GUESS: Gossiping Updates for Efficient Spectrum Sensing

Nabeel Ahmed
University of Waterloo
David R. Cheriton School of Computer Science
n3ahmed@uwaterloo.ca

David Hadaller
University of Waterloo
David R. Cheriton School of Computer Science
dthadaller@uwaterloo.ca

Srinivasan Keshav
University of Waterloo
David R. Cheriton School of Computer Science
keshav@uwaterloo.ca

ABSTRACT

Wireless radios of the future will likely be frequency-agile, that is, supporting opportunistic and adaptive use of the RF spectrum. Such radios must coordinate with each other to build an accurate and consistent map of spectral utilization in their surroundings. We focus on the problem of sharing RF spectrum data among a collection of wireless devices. The inherent requirements of such data and the time-granularity at which it must be collected makes this problem both interesting and technically challenging. We propose GUESS, a novel *incremental gossiping* approach to coordinated spectral sensing. It (1) reduces protocol overhead by limiting the amount of information exchanged between participating nodes, (2) is resilient to network alterations, due to node movement or node failures, and (3) allows *exponentially-fast* information convergence. We outline an initial solution incorporating these ideas and also show how our approach reduces network overhead by up to a factor of 2.4 and results in up to 2.7 times faster information convergence than alternative approaches.

Categories and Subject Descriptors

C.2.4 [**Distributed Systems**]: Distributed applications

General Terms

Algorithms, Performance, Experimentation

Keywords

Coordinated Spectrum Sensing, Gossip Protocols,
FM Aggregation, Incremental Algorithms

1. INTRODUCTION

There has recently been a huge surge in the growth of wireless technology, driven primarily by the availability of unlicensed spectrum. However, this has come at the cost of increased RF interference, which has caused the Federal Communications Commission (FCC) in the United States to re-evaluate its strategy on spectrum allocation. Currently, the FCC has licensed RF spectrum to a variety of public and private institutions, termed *primary users*. New spectrum allocation regimes implemented by the FCC use dynamic spectrum access schemes to either negotiate or opportunistically allocate RF spectrum to unlicensed *secondary users* that can use it when the primary user is absent. The second type of allocation scheme is termed *opportunistic spectrum sharing*. The FCC has already legislated this access method for the 5 GHz band and is also considering the same for TV broadcast bands [1]. As a result, a new wave of intelligent radios, termed *cognitive radios* (or software defined radios), is emerging that can dynamically re-tune their radio parameters based on interactions with their surrounding environment.

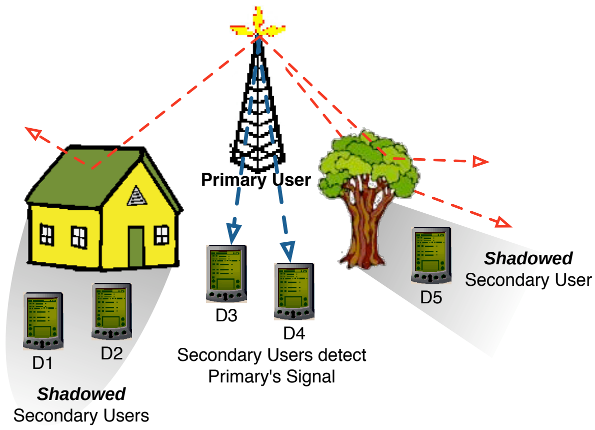

Figure 1: Without cooperation, shadowed users are not able to detect the presence of the primary user.

Under the new opportunistic allocation strategy, secondary users are obligated not to interfere with primary users (senders or receivers). This can be done by sensing the environment to detect the presence of primary users. However, local sensing is not always adequate, especially in cases where a secondary user is *shadowed* from a primary user, as illustrated in Figure 1. Here, coordination between secondary users is the only way for shadowed users to detect the primary. In general, cooperation improves sensing accuracy by an order of magnitude when compared to not cooperating at all [5].

To realize this vision of dynamic spectrum access, two fundamental problems must be solved: (1) Efficient and coordinated spectrum sensing and (2) Distributed spectrum allocation. In this paper, we propose strategies for coordinated spectrum sensing that are low cost, operate on timescales comparable to the agility of the RF environment, and are resilient to network failures and alterations. We defer the problem of spectrum allocation to future work.

Spectrum sensing techniques for cognitive radio networks [4, 17] are broadly classified into three regimes; (1) *centralized coordinated* techniques, (2) *decentralized coordinated* techniques, and (3) *decentralized uncoordinated* techniques. We advocate a decentralized coordinated approach, similar in spirit to OSPF link-state routing used in the Internet. This is more effective than uncoordinated approaches because making decisions based only on local information is fallible (as shown in Figure 1). Moreover, compared to centralized approaches, decentralized techniques are more scalable, robust, and resistant to network failures and security attacks (e.g. jamming).

Coordinating sensory data between cognitive radio devices is technically challenging because accurately assessing spectrum usage requires exchanging potentially large amounts of data with many radios at very short time scales. Data size grows rapidly due to the large number (i.e. thousands) of spectrum bands that must be scanned. This data must also be exchanged between potentially hundreds of neighboring secondary users at short time scales, to account for rapid changes in the RF environment.

This paper presents GUESS, a novel approach to coordinated spectrum sensing for cognitive radio networks. Our approach is motivated by the following key observations:

1. *Low-cost sensors collect approximate data*: Most devices have limited sensing resolution because they are low-cost and low duty-cycle devices and thus cannot perform complex RF signal processing (e.g. matched filtering). Many are typically equipped with simple energy detectors that gather only approximate information.

2. *Approximate summaries are sufficient for coordination*: Approximate statistical summaries of sensed data are sufficient for correlating sensed information between radios, as relative usage information is more important than absolute usage data. Thus, exchanging *exact* RF information may not be necessary, and more importantly, too costly for the purposes of spectrum sensing.

3. *RF spectrum changes incrementally*: On most bands, RF spectrum utilization changes infrequently. Moreover, utilization of a specific RF band affects only that band and not the entire spectrum. Therefore, if the usage pattern of a particular band changes substantially, nodes detecting that change can initiate an update protocol to update the information for that band alone, leaving in place information already collected for other bands. This allows rapid detection of change while saving the overhead of exchanging unnecessary information.

Based on these observations, GUESS makes the following contributions:

1. A novel approach that applies randomized gossiping algorithms to the problem of coordinated spectrum sensing. These algorithms are well suited to coordinated spectrum sensing due to the unique characteristics of the problem: i.e. radios are power-limited, mobile and have limited bandwidth to support spectrum sensing capabilities.

2. An application of in-network aggregation for dissemination of spectrum summaries. We argue that approximate summaries are adequate for performing accurate radio parameter tuning.

3. An extension of in-network aggregation and randomized gossiping to support incremental maintenance of spectrum summaries. Compared to standard gossiping approaches, incremental techniques can further reduce overhead and protocol execution time by requiring fewer radio resources.

The rest of the paper is organized as follows. Section 2 motivates the need for a low cost and efficient approach to coordinated spectrum sensing. Section 3 discusses related work in the area, while Section 4 provides a background on in-network aggregation and randomized gossiping. Sections 5 and 6 discuss extensions and protocol details of these techniques for coordinated spectrum sensing. Section 7 presents simulation results showcasing the benefits of GUESS, and Section 8 presents a discussion and some directions for future work.

2. MOTIVATION

To estimate the scale of the problem, In-stat predicts that the number of WiFi-enabled devices sold annually alone will grow to 430 million by 2009 [2]. Therefore, it would be reasonable to assume that a typical dense urban environment will contain several thousand cognitive radio devices in range of each other. As a result, distributed spectrum sensing and allocation would become both important and fundamental.

Coordinated sensing among secondary radios is essential due to limited device sensing resolution and physical RF effects such as shadowing. Cabric et al. [5] illustrate the gains from cooperation and show an order of magnitude reduction in the probability of interference with the primary user when only a small fraction of secondary users cooperate.

However, such coordination is non-trivial due to: (1) the limited bandwidth available for coordination, (2) the need to communicate this information on short timescales, and (3) the large amount of sensory data that needs to be exchanged.

Limited Bandwidth: Due to restrictions of cost and power, most devices will likely not have dedicated hardware for supporting coordination. This implies that both data and sensory traffic will need to be time-multiplexed onto a single radio interface. Therefore, any time spent communicating sensory information takes away from the device's ability to perform its intended function. Thus, any such coordination must incur minimal network overhead.

Short Timescales: Further compounding the problem is the need to immediately propagate updated RF sensory data, in order to allow devices to react to it in a timely fashion. This is especially true due to mobility, as rapid changes of the RF environment can occur due to device and obstacle movements. Here, fading and multi-path interference heavily impact sensing abilities. Signal level can drop to a deep null with just a $\lambda/4$ movement in receiver position (3.7 cm at 2 GHz), where λ is the wavelength [14]. Coordination which does not support rapid dissemination of information will not be able to account for such RF variations.

Large Sensory Data: Because cognitive radios can potentially use any part of the RF spectrum, there will be numerous channels that they need to scan. Suppose we wish to compute the average signal energy in each of 100 discretized frequency bands, and each signal can have up to 128 discrete energy levels. Exchanging complete sensory information between nodes would require 700 bits per transmission (for 100 channels, each requiring seven bits of information). Exchanging this information among even a small group of 50 devices each second would require (50 time-steps $\times$ 50 devices $\times$ 700 bits per transmission) = 1.67 Mbps of aggregate network bandwidth.

Contrast this to the use of a randomized gossip protocol to disseminate such information, and the use of FM bit vectors to perform in-network aggregation. By applying gossip and FM aggregation, aggregate bandwidth requirements drop to ($c{\cdot}logN$ time-steps $\times$ 50 devices $\times$ 700 bits per transmission) = 0.40 Mbps, since 12 time-steps are needed to propagate the data (with $c = 2$, for illustrative purpoes[1]). This is explained further in Section 4.

Based on these insights, we propose GUESS, a low-overhead approach which uses incremental extensions to FM aggregation and randomized gossiping for efficient coordination within a cognitive radio network. As we show in Section 7,

[1]Convergence time is correlated with the connectivity topology of the devices, which in turn depends on the environment.

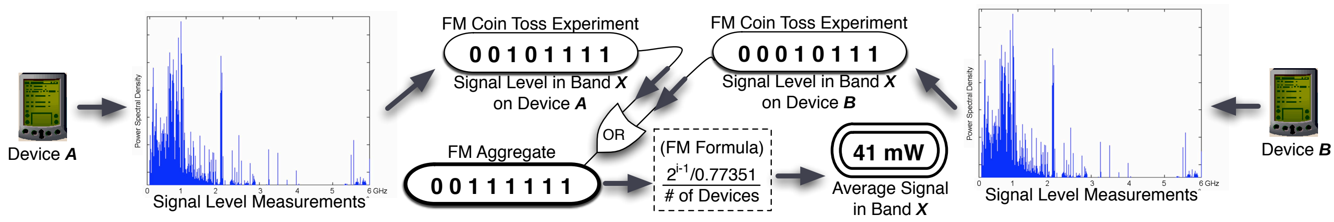

Figure 2: Using FM aggregation to compute average signal level measured by a group of devices.

these incremental extensions can further reduce bandwidth requirements by up to a factor of 2.4 over the standard approaches discussed above.

3. RELATED WORK

Research in cognitive radio has increased rapidly [4, 17] over the years, and it is being projected as one of the leading enabling technologies for wireless networks of the future [9]. As mentioned earlier, the FCC has already identified new regimes for spectrum sharing between primary users and secondary users and a variety of systems have been proposed in the literature to support such sharing [4, 17].

Detecting the presence of a primary user is non-trivial, especially a legacy primary user that is not cognitive radio aware. Secondary users must be able to detect the primary even if they cannot properly decode its signals. This has been shown by Sahai et al. [16] to be extremely difficult even if the modulation scheme is known. Sophisticated and costly hardware, beyond a simple energy detector, is required to improve signal detection accuracy [16]. Moreover, a shadowed secondary user may not even be able to detect signals from the primary. As a result, simple local sensing approaches have not gained much momentum. This has motivated the need for cooperation among cognitive radios [16].

More recently, some researchers have proposed approaches for radio coordination. Liu et al. [11] consider a centralized access point (or base station) architecture in which sensing information is forwarded to APs for spectrum allocation purposes. APs direct mobile clients to collect such sensing information on their behalf. However, due to the need of a fixed AP infrastructure, such a centralized approach is clearly not scalable.

In other work, Zhao et al. [17] propose a distributed coordination approach for spectrum sensing and allocation. Cognitive radios organize into clusters and coordination occurs within clusters. The CORVUS [4] architecture proposes a similar clustering method that can use either a centralized or decentralized approach to manage clusters. Although an improvement over purely centralized approaches, these techniques still require a setup phase to generate the clusters, which not only adds additional delay, but also requires many of the secondary users to be static or quasi-static. In contrast, GUESS does not place such restrictions on secondary users, and can even function in highly mobile environments.

4. BACKGROUND

This section provides the background for our approach. We present the FM aggregation scheme that we use to generate spectrum summaries and perform in-network aggregation. We also discuss randomized gossiping techniques for disseminating aggregates in a cognitive radio network.

4.1 FM Aggregation

Aggregation is the process where nodes in a distributed network combine data received from neighboring nodes with their local value to generate a combined *aggregate*. This aggregate is then communicated to other nodes in the network and this process repeats until the aggregate at all nodes has converged to the same value, i.e. the global aggregate. *Double-counting* is a well known problem in this process, where nodes may contribute more than once to the aggregate, causing inaccuracy in the final result. Intuitively, nodes can tag the aggregate value they transmit with information about which nodes have contributed to it. However, this approach is not scalable. Order and Duplicate Insensitive (ODI) techniques have been proposed in the literature [10, 15]. We adopt the ODI approach pioneered by Flajolet and Martin (FM) for the purposes of aggregation. Next we outline the FM approach; for full details, see [7].

Suppose we want to compute the number of nodes in the network, i.e. the COUNT query. To do so, each node performs a coin toss experiment as follows: toss an unbiased coin, stopping after the first "head" is seen. The node then sets the ith bit in a bit vector (initially filled with zeros), where i is the number of coin tosses it performed. The intuition is that as the number of nodes doing coin toss experiments increases, the probability of a more significant bit being set in one of the nodes' bit vectors increases.

These bit vectors are then exchanged among nodes. When a node receives a bit vector, it updates its local bit vector by bitwise OR-ing it with the received vector (as shown in Figure 2 which computes AVERAGE). At the end of the aggregation process, every node, with high probability, has the same bit vector. The actual value of the count aggregate is then computed using the following formula, $AGG_{FM} = 2^{j-1}/0.77351$, where j represents the bit position of the least significant zero in the aggregate bit vector [7].

Although such aggregates are very compact in nature, requiring only $O(logN)$ state space (where N is the number of nodes), they may not be very accurate as they can only approximate values to the closest power of 2, potentially causing errors of up to 50%. More accurate aggregates can be computed by maintaining multiple bit vectors at each node, as explained in [7]. This decreases the error to within $O(1/\sqrt{m})$, where m is the number of such bit vectors.

Queries other than count can also be computed using variants of this basic counting algorithm, as discussed in [3] (and shown in Figure 2). Transmitting FM bit vectors between nodes is done using randomized gossiping, discussed next.

4.2 Gossip Protocols

Gossip-based protocols operate in discrete time-steps; a time-step is the required amount of time for all transmissions in that time-step to complete. At every time-step, each node having something to send randomly selects one or more neighboring nodes and transmits its data to them. The randomized propagation of information provides fault-tolerance and resilience to network failures and outages. We emphasize that this characteristic of the protocol also allows it to operate **without relying on any underlying network structure**. Gossip protocols have been shown to provide exponentially fast convergence[2], on the order of $O(\log N)$ [10], where N is the number of nodes (or radios). These protocols can therefore easily scale to very dense environments.

[2] *Convergence* refers to the state in which all nodes have the most up-to-date view of the network.

Two types of gossip protocols are:

- **Uniform Gossip**: In uniform gossip, at each time-step, each node chooses a random neighbor and sends its data to it. This process repeats for $O(log(N))$ steps (where N is the number of nodes in the network). Uniform gossip provides exponentially fast convergence, with low network overhead [10].
- **Random Walk**: In random walk, only a subset of the nodes (termed designated nodes) communicate in a particular time-step. At startup, k nodes are randomly elected as designated nodes. In each time-step, each designated node sends its data to a random neighbor, which becomes designated for the subsequent time-step (much like passing a token). This process repeats until the aggregate has converged in the network. Random walk has been shown to provide similar convergence bounds as uniform gossip in problems of similar context [8, 12].

5. INCREMENTAL PROTOCOLS

5.1 Incremental FM Aggregates

One limitation of FM aggregation is that it does not support updates. Due to the probabilistic nature of FM, once bit vectors have been ORed together, information cannot simply be removed from them as each node's contribution has not been recorded. We propose the use of *delete vectors*, an extension of FM to support updates. We maintain a separate aggregate delete vector whose value is subtracted from the original aggregate vector's value to obtain the resulting value as follows.

$$AGG_{INC} = (2^{a-1}/0.77351) - (2^{b-1}/0.77351) \qquad (1)$$

Here, a and b represent the bit positions of the least significant zero in the original and delete bit vectors respectively.

Suppose we wish to compute the average signal level detected in a particular frequency. To compute this, we compute the SUM of all signal level measurements and divide that by the COUNT of the number of measurements. A SUM aggregate is computed similar to COUNT (explained in Section 4.1), except that each node performs s coin toss experiments, where s is the locally measured signal level. Figure 2 illustrates the sequence by which the average signal energy is computed in a particular band using FM aggregation.

Now suppose that the measured signal at a node changes from s to s'. The vectors are updated as follows.

- $s' > s$: We simply perform $(s' - s)$ more coin toss experiments and bitwise OR the result with the original bit vector.
- $s' < s$: We increase the value of the delete vector by performing $(s - s')$ coin toss experiments and bitwise OR the result with the current delete vector.

Using delete vectors, we can now support updates to the measured signal level. With the original implementation of FM, the aggregate would need to be discarded and a new one recomputed every time an update occurred. Thus, delete vectors provide a low overhead alternative for applications whose data changes incrementally, such as signal level measurements in a coordinated spectrum sensing environment. Next we discuss how these aggregates can be communicated between devices using incremental routing protocols.

5.2 Incremental Routing Protocol

We use the following incremental variants of the routing protocols presented in Section 4.2 to support incremental updates to previously computed aggregates.

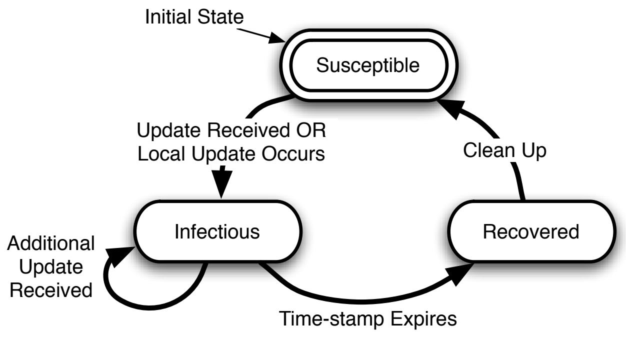

Figure 3: State diagram each device passes through as updates proceed in the system

- **Incremental Gossip Protocol (IGP)**: When an update occurs, the updated node initiates the gossiping procedure. Other nodes only begin gossiping once they receive the update. Therefore, nodes receiving the update become *active* and continue communicating with their neighbors until the update protocol terminates, after $O(log(N))$ time steps.
- **Incremental Random Walk Protocol (IRWP)**: When an update (or updates) occur in the system, instead of starting random walks at k random nodes in the network, all k random walks are initiated from the updated node(s). The rest of the protocol proceeds in the same fashion as the standard random walk protocol. The allocation of walks to updates is discussed in more detail in [3], where the authors show that the number of walks has an almost negligible impact on network overhead.

6. PROTOCOL DETAILS

Using incremental routing protocols to disseminate incremental FM aggregates is a natural fit for the problem of coordinated spectrum sensing. Here we outline the implementation of such techniques for a cognitive radio network. We continue with the example from Section 5.1, where we wish to perform coordination between a group of wireless devices to compute the average signal level in a particular frequency band.

Using either incremental random walk or incremental gossip, each device proceeds through three phases, in order to determine the global average signal level for a particular frequency band. Figure 3 shows a state diagram of these phases.

Susceptible: Each device starts in the susceptible state and becomes infectious only when its locally measured signal level changes, or if it receives an update message from a neighboring device. If a local change is observed, the device updates either the original or delete bit vector, as described in Section 5.1, and moves into the infectious state. If it receives an update message, it ORs the received original and delete bit vectors with its local bit vectors and moves into the infectious state.

Note, because signal level measurements may change sporadically over time, a smoothing function, such as an exponentially weighted moving average, should be applied to these measurements.

Infectious: Once a device is infectious it continues to send its up-to-date bit vectors, using either incremental random walk or incremental gossip, to neighboring nodes. Due to FM's order and duplicate insensitive (ODI) properties, simultaneously occurring updates are handled seamlessly by the protocol.

Update messages contain a time stamp indicating when the update was generated, and each device maintains a lo-

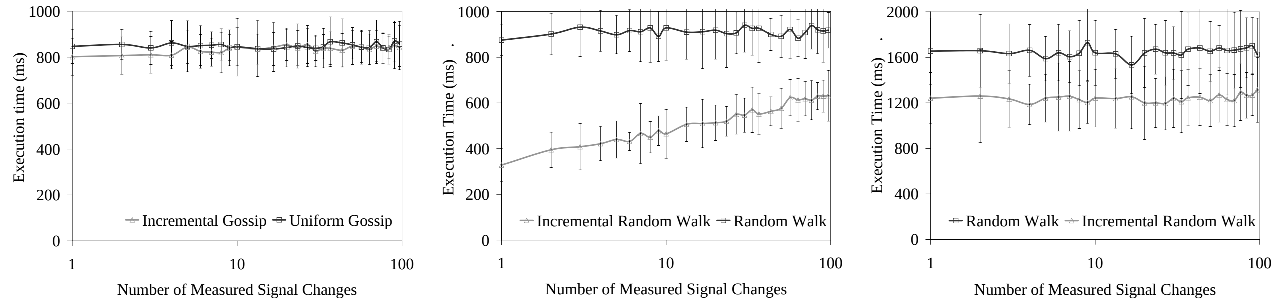

(a) Incremental Gossip and Uniform Gossip on Clique (b) Incremental Random Walk and Random Walk on Clique (c) Incremental Random Walk and Random Walk on Power-Law Random Graph

Figure 4: Execution times of Incremental Protocols

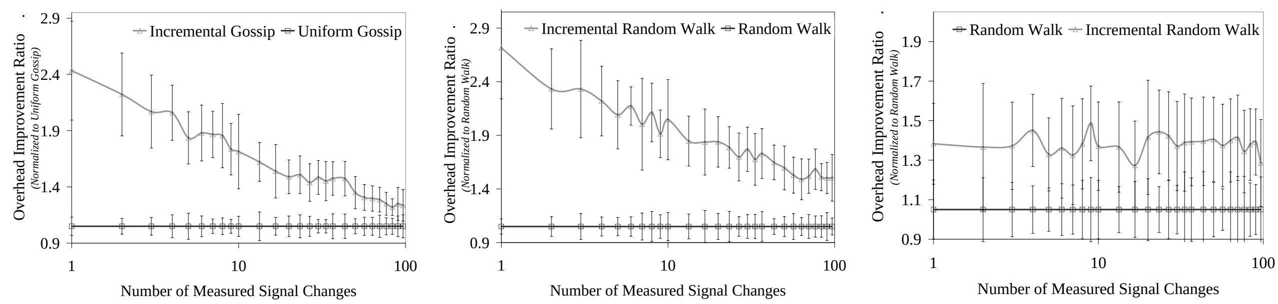

(a) Incremental Gossip and Uniform Gossip on Clique (b) Incremental Random Walk and Random Walk on Clique (c) Incremental Random Walk and Random Walk on Power-Law Random Graph

Figure 5: Network overhead of Incremental Protocols

cal time stamp of when it received the most recent update. Using this information, a device moves into the recovered state once enough time has passed for the most recent update to have converged. As discussed in Section 4.2, this happens after $O(log(N))$ time steps.

Recovered: A recovered device ceases to propagate any update information. At this point, it performs clean-up and prepares for the next infection by entering the susceptible state. Once all devices have entered the recovered state, the system will have converged, and with high probability, all devices will have the up-to-date average signal level. Due to the cumulative nature of FM, even if all devices have not converged, the next update will include all previous updates. Nevertheless, the probability that gossip fails to converge is small, and has been shown to be $O(1/N)$ [10].

For coordinated spectrum sensing, non-incremental routing protocols can be implemented in a similar fashion. Random walk would operate by having devices periodically drop the aggregate and re-run the protocol. Each device would perform a coin toss (biased on the number of walks) to determine whether or not it is a designated node. This is different from the protocol discussed above where only updated nodes initiate random walks. Similar techniques can be used to implement standard gossip.

7. EVALUATION

We now provide a preliminary evaluation of GUESS in simulation. A more detailed evaluation of this approach can be found in [3]. Here we focus on how incremental extensions to gossip protocols can lead to further improvements over standard gossiping techniques, for the problem of coordinated spectrum sensing.

Simulation Setup: We implemented a custom simulator in C++. We study the improvements of our incremental gossip protocols over standard gossiping in two dimensions: execution time and network overhead. We use two topologies to represent device connectivity: a clique, to eliminate the effects of the underlying topology on protocol performance, and a BRITE-generated [13] power-law random graph (PLRG), to illustrate how our results extend to more realistic scenarios. We simulate a large deployment of 1,000 devices to analyze protocol scalability.

In our simulations, we compute the average signal level in a particular band by disseminating FM bit vectors. In each run of the simulation, we induce a change in the measured signal at one or more devices. A run ends when the new average signal level has converged in the network.

For each data point, we ran 100 simulations and 95% confidence intervals (error bars) are shown.

Simulation Parameters: Each transmission involves sending 70 bits of information to a neighboring node. To compute the AVERAGE aggregate, four bit vectors need to be transmitted: the original SUM vector, the SUM delete vector, the original COUNT vector, and the COUNT delete vector. Non-incremental protocols do not transmit the delete vectors. Each transmission also includes a time stamp of when the update was generated.

We assume nodes communicate on a common control channel at 2 Mbps. Therefore, one time-step of protocol execution corresponds to the time required for 1,000 nodes to sequentially send 70 bits at 2 Mbps. Sequential use of the control channel is a worst case for our protocols; in practice, multiple control channels could be used in parallel to reduce execution time. We also assume nodes are loosely time synchronized, the implications of which are discussed further in [3]. Finally, in order to isolate the effect of protocol operation on performance, we do not model the complexities of the wireless channel in our simulations.

Incremental Protocols Reduce Execution Time: Figure 4(a) compares the performance of incremental gossip (IGP) with uniform gossip on a clique topology. We observe that both protocols have almost identical execution times. This is expected as IGP operates in a similar fashion to uniform gossip, taking $O(log(N))$ time-steps to converge.

Figure 4(b) compares the execution times of incremental random walk (IRWP) and standard random walk on a clique. IRWP reduces execution time by a factor of 2.7 for a small number of measured signal changes. Although random walk and IRWP both use k random walks (in our simulations k = number of nodes), IRWP initiates walks only from updated nodes (as explained in Section 5.2), resulting in faster information convergence. These improvements carry over to a PLRG topology as well (as shown in Figure 4(c)), where IRWP is 1.33 times faster than random walk.

Incremental Protocols Reduce Network Overhead: Figure 5(a) shows the ratio of data transmitted using uniform gossip relative to incremental gossip on a clique. For a small number of signal changes, incremental gossip incurs 2.4 times less overhead than uniform gossip. This is because in the early steps of protocol execution, only devices which detect signal changes communicate. As more signal changes are introduced into the system, gossip and incremental gossip incur approximately the same overhead.

Similarly, incremental random walk (IRWP) incurs much less overhead than standard random walk. Figure 5(b) shows a 2.7 fold reduction in overhead for small numbers of signal changes on a clique. Although each protocol uses the same number of random walks, IRWP uses fewer network resources than random walk because it takes less time to converge. This improvement also holds true on more complex PLRG topologies (as shown in Figure 5(c)), where we observe a 33% reduction in network overhead.

From these results it is clear that incremental techniques yield significant improvements over standard approaches to gossip, even on complex topologies. Because spectrum utilization is characterized by incremental changes to usage, incremental protocols are ideally suited to solve this problem in an efficient and cost effective manner.

8. DISCUSSION AND FUTURE WORK

We have only just scratched the surface in addressing the problem of coordinated spectrum sensing using incremental gossiping. Next, we outline some open areas of research.

Spatial Decay: Devices performing coordinated sensing are primarily interested in the spectrum usage of their local neighborhood. Therefore, we recommend the use of *spatially decaying aggregates* [6], which limits the impact of an update on more distant nodes. Spatially decaying aggregates work by successively reducing (by means of a decay function) the value of the update as it propagates further from its origin. One challenge with this approach is that propagation distance cannot be determined ahead of time and more importantly, exhibits spatio-temporal variations. Therefore, finding the optimal decay function is non-trivial, and an interesting subject of future work.

Significance Threshold: RF spectrum bands continually experience small-scale changes which may not necessarily be significant. Deciding if a change is significant can be done using a significance threshold β, below which any observed change is not propagated by the node. Choosing an appropriate operating value for β is application dependent, and explored further in [3].

Weighted Readings: Although we argued that most devices will likely be equipped with low-cost sensing equipment, there may be situations where there are some special *infrastructure* nodes that have better sensing abilities than others. Weighting their measurements more heavily could be used to maintain a higher degree of accuracy. Determining how to assign such weights is an open area of research.

Implementation Specifics: Finally, implementing gossip for coordinated spectrum sensing is also open. If implemented at the MAC layer, it may be feasible to piggy-back gossip messages over existing management frames (e.g. networking advertisement messages). As well, we also require the use of a control channel to disseminate sensing information. There are a variety of alternatives for implementing such a channel, some of which are outlined in [4]. The trade-offs of different approaches to implementing GUESS is a subject of future work.

9. CONCLUSION

Spectrum sensing is a key requirement for dynamic spectrum allocation in cognitive radio networks. The nature of the RF environment necessitates coordination between cognitive radio devices. We propose GUESS, an approximate yet low overhead approach to perform efficient coordination between cognitive radios. The fundamental contributions of GUESS are: (1) an FM aggregation scheme for efficient in-network aggregation, (2) a randomized gossiping approach which provides exponentially fast convergence and robustness to network alterations, and (3) incremental variations of FM and gossip which we show can reduce the communication time by up to a factor of 2.7 and reduce network overhead by up to a factor of 2.4. Our preliminary simulation results showcase the benefits of this approach and we also outline a set of open problems that make this a new and exciting area of research.

10. REFERENCES

[1] Unlicensed Operation in the TV Broadcast Bands and Additional Spectrum for Unlicensed Devices Below 900 MHz in the 3 GHz band, May 2004. Notice of Proposed Rule-Making 04-186, Federal Communications Commission.

[2] In-Stat: Covering the Full Spectrum of Digital Communications Market Research, from Vendor to End-user, December 2005. http://www.in-stat.com/catalog/scatalogue.asp?id=28.

[3] N. Ahmed, D. Hadaller, and S. Keshav. Incremental Maintenance of Global Aggregates. UW. Technical Report CS-2006-19, University of Waterloo, ON, Canada, 2006.

[4] R. W. Brodersen, A. Wolisz, D. Cabric, S. M. Mishra, and D. Willkomm. CORVUS: A Cognitive Radio Approach for Usage of Virtual Unlicensed Spectrum. Technical report, July 2004.

[5] D. Cabric, S. M. Mishra, and R. W. Brodersen. Implementation Issues in Spectrum Sensing for Cognitive Radios. In *Asilomar Conference*, 2004.

[6] E. Cohen and H. Kaplan. Spatially-Decaying Aggregation Over a Network: Model and Algorithms. In *Proceedings of SIGMOD 2004*, pages 707–718, New York, NY, USA, 2004. ACM Press.

[7] P. Flajolet and G. N. Martin. Probabilistic Counting Algorithms for Data Base Applications. *J. Comput. Syst. Sci.*, 31(2):182–209, 1985.

[8] C. Gkantsidis, M. Mihail, and A. Saberi. Random Walks in Peer-to-Peer Networks. In *Proceedings of INFOCOM 2004*, pages 1229–1240, 2004.

[9] E. Griffith. Previewing Intel's Cognitive Radio Chip, June 2005. http://www.internetnews.com/wireless/article.php/3513721.

[10] D. Kempe, A. Dobra, and J. Gehrke. Gossip-Based Computation of Aggregate Information. In *FOCS 2003*, page 482, Washington, DC, USA, 2003. IEEE Computer Society.

[11] X. Liu and S. Shankar. Sensing-based Opportunistic Channel Access. In *ACM Mobile Networks and Applications (MONET) Journal*, March 2005.

[12] Q. Lv, P. Cao, E. Cohen, K. Li, and S. Shenker. Search and Replication in Unstructured Peer-to-Peer Networks. In *Proceedings of ICS*, 2002.

[13] A. Medina, A. Lakhina, I. Matta, and J. Byers. BRITE: an Approach to Universal Topology Generation. In *Proceedings of MASCOTS conference*, Aug. 2001.

[14] S. M. Mishra, A. Sahai, and R. W. Brodersen. Cooperative Sensing among Cognitive Radios. In *ICC 2006*, June 2006.

[15] S. Nath, P. B. Gibbons, S. Seshan, and Z. R. Anderson. Synopsis Diffusion for Robust Aggregation in Sensor Networks. In *Proceedings of SenSys 2004*, pages 250–262, 2004.

[16] A. Sahai, N. Hoven, S. M. Mishra, and R. Tandra. Fundamental Tradeoffs in Robust Spectrum Sensing for Opportunistic Frequency Reuse. Technical Report UC Berkeley, 2006.

[17] J. Zhao, H. Zheng, and G.-H. Yang. Distributed Coordination in Dynamic Spectrum Allocation Networks. In *Proceedings of DySPAN 2005*, Baltimore (MD), Nov. 2005.