Multi-hop Greedy Gossip with Eavesdropping

Deniz Üstebay, Boris Oreshkin, Mark Coates, and Michael Rabbat
Department of Electrical and Computer Engineering
McGill University
Montréal, Québec, Canada
{deniz.ustebay,boris.oreshkin}@mail.mcgill.ca
{mark.coates,michael.rabbat}@mcgill.ca

Abstract – *Greedy Gossip with Eavesdropping (GGE) is a randomized gossip algorithm that computes the average consensus by exploiting the broadcast nature of wireless communications. Each node eavesdrops on its immediate neighbors to track their values so that when it comes time to gossip, a node can myopically exchange information with the neighbor that will give the greatest immediate improvement in local squared error. In previous work, we showed that the improvement achieved using GGE over standard randomized gossip (i.e., exchanging information equally often with all neighbors) is proportional to the maximum node degree. Thus, for network topologies such as random geometric graphs, where node degree grows with the network size, the improvements of GGE scale with network size, but for grid-like topologies, where the node degree remains constant, GGE yields limited improvement. This paper presents an extension to GGE, which we call "multi-hop GGE", that improves the rate of convergence for grid-like topologies. Multi-hop GGE relies on increasing artificially neighborhood size by performing greedy updates with nodes beyond one hop neighborhoods. We show that multi-hop GGE converges to the average consensus and illustrate via simulation that multi-hop GGE improves the performance of GGE for different network topologies.*

Keywords: Distributed signal processing, Gossip algorithms, Average consensus, Wireless sensor networks.

1 Introduction and Background

Distributed consensus is a fundamental problem in distributed control and signal processing (see, e.g., [5, 7–12] and references therein). The prototypical example of a consensus problem is computation of the *average consensus*: for a network of n nodes, initially each node has a scalar data value, y_i, and the goal is to asymptotically compute the average, $\bar{y} = \frac{1}{n}\sum_{i=1}^{n} y_i$ at every node i. An algorithm that computes the average consensus can further be used for computing linear functions of the data and can be generalized for averaging vectorial data.

In [13], we proposed a novel randomized gossip algorithm called *Greedy Gossip with Eavesdropping* (GGE). GGE exploits the broadcast nature of wireless communications to provide energy efficient average consensus computation. In a wireless sensor network setting, we assume all nodes are battery powered. We also assume that all the transmissions are wireless broadcasts, and therefore all nodes within range of a transmitting node can receive broadcasted messages. In this manner, we are able to perform greedy updates in contrast to random updates of previous gossip algorithms. That is, when a node becomes active, rather than gossiping with a randomly chosen neighbor, it greedily chooses to gossip with the neighbor that will result in the greatest decrease in local mean squared error. We have previously shown that GGE leads to faster rates of convergence than randomized gossip without the added overhead of geographic routing or node localization required by other fast gossiping algorithms. We showed that the extent to which GGE improves upon randomized gossip is a factor proportional to the maximum node degree. Hence, for topologies such as random geometric graphs, where node degree increases with network size, the improvements obtained by using GGE scale as network size grows. However, for grid-like topologies, where the maximum degree remains constant as network size increases, the improvement of GGE over randomized gossip is limited to a constant factor.

This paper presents multi-hop greedy gossip with eavesdropping, a variant of GGE designed to offer greater performance improvements on grid-like topologies. Like, GGE, all nodes eavesdrop on their immediate neighbors, in order to track their values. However, in k-hop GGE, when a node activates, it gossips with a node that is within k hops. This neighbor is determined by a sequence of greedy decisions made at each intermediate hop. Thus, multi-hop GGE also only requires nodes to store information about their immediate neighbors, and does not require any additional overhead to construct or maintain routes. We show that greedily gossiping over multiple hops in this fashion leads to faster convergence rates on topologies such as grids, and also leads to improved energy-accuracy tradeoffs over single-hop GGE. Although multi-hop GGE requires a mechanism to provide reliable information exchange over multiple wireless hops, the maximum number of hops is controlled. On moderate sized networks, we find that we can obtain improved rates of convergence without requiring nodes to exchange information across the entire network in a single iteration.

1.1 Background and Related Work

Randomized gossip algorithms have gained much attention in the wireless sensor networks community due to their simplicity and robustness. At iteration k, a node s_k is activated uniformly at random; it selects a neighbor, t_k, randomly; and this pair of nodes "gossips": s and t exchange values and perform the update $x_{s_k}(k) = x_{t_k}(k) = (x_{s_k}(k-1)+x_{t_k}(k-1))/2$, and all other nodes remain unchanged. It can be shown that under mild conditions on the way the random neighbor, t_k, is drawn, the values $x_i(k)$ converge to $\bar{y}$ at every node [15]. Relying on local information exchanges makes the algorithm simple and robust to changing topologies and wireless network conditions. However it causes slow information diffusion across the network for topologies such as grids and random geometric graphs [6]. Boyd et al. [4] show that for random geometric graphs, $O(n^2)$ transmissions are required for randomized gossip to approximate the average consensus well[1].

Slow convergence of standard randomized gossip motivated the development of new gossip algorithms. One such algorithm is geographic gossip [5]. Under the assumption that nodes know their own locations and the locations of their neighbors, geographic gossip operates with information exchange over multiple hops. It has been shown that long range information exchange enables information to diffuse faster compared to randomized gossip, and for random geometric graphs, only $O(n\sqrt{n/\log n})$ transmissions are required. In fact, if all nodes along the path average their values, rather than just the two endpoints, then the total number of transmissions required to gossip can be decreased to $O(n)$. However, geographic gossip and path averaging involve localization overhead. Furthermore, the network must supply geographic routing and reliable two-way transmission over many hops.

In a series of recent papers, Aysal et al. propose *broadcast gossip*, an algorithm that makes use of the broadcast nature of wireless networks [1, 2]. At each iteration, a node is chosen uniformly at random to broadcast its value. The nodes in the broadcast range of this node calculate a weighted average of their own value and the broadcasted value, and they update their value with this weighted average. In broadcast gossip, the value of the broadcasting node is independently incorporated at each neighbor. Broadcast gossip does not preserve the network average at each iteration. It achieves a low variance (i.e., rapid convergence), but introduces bias, the value to which broadcast gossip converges can be significantly different from the true average (see [13] for further discussion).

In previous work, we have proposed Greedy Gossip with Eavesdropping (GGE) which uses the broadcast medium to accelerate gossip [13]. In GGE, the selection procedure is done greedily such that the neighbor with most different value than the selecting node is chosen for gossip. We assume that all transmissions are wireless broadcasts and nodes eavesdrop on their neighbors' communication to keep track of their values. Accelerating convergence in this myopic way does not introduce bias to the computation and does not rely on geographic location information. Since GGE iterations depend on the values at each node, it is a data-driven algorithm. Therefore, the standard way of proving convergence of gossip algorithms (i.e., expressing updates in terms of a linear recursion and then imposing properties on this recursion) does not apply to GGE. Instead, we demonstrate that GGE updates correspond to iterations of a distributed randomized incremental subgradient optimization algorithm and we prove convergence. Similarly, analyzing rates of convergence of GGE is also a non-trivial task and the standard method for gossip algorithms (i.e., examining the mixing time of a related Markov chain) does not apply to data-driven nature of GGE. In [14], we prove that GGE always converges faster than standard randomized gossip and we develop a worst-case bound on its rate of convergence. We also show that the performance of GGE is closely related to the number of neighbors. Specifically, the maximum node degree in the network characterizes the improvement of GGE over randomized gossip.

In general, we are interested in topologies that are used for modeling wireless sensor and mesh networks, such as random geometric graphs (RGG) and grids. The maximum node degree in a RGG is in the order of $\log n$ and hence the improvement of GGE over randomized gossip is by a factor of $\log n$ (i.e. $O(n^2/\log n)$ transmissions are required for convergence to the average consensus). However, for grids the maximum number of neighbors is bounded by 4, therefore the improvement of GGE is only by a linear factor. Motivated by these results, we propose multi-hop GGE.

1.2 Paper Organization

The rest of the paper is organized as follows. Section 2 gives a detailed description of multi-hop GGE. In Section 2.3, we prove that multi-hop GGE converges to average consensus solution. Section 3 presents results from numerical simulations, and we conclude in Section 4.

2 Multi-hop GGE

We consider a network of n nodes and represent network connectivity as a graph, $G = (V, E)$, with vertices $V = \{1, \ldots, n\}$, and edge set $E \subset V \times V$ such that $(i, j) \in E$ if and only if nodes i and j directly communicate. We assume that communication relationships are symmetric and that the graph is connected. Let $\mathcal{N}_i = \{j : (i, j) \in E\}$ denote the

[1]Throughout this paper, when we refer to randomized gossip, we specifically mean the natural random walk version of the algorithm, where the node t_k is chosen uniformly from the set of neighbors at each iteration. For random geometric graph topologies, which are of most interest to us, Boyd et al. [4] prove that the performance of the natural algorithm scales identically to that of the optimal choice of transition probabilities, so there is no loss in generality.

set of neighbors of node i (not including i itself). Each node in the network has an initial value y_i, and the goal of the gossip algorithm is to use only local broadcast exchanges to arrive at a state where every node knows the average $\bar{y} = \frac{1}{n}\sum_{i=1}^{n} y_i$. To initialize the algorithm, each node sets its gossip value to $x_i(0) = y_i$.

2.1 One-hop GGE

We first describe the one-hop version of GGE. At the kth iteration of one-hop GGE, a node s_k is chosen uniformly at random from $\{1, \ldots, n\}$. (This can be accomplished using the asynchronous time model described in [3], where each node "ticks" according to a Poisson clock with rate 1.) Then, s_k identifies a neighboring node t_k satisfying

$$t_k \in \arg\max_{t \in \mathcal{N}_{s_k}} \left\{ \frac{1}{2}(x_{s_k}(k-1) - x_t(k-1))^2 \right\}, \quad (1)$$

where $\mathcal{N}_{s_k}$ is the set of neighbors of s_k. Hence, s_k identifies a neighbor that currently has the most different value from its own. This greedy selection is possible because each node i maintains not only its own local variable, $x_i(k-1)$, but also a copy of the current values at its neighbors, $x_j(k-1)$, for $j \in \mathcal{N}_i$. When s_k has multiple neighbors whose values are all equally (and maximally) different from s_k's, it chooses one of these neighbors at random. Then s_k and t_k perform the update

$$x_{s_k}(k) = x_{t_k}(k) = \frac{1}{2}\big(x_{s_k}(k-1) + x_{t_k}(k-1)\big), \quad (2)$$

while all other nodes $i \notin \{s_k, t_k\}$ hold their values at $x_i(k) = x_i(k-1)$. Finally, the two nodes, s_k and t_k, broadcast these new values so that their neighbors have up-to-date information.

2.2 Multi-hop GGE

In 2-hop GGE, at the kth iteration, a node s_k is chosen uniformly at random from $\{1, \ldots, n\}$, as in one-hop GGE. Then, s_k identifies a neighboring node t_k satisfying (1). However, instead of completing the update, t_k checks if any of its neighbors $u_k \in \mathcal{N}_{t_k}$ has a value even more different from s_k than its own; i.e., $\|x_{s_k} - x_{u_k}\| > \|x_{s_k} - x_{t_k}\|$. If such a neighbor u_k exists, then t_k facilitates a gossip update between s_k and u_k; otherwise, s_k and t_k gossip.

For p-hop GGE, when t_k finds a neighbor u_k, then u_k continues the process and checks its neighborhood, and so on, until the search has extended to at most p hops away from s_k. Note that this is not equivalent to searching the entire p hop neighborhood of s_k, since at each hop, only one option is explored in a greedy, myopic fashion. Nevertheless, we observe that gossiping in this fashion leads to improvements in number of messages that must be transmitted to achieve high-accuracy estimates of the average at every node. Of course, searching for the best node to gossip with over the entire p-hop neighborhood might lead to even faster rates of convergence, but we can motivate the proposed search over a restricted neighborhood using the following reasoning. The types of signals for which one-hop randomized gossip algorithms are slow to average are those signals that are correlated in some fashion with network structure. For example, if nodes are located in $\mathbb{R}^2$, suppose the initial value of a node located at coordinate (a, b) is $a + b$. Thus, nearby nodes have similar values, but distant nodes have very different values. In such a configuration, our greedy process of determining which pair of nodes gossips will track along the gradient, and information will be diffused through the network more quickly than if gossiping was restricted to be between immediate neighbors. On the other hand, if node values are initially i.i.d., then our greedy process of determining which neighbor to gossip with may not precisely identify the node in the p-hop neighborhood whose value is most different from s_k, but at the same time, since all initial values are i.i.d, any node should be able to average with its immediate neighbors and still obtain a reasonable estimate of the network average. Moreover, our greedy p-hop neighbor selection procedure requires at most p transmissions (fewer if we reach a node in fewer than p hops that doesn't have a neighbor whose value is more different than s_k than its own), which is much less costly than searching the entire p hop neighborhood, in general.

Calculating the greedy update in multi-hop gossip requires that each node know the values of its immediate neighbors. Similar to other randomized gossip algorithms, we assume that at the outset of gossip computation that each node i has already discovered its neighbor set, $\mathcal{N}_i$, but it does not know its neighbors' values. Instead, these values are learned during an initialization phase of multi-hop GGE. During this initialization phase, when s_k does not know the values of all its neighbors, it chooses t_k randomly from the subset of its neighbors whose values are currently unknow, rather than performing a GGE update. Since s_k and t_k broadcast their new values after averaging, the nodes in their neighborhoods overhear and acquire information accordingly. Once s_k has heard from all of its neighbors, the initialization process is complete for that particular node and it chooses t_k greedily for all subsequent iterations.

2.3 Convergence of Multi-hop GGE

Let $\{x(k)\}$ denote the sequence of iterates produced by multi-hop GGE, let $\{s_k\}$ denote the (random) sequence of nodes activated at each iteration, and let m_k denote the node within s_k's p-hop neighborhood with which it gossips at iteration k. The evolution of $x(k)$ can be expressed in recursive form as $x(k) = W(k)x(k-1)$, where $W(k)$ is a matrix with all diagonal entries equal to 1 except $W_{s_k,m_k}(k) = W_{m_k,s_k}(k) = W_{s_k,s_k}(k) = W_{m_k,m_k}(k) = 1/2$. It is clear that the only fixed point of this equation is a vector x whose entries are all identical (a consensus vector). Moreover, since at each iteration, the sum of the entries in $x(k)$ remains constant (i.e., $\mathbf{1}^T x(k) = \mathbf{1}^T x(0)$), the only fixed-point is the average consensus vector.

Now, let $M(k) = \|x(k) - \bar{y}\|^2$ denote the squared error after k iterations. Observe that we can write $\|x(k) - \bar{y}\|^2 = \|x(k-1) - \bar{y}\|^2 - \frac{1}{2}(x_{s_k}(k-1) - x_{m_k}(k-1))^2$. Let

$\Delta_k = \frac{1}{2}(x_{s_k}(k-1) - x_{m_k}(k-1))^2$. Then we have that $M(k) = M(k-1) - \Delta_k$ with probability 1. (Note that Δ_k is random, by virtue of the randomness in s_k at each iteration.) Conditioning on $x(k-1)$ and taking the expectation over s_k, it is clear that $\mathbb{E}[\Delta_k | x(k-1)] > 0$ unless $x(k-1) = \bar{y}$ is the average consensus solution. Thus, unless the algorithm has converged, we always make progress in expectation towards the consensus solution. Moreover, via repeated application of the recursion for $M(k)$, we have $M(k) = M(0) - \sum_{i=1}^{k} \Delta_i$. Since $M(k) \geq 0$, we have $\sum_{i=1}^{k} \Delta_i \leq M(0)$, which implies that $\Delta_k \to 0$ as $k \to \infty$. Since the only fixed point of the stochastic recursion is the average consensus solution, this implies that multi-hop GGE converges to the average consensus asymptotically as $k \to \infty$.

3 Numerical Simulations

To observe the effect of performing greedy updates over multiple hops, we conduct an experimental comparison between 1-hop, 2-hop, and 3-hop GGE. As a point of comparison, we also include curves for randomized gossip [4] and geographic gossip [5]. We examine the reduction they achieve in relative error, $\frac{||x(k)-\bar{x}||}{||x(0)-\bar{x}||}$, as a function of the number of transmissions. The number of transmissions is the cost of interest in wireless sensor network applications since each transmission consumes valuable battery resources, and experimental studies have shown that transmitting generally consumes significantly more energy than performing local computation or taking measurements from simple sensors. The initial values for our experiments are determined according to the linearly varying field example discussed above; each node is assigned coordinates in $\mathbb{R}^2$, and a node at location (a_i, b_i) has initial value $x_i(0) = a_i + b_i$. We consider two network topologies: the two-dimensional grid, and the family of random geometric graphs [6]. In a random geometric graph, nodes are assigned i.i.d. coordinates (uniformly) in the unit square, and two nodes are connected if they are separated by a distance of no more than $\sqrt{\frac{2 \log n}{n}}$. Both the grid and random geometric graph are commonly adopted as models for connectivity in wireless networks.

Figures 1 and 2 show simulation results for the 196-node grid and 200-node random geometric graph topologies. Observe that on random geometric graph topologies (Fig. 1), one-hop GGE performs comparably to geographic gossip, and multi-hop gossip improves upon this performance. However, on a grid topology of comparable size, (see Fig. 2), the performance of one-step GGE is much worse than that of geographic gossip. The reason for this disparity is that in the random geometric graph setup, most nodes have many more than four neighbors. Thus, in a single one-hop GGE iteration is is possible to find a neighbor that is quite different than the activated node, and thus spread information quickly. On the other hand, in the grid scenario, each node has at most four neighbors. Thus, all information exchange is highly local. By expanding the search radius to carry out 2-hop or 3-hop GGE iterations, we see a marked improvement.

We conjecture that the gain obtained by going from 1-hop GGE to p-hop GGE is roughly a factor of p. Consider the following heuristic explanation. After $k-1$ iterations, the square difference between values at nodes s_k and m_k, relative to the current squared error, $M(k-1)$, is on average at least p^2. (Again, think of a regular grid with the linearly varying field setup, in which case $(x_{s_k} - x_{m_k})^2 = p^2$ for nodes s_k and m_k that are p hops away from each other along a given axis.) The number of transmissions required to carry out this greedy p-hop update is also p, so the resulting gain is a factor of approximately $p^2/p = p$. This improvement is also evident when examining how performance scales with network size.

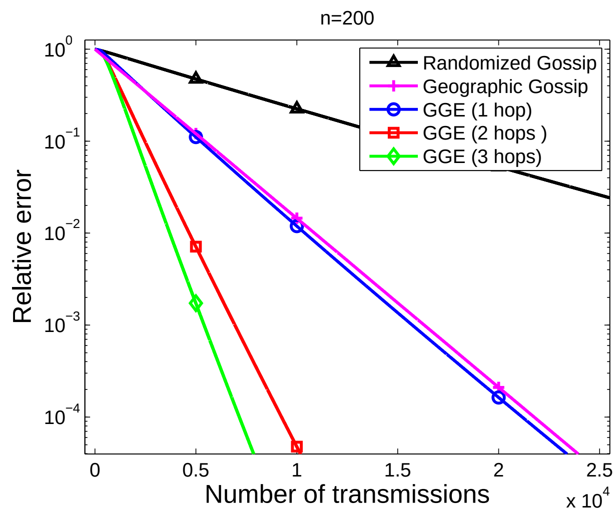

Figure 1: A comparison of the performance of randomized gossip, GGE (1-hop), multi hop GGE (2 and 3-hops) and geographic gossip for linearly-varying field initialization of $x(0)$ in RGG topology of 200 nodes. Results are averaged over 100 runs of the algorithm.

Next, we examine how the communication complexity scales with respect to the number of nodes in the network. The rate of convergence for gossip algorithms is typically quantified in terms of the ϵ-averaging time,

$$T_{ave}(\epsilon) = \sup_{x(0)\neq 0} \inf \left\{ k \; : \; \Pr\left(\frac{\|x(k) - \bar{x}\|}{\|x(0) - \bar{x}\|} \geq \epsilon \right) \leq \epsilon \right\}.$$

Figure 3 displays how the averaging time scales as a function of the number of nodes n, for the grid topology, for 1-hop, 2-hop, and 3-hop GGE. The averaging time has been estimated with simulating the gossip algorithms over the same grid network for 100 times. Note that the averaging time is shown in terms of the number of iterations per node. The slope of the curve for 2-hop gossip is roughly half of that of one-hop gossip, and the curve for 3-hop gossip has slope roughly 1/3 of that of one-hop gossip, providing further evidence to our conjecture that the improvement obtained by using p-hop gossip is roughly a factor of p.

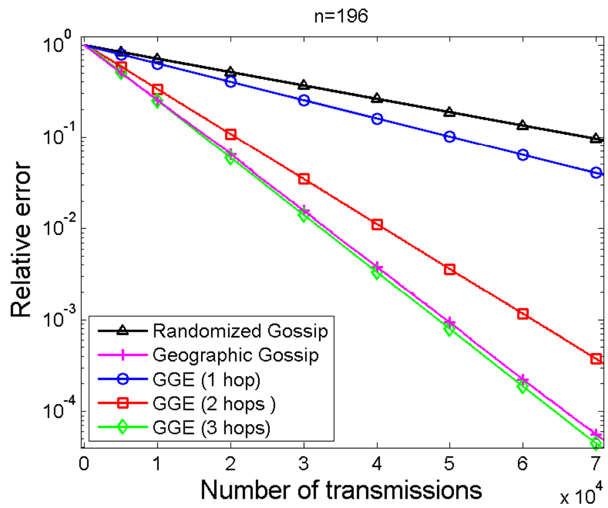

Figure 2: A comparison of the performance of randomized gossip, GGE (1-hop), multi hop GGE (2 and 3-hops) and geographic gossip for linearly-varying field initialization of $x(0)$ in grid topology of 196 nodes. Results are averaged over 100 runs of the algorithm.

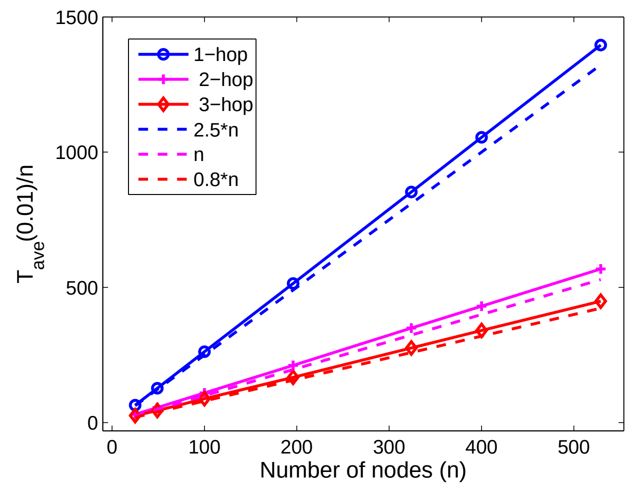

Figure 3: The averaging time $T_{ave}(\epsilon)$ for $\epsilon = 0.01$ as a function of number of nodes in the network. Results are averaged over 100 runs of the algorithm. The lines $2.5n$, n and $0.8n$ are shown for reference.

4 Discussion

In this paper, we proposed multi-hop greedy gossip with eavesdropping, an extension of GGE for faster convergence in grid-like topologies. Multi-hop GGE suggests a mechanism for practical fast gossiping in moderate sized networks without requiring the overhead entailed in localizing nodes so as to enable greedy geographic routing. Accelerated rates of convergence are achieved by having nodes eavesdrop on their neighbors (exploiting the broadcast nature of wireless communications), and then making greedy decisions about who to gossip, rather than selecting a node randomly. In the p-hop extension to GGE, we allow nodes to exchange information over up to p hops in each iteration. This serves to both facilitate faster convergence via longer range information spreading, while at the same time forcing a hard constraint on this distance. From a practical standpoint, constraining the number of hops over which one gossips has the advantage that reliable information exchange is more easily facilitated over fewer hops (less queueing, fewer opportunities for collisions or dropped packets, and thus fewer retransmissions). We provide theoretical arguments supporting the claim that multi-hop GGE converges to the average consensus solution, and we investigate the convergence rate of the algorithm via simulation.

We conjecture that performing p-hop GGE updates leads to improvements in rate of convergence by a factor of p. In our previous analysis, we found that one-step improves on the performance of randomized gossip by a factor of at most $d_{\max}$, the maximum degree of the network. The basic idea is that in a single one-hop GGE iteration, we get to search over a local neighborhood of at most $d_{\max}$ nodes. Since we greedily choose which of these nodes to gossip with, rather than drawing one at random, we obtain a speedup by at most that factor. In p-hop GGE, the size of the set we search over in each iteration is at most $pd_{\max}$, and so similar reasoning leads one to the conclusion that p-hop GGE leads to an improvement of at most this factor. Obtaining a more thorough theoretical characterization of the rates of convergence of multi-hop GGE is a topic of our future work.

References

[1] T. Aysal, M. Yildiz, A. Sarwate, and A. Scaglione. Broadcast gossip algorithms: Design and analysis for consensus. In *Proc. IEEE Conf. on Decision and Control*, Cancun, Mexico, Dec. 2008.

[2] T. Aysal, M. Yildiz, and A. Scaglione. Broadcast gossip algorithms. In *Proc. IEEE Information Theory Workshop*, Porto, Portugal, May 2008.

[3] D. Bertsekas and J. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific, Belmont, MA, 1997.

[4] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Trans. Info. Theory*, 52(6):2508–2530, June 2006.

[5] A. Dimakis, A. Sarwate, and M. Wainwright. Geographic gossip: Efficient aggregation for sensor networks. In *Proc. Int. Conf. Inf. Proc. in Sensor Networks (IPSN)*, Nashville, TN, Apr. 2006.

[6] P. Gupta and P. Kumar. The capacity of wireless networks. *IEEE Trans. Info. Theory*, 46(2):388–404, March 2000.

[7] A. Kashyap, T. Basar, and R.Srikant. Quantized consensus. *Automatica*, 43:1192–1203, July 2007.

[8] M. Rabbat, J. Haupt, A. Singh, and R. Nowak. Decentralized compression and predistribution via randomized gossiping. In *Proc. Information Processing in Sensor Networks*, Nashville, TN, Apr. 2006.

[9] M. Rabbat, R. Nowak, and J. Bucklew. Robust decentralized source localization via averaging. In *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, Phil., PA, Mar. 2005.

[10] S. Sundaram and C. Hadjicostis. Distributed function calculation and consensus using linear iterative strategies. *IEEE J. Selected Areas in Communications*, 26(4):650–660, May 2008.

[11] S. Sundhar Ram, V. Veeravalli, and A. Nedić. Distributed and recursive parameter estimation in parametrized linear state-space models. Submitted, Apr. 2008.

[12] J. Tsitsiklis. *Problems in decentralized decision making and computation*. PhD thesis, Massachusetts Institute of Technology, 1984.

[13] D. Üstebay, M. Coates, and M. Rabbat. Greedy gossip with eavesdropping. In *Proc. IEEE Int. Symp. on Wireless Pervasive Computing*, Santorini, Greece, May 2008.

[14] D. Üstebay, B. Oreshkin, M. Coates, and M. Rabbat. Rates of convergence for greedy gossip with eavesdropping. In *Proc. Allerton Conf. on Comm., Control, and Computing*, IL, USA, September 2008.

[15] L. Xiao and S. Boyd. Fast linear iterations for distributed averaging. *Systems and Control Letters*, 53(1):65–78, Sep. 2004.