

A (Truly) Local Broadcast Layer for Unreliable Radio Networks

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

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

ABSTRACT

In this paper, we implement an efficient *local broadcast* service for the dual graph model, which describes communication in a radio network with both reliable and unreliable links. Our local broadcast service offers probabilistic latency guarantees for: (1) message delivery to all reliable neighbors (i.e., neighbors connected by reliable links), and (2) receiving *some* message when one or more reliable neighbors are broadcasting. This service significantly simplifies the design and analysis of algorithms for the otherwise challenging dual graph model. To this end, we also note that our solution can be interpreted as an implementation of the *abstract MAC layer* specification—therefore translating the growing corpus of algorithmic results studied on top of this layer to the dual graph model. At the core of our service is a *seed agreement* routine which enables nodes in the network to achieve “good enough” coordination to overcome the difficulties of unpredictable link behavior. Because this routine has potential application to other problems in this setting, we capture it with a formal specification—simplifying its reuse in other algorithms. Finally, we note that in a break from much work on distributed radio network algorithms, our problem definitions (including error bounds), implementation, and analysis do not depend on global network parameters such as the network size, a goal which required new analysis techniques. We argue that breaking the dependence of these algorithms on global parameters makes more sense and aligns better with the rise of ubiquitous computing, where devices will be increasingly working locally in an otherwise massive network. Our push for locality, in other words, is a contribution independent of the specific radio network model and problem studied here.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless Communication; G.2.2 [Discrete Mathematics]: Graph Theory—*network problems*

Permission to make digital or hard copies of all or part 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
PODC'15, July 21–23, 2015, Donostia-San Sebastián, Spain.
Copyright © 2015 ACM 978-1-4503-3617-8/15/07 ...\$15.00.
DOI: <http://dx.doi.org/10.1145/2767386.2767411>.

Keywords

radio network; local broadcast; abstract MAC layer; dual graph; unreliability

1. INTRODUCTION

In this paper, we implement an efficient *local broadcast* service in the dual graph radio network model [14, 16, 3, 7, 10, 8], which captures wireless communication over both reliable and unreliable links. In more detail, the dual graph model describes a network with two graphs, one for each type of link. In each round, the network topology used for node communication is a combination of the edges from reliable link graph and some subset of the edges from the unreliable link graph, the latter chosen adversarially. As argued in our earlier studies of this setting, the inclusion of unpredictable links in formal models of radio communication is motivated by the ubiquity of such behavior in real networks (e.g., [23]).

Our local broadcast algorithm yields two types of probabilistic latency guarantees: a fast *progress bound*, which bounds the time for a node to receive *something* when one or more of its reliable neighbors (i.e., neighbors connected by reliable links) are transmitting, and a slower but still reasonable *acknowledgment bound*, which bounds the time for a sender to finish delivering a broadcast message to all of its reliable neighbors. The service we implement operates in an ongoing manner, which makes it suitable for use as an abstraction layer for developing higher-level distributed algorithms for unreliable radio networks. To this end, we note that our algorithm can be interpreted as an implementation of the *Abstract MAC Layer* specification [13, 15]. It follows that the growing corpus of results designed to run on top of this abstraction (e.g., [9, 20, 6, 12, 11, 5]) can be composed with our implementation, automatically porting these existing solutions to the dual graph model for the first time.

More generally speaking, we note that since 2009 [14] we have waged (algorithmic) battle with the complexities introduced by the presence of unpredictable link behavior in the dual graph model—describing upper and lower bounds for a variety of problems [14, 16, 3, 7, 10, 8]. This paper can be seen as the culmination of this half-decade of effort, in which we integrate what we have learned into a powerful abstraction that renders this realistic but difficult setting tractable to a wider community of algorithm designers.

True Locality. Before proceeding to the details of our results, we must first emphasize an important property of our service: its implementation is “truly local,” by which we mean that its specification, time complexity, and error

bounds are expressed independent of global parameters such as network size. To make this work: we define our correctness and performance properties locally, in terms of individual nodes, rather than globally, in terms of all the nodes; we express our time complexity bounds with respect to local properties, such as bounds on local neighborhood size, not the full network size, n ; and we capture our error probabilities as generic ϵ parameters, rather than the common approach of bounding the error in terms of $(1/n^c)$, for some constant $c \geq 1$. Though locality of this type has been well-studied in other network models (e.g., [19, 17]), it is less understood in the wireless setting—a deficiency we believe must be addressed. There are two justifications for this belief.

First, the common practice of seeking “high probability” (i.e., an error probability bound of the form n^{-c}) seems unnatural for most deployment scenarios—why should you have to grow your network size to decrease your error probability?¹ Of course, by instead introducing a generic error parameter, ϵ , as we do in this paper, we are *not* eliminating the possibility of high probability error bounds when useful, as you can simply set $\epsilon = (1/n^c)$ if needed. But we believe that one should try to postpone considering such dependence until it is really necessary. Second, as we see an increasing emphasis on an *Internet of Things* style ubiquity in the wireless setting [1], global properties such as network size can grow to a massive scale. In studying local algorithms for such scenarios, it is important that we separate time complexity and error guarantees from global parameters and instead reorient them toward the relevant local conditions. This paper provides an example of what is required for a reorientation of this type.

Results. Our local broadcast service is parametrized by two time bounds, t_{ack} and t_{prog} , and an error bound, ϵ . It assumes a dual graph model with an oblivious link scheduler (i.e., decisions about which unreliable links to include in the topology in each round are made at the beginning of the execution), and a natural geographic constraint that requires all nodes within distance 1 be connected by a reliable edge, and no nodes at distance more than $r \geq 1$ be connected by an unreliable edge. (We typically assume r is constant and therefore omit it in our result summary below. For the sake of completeness, however, we keep r as a parameter in our analysis during all steps leading to these results.) Our service guarantees: (1) for each broadcast of a message m by a node u , with probability at least $1 - \epsilon$, every reliable neighbor of u will receive m within t_{ack} rounds; and (2) for a given receiver v and span of t_{prog} rounds, such that at least one reliable neighbor of v is broadcasting throughout the entire span, the probability that v *fails* to receive at least one message during these rounds is no more than ϵ . We present an algorithm that takes ϵ as a parameter and implements this service for values of t_{ack} and t_{prog} bounded as follows:

- $t_{prog} = O\left(\log \Delta \log\left(\frac{\log^4 \Delta}{\epsilon}\right)\right)$
- $t_{ack} = O\left(\Delta \log(\Delta/\epsilon) \log \Delta \log\left(\frac{\log^4 \Delta}{\epsilon}\right)\left(\frac{1}{1-\epsilon}\right)\right)$

where Δ is an upper bound on node degree. We emphasize that these results are near optimal, as even in the absence of unreliable links: (1) any progress bound (which reduces to symmetry breaking among an unknown set of nodes) requires logarithmic rounds (e.g., [21]); and (2) any acknowledgement bound requires at least Δ rounds in the worst case (imagine a receiver u neighboring Δ broadcasters: u can only receive one message per round, delaying one of these broadcasters by at least Δ rounds).

Discussion. A core difficulty in solving local broadcast in the dual graph model is the presence of the unreliable links, which are included or excluded in the network topology according to an arbitrary link schedule (see the model section below for details). To understand this difficulty, recall that the standard strategy for broadcast in a radio network is to cycle through a fixed schedule of geometrically decreasing broadcast probabilities [2]. The intuition for this fixed schedule approach is that for each receiver, one of these probabilities will be well-suited for the local contention among its nearby broadcasters. In the dual graph model, however, there is no fixed amount of local contention: the link schedule can effectively change this amount at each receiver at each round by changing the edges included in the network topology. It is possible, for example, that the link schedule was constructed with the intent of thwarting this fixed schedule strategy by including many links (i.e., increasing contention) when the schedule selects high probabilities, and excluding many links (i.e., decreasing contention) when the schedule selects low probabilities.

To overcome this difficulty, we use as a starting point the general strategy we originally identified in [10]: permute the broadcast probability schedule *after* the execution begins (and therefore, *after* the link schedule has already been generated) to regain independence from the topology. The challenge in permuting a broadcast probability schedule, however, is coordinating the nodes sufficiently that they can apply the same permutation. This creates a chicken and egg problem: to share information among processes requires that we solve broadcast, but we are sharing this information to help solve broadcast. Our solution is to instead solve a form of loose coordination we call *seed agreement*. This problem assumes each participant generates a *seed* (for our purposes, this seed will be random bits for use in generating a probability permutation), and then attempts to convince nearby nodes to *commit* to its seed. A solution to this problem must guarantee that every node commits to some nearby seed (perhaps its own), and, crucially, that there are not too many unique seeds committed in any node’s neighborhood (in this paper, we achieve a bound on the order of $\log(1/\epsilon)$, for error probability ϵ). If nodes subsequently use these seeds to permute their broadcast probability schedules (thereby gaining independence from the link schedule), we are assured that there are not *too many* different schedules in a given neighborhood. An extra conflict resolution mechanism is then introduced to our broadcast strategy to help resolve contention among these competing schedules.

We note that the seed agreement subroutine provides a general strategy for taming adversarial link scheduling, and is therefore of independent interest. Accordingly, in this

¹The most likely answer for why these high probability bounds persist is that they make life easier for the algorithm designer. In particular, if a property holds with high probability in n , a basic union bound provides that the property holds for all nodes in a network for all rounds of any reasonable length execution, which greatly simplifies subsequent analysis. We believe, however, that the algorithm designer should do more work to make the treatment of error probability more natural to the practitioner using the algorithms.

paper we provide a standalone formal specification for the problem that is satisfied by our algorithm. This simplifies the process of subsequently integrating our seed agreement solution into other algorithms.

Finally, we note that to solve these problems in the absence of global parameters such as n requires the introduction of new and non-trivial proof techniques, also of independent interest for the general pursuit of true locality in radio network algorithms. For example, in analyzing our seed agreement subroutine, we could not simply assume that certain key parameters involving local contention hold for the whole network, as this would require a dependence on n . We instead established a *region of goodness* surrounding the target node in our analysis. Using a careful induction on algorithm steps, we showed that although the guaranteed radius of this region must contract as time advances (due to the influence of nodes outside the region for whom we make no assumptions), this contraction is slow enough that our target node safely completes its computation in time. Similarly, in analyzing the local broadcast routines that leverage the seed agreement bits, we had to leverage a new notion of “high probability” that is defined with respect to node degree, not network size (this is the source of the $\log(\Delta/\epsilon)$ factors in the t_{ack} bound described above).

Related Work. The dual graph model of unreliable radio communication was introduced by Clementi et al. [4], and subsequently given the name “dual graph” by Kuhn et al. [14]. It has since been well-studied [14, 16, 3, 7, 10, 22]. Under the pessimistic assumption of an adaptive adversary (as oppose to the oblivious adversary considered in this paper), we previously explored bounds for global broadcast [14, 16], local broadcast [7], and graph structuring algorithms [3, 22]. In [10], we studied the impact of different link scheduler models by proving that some of the more pessimistic bounds from our previous work depended on the assumption that the link schedule was constructed by an adaptive adversary. Of particular relevance to this paper, we proved in [10] that local broadcast with efficient progress is impossible with an adaptive link scheduler of this type, but is feasible with an oblivious link schedule. To establish this latter point, we designed and analyzed a one-shot local broadcast algorithm that offers a progress guarantee (i.e., every node that neighbors a broadcaster will get *some* message quickly) but no reliability guarantees (i.e., no particular message is guaranteed to be delivered). The algorithm in [10] introduced the basic ideas that we developed into the seed agreement specification and *SeedAlg* algorithm presented in this paper. We also note that all results [10] depend on global parameters, whereas here we invest significant effort in gaining true locality.

The abstract MAC layer [13, 15] is an approach to designing wireless network algorithms that defines an abstraction that captures the main guarantees of most wireless link layers, and then divides the task of algorithm design into two efforts: (1) describing and analyzing algorithms that run on top of the abstraction, and (2) implementing the abstraction in specific low-level wireless models. Our local broadcast problem was defined with the standard parameters of a (probabilistic) abstract MAC layer in mind. Our algorithm, therefore, can be interpreted as a strategy for implementing this layer in the dual graph radio network model, and therefore providing a way to translate to the dual graph low level model the growing corpus of algorithms designed and

analyzed onto of the abstract MAC layer [9, 20, 6, 12, 11, 5]. We note, however, that the translation from our algorithm to an abstract MAC layer implementation is not immediate, as some (presumably straightforward) work will be required to mediate between our definition, expressed in terms of low-level details like rounds and receiving messages, and the higher level specification of the abstract MAC layer, which is usually expressed only in terms of the ordering and timing of input and output events.

2. THE DUAL GRAPH MODEL

We use a radio network model based on dual graphs, which describes randomized algorithms executing in a synchronous multihop radio network with both reliable and unreliable links. The model describes the network topology with a dual graph (G, G') , where $G = (V, E)$, $G' = (V, E')$, and $E \subseteq E'$, where E describes reliable links and $E' \setminus E$ describes unreliable links. We use n to denote $|V|$, the number of vertices in the graphs. For $u \in V$, we write $N_G(u)$ ($N_{G'}(u)$) to denote u 's immediate neighbors in G (G'), not including u itself. We assume two degree bounds: Δ , an upper bound on $|N_G(u) \cup \{u\}|$, and Δ' , an upper bound on $|N_{G'}(u) \cup \{u\}|$, defined over every u .

An *embedding* of a (finite) set V of graph vertices in the Euclidean plane is simply a mapping $emb : V \rightarrow \mathbb{R}^2$; this provides a pair of (x, y) coordinates for each vertex V . If emb is an embedding of the vertices V of a dual graph (G, G') and r is a real number, $r \geq 1$, then we say that (G, G') is *r-geographic* with respect to emb provided that, for every $u, v \in V, u \neq v$, the following conditions hold (where d represents Euclidean distance):

1. If $d(emb(u), emb(v)) \leq 1$ then $\{u, v\} \in E$.
2. If $d(emb(u), emb(v)) > r$, then $\{u, v\} \notin E'$.

In other words, nearby vertices must be neighbors in G , and distant vertices cannot even be neighbors in G' , but vertices in the grey zone represented by the intermediate distances in $(1, r]$ might or might not be neighbors in G or G' . We say that (G, G') is *r-geographic* provided that there exists an embedding emb of the vertex set V such that (G, G') is *r-geographic* with respect to emb . We sometimes also say that (G, G') is *geographic* provided that there exists a real $r \geq 1$ such that (G, G') is *r-geographic*.

The dual graphs we consider in this paper are assumed *r-geographic*, for some particular r , which we fix for the rest of the paper. Moving forward, fix an space I . An *algorithm* is an injective mapping $proc()$ from I to some set of *processes*, which are some type of probabilistic timed automata that model wireless devices. Thus, $proc(i)$ denotes the process with id i . We assume that each process knows (e.g., has in a special component of its initial state) its own id, and also knows the quantities Δ , and Δ' . Notice, we do not assume processes know n (as is typical in such models) as we seek problem definitions and solutions that operate independently of the network size. A *process assignment* for a dual graph (G, G') and id space I is an injective mapping $id()$ from V to I , that assigns a different id to each graph vertex. The two mappings, $proc$ and id , together serve to assign a distinct process to each graph vertex. That is, $proc(id(u))$ is the process assigned to graph vertex u . To simplify terminology, we often write *node* u to indicate $proc(id(u))$ or

process i to indicate $proc(i)$. We assume that processes do not know the $id()$ mapping in advance.

An execution of an algorithm in a given dual graph network topology (G, G') proceeds in synchronous rounds $1, 2, \dots$. In each round t , each node decides whether to transmit a message or receive, based on its process definition; this might involve a random choice. The communication topology in round t consists of the edges in E plus an arbitrary subset of the edges in $E' \setminus E$. This subset, which can change from round to round, is determined by an adversary that we call a *link scheduler* (see below). Once the topology is fixed for a given round, we use the following standard collision rules to determine communication: node u receives a message m from node v in round t , if and only if: (1) u is receiving; (2) v is transmitting m ; and (3) v is the only node transmitting among the neighbors of u in the communication topology chosen by the link scheduler for round t . If node u does not receive a message, then we assume that it receives a special “null” indicator \perp : that is, we assume no collision detection.

We now formalize the notion of a *link scheduler*: the entity responsible for resolving the non-determinism concerning which edges from $E' \setminus E$ are added to the topology in each round. Formally, we define a link scheduler to be a sequence $\mathbb{G} = G_1, G_2, G_3, \dots$, where each G_t (also denoted $\mathbb{G}[t]$) is the graph used for the communication topology in round t . We assume each G_t is allowable given our above model definition.² We assume the link scheduler for a given execution is specified at the beginning of the execution. Notice, this definition implies oblivious behavior concerning the network dynamism, as all decisions on the topology are made at the beginning of an execution.

The other relevant source of non-determinism in our model is the *environment* which we use to provide inputs and receive outputs as required by a specific problem (when relevant). For example, in solving local broadcast, an environment provides the messages to broadcast, whereas for a problem like consensus, it provides the initial values. The details of what defines a *well-formed* environment is specified on a problem-by-problem basis. As with the scheduler, when analyzing an execution we first fix the environment for the execution. Though it is possible to conceive of an environment as a probabilistic entity, for the sake of simplicity, the environments we consider in this paper are all deterministic (i.e., once you fix an environment for an execution, all non-determinism regarding inputs is resolved). To formally model the interaction with an environment, we break down the synchronous steps each process takes within a round as follows: first all processes receive inputs (if any) from the environment, next all processes that decide to transmit do so, they then all receive, and finally, they generate outputs (if any) which are processed by the environment to end the round.

We call the combination of a dual graph, process assignment, link scheduler, and environment a *configuration*. Notice, a configuration resolves all relevant model and problem nondeterminism. It follows that a configuration combined with a probabilistic algorithm defines a distribution over possible executions (which we sometimes call the execution tree). When we say an algorithm satisfies a certain property with a given probability, we mean that for all allowable configurations, in the execution tree that results from com-

²That is, it is a graph that includes all the nodes and edges of G with (perhaps) some edges from $E' \setminus E$ also included.

binning the algorithm with the configuration, the property is satisfied with that probability.

3. SEED AGREEMENT

The *seed agreement* problem provides a loose form of coordination: each participating node u generates a *seed* s from some known seed domain S , and then eventually commits to a seed generated by a node in its neighborhood (perhaps its own). The safety goal is to bound the number of unique seeds committed in any given neighborhood by a sufficiently small factor δ , while the liveness goal is to do so in a minimum of rounds. In this section, we provide a dual graph algorithm that yields a bound δ that is roughly $O(\log(\frac{1}{\epsilon}))$, and that operates within time that is polynomial in $\log(\Delta)$ and $\log(\frac{1}{\epsilon})$, with (provided) error probability ϵ .

In Section 4, we use seed agreement as a crucial subroutine in our local broadcast service implementation. It is potentially useful, however, to any number of problems in the dual graph model (and elsewhere), so we take our time here to first provide a careful formal specification, which we then satisfy with a new dual graph algorithm. The analysis of our algorithm was rendered particularly tricky by our goal of avoiding dependence on global parameters such as n , and provides some of the main technical contributions of this paper. Due the long length of this analysis, we defer the details to the full version of this paper [18]. The strategy deployed in this section is to bound the rate at which a region of “goodness” (i.e., sufficiently bounded contention) surrounding our target node contracts as the node races toward termination.

3.1 The Seed Agreement Problem

Fix a finite *seed domain* S . We specify the problem as *Seed* (δ, ϵ) , where δ is a positive integer representing the seed partition bound, and ϵ is a small nonnegative real representing an error probability. This specification describes correctness for a system based on some arbitrary system configuration, running according to our execution definition. The specification has no inputs. Its outputs are of the form *decide* $(j, s)_u$, where $j \in I$, $s \in S$, and $u \in V$. This represents a decision by the node at graph vertex i to commit to the seed s proposed by the node with id j (in the following, we call j the *owner* of seed s). We begin with two basic non-probabilistic conditions on the outputs; these must hold in every execution:

1. *Well-formedness*: In every execution, for each vertex u , exactly one *decide* $(*, *)_u$ occurs.
2. *Consistency*: In every execution, for each pair of vertices u_1, u_2 , if *decide* $(j, s_1)_{u_1}$ and *decide* $(j, s_2)_{u_2}$ both occur, then $s_1 = s_2$.

That is, if outputs contain the same owners then they also contain the same seeds. The two remaining conditions are probabilistic. To talk about probabilities of events, we must first specify the probability distribution. As noted in Section 2, the combination of the system configuration fixed above and a given seed agreement algorithm defines a distribution on executions. We state our remaining properties in terms of this distribution. In more detail, we start by considering an agreement property. Let $B_{u,\delta}$ be the event

(in the probability space of executions) that at most δ distinct ids appear as seed-owners in *decide* outputs at nodes in $N_{G'}(u) \cup \{u\}$.

3. Agreement: For each vertex u , $\Pr(B_{u,\delta}) \geq 1 - \epsilon$.

Note that we state Condition 3 for each vertex u separately, rather than in terms of all vertices, as in [10]. This change is needed for expressing costs in terms of local parameters.

We now express independence of the choices of seed values corresponding to different owners. An *owner mapping* M_o is a mapping from V to I , that is, an assignment of an (owner) id to each vertex. A *seed mapping* M_s is a mapping from V to S , that is, an assignment of a seed to each vertex. We say that a seed mapping M_s is *consistent* with an owner mapping M_o provided that, if two vertices have the same owner, then also have the same seed. That is, if $M_o(u) = M_o(v)$ then $M_s(u) = M_s(v)$. Let Own_{M_o} be the event in the probabilistic execution that M_o is the owner mapping.

4. Independence: Suppose that M_o is an owner mapping and M_s is a seed mapping, where M_s is consistent with M_o . Suppose that $\Pr(Own_{M_o}) > 0$. Then, conditioned on Own_{M_o} , the probability that M_s is the seed mapping that appears in the execution is exactly $(\frac{1}{|S|})^{|range(M_o)|}$.

Condition 4 says that the probability of each consistent seed mapping is just what it would be if the seed mapping were determined in the following way: All processes first choose a seed from S , uniformly at random. Then after every process chooses a seed owner, it also adopts the associated seed value.

3.2 A Seed Agreement Algorithm

We now describe our seed agreement algorithm, *SeedAlg*, which takes as its sole parameter, an error bound, ϵ_1 . We will show, in Theorem 3.1, that this algorithm implements *Seed*(δ, ϵ) for values of δ and ϵ that depend on ϵ_1 . Its main strategy is to hold aggressive local leader elections that yield bounded safety violations (i.e., multiple nearby leaders). In the following description, we assume for simplicity that Δ is a power of 2. We also use a “sufficiently large” constant c_4 for the phase length.

Algorithm *SeedAlg*(ϵ_1), for process i at graph vertex u , where $0 < \epsilon_1 \leq \frac{1}{4}$:

The algorithm uses $\log \Delta$ *phases*, each consisting of $c_4 \log^2(\frac{1}{\epsilon_1})$ rounds.

Process i maintains a local variable containing its “initial seed” in S , which it chooses uniformly at random from the seed domain S . It also keeps track of its *status* $\in \{\text{“active”}, \text{“leader”}, \text{“inactive”}\}$, and the current phase number and round number.

Now we describe process i ’s behavior in any particular phase $h \in \{1, \dots, \log \Delta\}$. If *status* = *active* at the beginning of phase h , then process i becomes a leader; i.e., sets *status* to *leader*, with probability $2^{-(\log \Delta - h + 1)}$. Thus, it uses probabilities: $\frac{1}{\Delta}, \frac{2}{\Delta}, \dots, \frac{1}{4}, \frac{1}{2}$, as it progresses through the phases.

If process i becomes a leader at the start of phase h , it immediately outputs *decide*(i, s), where s

is its initial seed. Then, during the remaining rounds of the phase, process i broadcasts (i, s) with probability $\frac{1}{\log(\frac{1}{\epsilon_1})}$ in each round. At the end of the phase, it becomes inactive.

If process i is active but does not become a leader at the start of phase h , then it just listens for the entire phase. If it receives a message containing a pair (j, s) , then it immediately outputs *decide*(j, s) and becomes inactive. If it receives no messages during this phase, then it remains active.

If process i completes all phases and is still active, then it outputs *decide*(i, s), where s is its initial seed.

3.3 Correctness of *SeedAlg*

The analysis of *SeedAlg* contained in the full version [18] culminates with the main theorem below. For the following, recall that r is the value used in defining the r -geographic property assumed of our dual graph (Section 2), and Δ is the maximum node degree in G .

THEOREM 3.1. *SeedAlg*(ϵ_1) satisfies the *Seed*(δ, ϵ) specification, where:

- $\delta = O(r^2 \log(\frac{1}{\epsilon_1}))$, and
- $\epsilon = O(r^4 \log^4(\Delta)(\epsilon_1)^{c_4 r^2})$, for a constant c , $0 < c < 1$.

The algorithm takes $O((\log \Delta) \log^2(\frac{1}{\epsilon_1}))$ rounds.

4. LOCAL BROADCAST

We now define our local broadcast service, then describe and analyze an efficient solution which uses the *SeedAlg* algorithm from Section 3 as a key subroutine.

4.1 The Local Broadcast Problem

The local broadcast problem described here requires nodes to implement an ongoing probabilistic local communication service with timing and reliability guarantees. In more detail, we call the problem $LB(t_{ack}, t_{prog}, \epsilon)$, where $t_{ack} \geq t_{prog} \geq 1$ are integer round bounds, and ϵ is a small real representing an upper bound on the error probability. To define the problem, we must first fix the underlying dual graph network in which it is being solved: $(G = (V, E), G' = (V, E'))$. We then define a set \mathcal{M}_u of possible messages for each $u \in V$. For simplicity, we assume these sets are pairwise disjoint. Let $\mathcal{M} = \bigcup_u \mathcal{M}_u$ be the set of all possible messages. Every node $u \in V$ has a *bcast*(m) _{u} input and *ack*(m) _{u} output, for each $m \in \mathcal{M}_u$. Node u also has a *recv*(m') _{u} output for each $m' \in \mathcal{M}$.

We now restrict the behavior of the environments we consider for this problem. In more detail, we assume that (1) the environment generates each input at a given node u at most once per execution (i.e., each message it passes a node to broadcast is unique), and (2) if the environment generates a *bcast*(m) _{u} input at u at some round r , it must then wait until u subsequently generates a *ack*(m) _{u} output (if ever) before it can generate another *bcast* input at u . To simplify analysis, we restrict our attention to deterministic environments. Therefore, we can assume the environment

is modeled as synchronous deterministic automaton that receives the nodes' *ack* outputs as input, and generates their *bcast* inputs as its output.

In the following, we say a node u is *actively broadcasting* m in round r , if node u received a $bcast(m)_u$ input in some round $r' \leq r$, and through round r , node u has not yet generated a subsequent $ack(m)_u$ output. Similarly, we say u is *active* in a given round if there is some message that u is actively broadcasting during this round. The problem places deterministic and probabilistic constraints on the nodes' output behavior. We begin with the deterministic constraints.

In every execution, the following must always hold:

1. *Timely Acknowledgement.* If node u receives a $bcast(m)_u$ input in round r , it will generate a single corresponding $ack(m)_u$ output in the interval r to $r + t_{ack}$. These are the only *ack* outputs that u generates.
2. *Validity.* If node u performs a $recv(m)_u$ output in some round r , then there exists some $v \in N_{G'}(u)$ such that v is actively broadcasting m in round r .

Recall that if we fix some configuration and an algorithm, the combination yields a well-defined probability distribution (equiv., execution tree) over executions of the algorithm in this configuration. To aid our probabilistic constraint definitions, we first introduce some notation for discussing an execution tree determined by a configuration. When considering any execution from such a tree, we can partition time into *phases* of length t_{prog} starting in the first round. We number these 1, 2, We use the terminology *phase i prefix* to describe a finite execution prefix that includes every round up to the beginning of phase i (i.e., it does not include the first round of phase i but does include the last round—if any—before phase i begins). For a given execution tree, phase i prefix α in this tree, and node u , let A_α^u describe the set of t_{prog} -round extensions of α in which there is a G -neighbor of u that is active throughout every round of phase i , and let B_α^u describe the set of t_{prog} -round extensions where u receives at least one message $m_v \in M_v$ from a node v in a round r such that v is actively broadcasting m_v in r .

We now define two probabilistic constraints that must hold for every configuration:

1. *Reliability.* For every configuration, node u , and r -round execution prefix such that u receives a $bcast(m)_u$ input at the beginning of round r : the probability that every $v \in N_G(u)$ generates a $recv(m)_u$ output before u 's corresponding $ack(m)_u$ output, is at least $1 - \epsilon$.

(Notice: this property leverages the *timely acknowledgement* property which tells us that in every extension of this prefix, u generates an $ack(m)_u$ output within t_{ack} rounds.)

2. *Progress.* For every configuration, node u , phase i , and phase i prefix α in the resulting execution tree: $\Pr(B_\alpha^u \mid A_\alpha^u) \geq 1 - \epsilon$.

4.2 Constants

Below is a summary of the constants used in the algorithm description and analysis that follow. In the following, we assume Δ is a power of two.

- ϵ_1 is the notation used in the below algorithm description to describe the desired error probability.
- ϵ' is the maximum error probability bound that guarantees, given the constraints of Theorem 3.1, that the $SeedAlg(\epsilon')$ algorithm satisfies the $Seed(\delta, \epsilon)$ spec for an $\epsilon \leq \epsilon_1/2$.

Note: given the relationship between $SeedAlg$ and $Seed$'s error bounds, as established in Theorem 3.1, $\epsilon' = \Theta\left(\left(\frac{\epsilon_1}{r^4 \log^4 \Delta}\right)^{(1/(cr^2))}\right)$, where c , $0 < c < 1$, is the constant provided by Theorem 3.1. Because $c < 1$, we can rewrite the bound as $\Theta\left(\left(\frac{\epsilon_1}{r^4 \log^4 \Delta}\right)^{(\gamma/r^2)}\right)$, for some constant $\gamma > 1$.

- $\epsilon_2 = \min\{\epsilon', \epsilon_1\}$.³
- c_2 is a constant used in our analysis of successful receptions of messages.
- c_1 is a constant we use in defining the length of a phase in the algorithm (see T_{prog} below).
- $T_{prog} = \lceil c_1 \cdot r^2 \cdot \log\left(\frac{1}{\epsilon_1}\right) \cdot \log\left(\frac{1}{\epsilon_2}\right) \cdot \log \Delta \rceil = O\left(r^2 \log\left(\frac{1}{\epsilon_1}\right) \log\left(\frac{1}{\epsilon_2}\right) \log \Delta\right)$ is the number of rounds required by our algorithm to ensure progress.
- $\kappa = T_{prog} \cdot \lceil \log\left(r^2 \log\left(\frac{1}{\epsilon_2}\right)\right) \rceil \cdot \log \log \Delta$: the maximum number of bits consumed from a seed agreement seed in a single *phase* of length T_{prog} worth of broadcasting.
- $T_{ack} = \frac{r^2 12 \log(1/\epsilon_2) \log \Delta \ln\left(\frac{2\Delta}{\epsilon_1}\right) \Delta'}{c_2 T_{prog} (1 - \epsilon_1/2)} = \frac{12 \ln\left(\frac{2\Delta}{\epsilon_1}\right) \Delta'}{c_2 c_1 \log(1/\epsilon_1) (1 - \epsilon_1/2)} = O\left(\frac{\Delta \log(\Delta/\epsilon_1)}{(1 - \epsilon_1)}\right)$ is the number of phases a node will spend attempting to send a message that arrives as a *bcast* input.
- Let $T_s = O(\log \Delta \log^2\left(\frac{1}{\epsilon_2}\right))$ be the number of rounds required for the seed agreement algorithm $SeedAlg(\epsilon_2)$ (as provided by Theorem 3.1).

4.3 A Local Broadcast Algorithm

We now describe $LBAlg$: our solution to the local broadcast problem. This description makes use of several constants described in the previous section.

Algorithm $LBAlg(\epsilon_1)$, for process i at vertex u for some real $\epsilon_1, 0 < \epsilon_1 \leq \frac{1}{2}$.

Node u partitions rounds into *phases* of length $T_s + T_{prog}$ rounds. We label these phases 1, 2, 3, During each phase, u can be in one of two states:

³*Note:* For asymptotic concision, we want to ensure that the error probability we use for $SeedAlg$ is no more than ϵ_1 . We cannot simply claim that ϵ' , as defined above, satisfies this constraint because given its relationship to ϵ_1 from above, it is possible that if γ is sufficiently small compared to r^2 , and ϵ_1 is sufficiently small compared to r and $\log \Delta$, that this exponent will be a sufficiently small fraction to increase $\left(\frac{\epsilon_1}{r^4 \log^4 \Delta}\right)$ to something larger than ϵ_1 . This min statement handles this possibility.

receiving or sending. Node u begins the execution in the receiving state. After receiving a $bcast(m)_u$ input, u will spend the next T_{ack} full phases in the sending state (if it receives the $bcast$ input in the middle of a phase, it waits until the beginning of the next phase to switch to the sending state). At the end of the last round of the last of these T_{ack} phases, node u generates an $ack(m)_u$ output, and then returns to the receiving state.

We now define what happens during these phases. At the beginning of each phase, regardless of u 's state, it executes $SeedAlg(\epsilon_2)$ as a subroutine, using the seed set $S_\kappa = \{0, 1\}^\kappa$; i.e., the set describing every bit sequence of length κ . Let $s_u^{(j)}$ be the seed that node u commits in the beginning of phase j . We call the rounds spent at the beginning of a phase running $SeedAlg$ the *preamble* of the phase, and the remaining rounds the *body* of the phase.

Node u 's behavior during the body of a given phase j depends on its state. If it is in the receiving state, it simply receives during each of these rounds. If during one of these rounds, node u receives a message m' that it has not yet previously received, it generates a $recv(m')_u$ output.

On the other hand, if u is in the sending state for this phase, during each of the body rounds, it does the following:

1. Node u consumes $\lceil \log(r^2 \log(\frac{1}{\epsilon_2})) \rceil$ new bits from its seed $s_u^{(j)}$. If all of these bits are 0 (which occurs with probability $a \cdot \frac{1}{r^2 \log(\frac{1}{\epsilon_2})}$, for some $a \in [1, 2)$) it sets its status to *participant*, otherwise it sets its status to *non-participant*.
2. If u is a non-participant, it receives.
3. If u is a participant, it next consumes $\log \log \Delta$ new bits from $s_u^{(j)}$. Let b be the value in $\lceil \log \Delta \rceil$ specified by these bits. The node then uses an independent (with respect to the other processes) local source of randomness (i.e., *not* bits from $s_u^{(j)}$), to generate b bits with uniform randomness. If all b bits are 0 (which occurs with probability 2^{-b}), u broadcasts its message.

As with the receiving state, if during any of these body rounds, node u receives a message m' that it has not yet previously received, it generates a $recv(m')_u$ output.

The above algorithm divides rounds into phases and then runs a seed agreement algorithm at the beginning of each phase to synchronize shared random bits for the remainder of the phase. Notice that there is nothing fundamental about this frequency of seed agreements. In some settings, it might make sense to run the agreement protocol less frequently, and generate seeds of sufficient length to satisfy the demands of multiple phases. Such modifications do not change our worst-case time bounds but might improve an average case cost or practical performance.

4.4 Correctness of LBA_{lg}

Our goal here is to analyze LBA_{lg} to prove the following theorem:

THEOREM 4.1. $LBA_{lg}(\epsilon_1)$ solves the $LB(t_{ack}, t_{prog}, \epsilon_1)$ problem for parameters:

- $t_{prog} = T_s + T_{prog} = O\left(r^2 \log \Delta \log\left(\frac{r^4 \log^4 \Delta}{\epsilon_1}\right)\right)$
- $t_{ack} = (T_{ack} + 1)(T_s + T_{prog}) = O\left(r^2 \Delta \log(\Delta/\epsilon_1) \log \Delta \log\left(\frac{r^4 \log^4 \Delta}{\epsilon_1}\right)\left(\frac{1}{1-\epsilon_1}\right)\right)$

Below is the core lemma on which we build our proof of Theorem 4.1. This lemma bounds the behavior of LBA_{lg} within the scope of a single phase. To do so, we first introduce some useful notation. For a given phase i of an execution, let B_i be the set of nodes that are in sending status during phase i , and $R_i = N_G(B_i)$ be the set of nodes that neighbor B_i in G . Notice, because sending status is fixed for the duration of a given phase, both B_i and R_i are determined at the beginning of phase i and cannot change during the phase. Also recall from the model definitions that Δ' bounds the maximum degree in G' . Using this notation, we specify and prove the following key probabilistic behavior:

LEMMA 4.2. Fix some phase i and an execution prefix through the $(j-1)^{th}$ body round of this phase, for some $j \in \{2, \dots, T_{prog}\}$. Fix nodes u and v , where $u \in R_i$ and $v \in N_G(u) \cap B_i$. Assume the call to $SeedAlg$ at the beginning of phase i in this prefix satisfies $B_{u,\delta}$. Let p_u be the probability that u receives some message in the j^{th} round, and let $p_{u,v}$ be the probability that u receives a message from v in this round. It follows that:

- $p_u \geq \frac{c_2}{r^2 \log(\frac{1}{\epsilon_2}) \log \Delta}$
- $p_{u,v} \geq p_u / \Delta'$

PROOF. Fix some u, v, t and a prefix, as specified by the lemma statement. (Notice, we know that a node v satisfying the constraints of the statement exists due to the assumption that $u \in R_i$, which implies that $|N_G(u) \cap B_i| > 0$.) Let $\mathcal{S} = \{S_1, S_2, \dots, S_k\}$ be a minimum-sized partition of the nodes in $N_{G'}(u) \cap B_i$ such that all nodes in S_j committed to the same seed in the beginning of this phase. Given the lemma assumption that the preamble of this phase satisfies $B_{u,\delta}$, it follows: $k \leq \delta$. Finally, let S_{i_v} be the set from \mathcal{S} that contains v .

We now analyze the next broadcast round. In this round, nodes in B_i use their seeds to decide whether or not to become a participant. In particular, they become a participant with probability $\frac{1}{r^2 \log(\frac{1}{\epsilon_2})} = c/\delta$, for some constant $c > 0$, using bits from their seeds to resolve the random choice. For each $S_j \in \mathcal{S}$, all nodes in S_j make the same decision in each round because they are using bits from the same seed. Let p_{i_v} be the probability that set S_{i_v} decides to be a participant, and *all other* sets in \mathcal{S} decide to be non-participants. To bound p_{i_v} , we apply Lemmata B.17 and B.18 from the analysis of seed agreement in the full version of this paper [18] to obtain the uniformity and independence properties needed to prove the following:

$$\begin{aligned}
p_{i_v} &= (c/\delta)(1 - (c/\delta))^{k-1} \\
&> (c/\delta)(1/4)^{\frac{c(k-1)}{\delta}} \\
&\stackrel{(k \leq \delta)}{>} (c/\delta)(1/4)^c \\
&= \Theta(1/\delta)
\end{aligned}$$

Assume this event—that only S_{i_v} decides to participate from among the sets in \mathcal{S} —occurs. It follows that only nodes in S_{i_v} can potentially broadcast in this round. Let ℓ be the number of links from nodes in S_{i_v} to u included in the network topology for the round by the link scheduler included in our configuration definition. Because S_{i_v} contains the G -neighbor v (by definition), we know that v is connected to u and that therefore $\ell > 0$.

The next thing that happens in this round is that the nodes in S_{i_v} use more random bits from their shared seed to choose a value uniformly from $[\log \Delta]$. If $\ell = 1$, we define the *correct* choice of value from $[\log \Delta]$ to be 1. If $\ell > 1$, we define the *correct* choice to be $\lceil \log \ell \rceil$. By Lemma B.17 (from [18]), we know the nodes in S_{i_v} will choose a value from this set with uniform probability. The probability they choose a correct value with respect to ℓ is therefore at least $1/\log \Delta$.

Assume that this event also occurs. At this point, by assumption, only nodes in S_{i_v} are potential broadcasters. Each such node decides to broadcast with the correct probability, p_c , which, as defined above, is within a factor of 2 of $1/\ell$. Let us consider the possibilities. We first note that with probability at least $1/2$, u will decide to receive in this round (broadcast probability $1/2$, which corresponds to choosing the value 1 from $[\log \Delta]$, is the largest possible broadcast probability).

Assume this event occurs. The probability that *exactly one* neighbor among the ℓ neighbors connected to u subsequently decides to broadcast is constant (as there are ℓ neighbors, each deciding to broadcast with probability $p_c = \Theta(1/\ell)$). To calculate p_u we must now combine all three independent probabilities: the $\Theta(1/\delta)$ probability that S_{i_v} is the only set in \mathcal{S} to participate, the $1/\log \Delta$ probability that S_{i_v} nodes choose the correct value, and the $\Theta(1)$ probability that u decides to receive and exactly one neighbor of u in the topology for the round broadcasts. We combine the constants in these asymptotic expression to define a lower bound on the constant c_2 used in our definition of p_u from the lemma statement.

Now we step back to consider $p_{u,v}$. Whereas we just calculated that there is a constant probability that exactly one node from among ℓ nodes decides to broadcast using broadcast probability p_c , we must now ask the probability that a specific node—i.e., v —is this broadcaster. We can bound this probability as:

$$p_c(1 - p_c)^{\ell-1} = \frac{c'}{\ell} \left(1 - \frac{c'}{\ell}\right)^{\ell-1} \geq \frac{c'}{\ell} \left(\frac{1}{4}\right)^{c'} \geq \frac{1}{4\Delta'},$$

where c' is a constant of size at least 1 used in the definition of p_c . By replacing the constant probability for this final step used in the derivation of p_u above with this new $\approx 1/4\Delta'$ probability, we get the $p_{u,v}$ bound required by the lemma statement. (As a slight technicality, we omit the $1/4$ in the $\frac{1}{4\Delta'}$ calculation above in this final $p_{u,v}$ bound, as this can

be captured by adjusting the constant c_2 calculated for p_u by a factor of 4 to include this extra amount.) \square

We can now draw on Lemma 4.2 to prove the progress and reliability properties required by the LB specification. We begin with progress:

LEMMA 4.3. *LBAIlg(ϵ_1) solves the $LB(t_{ack}, t_{prog}, \epsilon_1)$ problem for: $t_{prog} = T_s + T_{prog}$.*

PROOF. Notice that this definition of t_{prog} is the same as the length used by the phases in our algorithm. It follows that the boundaries of the phases in the progress property align with the phase boundaries used by $LBAIlg$, so we can refer to both types of phases interchangeably. To prove progress, therefore, it is sufficient to show that for any node u and phase i such that u has an active G neighbor, the probability that u receives at least one broadcast message in this phase is at least $1 - \epsilon_1$.

To do so, fix some node u , phase i , and phase i prefix α such that A_α^u is non-empty; i.e., there is at least one G neighbor of u that is active throughout phase i . Let α' be the extension of α through the call to *SeedAlg* at the beginning of phase i . By Theorem 4.16, and the definition of the ϵ_2 error parameter passed to *SeedAlg*, this call to *SeedAlg* in this α' extension satisfies $B_{u,\delta}$ with probability at least $1 - (\epsilon_1/2)$.

Assume our above assumptions (including the assumption that $B_{u,\delta}$ is satisfied) hold. It follows that Lemma 4.2 applies with respect to u for all T_{prog} of the subsequent body rounds in phase i . Let $p_u^{(j)}$ be the (independent) probability that u receives a message in body round j of the phase. Notice, that Lemma 4.2 tells us that $p_u^{(j)} \geq p_u \geq \frac{c_2}{r^2 \log(\frac{1}{\epsilon_2}) \log \Delta}$ for each such round. We can now bound the probability that u fails to receive a message in all T_{prog} body rounds as:

$$\begin{aligned}
p_{fail} &= \prod_{j=1}^{T_{prog}} (1 - p_u^{(j)}) \leq \prod_{j=1}^{T_{prog}} (1 - p_u) < \\
&= (1/e)^{T_{prog} p_u} = (1/e)^{c_1 c_2 \cdot \log(1/\epsilon_1)}.
\end{aligned}$$

It is straightforward to show that for sufficiently large values of constants c_1 and c_2 , we get $p_{fail} \leq \epsilon_1/2$. To conclude the proof, we use a union bound to show that the probability that $B_{u,\delta}$ does not hold and/or the probability that u fails to receive a message when this property does hold, is less than ϵ_1 : providing the needed $1 - \epsilon_1$ probability for u receiving a least one message in phase i . \square

We now turn our attention to the reliability property of our local broadcast problem:

LEMMA 4.4. *LBAIlg(ϵ_1) solves the $LB(t_{ack}, t_{prog}, \epsilon_1)$ problem for: $t_{ack} = (T_{ack} + 1)(T_{prog} + T_s)$.*

PROOF. Fix some nodes u and v that are neighbors in G . Let $k = \ln(\frac{2\Delta}{\epsilon_1})/p$, for $p = \frac{c_2}{r^2 \log(\frac{1}{\epsilon_2}) \log \Delta}$ (i.e., p is the lower bound for $p_{u,v}$ from the statement of Lemma 4.2). In the following, we define a body round to be *useful* with respect to u , if it occurs in a phase such that $B_{u,\delta}$ holds for the preceding *SeedAlg* preamble. Let p_1 be the probability that u fails to receive a message m from v during k useful rounds in which v is active with m . Applying Lemma 4.2 to

lower bound the receive probability in each of these rounds, it follows:

$$p_1 \leq (1-p)^k < (1/e)^{pk} = \frac{\epsilon_1}{2\Delta}.$$

We now investigate the number of phases necessary to ensure that v experiences at least k useful rounds with a sufficiently high probability. To do so, we first fix $q = \lceil \frac{12k}{T_{prog}(1-\epsilon_1/2)} \rceil$. Consider an experiment where we run q consecutive phases. Let X_i , for $i \in [q]$, be a random variable that describes the number of useful rounds in phase i of the experiment. Notice, X_i either takes on the value T_{prog} (with probability at least $1 - \epsilon_1/2$) or 0 (else). Let $Y = X_1 + X_2 + \dots + X_q$ be the total number of useful rounds generated by the experiment. It follows:

$$\mathbb{E}[Y] = q \cdot T_{prog} \cdot (1 - \epsilon_1/2) = 12k.$$

We now apply a Chernoff Bound, to bound the probability that Y is more than a factor of 2 smaller than its expectation $\mu = \mathbb{E}[Y]$:

$$\Pr(Y < (1/2)\mu = 6k) < e^{-\frac{\mu}{12}} = e^{-k} \leq e^{-\ln(\frac{2\Delta}{\epsilon_1})} = \epsilon_1/(2\Delta).$$

(Notice, in the above we can bound $e^{-k} \leq e^{-\ln(\frac{2\Delta}{\epsilon_1})}$ because $k \geq \ln(\frac{2\Delta}{\epsilon_1})$.)

We have now established that with probability at least $1 - \epsilon_1/(2\Delta)$, u will experience at least $6k > k$ useful rounds in q phases. We earlier established that if u experiences at least k useful rounds during which v is broadcasting m , then u receives m from v with probability at least $1 - \epsilon_1/(2\Delta)$. Assume u broadcasts m for at least q consecutive phases. By a union bound, the probability that both events occur with respect to these phases, and u therefore receives v 's message m , is greater than $1 - \epsilon_1/\Delta$.

To conclude, we want to calculate, under the assumption that v broadcasts m for at least q phases, that every G neighbor of v succeeds in receiving m . Because there are at most Δ such neighbors, and each succeeds with probability at least $1 - \epsilon_1/\Delta$, a union bound says that every neighbor succeeds with probability at least $1 - \epsilon_1$, as required by the reliability property.

To satisfy reliability, therefore, it is sufficient for any node v receiving a $bcast(m)_u$ input to spend at least q full phases with sending status. Notice, this is exactly what $LBAIlg$ requires, as by definition it has v spend the next $T_{ack} = q$ full phases after a $bcast(m)_v$ input in sending status. The t_{ack} in the lemmas statement is defined to be long enough for v to wait up to a full phase length until the next phase boundary, plus the rounds required for an additional q phases. \square

We now pull together the pieces to prove Theorem 4.1:

PROOF PROOF (OF THEOREM 4.1). The definition of the local broadcast problem had four conditions, two deterministic and two probabilistic. We consider each in turn and argue that $LBAIlg$ satisfies the conditions for the parameter values specified in the theorem statement.

We first note that *timely acknowledgment* holds because $LBAIlg$, by definition, has each node generate an *ack* in response to a *bcast* within a fixed number of rounds that is strictly less than the t_{ack} factor from the theorem statement. Similarly, the *validity* condition holds as $LBAIlg$, by definition, only has nodes broadcast messages they received in a

bcast input, and nodes only *recv* messages that they actually received from another node. Moving on to the probabilistic properties, Lemma 4.3 tells us that $t_{prog} = T_s + T_{prog}$ rounds, and that $t_{ack} = (T_{ack} + 1)t_{prog}$ rounds. Notice, T_{ack} shows up in t_{ack} unchanged from its definition. The definition of t_{prog} , however, shows up in a form that is simplified as compared to the definition provided for T_s and T_{prog} (which contain both ϵ_1 and ϵ_2 factors). To match the bounds in the Theorem statement, therefore, it is sufficient to show that $T_s + T_{prog} = O(\log \Delta \log^2(\frac{1}{\epsilon_2}) + r^2 \log \Delta \log(\frac{1}{\epsilon_2}) \log(\frac{1}{\epsilon_1}))$ can be upper bounded by $O(r^2 \log \Delta \log(\frac{r^4 \log^4 \Delta}{\epsilon_1}))$. We dedicate the remainder of this proof to this effort.

By definition, $\epsilon_2 \leq \epsilon_1$. We can, therefore, substitute the former for the latter in our sum, yielding:

$$T_s + T_{prog} = O(r^2 \log \Delta \log^2(\frac{1}{\epsilon_2})).$$

We need a bound, however, that is expressed with respect to the problem parameter ϵ_1 . This requires us to dive deeper into the relationship between ϵ_2 and ϵ_1 . There are two cases to consider given our definition above that $\epsilon_2 = \min\{\epsilon', \epsilon_1\}$. The first case is that $\epsilon_2 = \epsilon_1$. If this is true, we can simply replace ϵ_2 with ϵ_1 in our above equation, and the result is clearly upper bounded by $O(r^2 \log \Delta \log(\frac{r^4 \log^4 \Delta}{\epsilon_1}))$ (which strictly increases the log factor by adding the $r^4 \log \Delta$ term).

The second case is that $\epsilon_2 = \epsilon' < \epsilon_1$. By definition of ϵ' , it would then follow that $\epsilon_2 = \Theta((\frac{\epsilon_1}{r^4 \log^4 \Delta})^{(\gamma/r^2)})$, for some constant $\gamma > 1$. The properties of γ allows us to simplify (asymptotically) the $\log^2(\frac{1}{\epsilon_2})$ term in our above equation as follows:

$$\log^2(\frac{1}{\epsilon_2}) = [\log(\frac{(r^4 \log^4 \Delta)^{(\gamma/r^2)}}{(\epsilon_1)^{(\gamma/r^2)}})]^2 =$$

$$[(\gamma/r^2) \log(\frac{(r^4 \log^4 \Delta)}{(\epsilon_1)})]^2 = O(\log^2(\frac{r^4 \log^4 \Delta}{\epsilon_1})).$$

Notice in the above we simply drop the $(1/r^2)$, as it too is bounded to be at least 1, so dropping it simply increases the value of the upper bound. Now to conclude our argument, we simply substitute this upper bound for $\log^2(\frac{1}{\epsilon_2})$ in our above sum to get the desired equation. \square

5. CONCLUSION

In this paper, we described and analyzed an ongoing local broadcast service for the dual graph model. This service hides the complexities introduced by unpredictable link behaviors and therefore has the potential to significantly simplify the development of distributed algorithms for this challenging setting. As noted in the introduction,

Our solution can also be adapted to implement the abstract MAC layer specification [13, 15], allowing existing results for this abstraction to translate to the dual graph model (e.g., [9, 20, 6, 12, 11, 5]). Though we leave the details of this adaptation to future work, we note that it would likely be straightforward, with the main effort focused on aligning our local broadcast problem definition—which depends on low level model details, such as rounds and receiving messages—with the higher level of the abstract MAC layer,

which is specified in terms of the timing and ordering of input and output events.

Finally, we emphasize that our commitment to *truly local* algorithms, which required us to avoid global parameters in problem definitions, algorithm strategies, and analysis, is of standalone interest to those studying radio network algorithms. As we argued at this paper's opening, a local perspective provides more flexibility to practitioners, and will become increasingly necessary as network sizes grows.

6. ACKNOWLEDGMENTS

This research is supported in part by: Ford Motor Company University Research Program, NSF Award CCF-1320279, NSF Award CCF-0937274, NSF Award CCF-1217506, NSF CCF-0939370, and AFOSR Award Number FA9550-13-1-0042.

7. REFERENCES

- [1] Luigi Atzori, Antonio Iera, and Giacomo Morabito. The Internet of Things: A Survey. *Computer Networks*, 54(15):2787–2805, 2010.
- [2] Reuven Bar-Yehuda, Oded Goldreich, and Alon Itai. On the Time-Complexity of Broadcast in Multi-Hop Radio Networks: An Exponential Gap between Determinism and Randomization. *Journal of Computer and System Sciences*, 45(1):104–126, 1992.
- [3] Keren Censor-Hillel, Seth Gilbert, Fabian Kuhn, Nancy Lynch, and Calvin Newport. Structuring unreliable radio networks. *Distributed Computing*, 27(1):1–19, 2014.
- [4] A. E. F. Clementi, A. Monti, and R. Silvestri. Round Robin is Optimal for Fault-Tolerant Broadcasting on Wireless Networks. *Journal of Parallel and Distributed Computing*, 64(1):89–96, 2004.
- [5] Alejandro Cornejo, Nancy Lynch, Saira Viqar, and Jennifer L Welch. Neighbor Discovery in Mobile Ad Hoc Networks Using an Abstract MAC Layer. In *Annual Allerton Conference on Communication, Control, and Computing*, 2009.
- [6] Alejandro Cornejo, Saira Viqar, and Jennifer L Welch. Reliable Neighbor Discovery for Mobile Ad Hoc Networks. In *Proceedings of the International Workshop on Foundations of Mobile Computing*, 2010.
- [7] Mohsen Ghaffari, Bernhard Haeupler, Nancy Lynch, and Calvin Newport. Bounds on contention management in radio networks. In *The International Symposium on Distributed Computing (DISC)*, 2012.
- [8] Mohsen Ghaffari, Erez Kantor, Nancy Lynch, and Calvin Newport. Multi-message broadcast with Abstract MAC layers and unreliable links. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, 2014.
- [9] Mohsen Ghaffari, Erez Kantor, Nancy Lynch, and Calvin Newport. Multi-message broadcast with abstract mac layers and unreliable links. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, 2014.
- [10] Mohsen Ghaffari, Nancy Lynch, and Calvin Newport. The cost of radio network broadcast for different models of unreliable links. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, 2013.
- [11] Majid Khabbazzian, Fabian Kuhn, Dariusz Kowalski, and Nancy Lynch. Decomposing Broadcast Algorithms Using Abstract MAC Layers. In *Proceedings of the International Workshop on Foundations of Mobile Computing*, 2010.
- [12] Majid Khabbazzian, Fabian Kuhn, Nancy Lynch, Muriel Medard, and Ali ParandehGheibi. MAC Design for Analog Network Coding. In *Proceedings of the International Workshop on Foundations of Mobile Computing*, 2011.
- [13] Fabian Kuhn, Nancy Lynch, and Calvin Newport. The Abstract MAC layer. In *Proceedings of the International Symposium on Distributed Computing (DISC)*, 2009.
- [14] Fabian Kuhn, Nancy Lynch, and Calvin Newport. Brief announcement: Hardness of broadcasting in wireless networks with unreliable communication. In *Proceedings of the ACM Symposium on the Principles of Distributed Computing (PODC)*, 2009.
- [15] Fabian Kuhn, Nancy Lynch, and Calvin Newport. The Abstract MAC layer. *Distributed Computing*, 24(3-4):187–206, November 2011. Special Issue for *DISC 2009: 23rd International Symposium on Distributed Computing*.
- [16] Fabian Kuhn, Nancy Lynch, Calvin Newport, Rotem Oshman, and Andrea Richa. Broadcasting in unreliable radio networks. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, 2010.
- [17] Fabian Kuhn, Thomas Moscibroda, and Rogert Wattenhofer. What cannot be computed locally! In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, 2004.
- [18] Nancy Lynch and Calvin Newport. A (truly) local broadcast layer for unreliable radio networks. Technical Report MIT-CSAIL-TR-2015-016, MIT Computer Science and Artificial Intelligence Laboratory, Cambridge, MA, May 2015.
- [19] Moni Naor and Larry Stockmeyer. What can be computed locally? *SIAM Journal on Computing*, 24(6):1259–1277, 1995.
- [20] Calvin Newport. Consensus with an abstract mac layer. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, 2014.
- [21] Calvin Newport. Lower bounds for radio networks made easy. In *Proceedings of the International Symposium on Distributed Computing (DISC)*, 2014.
- [22] Calvin Newport. Lower bounds for structuring unreliable radio networks. In *Proceedings of the International Symposium on Distributed Computing (DISC)*, 2014.
- [23] Calvin Newport, David Kotz, Yougu Yuan, Robert S Gray, Jason Liu, and Chip Elliott. Experimental Evaluation of Wireless Simulation Assumptions. *Simulation*, 83(9):643–661, 2007.