Toward Resource-Optimal Averaging Consensus over the Wireless Medium

Matthew Nokleby, Waheed U. Bajwa, Robert Calderbank, and Behnaam Aazhang

***Abstract*—We carry out a comprehensive study of the resource costs of distributed averaging consensus in wireless sensor networks. In particular, we consider two metrics appropriate to the wireless medium: total transmit energy and time-bandwidth product. Most previous approaches, such as gossip algorithms, suppose a graphical network, which abstracts away crucial features of the wireless medium, and measure resource consumption only in terms of the total number of transmissions required to achieve consensus. Under a path-loss dominated protocol interference model, we study the performance of several popular gossip algorithms, showing that they are nearly order-optimal with respect to transmit energy but strictly sub-optimal with respect to time-bandwidth product. We also propose a new scheme, termed *hierarchical averaging*, which is tailored to the wireless medium, and show that in general this approach is nearly order-optimal with respect to time-bandwidth product but strictly sub-optimal with respect to transmit energy. For the special case of free-space propagation, however, the proposed hierarchical scheme is approximately order-optimal with respect to both metrics.**

I. Introduction

Consider a wireless sensor network of N nodes, each of which has a measurement $z_n \in \mathbb{R}$. In *averaging consensus*, each node wishes to compute the average of the measurements:

$$z_{\text{ave}} = \frac{1}{N}\sum_{n=1}^{N} z_n. \tag{1}$$

This task, while conceptually simple, serves as a prototype for more complicated distributed signal processing tasks in sensor networks. The study of resource-efficient averaging algorithms is therefore an important research area.

Due to their simplicity, flexibility, and robustness, *gossip algorithms* have emerged as a popular approach to consensus. In gossip, the network is modeled by a graph. Nodes iteratively pair with neighbors, exchange estimates, and average those estimates together, eventually converging on the true average. A large body of excellent work on gossip has been developed, from the early *randomized gossip* of [1] to faster schemes such as *path averaging* [2] and *multi-scale gossip* [3]. Gossip is simple, requiring minimal processing and network knowledge, and it is robust to link and node failures.

However, the purpose of consensus strategies is typically to facilitate processing over *wireless* networks, and wireless affords possibilities that existing strategies do not fully exploit. For example, gossip algorithms are typically constructed over a fixed connectivity topology. In wireless, however, connectivity is adjustable dynamically by means of power control. In fact, one could trivially connect the entire network and achieve consensus in a single round, albeit at high energy costs. This suggests both that wireless permits flexibility that may improve performance, and that we must consider additional performance metrics—such as transmit power—that encompass more than just the number of transmissions.

A few works have addressed individually the wireless aspects we consider in this work. The broadcast nature of wireless is considered in [4], [5]; however, in these works broadcast does not significantly improve performance over randomized gossip. Multi-access interference is addressed—and in fact exploited—in [6], where lattice codes are used to compute sums of estimates "over the air." The notion that network connectivity can be optimized via power allocation is explored in [7], and in [8] the optimum graphical structure for consensus is derived.

By contrast, in this work we endeavor to address comprehensively the implications of the wireless medium on consensus. Accordingly, we study consensus using metrics more appropriate to the wireless medium: the total transmit energy and the time-bandwidth product required to achieve consensus. We first study the performance of gossip algorithms, observing that they are strictly suboptimal with respect to the required time-bandwidth product. We then present a new averaging algorithm, termed *hierarchical gossip*, which exploits the flexibility of wireless. Hierarchical gossip is nearly order optimal in terms of the required time-bandwidth product, and under free-space propagation it is also nearly order optimal in total transmit energy.

In Section II we detail the wireless model under consideration. In Section III we examine existing gossip algorithms with respect to the proposed metrics. In Section IV we present hierarchical averaging and prove scaling laws on its resource requirements. In Section V we present simulation results, and finally we conclude in Section VI. Proofs are omitted in Sections II and III due to space constraints. For details, as well as results not discussed herein, we direct the reader to the journal version of this work [9].

Matthew Nokleby and Behnaam Aazhang are with Rice University, Houston TX (email: {nokleby, aaz}@rice.edu). Waheed U. Bajwa is with Rutgers University, Piscataway, NJ (email: waheed.bajwa@rutgers.edu). Robert Calderbank is with Duke University, Durham, NC (email: robert.calderbank@duke.edu).

II. System Model

Each node n has a physical location $\mathbf{x}(n)$, which we take to be randomly and uniformly distributed across the unit square. We assume path-loss dominated propagation environment and calculate the channel gains between any two nodes m, n as

$$h_{mn} = \|\mathbf{x}(m) - \mathbf{x}(n)\|^{-\alpha/2}, \quad (2)$$

where $\alpha \geq 2$ is the path-loss exponent. We assume slotted discrete time. At each slot t, each node n transmits at power $P_n(t)$. We assume a simple transmission model: a signal arrives at a node if the receive signal-to-noise ratio is above an arbitrary threshold γ, but no interference is generated otherwise. Let the *neighborhood* of n be the nodes whose transmissions arrive at n at time t:

$$\mathcal{N}_n(t) = \{m : P_m(t)h_{mn}^2 \geq \gamma\}. \quad (3)$$

In order to manage multiple-access, we suppose that $\max_n |\mathcal{N}_n(t)|$ time-frequency resources are required for each time slot t.

Each node n maintains an estimate $z_n(t)$ of the average, where $z_n(0) = z_n$. After the transmission at time slot t, each node updates its estimates by computing a linear combination of its old estimate and the estimates of its neighbors at time t:

$$z_n(t+1) = \sum_{m \in \mathcal{N}_n(t)} a_{mn}(t) z_m(t), \quad (4)$$

for some collection of coefficients $a_{mn}(t)$. Note that the coefficient weights may vary with t along with the neighborhood $\mathcal{N}_n(t)$.

The ϵ-averaging time, denoted T_ϵ, is the number of time slots required to achieve consensus to within a specified tolerance:

$$T_\epsilon = \sup_{\mathbf{z}(0) \in \mathbb{R}^n} \inf \left\{ t : \Pr\left(\frac{\|\mathbf{z}(t) - z_{\mathrm{ave}}\mathbf{1}\|}{\|\mathbf{z}(0)\|} \geq \epsilon \right) \leq \epsilon \right\}, \quad (5)$$

where $\mathbf{z}(t)$ is the vector of estimates $z_n(t)$. The scaling law of T_ϵ is the primary focus of study for most gossip algorithms. However, it provides only a partial measure of resource consumption in wireless networks, so we consider other resources metrics.

An averaging algorithm requiring T_ϵ time slots and having power allocation $P_n(t)$ requires a *total transmit energy* of

$$E_\epsilon(N) = \sum_{t=1}^{T_\epsilon} \sum_{n=1}^{N} P_n(t), \quad (6)$$

and a required *time-bandwidth product* of

$$B_\epsilon(N) = \sum_{t=1}^{T_\epsilon} \max_n |\mathcal{N}_n(t)|, \quad (7)$$

In the sequel, we will require a lemma, proven in [10], about the approximate number of nodes in a particular region of the network.

Lemma 1 (Ozgur-Leveque-Tse, [10]): Let $A \subset [0,1] \times [0,1]$ be a region inside the unit square having area $|A|$, and let $\mathcal{C} = \{n : \mathbf{r}_n \in A\}$ be the nodes lying in A. Then, for any $\delta > 0$,

$$(1-\delta)|A|N \leq |\mathcal{C}| \leq (1+\delta)|A|N, \quad (8)$$

with probability greater than $1 - 1/|A|e^{-\Gamma(\delta)|A|N}$, where $\Gamma(\delta) > 0$ and is independent of N and $|A|$.

We can derive an asymptotic lower bound on the resources required by any averaging algorithm.

Theorem 1: For any distributed averaging algorithm over a random network of size N, the total transmit energy $E(N)$ and time-bandwidth product $B(N)$ scale[1], with high probability, as

$$E_\epsilon(N) = \Omega(N^{1-\frac{\alpha}{2}}) \quad (9)$$
$$B_\epsilon(N) = \Omega(1). \quad (10)$$

III. Gossip Algorithms

Nodes in gossip algorithms achieve consensus by means of pairwise interactions; nodes pair up, exchange current estimates, and form new estimates until the network converges. Under *randomized gossip* [1], nodes choose partners randomly and compute pairwise averages of estimates at each round. Randomized gossip is somewhat inefficient, requiring $\Theta(N^2)$ total transmissions in $\Theta(N)$ sequential rounds. Using these facts we can derive the performance with respect to our performance metrics.

Theorem 2: For randomized gossip over a random network of size N, the total transmit energy and time-bandwidth product scale, with high probability, as

$$E_\epsilon(N) = \Theta(N^{2-\frac{\alpha}{2}} \log^{1+\frac{\alpha}{2}}(N)) \quad (11)$$
$$B_\epsilon(N) = \Theta(N \log(N)). \quad (12)$$

Path averaging [11] is a more sophisticated gossip algorithm. Instead of exchanging estimates with neighbors, nodes select a geographically distant partner and exchange estimates via multi-hop routing. Furthermore, routing nodes contribute their estimates "along the way", allowing the entire route to average their estimates together. Path averaging requires only $\Theta(N \log(N))$ transmissions and is nearly order optimal in total transmit energy. However, this approach is strictly sub-optimal with respect to the required time-bandwidth product.

Theorem 3: For path averaging over a random network of size N, the total transmit energy and time-bandwidth product scale, with high probability, as

$$E_\epsilon(N) = \Theta(N^{1-\frac{\alpha}{2}} \log^{1+\frac{\alpha}{2}}(N)) \quad (13)$$
$$B_\epsilon(N) = \Omega(N^{\frac{1}{2}}). \quad (14)$$

IV. Hierarchical Averaging

Our proposed *hierarchical averaging* is similar to the multi-scale gossip of [3] in that we perform a recursive, geographical partition of the network into a hierarchy of clusters. At the highest level of the hierarchy the entire network is a single cluster, at the second-to-highest level the network is partitioned geographically into four equal-sized square clusters, and so on for approximately $\log(N)$ layers of hierarchy. The partitioning is depicted in Figure 1.

[1]We use the notation $f(x) = \Omega(x)$ to imply $f(x) \geq cg(x)$, $f(x) = O(g(x))$ to imply $f(x) \leq cg(x)$, and $f(x) = \Theta(g(x))$ to imply $cg(x) \leq f(x) \leq dg(x)$, all for arbitrary constants c and d and for large x.

A. Hierarchical Partitioning

We partition the network into T sub-network layers, one for each round of consensus, as depicted in Figure 1. At the top layer, which corresponds to the final round $t = T$ of consensus, there is a single cell. At the next-highest level $t = T - 1$, we divide the network into four equal-area square cells. Continuing, we recursively divide each cell into four smaller cells until the lowest layer $t = 1$, which corresponds to the first round of consensus. At each level t there are 4^{T-t} cells, formally defined as

$$\mathcal{C}_{jk}(t) = \{n : \mathbf{r} \in [(j-1)2^{t-T}, j2^{t-T}) \times [(k-1)2^{t-T}, k2^{t-T})\}, \quad (15)$$

where $1 \leq j, k \leq 2^{T-t}$ index the geographical location of the cell.

Let $\mathcal{C}(n, t)$ denote the unique cell at layer t containing node n. Using the Pythagorean theorem, we can easily bound the maximum distance between any two nodes:

$$M(t) = \sqrt{2} \cdot 4^{\frac{t-T}{2}} = \Theta(4^{\frac{t-T}{2}}), \quad (16)$$

where the maximum is achieved when two nodes lie on opposite corners of the cell.

We take $T = \lceil \log_4(N^{1-\kappa}) \rceil$, where $\kappa > 0$ is a constant smaller than unity.

B. Algorithm Description

Here we lay out the details of hierarchical averaging. We suppose that each node n knows the following information about the network: the total number of nodes N, its own location $\mathbf{r}_n$, and the number of layers T.

First, at time slot $t = 1$ each node broadcasts its initial estimate $z_n(0)$ to each member of its cluster $\mathcal{C}(n, t)$. In order to ensure that $n \in \mathcal{N}_m(t)$ for every $m \in \mathcal{C}(n, 1)$, each node transmits at power

$$P_n(1) = \gamma \max_{m \in \mathcal{C}(n,1)} h_{nm}^{\alpha} \leq \gamma M(1)^{\alpha} = O(N^{(\kappa-1)\alpha/2}). \quad (17)$$

Each node n takes a weighted average of the estimates in its cluster:

$$z_n(1) = \frac{1}{4^{1-T}N} \sum_{m \in \mathcal{C}(n,1)} z_m(0). \quad (18)$$

We use the approximate normalization factor $1/4^{1-T}N$ instead of the exact factor $1/|\mathcal{C}(n, 1)|$ so that nodes at higher levels of the hierarchy need not know the cardinality of the cells. As we shall see, this approximation introduces no error into the final estimate.

After time slot $t = 1$, each node in each cluster $\mathcal{C}_{jk}(1)$ has the same estimate, which we denote by $z_{\mathcal{C}_{jk}(1)}(1)$. At each subsequent time slot $2 \leq t \leq T$, a single representative node in $\mathcal{C}_{jk}(t-1)$ is chosen arbitrarily to transmit $z_{\mathcal{C}_{jk}(t-1)}(t-1)$ to its parent cluster at layer t. In order for its transmission to be received by every node in the parent cluster, the representative node n must transmit at power satisfying

$$P_n(t) = \gamma \max_{m \in \mathcal{C}(n,t)} h_{nm}^{\alpha} \leq \gamma M(t)^{\alpha} \quad (19)$$

$$= O(4^{\frac{(t-T)\alpha}{2}}). \quad (20)$$

After receiving estimates from the other sub-clusters, each node updates its estimate by taking the sum:

$$z_n(t) = \frac{1}{4} \sum_{\mathcal{C}(n,t-1) \subset \mathcal{C}(n,t)} z_{\mathcal{C}(n,t-1)}$$
$$= \frac{1}{4^{t-T}N} \sum_{m \in \mathcal{C}(n,t)} z_m(0),$$

where the second equality follows straightforwardly by induction. At time t, the identical estimate at each cluster is a weighted average of the measurements from within that cluster.

Consensus is achieved at round T, where the four sub-clusters at level $t = T - 1$ broadcast their estimates to the entire network. Evaluating (21) for $t = T$, we observe that hierarchical averaging achieves perfect consensus; there is no need for a tolerance parameter ϵ. This somewhat surprising result is the consequence of combining the flexibility of wireless, which allows us to adjust the network connectivity at will, with the simplifying assumption of infinite-rate links. In the next section we will revisit this assumption.

Theorem 4: For hierarchical averaging over a random network of size N, the total transmit energy $E(N)$ and time-bandwidth product $B(N)$ scale, with high probability, as

$$E_{\epsilon}(N) = \Theta(N^{\kappa}) \quad (21)$$
$$B_{\epsilon}(N) = \Theta(N^{\kappa}), \quad (22)$$

for path-loss exponents $2 \leq \alpha < 4$, and for any $\kappa > 0$.

Proof: We derive the bound on $B_{\epsilon}(N)$ by examining the cardinality of the neighborhoods for each node. At time slot $t = 1$, by (17) each node transmits at power $P_n(1) = O(N^{(\kappa-1)\alpha/2})$. The neighborhood size of each node therefore scales as the number of nodes in a circle of radius $O(N^{\kappa-1})$. By Lemma 1, this number is $|\mathcal{N}_n(1)| = O(N^{\kappa})$ with probability approaching 1 as $N \to \infty$.

For rounds $2 \leq t \leq T$, since only one node per cluster transmits, we need to bound the number of *clusters* in range of each node. We chose the transmit powers such that the clusters transmit to each node in a circle of area $\pi M^2(t) = O(4^t N^{1-\kappa})$. By construction, each cluster $\mathcal{C}(n, t)$ covers an area of $O(4^t N^{1-\kappa})$. Therefore, the number of clusters that can fit into the circle is constant, so we have $B(t) = O(1)$. Summing over all rounds, we get

$$B_{\epsilon}(N) = O(N^{\kappa}) + \sum_{t=2}^{T} O(1) = O(N^{\kappa}). \quad (23)$$

Finally, we derive the bounds on E_{ϵ}. Substituting (17) and (20) into the definition of E_{ϵ}, we obtain

$$E_{\epsilon}(N) = \sum_{t=1}^{T} \sum_{n=1}^{N} P_n(t) \quad (24)$$

$$= O(N^{1+(\kappa-1)\alpha/2}) + \sum_{t=2}^{T} O(4^{T-t} 4^{(t-T)\alpha/2}), \quad (25)$$

where the second term in (25) follows from the fact that only one node from each of the 4^{T-t} clusters transmits at each

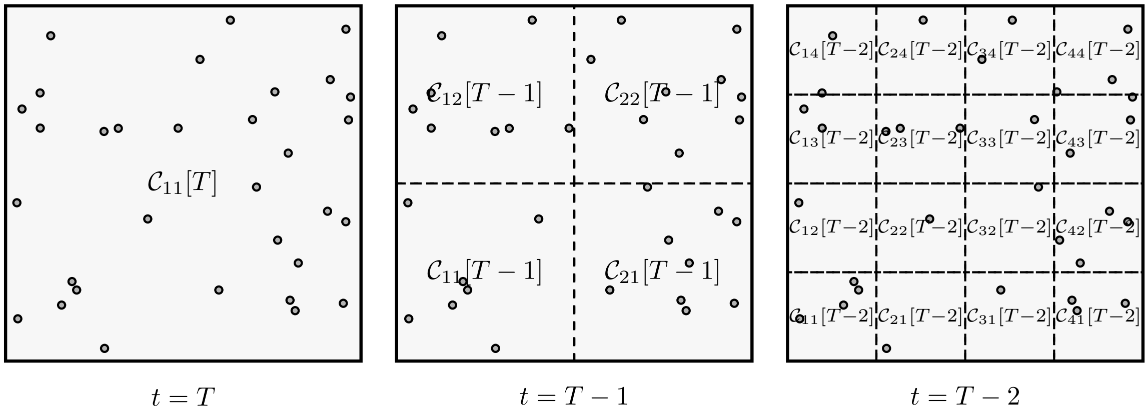

Fig. 1. Hierarchical partition of the network. Each square cell is divided into four smaller cells, which are each divided into four smaller cells, and so on.

round. Continuing, we get

$$E_\epsilon(N) = O(N^{1-\alpha/2+\kappa\alpha/2}) + 4^{T(1-\alpha/2)} \sum_{t=2}^{T} 4^{(\alpha/2-1)t} \quad (26)$$

$$= O(N^{1-\alpha/2+\kappa\alpha/2}) + O(4^{T(1-\alpha/2)} 4^{T(\alpha/2-1)}) \quad (27)$$

$$= O(N^{1-\alpha/2+\kappa\alpha/2} + O(1) \quad (28)$$

$$= O(N^{\kappa}), \quad (29)$$

where we have employed the finite geometric sum identity. ■

V. Numerical Results

We examine the empirical performance of the several consensus algorithms presented. First we choose $\gamma = 10$dB, $\alpha = 2$, $\epsilon = 10^{-4}$, and $\kappa = 0$. We let N run from 10 to 1000, averaging performance over 20 random initialization for each value of N. In Figure 3 we display the average transmit energy E_ϵ and the time-bandwidth product B_ϵ for both hierarchical averaging and path-averaging.

The simulations correspond to the theoretical results proven. With respect to the time-bandwidth product, hierarchical averaging performs best, the required number of sub-channel uses growing slowly and nearly on part with the constant lower bound. For path-averaging, on the other hand, the required sub-channel uses grows relatively quickly in N. With respect to transmit energy, the two algorithms perform similarly, the energy demands growing slowly in N. However, hierarchical averaging requires less energy on an absolute scale. As expected, randomized gossip performs worst with respect to both metrics.

Keeping the other parameters constant, we run another batch of simulations for $\alpha = 4$ and plot the results in Figure 3. Again hierarchical averaging performs best with respect to time-bandwidth product. With respect to total transmit energy, on the other hand, the relative performance depends on N. Path averaging achieves a better scaling law than hierarchical averaging, so for large N the energy required is smaller. For small N, however, hierarchical averaging requires less power on an absolute scale. Again randomized gossip performs worst.

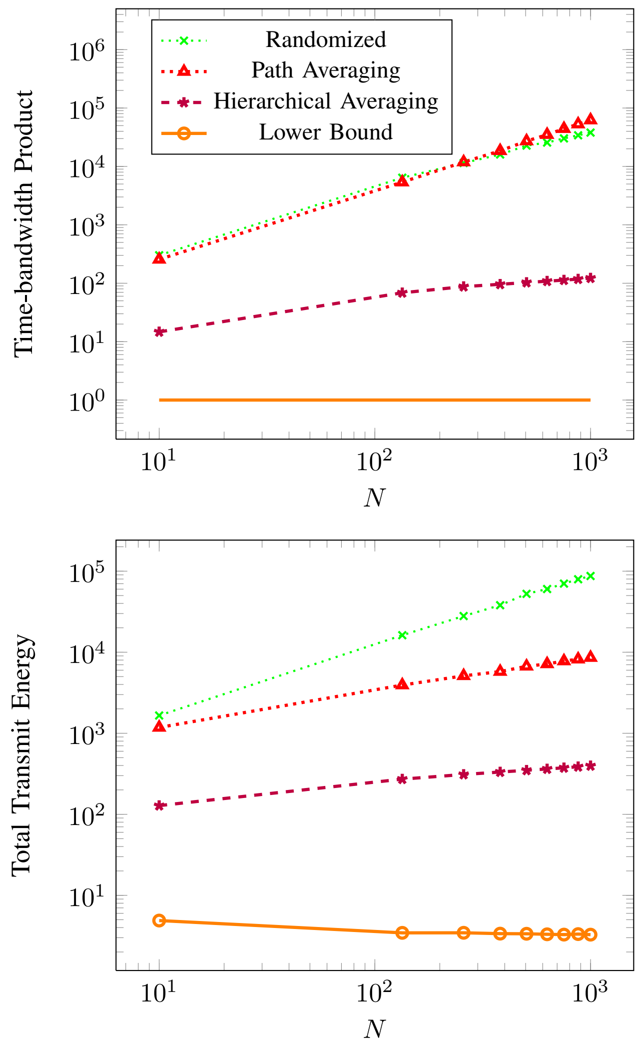

Fig. 2. Transmit energy E_ϵ and time-bandwidth product B_ϵ for a variety of consensus algorithms.

VI. Conclusion

Theorems 1-4 substantiate the claims made in Section I. For $\alpha > 2$, no scheme is order optimal in both metrics: Path averaging and multiscale gossip are nearly order-optimal in

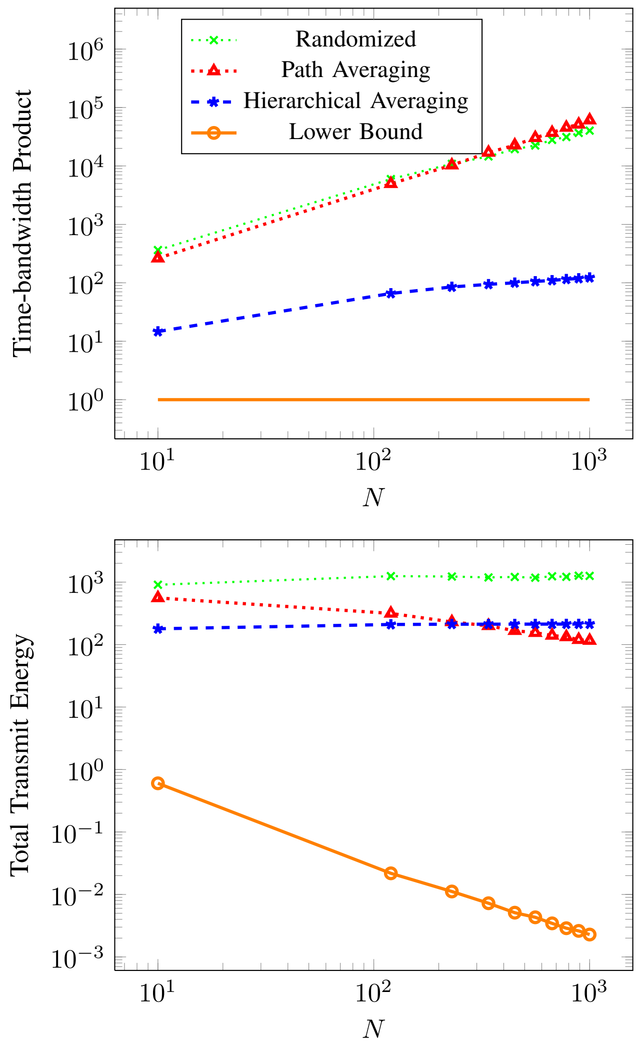

Fig. 3. Transmit energy E_ϵ and time-bandwidth product B_ϵ for a variety of consensus algorithms.

transmit energy, while hierarchical averaging is nearly order-optimal in time-bandwidth product. For $\alpha = 2$, however, hierarchical averaging achieves the lower bound in both metrics to within an arbitrarily small exponent. In the journal version [9], we present a cooperative version of hierarchical averaging which achieves better performance. We also explore the effects of quantization on consensus over the wireless medium.

References

[1] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Randomized gossip algorithms," *IEEE Trans. Info. Theory*, vol. 52, no. 6, pp. 2508 – 2503, June 2006.

[2] F. Benezit, A. Dimakis, P. Thiran, and M. Vetterli, "Order-optimal consensus through randomized path averaging," *Information Theory, IEEE Transactions on*, vol. 56, no. 10, pp. 5150 –5167, oct. 2010.

[3] K. I. Tsianos and M. G. Rabbat, "Multiscale gossip for efficient decentralized averaging in wireless packet networks," *submitted to IEEE Trans. Info. Theory*, 2010.

[4] D. Ustebay, B. Oreshkin, M. Coates, and M. Rabbat, "Rates of convergence for greedy gossip with eavesdropping," in *Proc. Allerton*, Monticello, IL, 2008.

[5] T. C. Aysal, M. E. Yildiz, A. D. Sarwate, and A. Scaglione, "Broadcast gossip algorithms for consensus," *IEEE Trans. Signal Processing*, vol. 57, no. 7, pp. 2748–2761, July 2009.

[6] B. Nazer, A. G. Dimakis, and M. Gastpar, "Local interference can accelereate gossip algorithms," *IEEE J. Selected Topics Signal Proc.*, vol. 5, no. 4, pp. 876–887, Aug. 2011.

[7] S. Vanka, V. Gupta, and M. Haenggi, "Power-delay analysis of consensus algorithms on wireless networks with interference," *Int. J. Systems, Control, and Communications*, vol. 2, no. 3, pp. 256–274, 2010.

[8] S. Kar, S. Aldosari, and J. Moura, "Topology for distributed inference on graphs," *Signal Processing, IEEE Transactions on*, vol. 56, no. 6, pp. 2609 –2613, june 2008.

[9] M. Nokleby, W. U. Bajwa, A. R. Calderbank, and B. Aazhang, "Toward resource-optimal consensus over the wireless medium," *submitted to IEEE Journal of Selected Topics in Signal Processing*, 2012. [Online]. Available: http://arxiv.org/abs/1208.3251

[10] A. Ozgur, O. Leveque, and D. Tse, "Hierarchical cooperation achieves optimal capacity scaling in ad hoc networks," *IEEE Trans. Info. Theory*, vol. 53, no. 10, pp. 3549 –3572, Oct. 2007.

[11] F. Benezit, A. G. Dimakis, P. Thiran, and M. Vetterli, "Gossip along the way: Order-optimal consensus through randomized path averaging," in *Proc. Allerton*, Sept. 2007.