Brief Announcement: A Shorter and Stronger Proof of an $\Omega(D \log(n/D))$ Lower Bound for Broadcast in Radio Networks

Calvin Newport*
Georgetown University, Washington DC
cnewport@cs.georgetown.edu

ABSTRACT

A seminal 1998 paper by Kushilevitz and Mansour [10] proved that for any randomized radio network broadcast algorithm, there exists a network in which the algorithm requires an expected time of $\Omega(D \log(n/D))$ rounds, for network size n and diameter D. In this study, we apply a new technique to generate a shorter and stronger version of this proof. In more detail, our new version fits in two pages, and it strictly strengthens the existing result by now allowing for active collision detection and an unlimited number of communication channels—assumptions which break the proof argument of [10].

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Wireless Communication

Keywords

Broadcast; Algorithms; Wireless; Theory

1. INTRODUCTION

Broadcast in radio networks has been studied from an algorithmic perspective for over 25 years (e.g., since [3]). Work on randomized distributed solutions to this problem build on three seminal papers written between 1987 and 1998: (1) the $O(D \log n + \log^2 n)$ randomized broadcast algorithm of Bar-Yehuda et al. [2], for network size n and network diameter D (a result which was later optimized slightly to $O(D \log(n/D) + \log^2 n)$ [9, 4]); (2) the $\Omega(\log^2 n)$ lower bound of Alon et al. [1], which proves the Bar-Yehuda bound tight for small D; and the (3) $\Omega(D \log(n/D))$ lower bound of Kushilevitz and Mansour [10], which proves (optimized) Bar-Yehuda optimal for larger D. Of these three important bounds, the proof argument by Kushilevitz and Mansour is the longest and arguably the most complicated (which perhaps explains the long gap between this bound and the $\Omega(\log^2 n)$ bound of [1] that preceded it).

In this (purposefully) brief paper, we demonstrate a surprising reality: the core technical argument of Alon et al. [1] can be used to prove a strengthened and much simplified version of the Kushilevitz and Mansour bound. (That is, the authors of [1] had, probably without knowing it, all the pieces necessary to prove the *full* Bar-Yehuda result optimal.) To support this claim, we leverage a core result of [1] to prove $\Omega(D \log(n/D))$ rounds are necessary to solve distributed broadcast. Our proof, which uses a different approach than [10], requires *only two pages* (including discussion). Furthermore, the result is *strictly* stronger than the original, as it now holds even if we assume active collision detection (i.e., active nodes can use collision detection) and provide the nodes access to an unlimited number of orthogonal communication channels. Both of these additional assumptions break the proof argument in [10]. In fact, before this paper, it was an open question whether $D \log(n/D)$ rounds are needed for broadcast in the presence of collision detection or multiple channels: both these assumptions, for example, have been shown to speed up leader election in radio networks [7, 5], so it stands to reason they would do the same for multihop broadcast.[1] We prove here—perhaps surprisingly—they do not.

Related Work.

We are not the first to simplify the complex bound of Kushilevitz and Mansour. This effort was previously undertaken by Liu and Prabhakaran [11], who simplified the result of [10] by first bounding a key deterministic behavior, then translating the result to the randomized setting using Yao's minimax principle. Though this approach differs from ours in its specifics, it shares the same general attack: finding a way to bound the progress of the message through a layered network without having to argue directly about the behavior of the randomized algorithm (the main source of complexity in the original proof). Whereas Liu and Prabhakaran leverage a connection to determinism in this effort, we instead work by reducing randomized broadcast to an easily bounded combinatorial game. We also note that their solution assumes a graph structure that does not satisfy ge-

*The author was supported in part by a Ford Motor Company University Research Program grant.

[1] In fairness, the leader election algorithms of [7, 5] assume all nodes start the execution active, whereas the broadcast problem requires nodes to remain inactive before receiving the message, so these strategies do not directly translate to the broadcast setting.

PODC'13, July 22–24, 2013, Montréal, Québec, Canada.
ACM 978-1-4503-2065-8/13/07.

ographic constraints such as the unit disk graph property, while our bound, as in the original bound of Kushilevitz and Mansour, still holds under such restrictions.

Model.

We model a multihop radio network as an undirected connected graph $G = (V, E)$, where the $n = |V|$ nodes in V correspond to the wireless devices, the edges in E indicate which devices are within communication range, and we use D to describe the graph diameter. An *algorithm* consists of n randomized processes. An execution begins with an adversary assigning the processes to nodes in V. It then proceeds in synchronous rounds. In each round, each node decides whether to transmit or receive, based on its corresponding process definition. A node u receives a message m in a given round r, if and only if : (a) u decides to receive in r; and (b) exactly one neighbor of u transmits in r and it transmits m. In the standard version of this model, a node cannot distinguish between silence and collision. Below, we define a type of *collision detection* for which our lower bound still holds.

We study the *broadcast* problem, which requires a designated *source* node to propagate a message to every node in the network. In this problem, a node is *inactive* until it first receives the broadcast message; at which point it becomes *active* and can participate in the algorithm. We say a graph satisfies the *unit disk graph* property if there is a way to assign nodes locations in a 2-dimensional cartesian plane such that $E = \{(u, v) \mid d(u, v) \leq 1\}$, for distance metric d. We say nodes have *active collision detection* if they can distinguish between silence and collisions once they become active (i.e., after they have received their first message). We say nodes have *access to multiple channels* to describe a generalization of the model where each node can choose a channel from a set of multiple orthogonal communication channels in each round, such that the message receive rules apply to each channel individually (e.g., u receives a message on channel i if and only if u chooses to receive on i and exactly one neighbor of u broadcasts on i). *Inactive* nodes can receive on some predetermined default channel.

2. LOWER BOUND

Alon et al. [1] proved the existence of a bipartite radio network (V_1, V_2) of size N, with $|V_1| = n$ and $|V_2| \approx n^c$, for some constant $c > 1$, where delivering a message from V_1 to every node in V_2 requires $\Omega(\log^2 n) = \Omega(\log^2 N)$ rounds. Their proof argument divided the nodes in V_2 into $\Theta(\log n)$ different *groups*, each corresponding to a different number of neighbors in V_1. Intuitively, each set of broadcasts $B \subset V_1$ can only help deliver messages to nodes in a small number of these groups. This observation is the starting point for their eventual argument that $\Omega(\log^2 N)$ rounds are necessary to get a message to every node in V_2. For our purposes, however, we do not need to follow this proof to its conclusion. We are instead content to make use of the following intermediate result (recently isolated and proved in [8]), which follows in a straightforward manner from the proof in [1]:

LEMMA 2.1 (ADAPTED FROM [8, 1]). *Fix any $n > 0$. There exists a constant $\alpha > 0$ and bipartite radio network $H = (V_1, V_2)$, with $|V_1| = n$ and $|V_2| > n$, such that in each round, regardless of which nodes in V_1 transmit, at most an $\alpha/\log n$ fraction of nodes in V_2 receive a packet.*

The above lemma provides the core technical result upon which we will now build our lower bound. At a high-level, our strategy will proceed as follows: (1) we will use Lemma 2.1 to bound an abstract game called *set isolation* (which we previously introduced in [6] to derive a shorter and stronger version of the $\Omega(\log^2 n)$ lower bound on the single-hop *wake-up* problem); (2) we will reduce set isolation to distributed broadcast, applying our bound from the preceding step to achieve a bound for distributed broadcast. The set isolation game, in other words, is the technical glue that connects the result of Alon et al. to the result of Kushilevitz and Mansour.

The Set Isolation Game.

The *k-set isolation game*, defined for some $k > 1$, has a *player* face off against an adversarial *referee*. At the beginning of the game, the referee secretly selects a *target set* $T \subseteq \{1, ..., k\}$. In each round, the player generates a *proposal* $P \subseteq \{1, ..., k\}$. If $|P \cap T| = 1$, then the player wins the game. Otherwise, the player moves on to the next round learning no information other than the fact that its proposal failed. We leverage Lemma 2.1 to prove a lower bound on the expected time to win this game:

LEMMA 2.2. *Fix some $k > 1$. There exists a referee strategy for the k-set isolation game, such that for every player strategy, the expected time to win the game is $\Omega(\log k)$ rounds.*

PROOF. We begin by discussing broadcast. Using $n = k$, fix the constant α and bipartite graph $H = (V_1, V_2)$ provided by Lemma 2.1. A consequence of Lemma 2.1 is that $\Omega(\log n)$ rounds of broadcasting in V_1 are needed before half or more of the nodes in V_2 have received the message. This follows because at most $(|V_2|\alpha)/\log n$ nodes in V_2 receive the message in each round, regardless of how we choose broadcasters (by Lemma 2.1). It takes, therefore, a minimum of $\beta = \lfloor \log n/(2\alpha) \rfloor$ rounds before at least half the nodes in V_2 can receive the message.

Now we connect this observation to set isolation. By construction, $k = |V_1|$. For use in the remainder of the proof, assign each node in V_1 a unique label from $\{1, ..., k\}$. When playing the k-set isolation game, we can interpret each proposal P from the player as the nodes $P \subseteq V_1$ broadcasting in H. Consider the referee strategy that chooses a node from $u \in V_2$ with uniform randomness, and then sets $T = N_H(u)$, where N_H is the neighbor function on H. It follows that the player wins the game during the first round when u receives a message in the corresponding broadcast simulation on H. By our above argument, the player's simulation will have delivered the message to less than half the nodes in V_2 by the end of round $\beta - 1$. Because the referee chooses u at random (without revealing this choice to the player), the probability that the player has won the game in $\beta - 1 = O(\log n)$ rounds is less than $1/2$. Therefore, regardless of the player strategy, this specific referee strategy yields an expected time of $\Omega(\log n)$ rounds to win the isolation game. □

Reducing Set Isolation to Distributed Broadcast.

We have used a lower bound regarding the existence of slow broadcast graphs to generate a lower bound for our abstract set isolation game. We will now use our set isolation game bound to generate a lower bound for the expected time for distributed broadcast. The following theorem statement is strictly stronger than the statement from [10]:

THEOREM 2.3. *For every broadcast algorithm $\mathcal{A}$, number of processors $n > 1$, and diameter $D > 0$: there exists a network in which the expected time to complete broadcast is $\Omega(D \log(n/D))$ rounds. This holds even if we restrict ourselves to network topologies that satisfy the unit disk graph property, and assume unique ids, active collision detection, and any number of available communication channels.*

PROOF. We proceed by reduction from a variant of set isolation. In more detail, let (k, k')-*multi-set isolation*, for $1 \leq k' \leq k$, be a variation of set isolation in which we run k' consecutive instances of $(\lfloor k/k' \rfloor)$-set isolation, requiring the player to win instance $i \in \{1, ..., k'-1\}$ before proceeding to instance $i+1$. Two technical points that aid the below argument: assume the referee selects all k' targets at the beginning of the game, and that the referee reveals the target for instance i in the round when the player wins that instance. For a given execution of a player and referee strategy for (k, k')-multi-set isolation, let X_i, for $i \in \{1, ..., k'\}$, be the time required to win instance i of the game, and let $Y = X_1 + X_2 + ... + X_{k'}$ be the time required to win the full multi-set game. By linearity of expectation and our result from Lemma 2.2, we note there is a referee strategy that allows us to bound $\mathbb{E}[Y]$ as follows: $\mathbb{E}[Y] = \mathbb{E}\big[\sum_i^{k'} X_i\big] = \sum_i^{k'} \mathbb{E}[X_i] = \Omega(k' \log(k/k'))$.

Assume that $\mathcal{A}$ solves broadcast in $f(n, D)$ rounds, in expectation, in networks of size n and diameter D. We now devise a (n, D)-multi-set isolation player strategy that simulates $\mathcal{A}$ to win the game in expected time $f(n, D)$ as well. In more detail, the player simulates $\mathcal{A}$ on a network consisting of $D+1$ layers, $L_1, L_2, ..., L_D, L_{D+1}$, where the first D layers each include $\lfloor n/D \rfloor$ nodes, and the last layer includes at least 1 node (if D divides n evenly, then we can add an extra node to the system to populate L_{D+1}, without affecting the asymptotic bounds below; otherwise we add the leftover nodes from the smaller layers). For the sake of construction, for each L_i, assign unique ids from $\{1, ..., k\}$ to the nodes in L_i. Let T_i be the target chosen by the referee for instance i. In our construction, we connect L_i and L_{i+1} by including an edge from every node in L_i corresponding to a value in T_i to every node in L_{i+1}. Notice, the player simulating this network does not know these T_i values, but we will now show this does not matter. Finally, the nodes within each layer are connected as a clique. (Notice that this graph clearly satisfies the unit disk graph property.)

The simulation begins with the player choosing some node in L_1 as the source. In each round r of the simulation, let $\hat{i}$ be the largest value of i such that the nodes in L_i are active (i.e., have the message). Let $B^r_{\hat{i}}$ be the nodes in $L_{\hat{i}}$ that broadcast in r, if any. If we are assuming multiple channels, let $B^r_{\hat{i}}$ be the nodes in $L_{\hat{i}}$ that broadcast on the default channel where inactive nodes listen. The player uses $B^r_{\hat{i}}$ as its proposal in this round of the mutli-set isolation game. The key insight of this reduction is that the player only needs to simulate communication between $L_{\hat{i}}$ and $L_{\hat{i}+1}$ if exactly one node connecting $L_{\hat{i}}$ to $L_{\hat{i}+1}$ is in $B^r_{\hat{i}}$. When this occurs, the player will learn of this fact, because its corresponding guess in the set isolation game will win this instance of the game (and once it wins instance i for the first time, it learns T_i, so it can, moving forward, successfully simulate all future communication between these two layers). Collision detection and multiple channels break the original proof of [10] because their argument requires that nodes in the same layer receive silence in all rounds before they advance the message. If the active nodes in a layer had collision detection, for example, they could quickly achieve some communication using collisions, at which point the argument of [10] fails. Our argument can tolerate such intra-layer communication as it focuses only on the externally observable behavior of the layer: i.e., which nodes broadcast and whether or not they help advance the message.

We conclude by noting that the player using this strategy will win the multi-set isolation game when the message arrives at L_{D+1}. By assumption, this occurs in expected time of $f(n, D)$ rounds. By our above bound on $\mathbb{E}[Y]$, it must follow: $f(n, D) = \Omega(D \log(n/D))$, as needed. □

3. ACKNOWLEDGEMENTS

The author thanks Mohsen Ghaffari for his helpful discussions regarding the arguments of Alon et al. [1].

4. REFERENCES

[1] N. Alon, A. Bar-Noy, N. Linial, and D. Peleg. A Lower Bound for Radio Broadcast. *Journal of Computer and System Sciences*, 43(2):290–298, 1991.

[2] R. Bar-Yehuda, O. Goldreigch, and A. Itai. On the Time-Complexity of Broadcast in Multi-Hop Radio Networks: An Exponential Gap between Determinism and Randomization. In *Proceedings of the ACM Conference on Distributed Computing*, 1987.

[3] I. Chlamtac and S. Kutten. On Broadcasting in Radio Networks–Problem Analysis and Protocol Design. *IEEE Transactions on Communications*, 33(12):1240–1246, 1985.

[4] A. Czumaj and W. Rytter. Broadcasting algorithms in radio networks with unknown topology. *Journal of Algorithms*, 60:115–143, 2006.

[5] S. Daum, S. Gilbert, F. Kuhn, and C. Newport. Leader Election in Shared Spectrum Radio Networks. In *Proceedings of the ACM Conference on Distributed Computing*. ACM, 2012.

[6] S. Daum, F. Kuhn, and C. Newport. Efficient Symmetry Breaking in Multi-Channel Radio Networks. In *Proceedings of the International Conference on Distributed Computing*, 2012.

[7] M. Ghaffari and B. Haeupler. Near Optimal Leader Election in Multi-Hop Radio Networks. 2013.

[8] M. Ghaffari, B. Haeupler, and M. Khabbazian. A Bound on the Throughput of Radio Networks. *Pre-print*, February 2013. ArXiv preprint arXiv:1302.0264.

[9] D. Kowalski and A. Pelc. Broadcasting in Undirected Ad Hoc Radio Networks. *Distributed Computing*, 18(1):43–57, 2005.

[10] E. Kushilevitz and Y. Mansour. An Ω(D\log(N/D)) Lower Bound for Broadcast in Radio Networks. *SIAM Journal on Computing*, 27(3):702–712, 1998.

[11] D. Liu and M. Prabhakaran. On Randomized Broadcasting and Gossiping in Radio Networks. In *Proceedings the International Computing and Combinatorics Conference*, 2002.