Brief Announcement: Hardness of Broadcasting in Wireless Networks with Unreliable Communication*

Fabian Kuhn
MIT CSAIL
Cambridge, MA, USA
fkuhn@csail.mit.edu

Nancy Lynch
MIT CSAIL
Cambridge, MA, USA
lynch@csail.mit.edu

Calvin Newport
MIT CSAIL
Cambridge, MA, USA
cnewport@csail.mit.edu

ABSTRACT

We prove two broadcast lower bounds for a wireless network model that includes unreliable links. For deterministic algorithms, we show $n-1$ rounds are required, where n is the number of processes. For randomized algorithms, $\epsilon(n-1)$ rounds are required for success probability ϵ. In both cases, the bounds are proved for a network in which constant-time broadcast is possible.

Categories and Subject Descriptors

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

General Terms

Algorithms, Theory

1. INTRODUCTION

We consider a wireless network model based on a pair of graphs, G and G', where G represents reliable communication links and G' represents both reliable and unreliable links. This model reflects a key aspect of the behavior of real wireless networks: some links are always reliable while others are unreliable—sometimes delivering messages and sometimes not.[1] (Such a *dual-graph* model is used, for example, in [3].) This model appears to be significantly less powerful than models based on single graphs, due to the fact that processes cannot easily discover which edges are reliable. It is also similar to the *quasi-unit-disk* model [1, 4]. The key difference, however, is that we allow links to fluctuate between delivering and not delivering messages throughout an execution, whereas the quasi-unit-disk model resolves the uncertainty permentantly at the beginning of the execution.

In this paper we use the dual-graph model to study the fundamental problem of broadcasting a single message. We assume processes can have unique ids and know the graph topology, but do not know *a priori* their location in the graph. We show lower bounds on the time required to complete the broadcast of a single message for both deterministic and randomized algorithms. We show these bounds to be significantly worse than the optimal time possible if the processes had full information about the network.

2. MODEL

Fix some $n > 2$. We define a network (G, G') to consist of two graphs $G = (V, E)$ and $G' = (V, E')$, where V is a set of n wireless nodes and $E \subseteq E'$. We define an algorithm $\mathcal{A}$ to be a collection of n processes. An execution of an algorithm $\mathcal{A}$ on network (G, G') first fixes some bijection *proc* from processes of $\mathcal{A}$ to V. It then proceeds in synchronous rounds $1, 2, \ldots$ Fix some node $v \in V$. In each slot, v may or may not send a message, as indicated by its process ($proc(v)$). To model communication, we use G and G'. A message sent by v *reaches* every neighbor of v in G and some subset of v's neighbors that are in G' but not G. *The choice of this subset is nondeterministic and can vary from round to round.* What node v *receives* is determined by the messages that reach it: if no message reaches v, it receives $\perp$; if exactly one message reaches v, it receives just that message; if two or more messages reach v, it receives $\top$, indicating a collision.

We use $[x]$ to refer to the set $\{1, ..., x\}$. For a fixed algorithm $\mathcal{A}$, we assign each process a unique label from $[n]$. Given an execution of $\mathcal{A}$ in some network $(G = (V, E), G' = (V, E'))$, we use ID_v, for every $v \in V$, to refer to the label of the process assigned to v by *proc* in this execution. This labelling is a convenience for our discussions, we do not necessarily assume that the labels are encoded in the processes definitions. When we refer to "process i" in a given execution, we mean the process assigned label i.

3. THE BROADCAST PROBLEM

The broadcast problem requires the dissemination of a message from a single *source* node to all nodes. To simplify our lower bounds, we assume that the processes in a broadcast algorithm treat the message like a *black box*; i.e., behave the same regardless of the message contents. We say

*This work has been support in part by Cisco-Lehman CUNY A New MAC-Layer Paradigm for Mobile Ad-Hoc Networks, AFOSR Award Number FA9550-08-1-0159, NSF Award Number CCF-0726514, and NSF Award Number CNS-0715397.

[1]In the systems literature, the phenomenon of links that are sometimes reliable and sometimes not, is described by the term *fading*, and is known to be generated by multi-path propagation and transient obstacles/interference, among other causes; c.f., [5, 2].

PODC'09, August 10–12, 2009, Calgary, Alberta, Canada.

a network (G, G') is *k-broadcastable* if and only if for every source there exists a sequence of broadcasts that disseminates the message to all nodes within k rounds, regardless of the resolution of the communication non-determinism.

4. LOWER BOUNDS

We begin with a deterministic lower bound that shows that even on a network where fast broadcast is possible, a linear number of rounds is required.

THEOREM 1. *There exists a 2-broadcastable network (G, G'), such that there does not exist a deterministic algorithm $\mathcal{A}$ that solves broadcast in less than or equal to $n-2$ rounds in (G, G').*

PROOF. Let G be a clique consisting of $n-1$ nodes, with one additional node r that is connected to a single node b in the clique. Specifically, $G = (V, E)$, $|T| = n-2$, $V = T \cup \{b, r\}$, $E = \{(u, v) | u, v \in (T \cup \{b\}), u \neq v\} \cup \{(b, r)\}$. Thus, node b is a "bridge" node that connects the clique to r. Assume G' is the complete graph over V. The network (G, G') is clearly 2-broadcastable: the source broadcasting followed by b broadcasting will always deliver the message to all nodes. In all executions involving this network, we assume that the adversary resolves the communication non-determinism as follows: If more than one node transmit a message in the same time slot, all messages reach all nodes and thus all nodes receive $\top$ from the collision detector. If a single node in T transmits a message, the message reaches only the nodes in $T \cup \{b\}$ and thus r receives $\perp$. If only r or only b transmits a message, the message reaches all nodes.

Assume for the sake of contradiction that there exists an algorithm $\mathcal{A}$ that solves broadcast in network (G, G') in $n-2$ rounds. For every $i \in [n-1]$, we fix an execution α_i of $\mathcal{A}$, with our specified adversary, in which $\mathrm{ID}_b = i$, $\mathrm{ID}_r = n$, and the node v with $\mathrm{ID}_v = 1$ is the source. We claim that for every $k \in [n-2]$, there exists a subset $I_k \subseteq [n-1]$ such that: (a) $|I_k| \geq n-1-k$; (b) for every $i \in I_k$, process i does not broadcast alone in the first k slots of α_i; and (c) for every $j \in [n]$, and every $i_1, i_2 \in I_k$, process j is in the same state after k rounds in both α_{i_1} and α_{i_2}.

We prove the claim by induction on k. The base case ($k = 0$) is trivial. Assume the hypothesis holds for some $k < n-2$. Let B be the set of processes that broadcast in round $k+1$ for every execution α_i where $i \in I_k$. (By condition (c), this set is the same in all such executions.) If $B = \{i\}$, $i \in I_k$, then we define $I_{k+1} = I_k \setminus \{i\}$. Otherwise, we set $I_{k+1} = I_k$. By construction, conditions (a) and (b) hold for I_{k+1}, we focus, therefore, on condition (c). Fix some process j. We show that j receives the same thing in round $k+1$ of every I_{k+1} execution and therefore is in the same state after this round in each of these executions. If $|B| > 1$, then j receives $\top$ in each execution, if B contains only a single process j' assigned to a node in T, then every process except n receives the message while n receives $\perp$, and, finally, if $B = \{n\}$ then all nodes receive its message. To conclude the proof, we consider some $i \in I_{n-2}$. By our above claim, the bridge process i does not broadcast alone in the first $n-2$ rounds of α_i, preventing process n from receiving the message during this interval. This contradicts the correctness of $\mathcal{A}$. $\square$

We now prove a similar result holds for randomized algorithms.

THEOREM 2. *There exists a 2-broadcastable network (G, G') such that there does not exist a randomized algorithm $\mathcal{A}$ and integer k, $1 \leq k \leq n-2$, where $\mathcal{A}$ solves broadcast in k rounds in (G, G') with probability greater than $k/(n-1)$.*

PROOF. Fix (G, G') as in Theorem 1, and some randomized algorithm $\mathcal{A}$ and integer k, $1 \leq k \leq n-2$. Assume for contradiction that $\mathcal{A}$ solves broadcast in k rounds in (G, G') with probability greater than $k/(n-1)$. Recall that for any deterministic algorithm, our proof of Theorem 1 exhibits a subset $I_k \subseteq [n-1]$ with $|I_k| \geq n-1-k$ such that, for every $i \in I_k$, i does not broadcast alone in the first k rounds in α_i. For the randomized algorithm $\mathcal{A}$, each way of fixing the random choices (using a predetermined choice sequence) yields a deterministic algorithm, and so also yields such a subset I_k with the same properties with respect to the α_i defined for these fixed choices. From the probability distribution of random choices we derive a probability distribution of subsets I_k, each with at least $(n-1-k)$-elements from $[n-1]$. There must exist some process i such that a subset I_k chosen by this distribution includes i with probability at least $(n-1-k)/(n-1)$. This i does not broadcast alone in the first k rounds in any of the executions associated with these subsets. Have the adversary fix $\mathrm{ID}_b = i$. It follows that with probability at least $(n-1-k)/(n-1)$ the random choices prevent process i from broadcasting alone in the first k rounds. Because i is the bridge, it follows that with this probability the message does not get to n within this time, contradicting the success probability assumption for $\mathcal{A}$. $\square$

5. CONCLUSION

Theorems 1 and 2 describe fundamental limitations on the power of wireless networks with unreliable communication. The uncertainty induced by this unreliability requires even randomized solutions to run slowly on networks for which fast broadcast is possible.

6. REFERENCES

[1] L. Barrière, P. Fraigniaud, and L. Narayanan. Robust position-based routing in wireless ad hoc networks with unstable transmission ranges. In *Proc. of the 5th Int. Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIALM)*, pages 19–27, 2001.

[2] D. Kotz, C. Newport, R. S. Gray, J. Liu, Y. Yuan, and C. Elliott. Experimental evaluation of wireless simulation assumptions. In *Proceedings of the ACM/IEEE International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, pages 78–82, October 2004.

[3] F. Kuhn, N. Lynch, and C. Newport. The abstract MAC layer. Technical report, MIT, 2009.

[4] F. Kuhn, R. Wattenhofer, and A. Zollinger. Ad hoc networks beyond unit disk graphs. *Wireless Networks*, 14(5):715–729, 2008.

[5] S. Obayashi and J. Zander. A body-shadowing model for indoor radio communication environments. *IEEE Transactions on Antennas and Propagation*, 46(6):920–927, 1998.