorithm that uses power control and requires no knowledge of n . First, we give some key ideas behind our algorithm. Then, we present a 2-round leader election algorithm followed by the analysis.

4.1. The essence of our algorithm

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

- (i) **Breaking symmetry:** Each node computes a geometric random variable k . Then, it picks an ID uniformly at random from a range that depends on k . We ensure that with high enough probability the largest ID is picked by exactly one node.
- (ii) **The loudest node wins:** Each broadcasting node v determines its transmission power by evaluating power function $f(ID) := P_{\min} \cdot ID^{\gamma ID}$ using its identification number, ID_v . Any listening nodes receive the message sent by the node with the highest ID, with a high enough probability.
- (iii) **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 of the two-round protocol. The listeners inform the competitors of the winner during the second round.

4.2. Leader election algorithm

The algorithm proceeds as follows. Initially, each node v flips a fair coin (a Bernoulli($\frac{1}{2}$) random variable) to determine its *role*, which is a *competitor* if heads are flipped, and *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 $J = g(k) := 2^k k^4$. Finally, the power P_v that v uses for broadcast is given by $f(ID) := P_{\min} \cdot ID^{\gamma ID}$, where P_{\min} is the minimum power needed to reach all nodes in the network (overcoming the ambient noise).

During round 1, each competitor v transmits its ID using power P_v , which is to be received by the listeners. In round 2, the roles are reversed, as each listener reports back the ID of the purported leader that it received. In the unlikely scenario that no message was received, the listener broadcasts an empty string.

We shall argue that, with high probability, a unique competitor succeeds in transmitting to all the listeners, and a unique listener succeeds in reporting back to all the competitors. The leader is then that successful competitor.

Algorithm 1 2-round leader election algorithm for node v .

```

1: Preprocessing:
2: Role: a boolean Bernoulli( $\frac{1}{2}$ ) random variable ('competitor' if heads, 'listener' if tails)
3:  $k$ : a Geometric( $\frac{1}{2}$ ) random variable,  $k \in \mathbb{Z}_{\geq 0}$ 
4:  $ID$ : chosen uniformly at random from  $[J, 2 \cdot J]$ , where  $J = g(k) := 2^k k^4$ ,  $ID \in \mathbb{Z}_{\geq 0}$ 
5:  $P$ : the transmission power,  $P = f(ID) := P_{\min} \cdot (ID)^{\gamma ID}$ ,  $P \in \mathbb{Z}_{\geq 0}$ 
6: Leader: a string denoting the identity of the leader, initially empty
7:
8: Round 1:
9: if Role = competitor then
10:   Broadcast  $ID$  using power  $P$ 
11: else
12:   If  $v$  receives message  $m$ , set Leader =  $m$ 
13:
14: Round 2:
15: if Role = competitor then
16:   If  $v$  receives message  $m$ , set Leader =  $m$ 
17: else
18:   Broadcast Leader using power  $P$ 

```

4.3. 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 logarithmic number of competitors have that large value. 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. 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.

We need the following version of Chernoff bounds.

Theorem 3 (Chernoff bound). Let X_1, X_2, \dots, X_n be independent Bernoulli random variables and $X = \sum_{i=1}^n X_i$. For $R \geq 4\mathbb{E}[X]$,

$$\Pr[X \geq R] \leq 2^{-0.55R}.$$

Proof. The standard Chernoff bound is that for any $\delta > 0$,

$$\Pr[X \geq (1 + \delta)\mathbb{E}[X]] \leq \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}}\right)^{\mathbb{E}[X]}.$$

Set δ be such that $R = (1 + \delta)\mathbb{E}[X]$, so $\delta \geq 3$. We get that

$$\Pr[X \geq R] \leq \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}}\right)^{\mathbb{E}[X]} \leq \left(\frac{e}{1 + \delta}\right)^{(1+\delta)\mathbb{E}[X]} = 2^{-\log((1+\delta)/e)R},$$

which is maximized when δ is minimized. Finally, observe that $\log(4/e) \geq 0.55$. \square

Lemma 4. Let $k_1 := \log n - \log \log n - 2$. With probability greater than $1 - \frac{1}{8n}$, for at least one and at most $32 \log n$ competitors v , it holds that $k_v \geq k_1$.

Proof. Let $t = \lceil k_1 \rceil = \lceil \log n - \log \log n - 2 \rceil$. Let A_v be the event that a given node v is a competitor and has $k_v \geq t$. The probability of A_v is $\Pr[A_v] = 2^{1-t} = 2^{1-\lceil k_1 \rceil}$. Thus,

$$\frac{4 \log n}{n} = 2^{-k_1} \leq \Pr[A_v] \leq 2^{1-k_1} = \frac{8 \log n}{n}.$$

The probability that no node satisfies A_v is then at most

$$\Pr\left[\bigwedge_v \overline{A_v}\right] \leq \left(1 - \frac{4 \log n}{n}\right)^n \leq e^{-4 \log n} \leq n^{-5.7} \leq \frac{1}{16n},$$

for n sufficiently large, establishing the first part of the claim.

Let X be the number of nodes v for which A_v holds. Then $\mathbb{E}[X] \leq 8 \log n$ and by Chernoff bound (Theorem 3) with $R = 32 \log n$,

$$\Pr[X \geq 32 \log n] \leq 2^{-0.55 \cdot 32 \log n} = n^{-17.6} \leq \frac{1}{16n},$$

for n large enough. I.e., at most $32 \log n$ nodes satisfy A_v , with probability greater than $1 - \frac{1}{16n}$.

Combined, with probability at least $1 - \frac{1}{8n}$, both of these claimed events hold. \square

The range from which the IDs are chosen is $[J, 2J]$, for $J \geq g(k_1)$, with high probability. Observe that $g(k_1) = 2^{k_1} k_1^4 \geq \frac{n \cdot \log^3 n}{8}$, for sufficiently large values of n .

Lemma 5. A sole competitor receives the highest ID with probability greater than $1 - \frac{1}{8n}$, given that at least one node calculated $k_v \geq k_1$.

Proof. The ranges of IDs assigned to nodes of different k_v values are disjoint. The competitor receiving the highest ID will therefore necessarily be one with a highest k_v value, which we denote by K . Let Z be the set of competitors with $k_v = K \geq k_1 (= \log n - \log \log n - 2)$. By Lemma 4, Z is non-empty and contains at most $32 \log n$ nodes.

The probability that a given pair of nodes in Z receive the same ID is inversely proportional to the range of IDs sampled from, or $1/J \leq \frac{1}{g(k_1)} \leq \frac{8}{n \cdot \log^3 n}$. The probability that some pair of nodes in Z are assigned the same ID is then, by the union bound, at most

$$\frac{\binom{|Z|}{2}}{J} \leq \frac{(32 \log n)^2}{\frac{n \cdot \log^3 n}{8}} = \frac{128}{n \log n} < \frac{1}{8n},$$

for large enough n . In particular, all nodes in Z receive different IDs with probability greater than $1 - \frac{1}{8n}$. \square

The highest ID received, ID_w , is at least $g(k_1) \geq n$, for sufficiently large values of n .

Lemma 6. *If a sole competitor receives the highest ID, then its transmission is received by all the listeners.*

Proof. Let w be the sole competitor with the highest ID. For any other competitor v it then holds that

$$\frac{P_w}{P_v} \geq \frac{f(ID_w)}{f(ID_w - 1)} \geq ID_w^\gamma \geq n^\gamma \geq \beta n^{c\alpha+1}. \tag{2}$$

Let u be a listener. We bound the noise and interference received by u in terms of the signal $S_u := P_w/d(w, u)^\alpha$ it receives from w . Recall that $d(w, u) \leq R \cdot d(v, u) \leq n^c \cdot d(v, u)$, and thus $d(w, u)^\alpha \leq n^{c\alpha} \cdot d(v, u)^\alpha$, for any competitor v . Hence, applying (2), the interference received from a competitor v is bounded by

$$I_v := \frac{P_v}{d(v, u)^\alpha} \leq \frac{P_w \cdot n^{c\alpha}}{\beta n^{c\alpha+1} \cdot d(w, u)^\alpha} = \frac{S_u}{\beta n}. \tag{3}$$

The definition of minimum power P_{min} ensures that $\frac{P_{min}/d(w, u)^\alpha}{N} \geq \beta$. Thus, we can use (2) to bound the noise term by

$$N \leq \frac{P_{min}}{d(w, u)^\alpha \cdot \beta} \leq \frac{P_w}{d(w, u)^\alpha \cdot n^\gamma \cdot \beta} = \frac{S_u}{\beta n^\gamma} \leq \frac{S_u}{\beta n}. \tag{4}$$

Combining (3) and (4), we get that the SINR of w 's signal at receiver u is bounded below by

$$\frac{S_u}{N + \sum_{v \in X} I_v} \geq \frac{\beta n}{1 + |X|} \geq \beta,$$

where X is the set of competitors other than w . Thus, w overpoweres all other competitors at all the listeners. \square

Theorem 7. *The 2-round leader election algorithm terminates with all nodes agreeing on a common leader, w.h.p.*

Proof. Using a union bound, we add up the error probabilities of Lemmas 4 and 5, we find that a sole competitor w receives the highest ID, with probability at least $1 - \frac{1}{4n}$. By Lemma 6, w then successfully informs all the receivers. All three lemmas work identically for the reporting process in round 2. Hence, with probability at least $1 - \frac{1}{2n}$, the algorithm succeeds. \square

Remark 8. Leader election can be achieved in a single round if simultaneous transmission and reception is possible. Such *full-duplex* radios operate by subtracting the transmitted signal from the received one. While they are still rare, being hard to implement, such technology has been progressing significantly in recent years and may well become a commodity feature. With full-duplex, our arguments apply unchanged to the success of reception by the other competitors, thus succeeding after only a single round.

5. Range of power needed for a 2-round leader election

Power control is the essential feature 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 $[P_{min}, \dots, X \cdot P_{min}]$. The basic question is then how the power range must grow as a function of n for leader election to work correctly.

5.1. Upper bound

Theorem 9. *Our 2-round leader election algorithm can be made to work correctly with a power range of $2^{\tilde{O}(n^{1.5})}$, w.h.p.*

Proof. The algorithm as is may select power assignments inducing a range of $2^{\tilde{O}(n^2)}$, since k_v is no larger than $2 \log n + 2$, with probability greater than $1 - \frac{1}{2n}$. However, if the 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} . We observe that this truncation can occur for at most one vertex, for the node with the highest ID to succeed. Namely, the probability that two or more nodes select a k_v value greater than $1.5 \log n$ is at most

$$\binom{n}{2} 2^{-3 \log n} \leq \frac{1}{2n}.$$

The bound on the maximum power now follows immediately. \square

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 at most $2(n \log^2 n)^{n \log^2 n} = 2^{\tilde{O}(n)}$.

Proposition 10. *When nodes know n , a power range of $2^{\tilde{O}(n)}$ suffices.*

5.2. Lower bound

We show that an exponential-size power range is actually necessary for any leader election protocol running in (at most) two rounds.

Theorem 11. *Every 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 if the nodes are located in a unit metric space (where all distances are equal).*

Proof. Consider n nodes located in a unit metric. In the unit metric, either a single message is received by all the listeners or none of them hear anything (assuming $\beta \geq 1$). Since the nodes don't operate full-duplex, two rounds are needed to inform the transmitting nodes of the winner, and the winner must be heard by all listeners in the first round.

We divide the available range of power into subranges, each within factor 2. Namely, if P_{max} is the maximum power available, then the i -th highest subrange is $[P_{max}/2^i, P_{max}/2^{i-1}]$. If the highest range used is used by two or more nodes, then the algorithm fails (assuming $\beta \geq 2$). We shall bound from below the probability that exactly two nodes use the highest subrange in use; this is clearly a lower bound on the failure probability of the algorithm.

Let X_i^v be the event that node v transmits in the first round using the i -th highest subrange. Since the nodes are identical, the same probability holds for them all, so let $p_i = \Pr[X_i^v]$. Observe that the probability that no node transmits in the round is at least $1 - n \sum_i p_i$, and since that can hold with probability at most $1/n$, it follows that

$$\sum_i p_i \geq \frac{1}{n} \left(1 - \frac{1}{n}\right). \tag{5}$$

Let q be the largest number such that

$$\sum_{i=1}^q p_i \leq \frac{1}{2n}. \tag{6}$$

So, a subrange of rank at least $q + 1$ is in use.

Let A_i be the event that at least two nodes use the i -th highest subrange, B_i be the event that no node transmits at subranges $1, 2, \dots, i - 1$, and $C_i = A_i \cap B_i$ be the event that both A_i and B_i occur, for $i = 1, 2, \dots$. Then, $C = \bigcup_i C_i$ is the event that at least two nodes use the highest subrange in use. Observe that $\Pr[A_i|B_i] \geq \Pr[A_i]$, since the non-use of the $i - 1$ highest subranges only makes the event A_i more likely. Then,

$$\Pr[C_i] = \Pr[A_i \cap B_i] = \Pr[A_i|B_i] \Pr[B_i] \geq \Pr[A_i] \Pr[B_i].$$

We bound the probability of A_i , $i \leq q$, by the first term of the binomial expansion:

$$\Pr[A_i] > \binom{n}{2} p_i^2 (1 - p_i)^{n-2} > \frac{n^2}{3} p_i^2 \left(1 - \frac{1}{2n}\right)^{n-2} > \frac{n^2}{3e} p_i^2.$$

Also, applying (6),

$$\Pr[B_i] \geq 1 - n \sum_{j=1}^{i-1} p_j \geq \frac{1}{2}.$$

Observe that the C_i 's are mutually exclusive and apply the Cauchy-Schwarz inequality followed by (5) to obtain:

$$\Pr[C] \geq \sum_{i=1}^q \Pr[C_i] \geq \frac{n^2}{3e} \sum_{i=1}^q p_i^2 \cdot \frac{1}{2} \geq \frac{n^2}{6e} \frac{(\sum_{i=1}^q p_i)^2}{q} \geq \frac{(1 - 1/n)^2}{6e \cdot q}.$$

The algorithm fails when C holds, and thus we may assume that $\Pr[C] \leq 1/n$, which implies that $q = \Omega(n)$. Hence, the claim. \square

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

Remark 12. We note that a construction can be given in the Euclidean plane that achieves the same result but with slightly weaker power tradeoffs. It consists of $n/2$ well-separated node-pairs that are internally close. It, however, does not avail itself to easy generalizations to protocols with greater number of rounds, and is therefore omitted.

6. Trading time for power range

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 $t \geq 2$.

6.1. 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 $f_t(ID_v) = P_{\min} \cdot ID_v^{\gamma(ID_v)^{1/t}}$. After each round-pair repetition, each competitor v updates its $leader_v$ value to the largest among those heard so far.

First, we observe that it suffices to succeed in one of the round-pairs.

Observation 13. *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.*

Proof. After this round-pair, all nodes have $leader_v$ value set as w . Thus, all broadcasts that follow use w for the value of $leader_v$. \square

We say that a node w is *strong* in a given round if $(ID_w)^{1/t} \geq (ID_v)^{1/t} - 1$, for all other nodes v . We first argue that a strong node will be heard by all.

Lemma 14. *A strong transmitting node is heard by all listeners.*

Proof. Recall that, as before, it holds that the highest ID, ID_w , is at least n . If a node w is strong in a round then it satisfies

$$\frac{f_t(ID_w)}{f_t(ID_v)} = \frac{ID_w^{\gamma(ID_w)^{1/t}}}{ID_v^{\gamma(ID_v)^{1/t}}} \geq ID_v^{\gamma(ID_w^{1/t} - ID_v^{1/t})} \geq ID_w^{\gamma} \geq n^{\gamma} \geq \beta n^{c\alpha+1}.$$

Thus, by the same argument as in Lemma 6, w overpowers all other transmitters at every receiver. \square

We next observe that a strong node is likely to emerge.

Lemma 15. *Let $U = n^{1/t}$. In a given round, some node is strong, with probability at least $1 - 1/(2U)$.*

Proof. Let Z be the set of competitors with the largest k_v -value. By Lemma 4, $|Z| \leq 32 \log n$. Recall that IDs are allocated uniformly at random, and for nodes in Z , the range is of size at least $g_t(k_1) = 2^{\log n - \log \log n - 2} (\log n - \log \log n - 2)^{3t+1} = \frac{n}{4 \log n} \left(\frac{\log n}{2}\right)^{3t+1} \geq \frac{1}{8 \cdot 2^{3t}} n \log^{3t} n$, for large enough n . The probability that a given pair of nodes u, v in Z receive nearly equivalent IDs, with $|(ID_u)^{1/t} - (ID_v)^{1/t}| \leq 1$, is at most $g_t(k_1)^{-1/t} \leq \frac{8^{1+1/t}}{n^{1/t} \log^3 n}$. Thus, the probability that some two nodes in Z receive nearly equivalent IDs is at most

$$\frac{\binom{|Z|}{2}}{g_t(k_1)^{1/t}} \leq \frac{8^{1+1/t} \cdot 32^2 \log^2 n}{n^{1/t} \log^3 n} < \frac{1}{2n^{1/t}},$$

for sufficiently large n . \square

The correctness of the algorithm follows from the above observations.

Theorem 16. *For each number $t = O(\log n / \log \log n)$, there is a $2t$ -round algorithm using a power range $2^{n^{O(1/t)}}$ that correctly elects a leader, w.h.p.*

Proof. A round-pair is *strong* if both of its round feature a strong (participating) node. In a strong round-pair, all the listeners hear from a strong competitor in the first round and all the competitors receive the acknowledgment from a strong listener, by Lemma 14. By Observation 13, the algorithm is therefore successful as long as there is some strong round-pair. We proceed to argue the probability of a strong round-pair occurring.

Let $U = n^{1/t}$. By Lemma 15, a strong node fails to emerge in a given round with probability at most $1/(2U)$. Thus, a given round-pair fails to be strong with probability at most $1/U$. By independence, the probability that none of the t round-pairs are strong is at most $1/U^t = 1/n$. \square

6.2. Lower bound for multi-round protocols

Theorem 17. 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 unit metric (where all distances are equal).

Proof. We consider n nodes located in a unit metric space. In this setting, after any round of the algorithm either all listening nodes receive a message, or no progress is made (assuming $\beta \geq 1$). Since the nodes do not operate full-duplex, any leader election algorithm requires at least two rounds, one round for the winner to broadcast its message, and one round to be informed of the victory.

Let A be a t -round leader election algorithm in the SINR model that runs correctly with probability greater than $1 - 1/n$. Since at least two rounds of successful communication are needed, Algorithm A fails when no listening node receives a message during the first $t - 1$ rounds. This happens with probability $\prod_{r=1}^{t-1} p_r$, where p_r denotes the probability that no listener receives a message in round r . Since algorithm A succeeds with probability greater than $1 - 1/n$,

$$\frac{1}{n} > \prod_{r=1}^{t-1} p_r.$$

Now, consider round r . Let q and C be as in Theorem 11. We can show by a similar argument that $\Pr[C] \geq \frac{1}{12e \cdot q}$, assuming $\beta \geq 2$. No listener receives a message in round r when C holds, and thus $\Pr[C] \leq p_r$, which implies that

$$q \geq \frac{1}{12e \cdot p_r}.$$

It follows that $1/n \geq (\frac{1}{12eq})^{t-1}$, and therefore $q \geq t^{-1/\sqrt{n}}/(12e) = \Omega(t^{-1/\sqrt{n}})$. Thus, algorithm A requires a power range $2^{\Omega(t^{-1/\sqrt{n}})}$. \square

7. 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 thank Hsin-Hao Su for helpful comments and discussions.

References

- [1] Norman Abramson, The ALOHA system: another alternative for computer communications, in: Proceedings of the November 17–19, 1970, Fall Joint Computer Conference, AFIPS '70 (Fall), ACM, New York, NY, USA, 1970, pp. 281–285.
- [2] Reuven Bar-Yehuda, Oded Goldreich, Alon Itai, On the time-complexity of broadcast in radio networks: an exponential gap between determinism and randomization, in: Proceedings of the Sixth Annual ACM Symposium on Principles of Distributed Computing, PODC '87, ACM, 1987, pp. 98–108.
- [3] S.A. Borbash, A. Ephremides, Wireless link scheduling with power control and SINR constraints, IEEE Trans. Inform. Theory 52 (11) (Nov 2006) 5106–5111.
- [4] John Capetanakis, Tree algorithms for packet broadcast channels, IEEE Trans. Inform. Theory 25 (5) (1979) 505–515.
- [5] Imrich Chlamtac, Shay Kutten, On broadcasting in radio networks—problem analysis and protocol design, IEEE Trans. Commun. 33 (12) (1985) 1240–1246.
- [6] R.L. Cruz, A.V. Santhanam, Optimal routing, link scheduling and power control in multihop wireless networks, in: INFOCOM, vol. 1, 2003, pp. 702–711.
- [7] Sebastian Daum, Seth Gilbert, Fabian Kuhn, Calvin Newport, Leader election in shared spectrum radio networks, in: Proceedings of the 2012 ACM Symposium on Principles of Distributed Computing, PODC '12, ACM, New York, NY, USA, 2012, pp. 215–224.
- [8] Jeremy T. Fineman, Seth Gilbert, Fabian Kuhn, Calvin Newport, Contention resolution on a fading channel, in: Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing, PODC '16, ACM, 2016, pp. 155–164.
- [9] Robert G. Gallager, A Perspective on Multiaccess Channels, Technical report, DTIC Document, 1985.
- [10] Thomas Moscibroda, Roger Wattenhofer, The complexity of connectivity in wireless networks, in: INFOCOM, 2006.
- [11] Calvin Newport, Radio network lower bounds made easy, in: International Symposium on Distributed Computing, Springer, 2014, pp. 258–272.
- [12] Dan E. Willard, Log-logarithmic selection resolution protocols in a multiple access channel, SIAM J. Comput. 15 (2) (May 1986) 468–477.