

# Multiscale Gossip for Efficient Decentralized Averaging in Wireless Packet Networks

Konstantinos I. Tsianos and Michael G. Rabbat

## Abstract

This paper describes and analyzes a hierarchical gossip algorithm for solving the distributed average consensus problem in wireless sensor networks. The network is recursively partitioned into subnetworks. Initially, nodes at the finest scale gossip to compute local averages. Then, using geographic routing to enable gossip between nodes that are not directly connected, these local averages are progressively fused up the hierarchy until the global average is computed. We show that the proposed hierarchical scheme with  $k$  levels of hierarchy is competitive with state-of-the-art randomized gossip algorithms, in terms of message complexity, achieving  $\epsilon$ -accuracy with high probability after  $O(n \log \log n \log \frac{kn}{\epsilon})$  messages. Key to our analysis is the way in which the network is recursively partitioned. We find that the optimal scaling law is achieved when subnetworks at scale  $j$  contain  $O(n^{(2/3)^j})$  nodes; then the message complexity at any individual scale is  $O(n \log \frac{kn}{\epsilon})$ , and the total number of scales in the hierarchy grows slowly, as  $\Theta(\log \log n)$ . Another important consequence of hierarchical construction is that the longest distance over which messages are exchanged is  $O(n^{1/3})$  hops (at the highest scale), and most messages (at lower scales) travel shorter distances. In networks that use link-level acknowledgements, this results in less congestion and resource usage by reducing message retransmissions. Simulations illustrate that the proposed scheme is more message-efficient than existing state-of-the-art randomized gossip algorithms based on averaging along paths.

## I. INTRODUCTION

Distributed signal and information processing applications arise in a variety of contexts including wireless sensor networks, the smart-grid, large-scale unmanned surveillance, and mobile social networks. Large-scale applications demand protocols and algorithms that are robust, fault-tolerant, and scalable. Energy-efficiency is also an increasingly important design factor. When a system is comprised of battery-powered nodes or agents equipped with wireless radios for transmission—such as in wireless sensor networks—energy-efficiency equates to requiring few transmissions since in addition to consuming bandwidth, each wireless transmission dissipates battery resources.

Gossip algorithms [2], [3], [4], [5], [6] are an attractive paradigm for decentralized, in-network processing, and have received much attention in the computer science, systems and control, information theory, and signal processing research communities of late. Gossip algorithms are frequently posed and studied as solutions to the *distributed averaging* problem: in a network of  $n$  nodes whose topology is described by a graph  $G = (V, E)$  with  $|V| = n$ , each node initially has a scalar value  $x_i(0)$ , and the goal is to approximate the average,  $x_{ave} = \frac{1}{n} \sum_{i=1}^n x_i(0)$  at every node. Nodes iteratively and asynchronously exchange estimates with a small subset of the entire network, updating their local estimate after each exchange. These protocols have a number of attractive properties. The simplicity of the protocol (exchange information, update, repeat) makes it extremely robust; since there is no fixed routing of information to a fusion center and since all nodes compute a solution, there is no single point of failure or bottleneck. Furthermore, past studies have demonstrated that gossip algorithms converge even under unreliable or dynamic networking conditions; see, e.g., [6] and references therein.

However, the standard gossip algorithms for distributed averaging [2], [3], [4], [5] constrain information to only be exchanged between neighboring nodes and exhibit poor scaling and energy-efficiency in topologies frequently used to model connectivity in wireless networks, such as grids and random geometric graphs [7]. Roughly speaking, the number of messages transmitted per node depends linearly on  $n$ , the size of the network. Since the  $n$  values required to compute the average are initially stored at different nodes, any distributed averaging algorithm requires that each node perform at least one transmission. This discrepancy between constant and linear transmissions per

K.I. Tsianos and M.G. Rabbat are with the Department of Electrical and Computer Engineering, McGill University, Montréal, Québec, H3A 2A7 Canada, e-mail: konstantinos.tsianos@mail.mcgill.ca, michael.rabbat@mcgill.ca.

Parts of this work were presented at the 2010 International Conference on Distributed Computing in Sensor Systems; see [1].

node, has motivated the development of a number of variants of gossip algorithms specifically aimed at improving the efficiency of gossip on grid and random geometric graph topologies (see Section II for more).

The principle of hierarchical (multiscale) decomposition, or divide-and-conquer, arises in a variety of settings as a mechanism which yields efficient information processing procedures. In the signal processing and coding communities, multiscale analysis is frequently associated with wavelet-based methods, e.g., for signal and image denoising, edge detection, and transform coding [8]. A hierarchical approach to communication over wireless networks was shown to achieve the optimal capacity scaling law [9]. A recent study also found that flocks of birds exhibit hierarchical organization and suggested that hierarchical behavior has been selected (in the evolutionary sense) because it is more efficient than democratic or individualistic strategies [10].

This paper describes and analyzes a multiscale gossip algorithm for distributed averaging in grid and random geometric graph topologies. The network is recursively partitioned into smaller subnetworks. The size of subnetworks at each scale and the number of scales of the partition depend on the number of nodes in the network. Multiscale gossip operates over this partition in a bottom-up fashion. First, all nodes within each subnetwork at the finest scale gossip until computing a suitably accurate local average. Then, a representative node is elected for each subnetwork; an overlay grid is formed among all representatives within the same subnetwork at the next higher scale, and the representatives gossip over the overlay grid. This procedure is repeated until representatives at the coarsest scale have computed an accurate approximation to the network average. At that point the representatives disseminate their estimate to all of their children in the hierarchy. Multi-hop communication between representatives at coarser scales is accomplished using geographic routing [11], [12].

Our main contribution is the analysis of multiscale gossip. In particular, for a carefully designed multiscale partition, we show that the total number of single-hop transmissions required to reach a desired level of accuracy  $1/n$  scales nearly-linearly, requiring  $O(n \log \log n \log n)$  total transmissions as  $n \rightarrow \infty$  on random geometric graph and grid topologies. Consequently, the average number of transmissions per node is  $O(\log \log n \log n)$ . Since information dissemination (randomized broadcast) is much more efficient than gossip (which is a form of information diffusion) in these topologies, representatives at all scales can optionally disseminate intermediate results to other nodes in their subnetwork, thereby improving robustness and fault-tolerance of the scheme, without affecting the order-wise scaling law. In contrast to *geographic gossip with path averaging* [13], a randomized gossip scheme with a linear scaling that also uses geographic routing to exchange information over multiple hops, multiscale gossip requires fewer and shorter multi-hop transmissions; for example, in a  $\sqrt{n} \times \sqrt{n}$  grid topology, path averaging requires relaying messages over  $O(n^{1/2})$  hops, whereas multiscale gossip messages at the coarsest scale are relayed over at most  $O(n^{1/3})$  hops, and messages at finer scales travel significantly shorter distances. This has advantages when reliable transmission (i.e., handshaking, forward error-correcting, and/or retransmission) protocols are used at the link-level to ensure accurate reception over each link of a multi-hop path. Moreover, at each iteration of multiscale gossip, information is only exchanged between one pair of nodes, as opposed to all nodes along a path.

The remainder of this paper is organized as follows. Section II covers background, and related work. Section III describes the procedure for recursively constructing the hierarchical network partition and for carrying out multiscale gossip. Then, our main results are presented in Section IV, with analysis and proofs provided in Section V. A numerical evaluation of the proposed algorithm is presented in Section VI. Some practical considerations are discussed in Section VII, and we conclude in Section VIII.

## II. BACKGROUND AND PROBLEM DEFINITION

Our primary measure of performance is *communication cost*—the number of messages (transmissions over a single hop) required to compute an estimate to  $\epsilon$  accuracy—which is also considered in [11], [13]. Moreover, we are interested in characterizing scaling laws, or the rate at which the communication cost increases as a function of network size. In the analysis of scaling laws for gossip algorithms, a commonly studied measure of convergence rate is the  $\epsilon$ -averaging time, denoted  $T_\epsilon(n)$  and defined as [2]

$$T_\epsilon(n) = \sup_{x(0)} \inf \left\{ t : \Pr \left( \frac{\|x(t) - x_{ave}\|}{\|x(0)\|} \geq \epsilon \right) \leq \epsilon \right\}, \quad (1)$$

which is the number of iterations required to reach an estimate with  $\epsilon$  accuracy with high probability. The  $\epsilon$ -averaging time  $T_\epsilon(n)$  reflects the idea that the complexity of gossiping on a particular class of network topologies

should depend both on the final accuracy and the network size. When only neighbouring nodes communicate at each iteration,  $T_\epsilon(n)$  and communication cost are identical up to a constant factor. Otherwise, communication cost can generally be bounded by the product of  $T_\epsilon(n)$  and a bound on the number of messages required per iteration.

In wireless sensor network applications, *random geometric graphs* are a typical model for connectivity since communication is restricted to nearby nodes. In the 2-dimensional random geometric graph model,  $n$  nodes are randomly assigned coordinates uniformly in the unit square, and two nodes are connected with an edge when their Euclidean distance is less than or equal to a connectivity radius,  $r(n)$  [7], [14]. In [7] it is shown that if the connectivity radius scales as  $r_{\text{con}}(n) = \Theta(\sqrt{\frac{\log n}{n}})$  then the network is connected with high probability. Throughout this paper when we refer to a random geometric graph, we mean one with the connectivity  $r_{\text{con}}(n)$ .

Although the standard neighbor gossip algorithms are known to be efficient on complete graphs and expander-like topologies, they are also known to converge slowly on grids and random geometric graphs, two topologies commonly used to model wireless networks [3], [2]. Kempe, Dobra, and Gehrke [3] initiated the study of scaling laws for gossip algorithms and showed that gossip requires  $\Theta(n \log \epsilon^{-1})$  total messages to converge on complete graphs. Boyd, Ghosh, Prabhakar, and Shah [2] studied scaling laws for standard randomized gossip on random geometric graphs and found that communication cost scales as  $\Theta(\frac{n^2}{\log n} \log \epsilon^{-1})$  messages even if the algorithm is optimized with respect to the topology. This finding motivated the pursuit of efficient gossip algorithms for wireless networks in a number of interesting directions. For a complete overview of this line of work, we refer the reader to the recent survey [6]. Here we briefly discuss different approaches, focusing on advances most closely related to the present article.

A number of approaches seek more efficient computation while enforcing the constraint that information only be exchanged between neighboring nodes at each iteration. Most of these approaches introduce memory at each node, creating higher-order updates similar to shift-registers or polynomial filters [15], [16]. Scaling laws for a deterministic, synchronous variant of this approach are presented in [17], leading to  $\Theta(\frac{n^{1.5}}{\sqrt{\log n}} \log \epsilon^{-1})$  communication cost. Related asynchronous gossip algorithms based on lifted Markov chains have been proposed that achieve similar scaling laws [18], [19]. Recent work [20] suggests that no gossip algorithm on grids and random geometric graphs can achieve better than  $O(n^{1.5} \log \epsilon^{-1})$  scaling while constraining information exchange to be solely between neighboring nodes.

A variant called *geographic gossip*, proposed by Dimakis, Sarwate, and Wainwright [11], achieves a communication cost of  $\Theta(\frac{n^{1.5}}{\sqrt{\log n}} \log \epsilon^{-1})$  by allowing distant (non-neighbouring) pairs of nodes to gossip at each iteration. Assuming that each node knows its own coordinates and the coordinates of its neighbours in the unit square, communication between arbitrary pairs of nodes is made possible using *greedy geographic routing*. Rather than addressing nodes directly, a message is sent to a randomly chosen target  $(x, y)$ -location, and the recipient of the message is the node closest to that target. To reach the target, a message is forwarded from a node to its neighbour who is closest to the target. If a node is closer to the target than all of its neighbours, this is the final message recipient. It is shown in [11] that for random geometric graphs with connectivity radius  $r(n) = r_{\text{con}}(n)$ , greedy geographic routing succeeds with high probability. For an alternative form of greedy geographic routing, which may be useful in implementations see [12]. The main contribution of [11] is to illustrate that allowing nodes to gossip over multiple hops can lead to significant improvements in message cost. In follow-up work, Benezit, Dimakis, Thiran, and Vetterli [13] showed that a modified version of geographic gossip, called *path averaging*, can achieve  $\Theta(n \log \epsilon^{-1})$  message cost on random geometric graphs. To do this, all nodes along the path from the source to the target participate in a gossip iteration. If geographic routing finds a path of nodes  $S = \{x_i, \dots, x_j\}$  to deliver a message from  $x_i$  to  $x_j$ , the estimates of all nodes in  $S$  are accumulated on the way to  $x_j$ . Then  $x_j$  computes the average of all  $|S|$  values and sends the average back down the same path towards  $x_i$ , and all nodes in  $S$  update their estimates.

The multiscale approach considered in this paper also assumes that the network is capable of geographic routing in order to gossip among representative nodes at each scale. Below, we show that, asymptotically, the communication cost of multiscale gossip is  $O(n \log \log n \log \epsilon^{-1})$  messages, which is equivalent to that of path averaging up to a logarithmic factor. However, in multiscale gossip, information is only exchanged between pairs of nodes, and there is no averaging along paths. If an adversary wishes to disrupt gossip computation by forcing the network to drop a particular message or by deactivating a node in the middle of an iteration, a substantial amount of information can be lost in path averaging since each iteration involves  $O(\sqrt{\frac{n}{\log n}})$  nodes on average. In addition, the longest distance

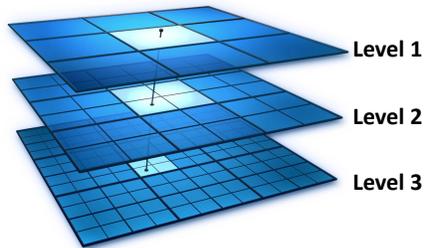


Fig. 1. Hierarchical multiscale subdivision of the unit square. At each level, each cell is split into equal numbers of smaller cells. Before the representatives of the cells can gossip on a grid graph, we run gossip on each cell.

a message travels in our multiscale approach is  $O(n^{1/3})$  hops in comparison to  $O(n^{1/2})$  hops for geographic gossip or path averaging.

We note that we are not the first to propose gossiping in a multiscale or hierarchical manner. Sarkar et al. [21] describe a hierarchical approach for computing aggregates, including the average. However, because their algorithm uses order and duplicate insensitive synopses to estimate the desired aggregate, the size of each message exchanged between a pair of nodes must scale with the size of the network. Other hierarchical distributed averaging schemes that have been proposed in the literature focus on the synchronous form of gossip, and they do not prove scaling laws for communication cost neither do they provide rules for forming the hierarchy (i.e. assume the hierarchical decomposition is given) [22], [23], [24].

A preliminary version of this work appears in the conference paper [1]. That paper describes multiscale gossip, including our construction of the multiscale network partition. The present manuscript builds upon [1], providing a more detailed analysis, a discussion of alternative partition schemes, and an evaluation of multiscale gossip in wireless packet networks when a link-layer acknowledgement scheme is used to ensure reliable delivery of messages over each hop.

### III. MULTISCALE GOSSIP

Multi-scale gossip performs averaging in a hierarchical manner. At each moment only nodes in the same level of hierarchy do computations at a local scale and computation at one level begins after the previous level has finished. By hierarchically decomposing the initial graph into subgraphs, we impose an order in the computation. As shown in the next section, for a specific decomposition it is possible to divide the overall work into a small number of linear sub-problems and thus obtain very close to linear complexity in the size of the network overall.

Assume we have a random geometric graph  $G = (V, E)$  and each node knows its own coordinates in the unit square and the locations of its immediate neighbours. Each node also knows the total number of nodes  $n$  in the network, and  $k$ , the desired number of hierarchy levels<sup>1</sup>. Figure 1 illustrates an example with  $k = 3$ . We use the convention that level  $k$  is the lowest level where the unit square is split into many small cells. Level 1 is the top level where we only have few big cells. All cells at the same level have the same area. The way we split each cell into subcells is directed by a subdivision constant  $a = \frac{2}{3}$  whose value is justified in Section V. If a cell contains  $n$  nodes, it is split into  $n^{1-a}$  cells of dimensions  $n^{\frac{a-1}{2}} \times n^{\frac{a-1}{2}}$  each.

At level  $k$  the unit square is split into  $d_k$  cells  $C_{(k,1)}, \dots, C_{(k,d_k)}$ . Each cell  $C_{(k,\cdot)}$  contains nodes forming a subgraph of  $G$ . Each subgraph  $G_{(k,\cdot)}$  runs standard randomized gossip until convergence. Then, in each subgraph  $G_{(k,\cdot)}$  we elect one representative node  $L_{(k,\cdot)}$ . The representative selection can be randomized or deterministic as explained in Section VII. Generally, not all subgraphs have the same number of nodes. For this reason, the value of each representative has to be reweighed proportionally to its subgraph's size. At level  $k - 1$ , the unit square is split into  $d_{k-1}$  cells  $C_{(k-1,1)}, \dots, C_{(k-1,d_{k-1})}$ . Each cell  $C_{(k-1,\cdot)}$  contains the same number of  $C_{(k,\cdot)}$  cells. The

<sup>1</sup>As explained in Section V, given  $n$ , the number of levels  $k$  can be computed automatically.

representatives of the  $G_{(k,\cdot)}$  graphs form grid graphs  $G_{(k-1,1)}, \dots, G_{(k-1,d_{k-1})}$  with two representatives  $L_{(k,i)}$  and  $L_{(k,j)}$  sharing an edge in a  $G_{(k-1,\cdot)}$  if cells  $C_{(k,i)}$  and  $C_{(k,j)}$  are adjacent and contained in the same  $C_{(k-1,\cdot)}$  cell. Note that representatives can determine which cells they are adjacent to given the current level of hierarchy and  $n$  since the cell construction is deterministic. Next, we run randomized gossip simultaneously on all  $G_{(k-1,\cdot)}$  grid graphs. Finally, we select a representative node  $L_{(k-1,\cdot)}$  for each  $G_{(k-1,\cdot)}$  grid graph and continue the next hierarchy level. The process repeats until we reach level 1 at which point we have only one grid graph  $G_{(1,1)}$  contained in the single cell  $C_{(1,1)}$ . By construction  $C_{(1,1)}$  coincides with the unit square. Once randomized gossip on  $G_{(1,1)}$  is over, each node of  $G_{(1,1)}$  as a representative  $L_{(2,\cdot)}$  disseminates its final value to all the nodes in its  $C_{(2,\cdot)}$  cell.

---

**Algorithm 1** MultiscaleGossip( $x_{init}, C, n, q, k, \epsilon$ )

---

```

1:  $a = \frac{2}{3}$ 
2: if  $q < k$  then
3:   Split  $C$  into  $m_{q+1} = n^{1-a}$  cells:  $C_{(q+1,1)}, \dots, C_{(q+1,m_{q+1})}$ 
4:   Select a representative node  $L_{(q+1,i)}$  for each cell  $C_{(q+1,i)}, i \in \{1, \dots, m_{q+1}\}$ 
5:   for all cells  $C_{(q+1,i)}$  do
6:     call HierarchicalGossip( $x_{init}(C_{(q+1,i)}), C_{(q+1,i)}, n^a, q+1, tol$ )
7:   end for
8:   Form grid graph  $G_{(q,\cdot)}$  of representatives  $L_{(q+1,i)}$ 
9:   call RandomizedGossip( $x_{init}(L_{(q+1,1:m_{q+1})}), G_{(q,\cdot)}, \epsilon$ )
10:  if  $q = 1$  then
11:    Spread value of  $L_{(2,i)}$  to all nodes in each  $C_{(2,i)}$ 
12:  end if
13: else
14:   Form graph  $G_{(k,\cdot)}$  only of nodes in  $V(G)$  contained in  $C$ 
15:   call RandomizedGossip( $x_{init}, G_{(k,\cdot)}, \epsilon$ )
16:   Reweight representative values as :  $x(L_{(k,i)}) = x(L_{(k,i)}) \frac{|V(G_k)| \cdot m_{k-1}}{|V(G)|}$ 
17: end if

```

---

Algorithm 1 describes multi-scale gossip in a recursive manner. The initial call to the algorithm has as arguments, the vector of initial node values ( $x_{init}$ ), the unit square ( $C = [0, 1] \times [0, 1]$ ), the network size  $n$ , the top level  $q = 1$ , the desired number of hierarchy levels  $k$  and the desired error tolerance  $\epsilon$  to be used by each invocation of randomized gossip. In a down-pass the unit square is split into smaller and smaller cells all the way to the  $C_{(k,\cdot)}$  cells. After gossiping in the  $G_{(k,\cdot)}$  subgraphs in *Line 15*, the representatives adjust their values (*Line 16*). As explained in the next section, if  $k$  is large enough, each  $G_{(k,\cdot)}$  is a complete graphs. Since each node knows the locations of its immediate neighbours (needed for geographic routing), at level  $k$  it is easy to also compute the size of each  $G_{(k,\cdot)}$  graph which is needed for the reweighting. The up-pass begins with the  $L_{(k,\cdot)}$  representatives forming the  $G_{(k-1,\cdot)}$  grid graphs (*Line 8*) and then running gossip in all of them in parallel. Between consecutive levels we use a parameter  $a = \frac{2}{3}$  to decide how many  $C_{(r+1,\cdot)}$  cells fit in each  $C_{(r,\cdot)}$  cell. As mentioned earlier, the motivation for this parameter and its specific value is explained in the following section. Notice the pseudocode mimics a sequential single processor execution which is in line with the analysis that follows in Section V. However, it should be emphasized that the algorithm is intended for and can be implemented in a distributed fashion. The notation  $x_{init}(C)$  or  $x_{init}(L)$  indicates that we only select the entries of  $x_{init}$  corresponding to nodes in cell  $C$  or representatives  $L$ .

#### IV. MAIN RESULTS

Before proceeding with the details we present here our main results. Proofs are provided in Section V below. Section III above describes multiscale gossip, an algorithm for distributed average consensus on random geometric graphs which uses randomized gossip as a black box. If each invocation of randomized gossip runs up to  $\epsilon$  accuracy, the total number of messages used by multiscale gossip is given in Theorem 1.

**Theorem 1.** *Let a random geometric graph  $G$  of size  $n$  and constant  $\epsilon > 0$  be given. As the graph size  $n \rightarrow \infty$ , the communication cost of the multiscale gossip scheme described above with scaling constant  $\alpha = \frac{2}{3}$  behaves as follows:*

- 1) *If the number of hierarchy levels  $k$  remains fixed as  $n \rightarrow \infty$ , then the communication cost of multiscale gossip is  $O(((k-1)n + n^{1+(\frac{2}{3})^{k-1}}) \log \epsilon^{-1})$  messages.*
- 2) *If the number of levels grows according to  $k = \Theta(\log \log n)$  as  $n \rightarrow \infty$ , then the communication cost of multiscale gossip is  $O(n \log \log n \log \epsilon^{-1})$  messages.*

Note that  $\epsilon$  in the theorem above is the target level of relative accuracy used each time we run randomized gossip on one overlay network, and not to the level of accuracy of the final average. Errors at intermediate levels can accumulate, but this accumulation is not catastrophic. An upper bound on this final accuracy is given in Theorem 2 below.

**Theorem 2.** *Let a random geometric graph  $G$  with  $n$  nodes and initial values on the nodes  $\bar{x}$  be given. If we run multiscale gossip on  $G$  using  $k$  levels of hierarchy and demand  $\epsilon$ -accuracy for randomized gossip at each subgraph, then with probability at least  $(1 - \epsilon)^g$ , where  $g$  is the total number of invocations of randomized gossip, the final error is not more than  $n\epsilon$  i.e.,*

$$\frac{\|\bar{x}_{final} - \bar{x}_{av}\|}{\|\bar{x}\|} \leq n\epsilon, \quad (2)$$

where  $\bar{x}_{final} \in \mathbb{R}^V$  denotes the vector of final estimates at each node, and  $\bar{x}_{av}$  denotes a  $n$ -vector whose entries are all set to the average of the initial values at each node.

The bound in Theorem 2 is loose in two senses as explained in the end of its proof; namely, both the error bound is loose, and the probability with which the result holds is loose. The recursive partitioning scheme produces  $g = 1 + n^{1-a} + n^{1-a^2} + \dots + n^{1-a^{k-1}} = O(kn)$  graphs and one invocation of randomized gossip for each. We can control both the accuracy and the probability of success by carefully setting the value of  $\epsilon$ , the accuracy used each time we invoke randomized gossip within the overall multiscale gossiping procedure. If we want the final accuracy of multiscale gossip to be  $\delta$  with high probability, we set the required accuracy for each randomized gossip call to  $\epsilon = \frac{\delta}{kn}$ . This will yield final accuracy  $\frac{\delta}{k}$  which is in fact better than required. Moreover, the probability of achieving this accuracy will be at least  $(1 - \frac{\delta}{kn})^g \geq (1 - \frac{\delta}{kn})^{kn} \rightarrow e^{-\delta} \geq 1 - \delta$  as  $n \rightarrow \infty$ . The adjustment in  $\epsilon$  also affects the total number of transmissions, as per Theorem 1. Specifically, multiscale gossip requires  $O(n \log \log n \log \frac{1}{\epsilon}) = O(n \log \log n \log \frac{kn}{\delta}) = O(n \log \log n \log(kn) + n \log \log n \log \frac{1}{\delta})$ . As we see the transmissions are only increased by a logarithmic factor. In particular, letting the number of levels of hierarchy scale as  $k = \Theta(\log \log n)$  and taking  $\delta = \frac{1}{n}$  yields an overall message complexity of  $O(n \log \log n \log n)$ .

Besides the above main theoretical results, we have compared multiscale gossip to path averaging which is a recent state-of-the-art linear complexity algorithm. The experiments presented in Section VI suggest that multiscale gossip has superior performance for graphs of up to many thousands or nodes. We also include an evaluation in scenarios with unreliable transmissions.

## V. ANALYSIS

### A. Proof of Theorem 1

Suppose we run multiscale gossip (Algorithm 1) on a random geometric graph  $G = (V, E)$  with  $|V| = n$  and transmission radius  $r(n) = \sqrt{\frac{c \log n}{n}}$ . Call the unit square cell  $C_{(1,1)}$ , the topmost cell in the partition hierarchy, and suppose we partition down to  $k$  levels. At the highest level, we split the unit square into  $m_2 = n^{1-a}$  cells  $C_{(2,\cdot)}$  each of dimensions  $n^{\frac{a-1}{2}} \times n^{\frac{a-1}{2}}$  where  $a = \frac{2}{3}$  as explained below. The representative nodes at this level form an overlay grid, so that representatives only exchange gossip messages with other representatives in adjacent cells. In general, on a 2-D grid of  $p$  nodes, randomized gossip requires  $O(p^2 \log \epsilon^{-1})$  messages to achieve accuracy  $\epsilon$  with probability  $1 - \epsilon$  (e.g. see [2]). The grid graph  $G_{(1,1)}$  formed by the representatives of the  $C_{(2,\cdot)}$  cells has  $n^{1-a}$  nodes. Moreover, by using appropriately large constant  $c$  in the transmission radius (e.g.  $c = 3$ ), the random geometric graph  $G$  is *geo-dense* [25] which means that a patch of area  $n^{a-1}$  contains  $\Theta(n^{a-1} \cdot n) = \Theta(n^a)$  nodes with high probability. The maximum distance between two nodes of  $G_{(1,1)}$  will be  $\sqrt{5}n^{\frac{a-1}{2}} = O(n^{\frac{a-1}{2}})$ . To see this compute the maximum possible distance between two nodes in adjacent  $C_{(2,\cdot)}$  cells using the Pythagorean theorem. If we divide by  $r(n)$ , we

get a worst case estimate of the cost for multi-hop messages between representatives at level 1:  $MsgCost_1 = O(n^{\frac{a}{2}})$  hops. Now the total number of transmissions on  $G_{(1,1)}$  will be  $O((n^{1-a})^2 \log \epsilon^{-1}) \cdot O(n^{\frac{a}{2}}) = O(n \log \epsilon^{-1})$  if  $a = \frac{2}{3}$ .

At the next level of hierarchy, we subdivide each cell  $C_{(2,\cdot)}$  of  $q = \Theta(n^a)$  nodes into  $q^{1-a'}$  cells in a recursive manner. Each cell  $C_{(3,\cdot)}$  will have dimensions  $q^{\frac{a'-1}{2}} \times q^{\frac{a'-1}{2}}$ . Using the exact same analysis as above, we will have  $n^{1-a}$  grid graphs  $G_{(2,\cdot)}$  and each has  $q^{1-a'}$  (representative) nodes. The message cost between  $L_{(3,\cdot)}$  representatives is  $MsgCost_2 = O(q^{\frac{a'}{2}})$ . To make the total number of messages at level 2 linear, we get  $n^{1-a} \cdot O((q^{1-a'})^2 \log \epsilon^{-1}) \cdot O(q^{\frac{a'}{2}}) = n \Rightarrow a' = \frac{2}{3}$  as well. In general, at any intermediate level  $j$  the total number of transmissions is going to be

$$n^{1-a^{j-1}} \cdot O\left(\left(\left(n^{a^{j-1}}\right)^{1-a}\right)^2 \log \frac{1}{\epsilon}\right) \cdot O\left(\left(n^{a^{j-1}}\right)^{\frac{a}{2}}\right) \quad (3)$$

which is linear in  $n$  if  $a = \frac{2}{3}$ . Finally, at the last level  $k$ , the algorithm runs gossip on each subgraph of  $G$  with nodes contained inside each of the  $C_{(k,\cdot)}$  cells. We have  $O(n^{1-(\frac{2}{3})^{k-1}})$   $C_{(k,\cdot)}$  cells, each containing  $n^{(\frac{2}{3})^{k-1}}$  nodes which can communicate via single hop messages. Since we run randomized gossip on each subgraph, the total number of messages at the last level is  $O(n^{1+(\frac{2}{3})^{k-1}})$ . Summing up all levels, plus  $n$  messages to spread the final result back to all nodes, the total number of messages for multiscale gossip is  $O(((k-1)n + n^{1+(\frac{2}{3})^{k-1}}) \log \epsilon^{-1} + n) = O(((k-1)n + n^{1+(\frac{2}{3})^{k-1}}) \log \epsilon^{-1})$ .

For the second part of the theorem, observe that at level  $k$  each cell contains a subgraph of  $n^{(\frac{2}{3})^{k-1}}$  nodes in expectation. For constants  $m \geq 2$  and  $M \geq m$ , we can choose  $k$  so that each cell at the finest scale contains between  $m$  and  $M$  nodes with high probability, so that the cost per cell is bounded by  $M^2 \log \epsilon^{-1}$ . In other words, choose  $k$  such that  $m \leq n^{(\frac{2}{3})^{k-1}} \leq M$ , implying that  $k = \Theta(\log \log n)$ . Since the cost per cell at level  $k$  is now bounded by a constant for  $k = \Theta(\log \log n)$ , the total level  $k$  cost is  $O(n^{1-(\frac{2}{3})^{k-1}} \log \epsilon^{-1})$  and the overall cost is  $O(((k-1)n + n^{1+(\frac{2}{3})^{k-1}}) \log \epsilon^{-1} + n) = O(n \log \log n \epsilon^{-1})$ , completing the proof.

In Section VI we observe experimentally that we only need few levels in practice for networks of a few thousand nodes.

## B. Proof of Theorem 2

To simplify the discussion, we assume that each subgraph at each level has the same number of nodes. By geo-density of the random geometric graph [25], the number of nodes in each cell concentrates quickly as  $n$  grows. The same proof technique can be generalized, at the expense of much more cumbersome notation.

If each subgraph at level  $i$  has  $L_i$  nodes then we have  $n = \prod_{j=1}^k L_j$ . We introduce some notation to analyze the procedure or error propagation in its general form. The initial vector of values is  $(x_1, \dots, x_n)$ . It is convenient to rewrite each element in  $x$  as  $x_{l_1 l_2 \dots l_k}$  where  $1 \leq l_i \leq L_i$ . We will write the whole vector as  $\{x_{l_1 l_2 \dots l_k}\}$  using brackets. We overload our notation to describe the values of nodes that gossip at any level. For example at level  $j$  we have the node values  $x_{l_1 l_2 \dots l_j}$  where the number of indices is indicative of the level. We will use  $*$  notation to signify the converged values after gossiping at any level e.g. After gossiping at the level  $j$ , the node values are transformed to  $x_{l_1 l_2 \dots l_j}^*$ . Moreover, to advance from level  $j+1$  to level  $j$  we need to select one node at each subgraph at level  $j+1$  as a representative and promote its value to the next level. This means that  $x_{l_1 l_2 \dots l_j} = x_{l_1 l_2 \dots l_j c}^*$  for some  $c$  in the range  $1 \leq c \leq L_{j+1}$ .

Let us also write  $m_{l_1 l_2 \dots l_{j-1}}^{L_j}$  to denote the mean of the values in a subgraph at level  $j$  i.e.,

$$m_{l_1 l_2 \dots l_{j-1}}^{L_j} = \frac{\sum_{s_j=1}^{L_j} x_{l_1 l_2 \dots l_{j-1} s_j}}{L_j}. \quad (4)$$

To begin the proof we first state the error bounds for each intermediate subgraph after running randomized gossip to  $\epsilon$ -accuracy. At level  $j$  for each subgraph we have

$$\frac{\|\{x_{l_1 l_2 \dots l_{j-1} 1: L_j}^*\} - \overline{m}_{l_1 l_2 \dots l_{j-1}}^{L_j}\|}{\|\{x_{l_1 l_2 \dots l_{j-1} 1: L_j}\|} \leq \epsilon. \quad (5)$$

where  $\bar{m}$  signifies a vector with all its element equal to  $m$ . Using the definition of the 2-norm, the inequality should hold for each summand, and so,

$$\epsilon \|\{x_{l_1 l_2 \dots l_{j-1} 1 : L_j}\}\| \geq |x_{l_1 l_2 \dots l_{j-1} c}^* - m_{l_1 l_2 \dots l_{j-1}}^{L_j}| \quad (6)$$

$$= |x_{l_1 l_2 \dots l_{j-1}} - m_{l_1 l_2 \dots l_{j-1}}^{L_j}|, \quad (7)$$

where the last equality follows since  $x_{l_1 l_2 \dots l_{j-1} c}^* = x_{l_1 l_2 \dots l_{j-1}}$ . At the top level 1 we have:

$$\frac{\|\{x_{1:L_1}^*\} - \bar{m}^{L_1}\|}{\|\{x_{1:L_1}\}\|} \leq \epsilon. \quad (8)$$

We can obtain a bound by squaring both sides and using the definition of the norm, and a second bound is again obtained by observing that each summand must be less than the right hand side:

$$\sum_{s_1=1}^{L_1} (x_{s_1}^* - m^{L_1})^2 \leq \epsilon^2 \|\{x_{1:L_1}\}\|^2 \quad (9a)$$

$$|x_{s_1}^* - m^{L_1}| \leq \epsilon \|\{x_{1:L_1}\}\|. \quad (9b)$$

Once level 1 is finished, the final values are distributed to all the nodes that each node represents. We are interested to bound the following :

$$\frac{\|(x_1^*, \dots, x_1^*, \dots, x_{L_1}^*, \dots, x_{L_1}^*) - \bar{x}_{av}\|}{\|x\|} \quad (10)$$

where in the above expression each value  $x_{l_1}^*$  is repeated  $L_k \times \dots \times L_2$  times.

We start by bounding the squared expression for simplicity. Using the definition of the 2-norm, adding and subtracting the mean at level 1, expanding the quadratic term, and using the bound (9a),

$$\frac{\|(x_1^*, \dots, x_1^*, \dots, x_{L_1}^*, \dots, x_{L_1}^*) - \bar{x}_{av}\|^2}{\|x\|^2} \quad (11)$$

$$= \frac{L_k \dots L_2 \sum_{s_1=1}^{L_1} (x_{s_1}^* - x_{av})^2}{\|x\|^2} \quad (12)$$

$$= \frac{L_k \dots L_2 \sum_{s_1=1}^{L_1} (x_{s_1}^* - m^{L_1} + m^{L_1} - x_{av})^2}{\|x\|^2} \quad (13)$$

$$= \frac{L_k \dots L_2 \left[ \sum_{s_1=1}^{L_1} (x_{s_1}^* - m^{L_1})^2 + 2(m^{L_1} - x_{av}) \sum_{s_1=1}^{L_1} (x_{s_1}^* - m^{L_1}) + L_1 (m^{L_1} - x_{av})^2 \right]}{\|x\|^2} \quad (14)$$

$$\leq \frac{L_k \dots L_2 \left[ \epsilon^2 \|\{x_{1:L_1}\}\|^2 + 2(m^{L_1} - x_{av}) \sum_{s_1=1}^{L_1} (x_i^* - m^{L_1}) + L_1 (m^{L_1} - x_{av})^2 \right]}{\|x\|^2} \quad (15)$$

At this point we focus on bounding the parts of the above numerator separately. For details of the derivation please see the appendix. The final results are:

$$\frac{\epsilon^2 \|\{x_{1:L_1}\}\|^2}{\|x\|^2} \leq \epsilon^2 \quad (16)$$

$$\frac{m^{L_1} - x_{av}}{\|x\|} \leq (k-1)\epsilon \quad (17)$$

$$\frac{\sum_{s_1=1}^{L_1} (x_{s_1}^* - m^{L_1})}{\|x\|} \leq L_1 \epsilon^2 + \sqrt{\frac{L_1}{L_2}} L_1 \epsilon \quad (18)$$

$$\frac{L_1 (m^{L_1} - x_{av})^2}{\|x\|^2} \leq L_1 (k-1)^2 \epsilon^2 \quad (19)$$

Using these bounds we obtain the overall bound:

$$\frac{\|(x_1^*, \dots, x_1^*, \dots, x_{L_1}^*, \dots, x_{L_1}^*) - \bar{x}_{av}\|^2}{\|x\|^2} \quad (20)$$

$$\leq \frac{L_k \dots L_2 \left( \epsilon^2 \|\{x_{1:L_1}\}\|^2 + 2(m^{L_1} - x_{av}) \sum_{s_1=1}^{L_1} (x_{s_1}^* - m^{L_1}) + L_1(m^{L_1} - x_{av})^2 \right)}{\|x\|^2} \quad (21)$$

$$\leq L_k \dots L_2 \left( \epsilon^2 + 2(k-1)\epsilon(L_1\epsilon^2 + \sqrt{\frac{L_1}{L_2}}L_1\epsilon) + L_1(k-1)^2\epsilon^2 \right) \quad (22)$$

$$\leq L_k \dots L_2 \epsilon^2 + 2(k-1)n\epsilon^3 + 2(k-1)n\sqrt{\frac{L_1}{L_2}}\epsilon^2 + n(k-1)^2\epsilon^2 \quad (23)$$

Coarsely bounding all the terms related to  $k$  by  $n$  gives

$$\frac{\|(x_1^*, \dots, x_1^*, \dots, x_{L_1}^*, \dots, x_{L_1}^*) - \bar{x}_{av}\|^2}{\|x\|^2} \quad (24)$$

$$\leq n\epsilon^2 + 2n^2\epsilon^3 + 2n^2\epsilon^2 + n^2\epsilon^2 \leq n\epsilon^2 + 2n^2\epsilon^2 + 2n^2\epsilon^2 + n^2\epsilon^2 \leq 6n^2\epsilon^2, \quad (25)$$

and so we arrive at the bound

$$\frac{\|(x_1^*, \dots, x_1^*, \dots, x_{L_1}^*, \dots, x_{L_1}^*) - \bar{x}_{av}\|}{\|x\|} \leq \sqrt{6}n\epsilon. \quad (26)$$

This bound will hold whenever all randomized gossip operations at intermediate subgraphs achieve  $\epsilon$  accuracy. Any invocation of randomized gossip achieves  $\epsilon$  accuracy with probability at least  $1 - \epsilon$  and all randomized gossip operations are independent of each other. If we have  $g$  subgraphs total appearing during a run of multiscale gossip, the probability that we achieve final error  $cn\epsilon$  is at least  $(1 - \epsilon)^g$ .

Notice that this bound is relatively loose. This should be expected given it was obtained using very loose bounds for worst case errors at all levels through equations 5 and 9. Moreover, if the number of subgraphs  $g$  is large, the final probability of success is low. As explained in section IV however, we can select an  $\epsilon$  to control both the final accuracy and the probability of success at the expense of logarithmically more transmissions.

### C. Is $a = \frac{2}{3}$ optimal?

We have selected  $a = \frac{2}{3}$  in the previous section to get linear cost at each intermediate hierarchy level. One could ask, whether this is the best we can do. Maybe a different choice of  $a$  could yield even smaller communication cost. We investigate this question here. As we will see, although  $a = \frac{2}{3}$  is not the unique optimal option, it is a well justified choice.

For convenience in the analysis we change the notation a little bit. In Sections III and V we use  $a = \frac{2}{3}$  as a rule for subdividing each cell at one level to its subcells. Here, let's assume that we have subdivision parameters  $b_1, b_2, \dots, b_{k-1}$  with a slightly different meaning. While  $a$  is "local" and allows transitioning from one level to the next,  $b_j$  directly specifies exactly how many cells and nodes in each cell we have at level  $j$ . Specifically, at level  $j$  there will be a total of  $n^{1-b_{j-1}}$  cells and each cell will contain  $n^{b_{j-1}}$  nodes and  $n^{b_{j-1}-b_j}$  leaders communicating over distances of  $n^{\frac{b_j}{2}}$  hops. There is a connection between the  $b$ 's since, if  $a$  is constant for all levels then  $b_j = a^j$ . We also require that  $b_j \geq b_{j-1}$ ; otherwise, there would be more nodes in a  $C_{j+1}$  cell than in a  $C_j$  cell which is not consistent with our notion of refining the hierarchical partition.

Recall from the complexity analysis that the total number of messages  $M_j$  at some level  $j$  are

$$M_j = (\#C_j \text{ cells}) \times (\# \text{ leaders per } C_j)^2 \times (\# \text{ hops per message}) \quad (27)$$

It is useful to write the exact expression for the most important cases so:

- At level 1 we only have one cell so:

$$M_1 = 1 \cdot (n^{1-b_1})^2 \cdot n^{\frac{b_1}{2}} = n^{2-\frac{2}{3}b_1} \text{ messages} \quad (28)$$

- At level  $1 < j < k$ :

$$M_j = n^{1-b_{j-1}} \cdot (n^{b_{j-1}-b_j})^2 \cdot n^{\frac{b_j}{2}} = n^{1+b_{j-1}-\frac{2}{3}b_j} \text{ messages} \quad (29)$$

- At last level  $k$  all messages can be delivered in one hop so:

$$M_k = n^{1-b_{k-1}} \cdot (n^{b_{k-1}})^2 \cdot 1 = n^{1+b_{k-1}} \text{ messages} \quad (30)$$

Notice that even if we select all the  $b_j$ 's so that  $M_j = O(n)$  the  $k$ -th level will dominate with superlinear complexity since  $b_{k-1} > 0$ . Now, if we take a fixed number of levels  $k$  we are interested to choose the  $b_j$ s that minimize the total number of messages  $\sum_{j=1}^k M_j$  as  $n \rightarrow \infty$ . This is equivalent to the optimization problem:

$$\text{minimize}_{b_1, \dots, b_{k-1}} \quad \max\{2 - \frac{2}{3}b_1, 1 + b_1 - \frac{2}{3}b_2, \dots, 1 + b_{k-1}\} \quad (31)$$

$$\text{subject to} \quad 1 \geq b_1 \geq b_2 \quad (32)$$

$$b_2 \geq b_3 \quad (33)$$

$$\vdots \quad (34)$$

$$b_{k-1} \geq 0. \quad (35)$$

In general the solution to this problem does not yield  $a = \frac{2}{3}$  at each level.

As discussed in section V to get near linear (e.g.  $O(n \log \log n \log \frac{kn}{\epsilon})$ ) complexity we can allow the number of levels  $k$  to depend on  $n$  so that  $k(n) \rightarrow \infty$  as  $n \rightarrow \infty$ . Notice that even if we use enough levels to have a fixed number of nodes at the finest level, we end up with a linear  $O(n)$  number of cells at level  $k$  and require constant time to gossip in each. As a result the finest level's complexity can be linear at best.

If the number of levels is variable, it depends on the values of the subdivision parameters  $b_j$ . If the  $b_j$ 's are large then each cell contains a lot of nodes and we need to use many levels until we reach the finest level. If on the other hand  $b_j$ 's are small, we create a lot of small cells and few levels. If the number of cells is too large however, we can no longer have  $O(n)$  messages at an intermediate level. Consequently, we need to have as few levels as possible. This principle together with the desire to have as few messages as possible can justify the selection of  $a = \frac{2}{3}$  at each level. To see this, let's demand that  $M_j = O(n)$ . For the first level this is true if  $2 - \frac{2}{3}b_1 \leq 1$  which gives  $b_1 \geq \frac{2}{3}$ . Similarly for any other level we need  $b_j \geq \frac{2}{3}b_{j-1}$ . Obviously the smallest possible  $b_j$ s that are still large enough to admit linear complexity at each level are such that  $b_j = (\frac{2}{3})^j$ . This is the same as using the same  $a = \frac{2}{3}$  to subdivide each cell at each level.

## VI. EXPERIMENTAL EVALUATION

The formal analysis of the algorithm complexity is given in section V. In this section we evaluate multiscale gossip in simulation and study its behaviour in practical scenarios. First we investigate the effect of using few versus many levels. Then we show that multiscale gossip performs very well against path averaging [13], the current state-of-the-art gossip algorithm that requires linear number messages in the size of the network to converge to the average with  $\epsilon$  accuracy. Finally, we investigate scenarios where transmissions do not always succeed and messages are either retransmitted or lost.

### A. Varying levels of Hierarchy

In the analysis we concluded that we can select the number of levels  $k = \log \log n$  i.e. we don't need too many levels. This can be verified in practice. Figure 2 investigates the effect of increasing the levels of hierarchy. The figure shows the number of messages until convergence within 0.0001 error, averaged over ten graphs of 5000 nodes. More levels yield a diminishing reward and we do not need more than 4 or 5 levels. As discussed in the next subsection these observations lead us to try a scheme with only two levels of hierarchy which still produces an efficient algorithm.

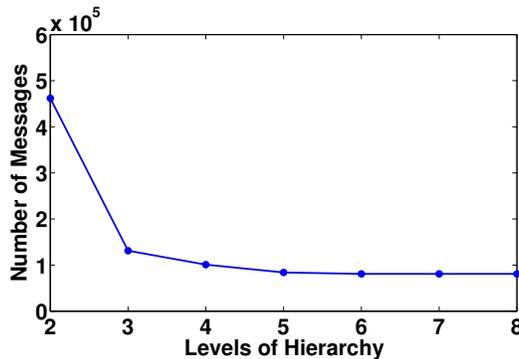


Fig. 2. Increasing the levels of Hierarchy yields a diminishing reward. Results are averages over 10 random geometric graphs with 5000 nodes and final desired accuracy  $\epsilon = 0.0001$ . All graphs are created with radius  $r = \sqrt{\frac{3 \log n}{n}}$ .

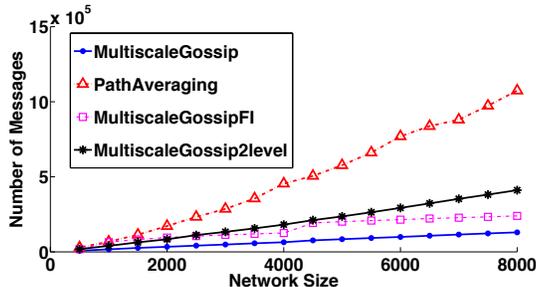


Fig. 3. Comparison of multiscale gossip to path averaging on random geometric graphs of increasing sizes. MultiscaleGossip used with 5 levels of hierarchy. MultiscaleGossipFI is the version using a fixed number of iterations for gossiping at a specific level. MultiscaleGossip2level is a version using only two levels of Hierarchy and is explained in Section VII. Results are averages over 20 random geometric graphs with 8000 nodes and final desired accuracy  $\epsilon = 0.0001$ . All graphs are created with radius  $r(n) = \sqrt{\frac{3 \log n}{n}}$ .

### B. Multiscale Gossip vs Path Averaging

We compare multiscale gossip against path average which is in theory the fastest algorithm for gossiping on random geometric graphs. Figure 3 shows the number of messages needed to converge within  $\epsilon = 0.0001$  error for graphs of sizes 500 to 8000. The bottom curve tagged *MultiscaleGossip* shows the ideal case where computation inside each cell stops automatically when the desired accuracy is reached. The curve labeled *MultiscaleGossipFI* was generated using fixed number of iterations per level based on worst case graph sizes and the curve labeled *MultiscaleGossip2level* was generated using only two levels of hierarchy and an  $a = \frac{1}{2}$  instead of  $\frac{2}{3}$ . Both of these variants are explained below. For path averaging we also simulate the ideal scenario where nodes stop transmitting automatically when achieving the targeted accuracy. As we see all variants of multiscale gossip use noticeably fewer transmissions than path averaging. One reason why path averaging seems to be slower than in [13] is because we use a smaller connectivity radius for our graphs ( $r = \sqrt{\frac{3 \log n}{n}}$  instead of  $r = \sqrt{\frac{10 \log n}{n}}$ ).

Figure 4 which depicts the cumulative density functions of transmissions for multiscale gossip and path averaging. Specifically, we plot the fraction of nodes that transmitted  $t$  times or less as a function for a random geometric graph with 2000 nodes. Both path averaging and multiscale gossip were stopped as soon as the desired error level is reached:  $\|x(t) - \bar{x}_{av}\| \leq 0.0001 \|x(0)\|$ . As we see, the node with most transmissions in multiscale gossip still sends less messages than about 22% of the nodes in path averaging.

Multiscale gossip has several advantages over Path Averaging. All the information relies on pairwise messages. In contrast, averaging over paths of length more than two has two main disadvantages. First, if a message is lost, a large number of nodes (potentially  $O(\sqrt{\frac{n}{\log n}})$ ) are affected by the information loss. Second, when messages are sent to a remote location over many hops, they increase in size as the message body accumulates the information of all the intermediate nodes. Besides being variable, the message size now depends on the length of the path and ultimately on the network size. Our messages are always of constant size and independent of the hop distance or

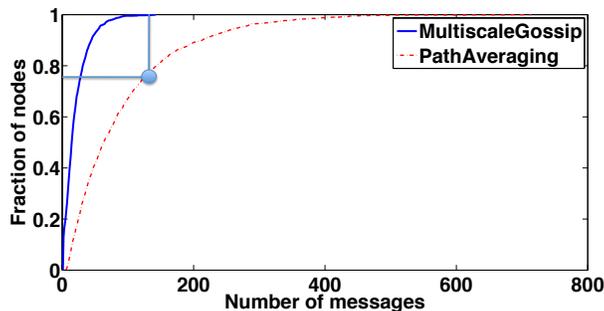


Fig. 4. Cumulative density function of the probability that a node sends less of equal to  $t$  messages as a function of  $t$  on a random geometric graph with 2000 nodes. The node with maximum number of transmissions for multiscale gossip has less transmissions than about 22% of the nodes in path averaging.

network size. Moreover, the maximum number of hops any message has to travel is  $O(n^{\frac{1}{3}})$  at worst<sup>2</sup>. This should be compared to distance  $O(\sqrt{n})$  which is necessary for path averaging to achieve linear scaling. Finally, multiscale gossip is relatively easy to analyze and implement using standard randomized gossip as a building block for the averaging computations.

*Fixed Number of Iterations per Level:* The ideal scenario for multiscale gossip is if computation inside each cell stops automatically when the desired accuracy is reached. This way no messages are wasted. However in practice for cells at the same level may need to gossip on graphs of different sizes that take different numbers of messages to converge. This creates a need for node synchronization so that all computation in one level is finished before the next level can begin. To alleviate the synchronization issue, we can fix the number of randomized gossip iterations per level so that all computation between different subgraphs at the same level take practically the same amount of time. However, we need to be careful not to do fewer iterations than needed for the desired accuracy. Given that nodes are deployed uniformly at random in the unit square, we can make a worst case estimate of how many nodes are expected to be in a cell of a certain area. Since by construction all cells at the same level have equal area, we gossip on all graphs at that level for a fixed number of iterations. Moreover, as seen in the previous section, we can use enough levels of hierarchy to only have  $m \leq k \leq M$  nodes at the last level. This can ensure that we will not do less iterations than necessary. In practice, usually at level  $k$ , we have less nodes than expected so we end up wasting messages running gossip for longer than necessary.

*Two-level Gossip:* multiscale gossip is a synchronized algorithm where computation in one level begins after the previous level has converged. Synchronization can be complicated or inefficient if we have too many levels. This motivates trying an algorithm with only two levels. In this case, for graphs of size a few thousand nodes, splitting the unit square into  $n^{1-a}$  cells with  $a = \frac{2}{3}$  is not a good choice as it produces a very small grid of representatives and quite large level-1 cells. To achieve better load balancing between the two levels, we use  $a = \frac{1}{2}$ . This choice has the advantage that the maximum number of hops any message has to travel is  $O(n^{\frac{1}{4}})$ . To see this is true, observe that each  $C_{(2,\cdot)}$  cell has area  $\frac{1}{n^a} = n^{-\frac{1}{2}} = n^{-\frac{1}{4}} \times n^{-\frac{1}{4}}$ . Thus the maximum distance between representatives is  $O(n^{-\frac{1}{4}})$ . If we divide by the connecting radius  $r(n) = \sqrt{\frac{c \log n}{n}}$  we get the result. Another interesting finding is that for moderate sized graphs, using cells of area  $n^{-\frac{1}{2}}$  produces subgraphs which are very well connected. Since nodes are deployed uniformly at random, an area  $n^{-\frac{1}{2}}$  is expected to contain  $n^{\frac{1}{2}}$  nodes. A subgraph inside a  $C_{(2,\cdot)}$  cell is still a random geometric graph with  $t = n^{\frac{1}{2}}$  nodes, but for which the radius used to connect nodes is not  $\sqrt{\frac{c \log t}{t}}$ . It is  $\sqrt{\frac{c \log n}{n}}$ . This is equivalent to creating a random geometric graph of  $t$  nodes in the unit square but with a scaled up radius of  $r_t = \sqrt{\frac{c \log n}{t}}$ . From [25] we know that a random geometric graph of  $t$  nodes is rapidly mixing (i.e. linear number of messages for convergence) if the connecting radius is  $r_{rapid} = \frac{1}{poly(\log t)}$ . Now, e.g. for  $c = 3$  and  $n \leq 9 * 10^6$ , we get  $r_t \geq \frac{1}{\log t} \geq r_{rapid}$  for  $t = \sqrt{n} \leq 3000$ . Consequently, the  $C_{(2,\cdot)}$  cells are rapidly mixing for networks of less than a few millions of sensors. In Figure 3 verifies this analysis. For graphs from 500 to 8000 nodes and final error 0.0001, we see that *MultiscaleGossip2level* performs very close to multiscale gossip

<sup>2</sup>At level 1 the distance in hops between leaders is at worst  $O(n^{\frac{a}{2}}) = O(n^{\frac{1}{3}})$  for  $a = \frac{2}{3}$ .

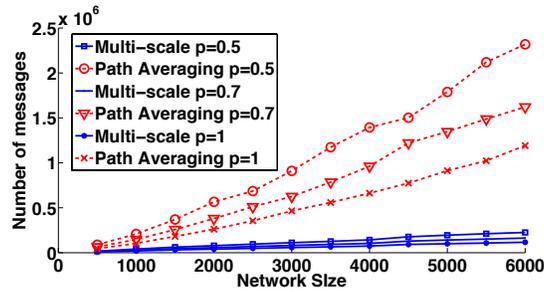


Fig. 5. Multiscale gossip and Path averaging on graphs of various sizes using hand shakes to overcome transmission failures. The probability of successful transmission  $p$  varies from 0.5 to 1. All results are averaged over 25 graphs of each size. Final accuracy is  $\epsilon = 0.0001$  and the random geometric graphs use  $r(n) = \sqrt{\frac{3 \log n}{n}}$ .

with more levels of hierarchy and better than path averaging.

### C. Operating under Transmission Failures

As explained in the previous section, multiscale gossip needs to send messages to shorter distances across the network than Path Averaging. It is important to see what effect this has on the robustness of the algorithms against transmission failures. Two different but general scenarios are considered. In the first scenario, no message is truly lost. There is a non-zero probability that a message over a network edge is not sent successfully, but the nodes communicate via a hand shake mechanism so messages are eventually delivered after a number of attempts. If the probability of successful transmission is  $p$ , then the cost for a single message transmission over an edge is geometrically distributed:  $P[\text{Cost} = m] = \text{Geo}(p, m)$ . The second scenario is more extreme. Each message is delivered with probability  $p$  over each edge, otherwise it is lost. This model has severe consequences. Depending on where in its path the message is lost, a number of nodes will not update their values properly so besides the overall delay in convergence, part of the signal energy is lost and the final estimate of the average is no longer guaranteed to be close to the true average.

1) *Hand Shake Model*: In this scenario we don't have to worry about convergence. All messages are eventually delivered and it is just a matter of time. Figure 5 shows the results of multiscale gossip against path averaging on networks of different sizes. The probability of successful transmission also varies from  $p = 0.5$  to  $p = 1$ . As we see multiscale gossip is significantly less affected by such failures. This example clearly illustrates the importance of not having to send messages in long distances of the network. Since each individual link introduces some delay, the fact that messages in multiscale gossip usually don't need to travel far and need to go  $O(n^{\frac{1}{3}})$  hops at most, allow the algorithm to converge using much fewer messages than path averaging.

2) *Message Loss Model*: In this scenario, a message is delivered with probability  $p$  or lost forever. This severely impedes the algorithms from converging fast. Moreover, depending on where along a path on the graph the network happens, information is lost permanently distorting the final result and making it impossible to meet the desired final accuracy. For example in path averaging, if the message is lost at the first transmission on its way back, all the nodes along the path except the last will have distorted information. Similar situations occur in multiscale gossip between leader communications where messages need to travel across multiple hops. Since in this scenario we have no guarantee that the desired final accuracy criterion will be met, it is hard to draw conclusions whether multiscale gossip is better than path averaging. Our observations showed that both algorithms can only reach an accuracy in the order of 0.01 when targeting at  $\epsilon = 0.0001$ . Specifically, multiscale gossip would only reach up to 0.06 accuracy while path averaging could not improve beyond 0.02. At the same time the total number of messages still increases linearly for multiscale gossip while it seem to blow up exponentially for path averaging.

## VII. PRACTICAL CONSIDERATIONS

There is a number of practical considerations that we would like to bring to the reader's attention. We list them in the form of questions below:

**How can we detect convergence in a subgraph or cluster? Do the nodes need to be synchronized?** At each hierarchy level, representatives know how big the grid that they are gossiping over is (function of  $n$  and  $k$  only). Moreover, all grids at the same level are of the same size and we have tight bounds on the number of messages needed to obtain  $\epsilon$  accuracy on grids w.h.p. We can thus gossip on all grids for a fixed number of rounds and synchronization is implicit. At level  $k$  however, in general we need to gossip on random geometric subgraphs which are not of exactly the same size. As  $n$  gets large though, random geometric graphs tend to become regular and uniformly spaced on the unit square. Therefore, the subgraphs contained in cells at level  $k$  all have sizes very close to the expected value of  $n^{(\frac{2}{3})^k}$ . Thus, we run gossip for a fixed number of rounds using the theoretical bound for graphs of the size  $n^{(\frac{2}{3})^k}$ . As discussed in Section III, fixing the number of iterations leads to redundant transmissions, however the algorithm is still very efficient.

**What happens with disconnected subgraphs or grids due to empty grid cells?** This is possible since the division of the unit square into grid cells does not mean that each cell is guaranteed to contain any nodes of the initial graph. Representatives use multi-hop communication and connected grids can always be constructed as long as the initial random geometric graph is connected. At level  $k$  the subgraphs of the initial graph contained in each cell could still be disconnected if edges that go outside the cell are not allowed. However, as explained in Section V we can use enough hierarchy levels so that each  $C_{(k,\cdot)}$  cell is a complete graph and the probability of getting disconnected  $C_{(k,\cdot)}$  cells tends to zero.

**How can we select representatives in a natural way?** The easiest solution is to pick the point  $p_c$  that is geographically at the center of each cell. Again, knowledge of  $n, k$  uniquely identifies the position of each cell and also  $p_c$ . By sending all messages to  $p_c$ , geographic routing will deliver them to the unique node that is closest to that location w.h.p. To change representatives, we can deterministically pick a location  $p_c + u$  which will cause a new node to be the closest to that location. A more sophisticated solution would be to employ a randomized auction mechanism. Each node in a cell generates a random number and the largest number is the representative. Once a new message enters a cell, the nodes knowing their neighbours values, route the message to the cell representative. Notice that determining cell leaders this way does not incur more than linear cost.

**Are representatives bottlenecks and single points of failure?** This is not an issue. There might be a small imbalance in the amount of work done by each node, but it can be alleviated by selecting different representatives at each hierarchy level. Moreover, for increased robustness, at a linear cost we can disseminate the representative's values to all the nodes in its cell. This way if a representative dies, another node in the cell can take its place. The new representative will have a value very similar (within  $\epsilon$ ) to that of the initial representative at the beginning of the computation at the current level. Thus node failure is expected to only cause small delay in convergence at that level. We should emphasize however that the effect of node failures has received little attention so far and still asks for a more systematic investigation.

**How much extra energy do the representatives need to spend?** This question is difficult to answer analytically. We use simulation to get a feel for it. Figure 6 shows the number of messages sent by each of the 5000 nodes in a random geometric graph. For this case we used five levels of hierarchy. We expect that some nodes will transmit more messages since, as we move down the hierarchy, cells get smaller and there are fewer nodes from which to draw representatives. In this example, by randomly selecting representatives at each level, no node was a representative more than 3 times. We show the number of transmissions for nodes of each type in table VII, including messages relayed by intermediate nodes using geographic routing. As we see, most of the nodes use very few messages. Moreover, the average degree of this particular example is 26, and thus, on average each node sends fewer messages than it has neighbors.

Figure 7 illustrates the typical fraction of transmissions per node as a function of location in the unit square normalized as a probability distribution. Each pixel is assigned an intensity between proportional to the number of messages that nodes located in that region will typically transmit. Thus this figure reveals which nodes suffer from the heaviest traffic. The figure is the result of averaging over 100 realizations of gossip for each of 100 different 1000-node random geometric graphs. As the figure shows, multiscale gossip tends to distribute traffic almost uniformly in all the nodes. The observed color pattern is consistent with the hierarchical nature of the algorithm and although nodes that become representatives send more messages, no node is heavily used. On the other hand, for path averaging we observe that the nodes at the center (red region) send many more messages than nodes around the perimeter. This should be expected since geographic routing (which path averaging relies on) is greedy when trying to deliver messages to remote locations across the network.

TABLE I  
MEAN AND STANDARD DEVIATION OF NUMBER OF TRANSMISSIONS FOR DIFFERENT TYPES OF NODES RUNNING MULTISCALE GOSSIP ON A RANDOM GEOMETRIC GRAPH WITH 5000 NODES WITH 5 LEVELS OF HIERARCHY (SEE ALSO FIGURE 6).

Node type	Mean #msg	Std
Three times representatives	58.63	27.05
Two times representatives	31.23	21.11
One time representatives	18.77	15.21
Never representatives	8.65	10.23
All nodes	16.85	16.67

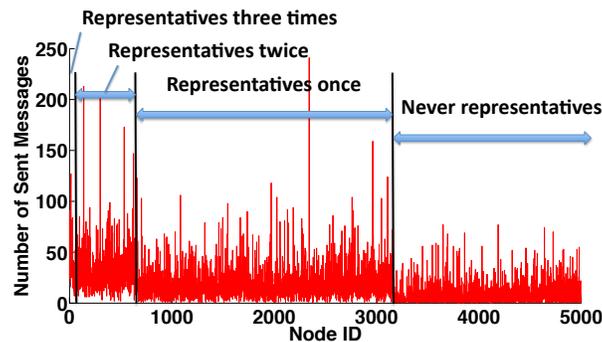


Fig. 6. Node utilization on a random geometric graphs with 5000 nodes and final desired accuracy  $\epsilon = 0.0001$ . For each node the total number of sent messages is shown. Nodes 1 to 30 have been representatives three times. Nodes 31 to 662 have been representatives two times. Nodes 663 to 3157 have been representatives once. The rest of the nodes never had a representative role. All nodes may have participated as intermediates in multi-hop communication.

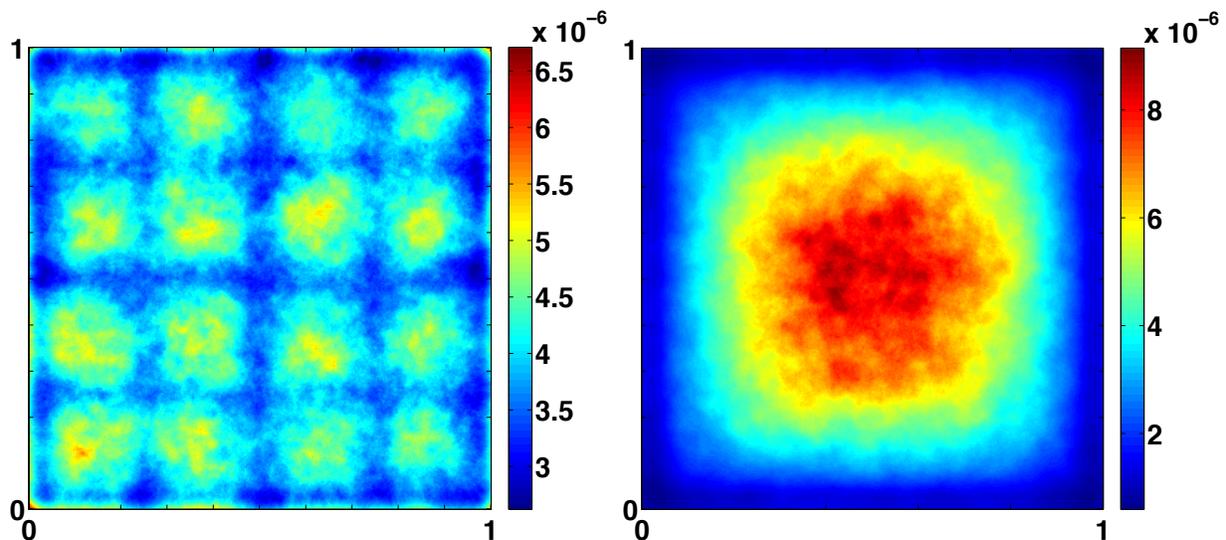


Fig. 7. Comparison of communication load between multiscale gossip (left) and path averaging (right). Color intensity is proportional to amount of transmissions through that part of the unit square. More blue (lower intensity) means little traffic while more red (higher intensity) means heavy traffic. Multiscale gossip distributes traffic relatively uniformly on all nodes while in path averaging the nodes at the center handle much more traffic than nodes around the perimeter. Results are averages over 100 graphs of 1000 nodes each.

### VIII. DISCUSSION AND FUTURE WORK

We have presented a new algorithm for distributed averaging exploiting hierarchical computation. Multiscale gossip separates the computation in linear phases and achieves very close to linear complexity overall. The key to achieving nearly-linear scaling lies in the way the recursive network partition is constructed. In particular, we argue that refining the network so that subnetworks at scale  $j$  contain  $O(n^{(2/3)^j})$  nodes provides optimal message-scaling laws with minimal number of levels in the hierarchy; although other hierarchical partitioning schemes could be constructed to achieve nearly-linear message complexity, they would require deeper hierarchies and, consequently, additional overhead for management. Our analysis focused on network topologies modeled as random geometric graphs, but it should be clear that the results translate directly to grids (2-dimensional lattices) embedded in the unit square.

Another feature of the proposed scheme is that the maximum distance any message has to travel is  $O(n^{\frac{1}{3}})$  hops, which is shorter than  $O(\sqrt{n/\log(n)})$  hops needed by path averaging [13], where each iteration involves averaging along a path of nodes potentially spanning the diameter of the network. Requiring transmissions over shorter distances is advantageous when transmitting over unreliable links that use acknowledgements and retransmission to ensure reliable communication at the link-layer, as is common practice in many existing systems. Intuitively, shorter paths translates directly to fewer retransmissions, and we illustrate this via simulation.

There is a number of interesting future directions that we see. In our present algorithm, gossip at higher levels happens on overlay grids which are known to require a number of messages which scales quadratically in the size of the grid. Since these grids already use multi-hop communication, it may be possible to further increase performance by devising other overlay graphs between representatives with better convergence properties, i.e. expander graphs [26]. Moreover, the subdivision of the unit square into a grid cell is not necessarily natural with respect to the topology of the graph, and one could use other methods for constructing hierarchical partitions which are tuned to the network topology. Our preliminary results with using hierarchical spectral clustering appear promising in simulation. It is, however, not clear how to carry out spectral clustering in a decentralized manner way and in linear number of messages. Another possibility is to combine the multiscale approach with the use of more memory at each node to get faster mixing rates. Notice however that how to use memory to provably accelerate asynchronous gossip is still an open question. Current results only consider synchronous algorithms [17]. Finally, an important advantage of gossip algorithms in general is their robustness. However, the general question of modeling and reacting to node failures has not been formally investigated in the literature. It would be very interesting to introduce failures and see the effect on performance for different gossip algorithms.

### APPENDIX A

#### COMPLEMENTARY DERIVATIONS FOR MULTI-SCALE GOSSIP ERROR BOUND

Looking at inequality 11, we have three terms that we need to bound.

- $T_1 = \frac{\epsilon^2 \|\{x_{1:L_1}\}\|^2}{\|x\|^2}$  : We can bound the numerator if we consider the following; At some level  $j$  we have  $x_{s_1 \dots s_j} = x_{s_1 \dots s_j c_{j+1}}^* \leq \max_{1 \leq i_{j+1} \leq L_{j+1}} (x_{s_1 \dots s_j i_{j+1}})$  for some  $1 \leq c_{j+1} \leq L_{j+1}$ . We can keep bounding all the way down to level  $k$  in the exact same fashion since  $x_{s_1 \dots s_j i_{j+1}} \leq x_{s_1 \dots s_j i_{j+1} c_{j+2}}^* \leq \max_{1 \leq i_{j+2} \leq L_{j+2}} (x_{s_1 \dots s_j i_{j+1} i_{j+2}})$  for  $1 \leq c_{j+2} \leq L_{j+2}$  and so on. If we do this for all terms appearing in the vector  $\{x_{1:L_1}\}$ , we get a numerator as a norm of elements from the initial  $x$ -vector. Since we need to divide expression by  $\|x\|$ , the ratio is less than one simply since the denominator has more terms. This means  $T_1 \leq \epsilon^2$ .
- $T_2 = \frac{m^{L_1} - x_{av}}{\|x\|}$  : Using definition (4) and the fact that an average such as  $x_{av}$  can be written as the average of averages,

$$\frac{m^{L_1} - x_{av}}{\|x\|} = \frac{1}{\|x\|^2} \left( \frac{\sum_{s_1=1}^{L_1} x_{s_1}}{L_1} - \frac{\sum_{s_1=1}^{L_1} \dots \sum_{s_k=1}^{L_k} x_{s_1 \dots s_k}}{L_1 L_2 \dots L_k} \right). \quad (36)$$

Pulling  $\frac{1}{L_1}$  and the summation over  $L_1$  out, adding and subtracting the  $L_2$  mean, taking the absolute value

and using triangular inequality

$$T_2 \leq \frac{1}{\|x\|^2} \frac{1}{L_1} \sum_{s_1=1}^{L_1} \left( |x_{s_1} - m_{s_1}^{L_2} + m_{s_1}^{L_2} - \frac{\sum_{s_2=1}^{L_2} \cdots \sum_{s_k=1}^{L_k} x_{s_1 \dots s_k}}{L_2 L_3 \cdots L_k}| \right) \quad (37)$$

$$\leq \frac{1}{\|x\|^2} \frac{1}{L_1} \sum_{s_1=1}^{L_1} \left( |x_{s_1} - m_{s_1}^{L_2}| + |m_{s_1}^{L_2} - \frac{\sum_{s_2=1}^{L_2} \cdots \sum_{s_k=1}^{L_k} x_{s_1 \dots s_k}}{L_2 L_3 \cdots L_k}| \right). \quad (38)$$

Using bound (5)

$$T_2 \leq \frac{1}{\|x\|^2} \frac{1}{L_1} \sum_{s_1=1}^{L_1} \left( \epsilon \|\{x_{s_1 1:L_2}\}\| + |m_{s_1}^{L_2} - \frac{\sum_{s_2=1}^{L_2} \cdots \sum_{s_k=1}^{L_k} x_{s_1 \dots s_k}}{L_2 L_3 \cdots L_k}| \right). \quad (39)$$

Using definition (4)

$$T_2 \leq \frac{1}{\|x\|^2} \frac{1}{L_1} \sum_{s_1=1}^{L_1} \left( \epsilon \|\{x_{s_1 1:L_2}\}\| + \left| \frac{\sum_{s_2=1}^{L_2} x_{s_1 s_2}}{L_2} - \frac{\sum_{s_2=1}^{L_2} \cdots \sum_{s_k=1}^{L_k} x_{s_1 \dots s_k}}{L_2 L_3 \cdots L_k} \right| \right). \quad (40)$$

Pulling the summation and the term  $\frac{1}{L_2}$  outside

$$T_2 \leq \frac{1}{\|x\|^2} \frac{1}{L_1} \sum_{s_1=1}^{L_1} \left( \epsilon \|\{x_{s_1 1:L_2}\}\| + \frac{1}{L_2} \sum_{s_2=1}^{L_2} \left| x_{s_1 s_2} - \frac{\sum_{s_3=1}^{L_3} \cdots \sum_{s_k=1}^{L_k} x_{s_1 \dots s_k}}{L_3 L_4 \cdots L_k} \right| \right). \quad (41)$$

We repeatedly pull out terms and add and subtract means to reach

$$T_2 \leq \frac{1}{\|x\|^2} \frac{1}{L_1} \sum_{s_1=1}^{L_1} \left( \epsilon \|\{x_{s_1 1:L_2}\}\| + \frac{1}{L_2} \sum_{s_2=1}^{L_2} (\epsilon \|\{x_{s_1 s_2 1:L_3}\}\| + \dots \right) \quad (42)$$

$$\dots + \frac{1}{L_{k-1}} \sum_{s_{k-1}=1}^{L_{k-1}} (\epsilon \|\{x_{s_1 \dots s_{k-1} 1:L_k}\}\|) + |m_{s_1 \dots s_{k-1}}^{L_k} - \frac{\sum_{s_k=1}^{L_k} x_{s_1 \dots s_k}}{L_k}| \dots \quad (43)$$

$$= \frac{1}{\|x\|^2} \frac{1}{L_1} \sum_{s_1=1}^{L_1} \left( \epsilon \|\{x_{s_1 1:L_2}\}\| + \frac{1}{L_2} \sum_{s_2=1}^{L_2} (\epsilon \|\{x_{s_1 s_2 1:L_3}\}\| + \dots \right) \quad (44)$$

$$\dots + \frac{1}{L_{k-1}} \sum_{s_{k-1}=1}^{L_{k-1}} (\epsilon \|\{x_{s_1 \dots s_{k-1} 1:L_k}\}\|) + 0 \dots \quad (45)$$

The comment from bounding  $T_1$  helps us here as well. Each term in the numerator can be bounded by terms from the initial vector and dividing by  $\|x\|$  we get

$$\frac{m^{L_1} - x_{av}}{\|x\|} \leq \frac{1}{L_1} \sum_{s_1=1}^{L_1} \left( \epsilon + \frac{1}{L_2} \sum_{s_2=1}^{L_2} \left( \epsilon + \dots + \frac{1}{L_{k-1}} \epsilon \right) \dots \right) = (k-1)\epsilon. \quad (46)$$

This is true since each summation cancels out with the corresponding denominator leaving only an  $\epsilon$  term and we have  $k-1$  such terms.

- $T_3 = \frac{\sum_{s_1=1}^{L_1} (x_{s_1}^* - m^{L_1})}{\|x\|}$  : Using bound (9b)

$$\frac{\sum_{s_1=1}^{L_1} (x_{s_1}^* - m^{L_1})}{\|x\|} \leq \frac{1}{\|x\|} \sum_{s_1=1}^{L_1} \epsilon \|\{x_{1:L_1}\}\| = \frac{1}{\|x\|} L_1 \epsilon \|\{x_{1:L_1}\}\|. \quad (47)$$

Adding and subtracting the vector of means at level  $L_2$

$$\frac{1}{\|x\|} L_1 \epsilon \|\{x_{1:L_1}\}\| = \frac{1}{\|x\|} L_1 \epsilon \|(x_1, \dots, x_{L_1}) - (m_1^{L_2}, \dots, m_{L_1}^{L_2}) + (m_1^{L_2}, \dots, m_{L_1}^{L_2})\|. \quad (48)$$

By triangle inequality

$$T_3 \leq \frac{1}{\|x\|} L_1 \epsilon (\|(x_1, \dots, x_{L_1}) - (m_1^{L_2}, \dots, m_{L_1}^{L_2})\| + \|(m_1^{L_2}, \dots, m_{L_1}^{L_2})\|). \quad (49)$$

Using the definition of a 2-norm

$$T_3 \leq \frac{1}{\|x\|} L_1 \epsilon \left( \sqrt{\sum_{s_1=1}^{L_1} (x_{s_1} - m_{s_1}^{L_2})^2} + \sqrt{\sum_{s_1=1}^{L_1} (m_{s_1}^{L_2})^2} \right). \quad (50)$$

Using bound (5) and definition (4)

$$T_3 \leq \frac{1}{\|x\|} L_1 \epsilon \left( \sqrt{\sum_{s_1=1}^{L_1} \epsilon^2 \|x_{s_1:1:L_2}\|^2} + \sqrt{\sum_{s_1=1}^{L_1} \left( \frac{\sum_{s_2=1}^{L_2} x_{s_1 s_2}}{L_2} \right)^2} \right). \quad (51)$$

Using the definition of a 2-norm

$$T_3 \leq \frac{1}{\|x\|} L_1 \epsilon \left( \epsilon \sqrt{\sum_{s_1=1}^{L_1} \sum_{s_2=1}^{L_2} x_{s_1 s_2}^2} + \frac{1}{L_2} \sqrt{\sum_{s_1=1}^{L_1} \left( \sum_{s_2=1}^{L_2} x_{s_1 s_2} \right)^2} \right) \quad (52)$$

Bringing the  $\|x\|$  inside using the definition of a 2-norm

$$T_3 \leq L_1 \epsilon \left( \epsilon \sqrt{\frac{\sum_{s_1=1}^{L_1} \sum_{s_2=1}^{L_2} x_{s_1 s_2}^2}{\sum_{i=1}^n x_i^2}} + \frac{1}{L_2} \sqrt{\frac{\sum_{s_1=1}^{L_1} \left( \sum_{s_2=1}^{L_2} x_{s_1 s_2} \right)^2}{\sum_{i=1}^n x_i^2}} \right). \quad (53)$$

At this point we can bound each element in the numerators using the same arguments as above in term  $T_4$ . For the first term we easily see that it is less or equal to one. For the second term we use the fact that for any  $n$  real numbers  $a_1, \dots, a_n$  we have  $\frac{(\sum_{i=1}^n a_i)^2}{\sum_{i=1}^n a_i^2} \leq n$  which is a direct consequence of the Cauchy-Schwartz inequality. In the second term, the numerator under the square root has  $L_1$  squares of sums of  $L_2$  values. The denominator has sums of  $n$  squares and can be bounded by sums of  $L_2$  squares. This gives rise to  $L_1$  ratios each of which is less than  $L_2$  according to the just mentioned inequality. As a result we get the bound

$$T_3 \leq L_1 \epsilon \left( \epsilon \sqrt{\frac{\sum_{s_1=1}^{L_1} \sum_{s_2=1}^{L_2} x_{s_1 s_2}^2}{\sum_{i=1}^n x_i^2}} + \frac{1}{L_2} \sqrt{\frac{\sum_{s_1=1}^{L_1} \left( \sum_{s_2=1}^{L_2} x_{s_1 s_2} \right)^2}{\sum_{i=1}^n x_i^2}} \right) \quad (54)$$

$$\leq L_1 \epsilon \left( \epsilon + \frac{1}{L_2} \sqrt{L_1 L_2} \right) = L_1 \epsilon^2 + \sqrt{\frac{L_1}{L_2}} L_1 \epsilon. \quad (55)$$

- $T_4 = \frac{L_1(m^{L_1} - x_{av})^2}{\|x\|^2}$  : After having a bound for term  $T_2$  it is easy to see that  $T_4 \leq L_1(k-1)^2 \epsilon^2$ .

#### ACKNOWLEDGMENT

This work is supported by grants from the *Fonds québécois de la recherche sur la nature et les technologies*, and the Mathematics of Information Technology and Complex Systems (MITACS) Canadian network of centres of excellence.

The authors would like to thank Marius Sucan for providing the multi-scale grid figure used in Section III.

#### REFERENCES

- [1] K. Tsianos and M. Rabbat, "Fast decentralized averaging via multi-scale gossip," in *Proc. Intl. Conf. on Distributed Computing in Sensor Systems*, Santa Barbara, CA, June 2010.
- [2] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Randomized gossip algorithms," *IEEE Trans. Inf. Theory*, vol. 52, no. 6, pp. 2508–2530, Jun. 2006.
- [3] D. Kempe, A. Dobra, and J. Gehrke, "Gossip-based computation of aggregate information," in *Proc. IEEE Foundations of Computer Science*, Cambridge, MA, Oct. 2003.
- [4] J. Tsitsiklis, D. Bertsekas, and M. Athans, "Distributed asynchronous deterministic and stochastic gradient optimization algorithms," *IEEE Trans. Automatic Control*, vol. AC-31, no. 9, pp. 803–812, Sep. 1986.

- [5] J. Tsitsiklis, "Problems in decentralized decision making and computation," Ph.D. dissertation, Massachusetts Institute of Tech., Nov. 1984.
- [6] A. Dimakis, S. Kar, J. Moura, M. Rabbat, and A. Scaglione, "Gossip algorithms for distributed signal processing," 2010, to appear in *Proceedings of the IEEE*.
- [7] P. Gupta and P. R. Kumar, "Critical power for asymptotic connectivity in wireless networks," in *Stochastic Analysis, Control, Optimization, and Applications*, Boston, 1998, pp. 1106–1110.
- [8] S. Mallat, *A Wavelet Tour of Signal Processing*. Academic Press, 1999.
- [9] A. Özgür, O. Lévêque, and D. Tse, "Hierarchical cooperation achieves optimal capacity scaling in *ad hoc* networks," *IEEE Trans. Inform. Theory*, vol. 53, no. 10, pp. 3549–3572, Oct. 2007.
- [10] M. Nagy, Z. Akos, D. Biro, and T. Vicsek, "Hierarchical group dynamics in pigeon flocks," *Nature*, vol. 464, no. 7290, pp. 890–893, Apr. 2010.
- [11] A. Dimakis, A. Sarwate, and M. Wainwright, "Geographic gossip: Efficient averaging for sensor networks," *IEEE Trans. Signal Processing*, vol. 56, no. 3, pp. 1205–1216, Mar. 2008.
- [12] R. Sarkar, X. Yin, J. Gao, F. Luo, and X. D. Gu, "Greedy routing with guaranteed delivery using ricci flows," in *Proc. Information Processing in Sensor Networks*, San Francisco, April 2009.
- [13] F. Benezit, A. Dimakis, P. Thiran, and M. Vetterli, "Gossip along the way: Order-optimal consensus through randomized path averaging," in *Proc. Allerton Conf. on Comm., Control, and Comp.*, Urbana-Champaign, IL, Sep. 2007.
- [14] M. Penrose, *Random Geometric Graphs*. Oxford University Press, 2003.
- [15] M. Cao, D. A. Spielman, and E. M. Yeh, "Accelerated gossip algorithms for distributed computation," in *Proc. 44th Annual Allerton Conf. Comm., Control, and Comp.*, Monticello, IL, Sep. 2006.
- [16] E. Kokiopoulou and P. Frossard, "Polynomial filtering for fast convergence in distributed consensus," *IEEE Trans. Signal Processing*, vol. 57, no. 1, pp. 342–354, Jan. 2009.
- [17] B. Oreshkin, M. Coates, and M. Rabbat, "Optimization and analysis of distributed averaging with short node memory," To appear *IEEE Trans. Signal Processing*, 2010.
- [18] W. Li and H. Dai, "Location-aided fast distributed consensus," in *IEEE Transactions on Information Theory*, submitted, 2008.
- [19] K. Jung, D. Shah, and J. Shin, "Fast gossip through lifted Markov chains," in *Proc. Allerton Conf. on Comm., Control, and Comp.*, Urbana-Champaign, IL, Sep. 2007.
- [20] —, "Distributed averaging via lifted Markov chains," Aug. 2009, to appear in *IEEE Trans. on Info. Theory*, and available online as arxiv:0908:4073v1.
- [21] R. Sarkar, X. Zhu, and J. Gao, "Hierarchical spatial gossip for multi-resolution representations in sensor networks," in *Proc. of the International Conference on Information Processing in Sensor Networks (IPSN'07)*, April 2007, pp. 420–429.
- [22] J.-H. Kim, M. West, S. Lall, E. Scholte, and A. Banaszuk, "Stochastic multiscale approaches to consensus problems," in *Proc. IEEE Conf. on Decision and Control*, Cancun, Dec. 2008.
- [23] M. Epstein, K. Lynch, K. Johansson, and R. Murray, "Using hierarchical decomposition to speed up average consensus," in *Proc. IFAC World Congress*, Seoul, Jul. 2008.
- [24] F. Cattivelli and A. Sayed, "Hierarchical diffusion algorithms for distributed estimation," in *Proc. IEEE Workshop on Statistical Signal Processing*, Wales, Aug. 2009.
- [25] C. Avin and G. Ercal, "On the cover time and mixing time of random geometric graphs," *Theoretical Computer Science*, vol. 380, pp. 2–22, June 2007.
- [26] G. Margulis, "Explicit group-theoretical constructions of combinatorial schemes and their application to the design of expanders and concentrators," *J. Probl. Inf. Transm.*, vol. 24, no. 1, pp. 39–46, 1988.