

Effect of Network Geometry and Interference on Consensus in Wireless Networks

Sundaram Vanka, Vijay Gupta, and Martin Haenggi

Abstract We study the convergence of the average consensus algorithm in wireless networks in the presence of interference. It is well-known that convergence of the consensus algorithm improves with network connectivity. However, from a networking standpoint, highly connected wireless networks may have lower throughput because of increased interference. This raises an interesting question: What is the effect of increased network connectivity on the convergence of the consensus algorithm, given that this connectivity comes at the cost of lower network throughput? We address this issue for two types of networks: regular lattices with periodic boundary conditions, and a hierarchical network where a backbone of nodes arranged as a regular lattice supports a collection of randomly placed nodes. We characterize the properties of an optimal TDMA protocol that maximizes the speed of convergence on these networks and provide analytical upper and lower bounds for the achievable convergence rate. Our results show that in a interference-limited scenario the fastest converging interconnection topology for the consensus algorithm crucially depends on the geometry of node placement. In particular, we prove that asymptotically in the number of nodes, forming long-range interconnections improves the convergence rate in one-dimensional tori, while it has the opposite effect in higher dimensions.

1 Introduction

Consensus has become an area of increasing research focus in recent years (e.g. see [1, 9, 10, 15, 19, 20, 2, 16] and the references therein). Given n nodes each with a scalar value and a possibly time-varying interconnection graph defined on these nodes, a consensus algorithm specifies the updating rule that every node should follow. The updated value of each node at every time step is a function of the value

The authors are with the Department of Electrical Engineering, University of Notre Dame, Notre Dame, IN 46556. (e-mail: {svanka, vgupta2, mhaenggi}@nd.edu)

held by itself and its neighbors at the previous time step. The conditions on interconnectivity graphs that permit convergence to a common value have been fairly well characterized. The focus has now shifted to analyzing the convergence properties in the face of communication constraints such as quantization [14], packet erasures [2, 6], additive channel noise [7, 8], and delays [13].

Such works typically assume that the communication channels between each pair of nodes are uncoupled. However, consensus algorithms are often employed over wireless networks, where channels are inherently coupled due to their broadcast nature and the presence of interference. Long range interconnections lead to a smaller graph diameter, but also to decreased spatial re-use. The effect of long-range interconnections on the rate of convergence of the consensus algorithm is thus, not clear. Moreover, in a wireless network, a communication graph cannot be considered to be given a priori, since any two nodes can communicate by spending enough energy. The communication topology in wireless networks thus depends on the network protocols and is, in fact, a design parameter.

In this work, we consider the rate of convergence of the average consensus algorithm while explicitly accounting for interference. We analyze the performance of scheduling algorithms that are optimal with respect to the rate of convergence. We also provide an analytical understanding of the impact of transmission power on the rate of convergence.

The paper is organized as follows. We begin in Section 2 by formulating the problem and introducing our notation. We concentrate on two geographical placements of the nodes: (i) nodes that are physically placed on a grid with periodic boundary conditions (considered in Section 3), and (ii) hierarchical networks with randomly placed sensor nodes and a regular communication backbone (considered in Section 4). Some avenues for future work are presented in Section 5.

2 Problem Formulation

Average Consensus Algorithm: Consider n nodes that aim to reach consensus with the final value being the average of their initial scalar values. Denote the value held by the i -th node at time k as $x_i(k)$. Also denote by $x(k)$ the n -dimensional vector obtained by stacking the values of all the nodes in a column vector. Let the nodes be connected according to a given interconnection topology at every iteration step. The topology can be described by a *consensus graph*, with an edge present between two nodes if and only if they can exchange information. Denote the neighbor set of node i at time k by $\mathcal{N}_i(k)$. An iteration consists of every node i exchanging its state variable $x_i(k)$ with all nodes in $\mathcal{N}_i(k)$, that is assumed to happen in a single packet transmission interval (also referred to as a *time slot* and normalized to 1). Then, the state of the system evolves as

$$x_i(k+1) = x_i(k) - h \sum_{j \in \mathcal{N}_i(k)} (x_i(k) - x_j(k)), \quad (1)$$

where h is a scalar constant designed to ensure convergence of the algorithm. In this case, from the perspective of the consensus algorithm, we say that the *iteration time* is one time slot.

Denote the interconnection graph at time k by $\mathcal{G}(k)$. The system thus evolves according to the discrete time equation

$$x(k+1) = (I - hL(k))x(k), \quad x(0) = x_0, \quad (2)$$

where $L(k)$ denotes the Laplacian matrix of the graph $\mathcal{G}(k)$. It can be shown (see, e.g., [15]) that under proper connectivity assumptions, if the parameter h is small enough, consensus is achieved with each node assuming the average value $x_{av} = \frac{1}{n} \sum_i x_i(0)$. Throughout our presentation, we will assume that h is fixed and has a value $h < \frac{1}{2d_{\max}}$ where d_{\max} is the maximum degree corresponding to any node in the consensus graph over all time. In other words, $d_{\max} = \max_{i,k} |\mathcal{N}_i(k)|$. To ensure that the nodes converge to the average of x_0 , it is also essential that the graph at every time step be balanced. The protocols we consider below will ensure that the graph is symmetric, which satisfies this condition.

The rate of convergence of the value of the nodes is a function of the graph topology. In the case of a static graph topology (i.e., $\mathcal{G}(k) = \mathcal{G}$ for all time k), it can be shown (see, e.g., [5, 15, 17]) that the convergence of the consensus protocol is geometric, with the rate being governed by the second largest eigenvalue modulus (SLEM) of the matrix $I - hL$. In general, a consensus algorithm on a graph with smaller SLEM converges more quickly.

However, in practice, a number of transmissions, each occupying a single time slot, may be necessary for this information exchange among these nodes to occur. This scenario is common in wireless networks, where concurrent transmissions in the same frequency band can interfere at a node, and hence may not be decodable. Therefore if each node can receive data from say at most one neighbor at any given time, the exchange of information necessary for iteration k will require at least $1 + \max_i |\mathcal{N}_i(k)|$. This idea is developed further in this paper.

Communication Protocols: We consider a situation in which the physical locations of the nodes are given in a d -dimensional space. Every node then decides on the power with which it transmits. This power determines the communication radius of the node in each axial direction according to the relation

$$P = P_0 r_c^\alpha,$$

where P_0 is a normalization constant, α is the path-loss exponent (typically $2 \leq \alpha \leq 5$), P is the transmission power and r_c is the communication radius along each axial direction. All nodes at a distance smaller than r_c from the transmitter can receive the transmitted message. Note that we assume that communication is possible only along axial directions. This is called a Manhattan connectivity model, and is suitable for a situation when only line of sight (LOS) communication is possible. In 2-dimensions this is a good model for urban environments where the presence of buildings inhibits most non-line of sight links. In an 1-dimensional node arrange-

ment, this model is identical to other connectivity models such as those based on communication disks centered at the nodes.

Similar to the communication radius, we can also define an interference radius r_i . A node at position x can receive a message successfully from a node at position y only if $\|y - x\| < r_c$, and there is no node at position z that is simultaneously transmitting, such that $\|z - x\| < r_i$ (interference constraint). The above equations should be interpreted in the framework of the Manhattan connectivity model, i.e., the distances should be measured only along the axial directions. In this paper, for simplicity, we assume $r_c = r_i$. The results can be generalized to other cases.

Given the above condition for successful transmission, we require a medium access control (MAC) protocol for the nodes. We focus on Time Division Multiple Access (TDMA)-based MAC protocols in this paper rather than random access protocols. These protocols assure successful communication by scheduling transmissions in time such that messages do not interfere. They demonstrate better throughput than collision-based MAC protocols, at the expense of greater synchronization and co-ordination requirements among the nodes [11, 21].

Problem Formulation The operation of the average consensus protocol can be divided into two phases that are repeated at every update of the node values. In the first phase, the nodes exchange their values through possibly multiple transmissions. We consider each transmission to consume one time slot. The effective communication graph at each update is composed of edges (i, j) such that node j has received the value of node i during the previous communication phase. In the second phase, the nodes update their values according to (1). As in the standard model, this step is assumed to be instantaneous. Therefore, due to multiple transmissions to set up the consensus graph, in our model, the state update does not occur at every time slot. In fact, assuming that each communication phase is completed in T time slots, the k^{th} update can be expressed as

$$x(kT + T) = (I - hL)x(kT). \quad (3)$$

Therefore the effect of finite communication time, possibly due to interference, is to slow down the convergence rate.

We are interested in the following problem: *Given a set of nodes at known locations, what is the effect of increasing transmit power on the convergence rate of the consensus algorithm when the channel-access mechanism accounts for interference?* In this context, we characterize the convergence of the consensus algorithm for the optimal MAC protocol that minimizes the number of time slots needed for communication in order to form a desired consensus graph \mathcal{G} (thus maximizing the update rate). We analyze this problem for two physical distributions of the nodes on a torus:

1. A regular grid of sensor nodes.
2. A regular grid of nodes that form a communication network for sensor nodes that are distributed as a binomial point process.

The transmit power at each node determines its neighbors in the consensus algorithm. The periodic boundary condition is chosen for analytical tractability, and our analysis becomes accurate as the number of nodes becomes large.

For analytical tractability, we make the following assumptions:

- We assume equal transmission power for all nodes.
- We limit the transmission policy to be time-invariant.
- At the time of an update of the values of the nodes, we require that the effective communication graph be undirected, i.e., for any two nodes i, j in the network, $j \in \mathcal{N}_i \Leftrightarrow i \in \mathcal{N}_j$. Note that this is slightly stronger than the necessary and sufficient condition for convergence of the average consensus algorithm that the graph be balanced [15].
- We do not assume explicit routing of values through nodes since the consensus algorithm itself incorporates implicit routing and in-network computation.
- We assume half-duplex operation, and further assume that packets that suffer collisions cannot be decoded.

Under these assumptions, the following are the main results of the paper:

1. We characterize the rate of convergence for the optimal MAC scheduling protocol for the average consensus algorithm for tori in n dimensions.
2. We show that network geometry plays a key role in identifying the optimum power allocation that maximizes the speed of convergence. In particular, while the convergence rate increases with the transmission power in 1-dimensional tori, the opposite is true in higher dimensions.
3. In hierarchical networks, we show that a positive fraction of nodes can always achieve consensus for certain scalings of backbone node density.

In the next section, we begin by studying the convergence properties of MAC protocols that maximize the speed of convergence for a given consensus graph \mathcal{G} . We begin by considering nodes placed on a regular grid with periodic boundary conditions.

3 Analysis of a Ring and a 2D Torus

3.1 The 1-D Case: Nodes on a Ring

Consider n nodes numbered $\{0, 1, \dots, n-1\}$ placed uniformly on a circle of radius r centered at the origin, as shown in Figure 1. Suppose that the transmission power is such that every node can transmit information to m of its nearest neighbors on either side. As an example, in Figure 1, $m = 1$. Define P_m , $m \leq \lfloor \frac{n}{2} \rfloor$ as the transmit power that provides a communication radius $r_c = 2r \sin\left(\frac{m\pi}{n}\right)$. Hence

$$P_m \propto \left(2r \sin\left(\frac{m\pi}{n}\right)\right)^\alpha, \quad (4)$$

where $\alpha \geq 2$ is the path-loss exponent. As stated above, for simplicity, we will assume that the interference radius $r_i = r_c$.

We note here that an alternative interpretation of this geometry of node placement is to consider the n nodes placed on a regular one-dimensional torus $[0, 1]$ (hereafter called a “1-torus” or $\mathcal{T}_1(n)$). This interpretation is useful in connecting these results with those for higher dimensional tori, that are discussed later. In this case, if choose the location of the first node as the origin, the position of the k^{th} node is given by $\frac{k}{n}$, $0 \leq k \leq n-1$, with a periodic boundary condition. In this geometry, the expression for P_m would appear as $P_m \propto \left(\frac{m}{n}\right)^\alpha$.




Fig. 1 Schematic of nodes placed along a ring.

If the wireless channel could support simultaneous transmissions by every node, the system would evolve according to (2), with $I - hL$ being an $n \times n$ circulant matrix with the first row given by

$$[1 - 2mh \mathbf{1}_m^* \ 0 \ 0 \ \dots \ 0 \ \mathbf{1}_m^*],$$

where

$$\mathbf{1}_m^* = [1 \ 1 \ \dots \ 1]_{1 \times m}.$$

For future reference, denote by $\mathcal{G}_{1,m}$, $L_{1,m}$ and $F_{1,m}$ the consensus graph, the Laplacian and the update matrix respectively for such a situation. Given the nodes placed on a ring, $\mathcal{G}_{1,m}$ is the consensus graph with the highest connectivity that can be formed for a given P_m , and therefore will have the fastest convergence. The MAC protocol that we propose guarantees that the system evolves according to this matrix. However, the communication phase occurs over multiple steps.

We begin by bounding the number of time slots required to form $\mathcal{G}_{1,m}$. Denote the smallest number of time slots used to form $\mathcal{G}_{1,m}$ by $T_1^*(m)$, or more compactly, as T_1^* . Also denote by $\mathcal{P}^*(m)$ the optimal TDMA protocol that forms the graph $\mathcal{G}_{1,m}$ in T_1^* number of steps. Observe that for all $m_1 \leq m_2$, $\mathcal{G}_{1,m_1} \subseteq \mathcal{G}_{1,m_2}$. Then $\mathcal{G}_{1,m}$ can always be formed in at most $T_1^*(m_1) \leq T_1^*(m_2)$ slots. This implies that

$m_1 \leq m_2 \implies T_1^*(m_1) \leq T_1^*(m_2)$. We say a *link* is formed from node v to node u whenever the message from v is successfully decoded at u . Since $\mathcal{G}_{1,m}$ is undirected and balanced, an edge $e \in \mathcal{E}_{1,m}$ connecting v and u is formed iff both v and u form links with each other. The following result bounds the length of the shortest TDMA schedule that forms the consensus graph $\mathcal{G}_{1,m}$ (and hence the smallest time T in the update equation (3)) and was proven in [18].

Lemma 1 (From [18]). *Consider the set-up described above, where the consensus graph $\mathcal{G}_{1,m}$ is to be formed in the smallest number of time slots. The optimal TDMA protocol forms $\mathcal{G}_{1,m}$ in the smallest possible number of time slots T_1^* where*

$$2m + 1 \leq T_1^*(m) \leq 4m + 1.$$

Using this result in conjunction with the spectral properties of $\mathcal{G}_{1,m}$ yields the following characterization of the fastest convergence rate that is possible for a given $\mathcal{G}_{1,m}$.

Theorem 1 (From [18]). *Consider the problem set-up described above. If the optimal TDMA protocol is used to construct $\mathcal{G}_{1,m}$ for each iteration, the error vector $\varepsilon(k) = x(k) - \mathbf{1}_n x_{av}$ converges geometrically to zero with the rate of decay β bounded as*

$$\rho_1^{\frac{1}{2m+1}} \leq \beta \leq \rho_1^{\frac{1}{4m+1}}$$

where

$$\begin{aligned} \rho_1 &= 1 - h(2m + 1) + hS_1^{(m,n)} \\ S_p^{(m,n)} &= \frac{\sin\left(\frac{(2m+1)\pi p}{n}\right)}{\sin\left(\frac{\pi p}{n}\right)}, \quad p = 0, 1, \dots, n-1. \end{aligned} \quad (5)$$

Remarks

1. For any given transmission power P_m , we see that the MAC constraints reduce the rate by a factor of T where $2m + 1 \leq T \leq 4m + 1$.
2. The speed of convergence is an increasing function in m , and hence in P_m . An illustration of this fact is provided in Figure 2. For the purpose of the plot, we show the time taken for the error norm to become half, termed the ‘‘half-value period’’, as a function of transmission power for 31 nodes arranged regularly on a ring of radius 1 unit. We have assumed $\alpha = 2$, and the constant of proportionality in (4) to be unity. For each P_m , we chose $h \propto \frac{1}{2m+1}$. The results are somewhat counter-intuitive since the rate reduction due to a larger number of steps in the communication phase is always compensated by the increase in rate due to higher connectivity. That forming long range communication links would lead to faster convergence even in networks with interference was not evident *a priori*.
3. The effect of increasing the transmission power are the most prominent at small P_m . This can again be seen from Figure 2. If $\theta = p\pi/n$ and $p \ll n$,

$$\sin \theta \approx \theta - \theta^3/3. \quad (6)$$

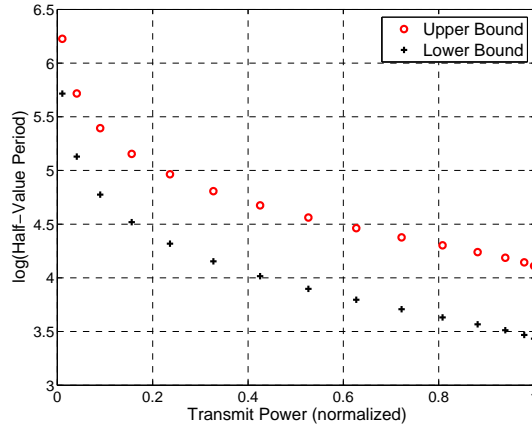


Fig. 2 Variation of the convergence rate with the transmission power for a ring of $n = 31$ nodes.

We use (6) to express the spectral gap $SG \triangleq 1 - \rho_1^{\frac{1}{T}}$ when $m \ll n$ as

$$SG = 1 - \left(1 - h(2m+1) + h \frac{\sin\left(\frac{(2m+1)\pi}{n}\right)}{\sin\left(\frac{\pi}{n}\right)} \right)^{\frac{1}{T}}$$

$$\approx \eta m h (m+1) (2m+1) T^{-1} n^{-2}.$$

where $\eta \triangleq \frac{4\pi^2}{3}$. Since $h \propto \frac{1}{2m+1}$ and $T = \Theta(m)$ for the optimal schedule, the spectral gap scales as $\frac{m}{n^2}$.

3.2 Nodes on a Two-Dimensional Torus

We now generalize our results to higher dimensional tori. Consider a set $\mathcal{T}_d(n)$ of $n = l^d$ regularly spaced points on a d -dimensional torus located at $[0, 1]^d$. An example when $d = 2$ is shown in Figure 3. Choose a node as the origin, and label each node using its displacements along each of the d axes (referred to as the d axial directions of the torus in this paper). For example, in Figure 3, $n = 9$ and the node $(1, 1)$ is located at $r_{11} \equiv r_{(1,1)} \triangleq \left(\frac{1}{3}, \frac{1}{3}\right)$. An alternative interpretation of a toroidal arrangement in the two-dimensional case is shown in Figure 4. Both these interpretations yield similar results in the limiting case of a large torus in which case local distances are not significantly affected by the curvature. We will focus on the former interpretation in this paper.

Suppose all nodes on a torus $\mathcal{T}_d(n)$ participate in an average consensus algorithm of the form (2) with a power allocation of P_m per node. The results for a disc connectivity model were provided in [18]. Here we present results for the Manhattan

connectivity model. Figure 5 describes the connectivity model for a set of transmitting nodes on a two-dimensional torus of $n = 25$ nodes and $m = 1$.

We now formally define the desired consensus graph $\mathcal{G}_{d,m} = (\mathcal{V}, \mathcal{E}_{d,m})$. The vertex set

$$\mathcal{V} = \{0, 1, \dots, l-1\}^d,$$

is the set of all points in $\mathcal{T}_2(n)$. We note that if each node can transmit at power P_m , the Manhattan connectivity model places edges between two distinct nodes v and u iff:

1. The locations of v and u must differ in at most one co-ordinate; and
2. At the co-ordinate where they differ, the absolute difference between the respective values can be at most m/l .

Thus the edge set $\mathcal{E}_{d,m}$ from the formed by connecting all pairs of distinct nodes whose positions satisfy these conditions.

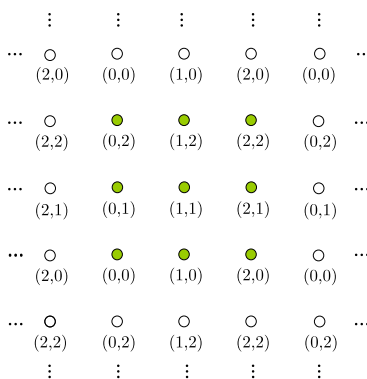


Fig. 3 The toroidal lattice $\mathcal{T}_2(9)$. Colored nodes indicate those physically placed in $[0, 1]^2$. Node $(0,0)$ represents the node at the origin, with each of the colored nodes (i, j) being placed at $(i/3, j/3)$. Nodes that left unfilled are the image nodes that arise due to the periodic boundary condition.

In keeping with the notation developed for the one-dimensional case, we will denote the Laplacian and the update matrix for $\mathcal{G}_{d,m}$ by $L_{d,m}$ and $F_{d,m} \triangleq I - hL_{d,m}$ respectively. Assuming as before that each transmission occupies one time slot, we now study the convergence properties of the optimal MAC protocol that will form $\mathcal{G}_{d,m}$ in the smallest number of time slots. In this paper, we set $d = 2$; the results can be generalized to tori of higher dimensions.

Using similar arguments as in the one dimensional case, we can show that the optimal TDMA protocol assigns an equal power P_m to each chosen transmitting node i in any time slot. Denote by $T_2^*(m)$, or more compactly by T_2^* the number of time slots required by an optimal schedule to construct $\mathcal{G}_{2,m}$. The optimal MAC schedule places the *maximum* number $N_2^*(t)$ of non-interfering transmitters on the torus in every time slot $t = 1, 2, \dots, T_2^*$.

Note that increased network dimensionality plays an important role in finding N_2^* (and consequently, T_2^*). This makes the problem of analytically finding T_2^* non-trivial. This effect is illustrated in Figure 5 for a 2-torus of $n = 25$ nodes and $m = 1$. The transmitters are chosen from the entire two-dimensional lattice. With power P_1 , each node can reach its 4 nearest neighbors as shown. Since each of these nodes must transmit at least once, at least $4 + 1 = 5$ slots are necessary to form the consensus graph $\mathcal{G}_{2,1}$. Figure 5 shows the optimal transmitting set in the first time slot.

As before, we begin by characterizing the length T_2^* of the shortest TDMA schedule that constructs $\mathcal{G}_{2,m}$. Thereafter, we exploit the properties of the consensus algorithm along with optimality of the MAC protocol and constraints imposed by our problem to bound the convergence rates for a 2-torus.




Fig. 4 Schematic of nodes placed along a 2-dimensional torus. The periodic square grid being considered can be considered a limiting case of a large torus, so that the effect of its curvature on small distances is not important.

As before, define P_m to be the transmit power that enables a node to form error-free links with m neighbors in the axial directions. Given that there are \sqrt{n} nodes in either of the axial directions,

$$P_m \propto \left(\frac{m}{\sqrt{n}} \right)^\alpha,$$

where $\alpha \geq 2$ is the path-loss exponent.

We now define the update matrix $F_{2,m}$. Define circulant matrix Q_m to be a circulant matrix with the first row $[1 - 4mh \ h\mathbf{1}_m^* \ 0 \cdots 0 \ h\mathbf{1}_m^*]_{1 \times l}$. Further define $R_m \triangleq h\mathbf{1}_m^* \otimes I = h[I \ I \ \cdots \ I]_{l \times lm}$

If each node uses power P_m , the Manhattan connectivity model results in a block circulant update matrix $F_{2,m}$ with its first row being

$$[Q_m \ R_m \ 0 \ \cdots \ 0 \ R_m]_{l \times n}. \quad (7)$$

We now bound the number of time slots required to form $\mathcal{G}_{2,m}$.




Fig. 5 Effect of network geometry on choosing the transmitter set, shown for $\mathcal{F}_2(25)$ and $m = 1$. The transmitters are shown in red (e.g., (0,0)). Note that the dimensionality is exploited to allow more concurrent transmissions. The nodes in light-blue (e.g., (4,0)) are covered by transmitters via their images (shown unfilled).

Lemma 2. *If each node transmits at power P_m and the optimal schedule over the 2-torus constructs $\mathcal{G}_{2,m}$ in T_2^* time slots, for $1 < m < \lfloor l/2 \rfloor$, T_2^* always satisfies*

$$T_l \leq T_2^* \leq T_u$$

where

$$T_l = m^2 + 2m + 2$$

and

$$T_u = 16m^2 + 8m + 1.$$

Proof. Using similar arguments as in Lemma 1 it is easy to show that power allocation for any node should be at least P_m . Define a *feasible* TDMA schedule for a power allocation P_m per node as one that which constructs $\mathcal{G}_{2,m}$ while satisfying the half-duplex and interference constraints described in Section 2.

Without loss of generality, suppose node (0,0) transmits in the first time slot with power P_m . From the definition of a feasible schedule, during this time there cannot exist a transmitter inside the square with vertices (1,1), (1,m), (m,m), (m,1). As a result, each node that is contained in the square requires time slot each. It is easy to see that there are $(m-1)^2$ such nodes. Moreover, any of the $4m$ nodes that are currently receiving a message from (0,0) cannot transmit at this time. Therefore, accounting for the current slot, the length of a feasible schedule cannot be shorter than $T_l = (m-1)^2 + 4m + 1 = m^2 + 2m + 2$ slots. In particular, since T_2^* is the length of the shortest feasible schedule, $T_2^* \geq T_l$.

If T_u is number of time slots taken by any feasible schedule to form $\mathcal{G}_{2,m}$, $T_2^* \leq T_u$. Consider the following schedule: in the first time slot, choose $(0,0)$ as a transmitter, and schedule nodes $(0 + p(2m+1), 0 + q(2m+1))$ for $p, q = 1, \dots, \lfloor \frac{l}{2m+1} \rfloor$ to transmit. In other words, we attempt to tile the torus with squares of side $\frac{2m+1}{\sqrt{n}}$. Clearly this is feasible, since each node receives from at most one transmitter. In each subsequent time slot, repeat this process by choosing some other node (i, j) inside the square of side $(2m+1)/\sqrt{n}$ centered at the origin and schedule nodes $(i + p(2m+1), j + q(2m+1))$ for transmission. Repeat this process until all the $(2m+1)^2$ nodes in this square have been chosen once. Using arguments similar to those used in Lemma 1, a maximum of $\lfloor \frac{l}{2m+1} \rfloor^2$ simultaneous transmissions can be scheduled per time slot. After $(2m+1)^2$ time slots,

$$\begin{aligned} & n - \lfloor \frac{l}{2m+1} \rfloor^2 (2m+1)^2 \\ &= 2 \lfloor \frac{l}{2m+1} \rfloor (2m+1) \text{rem}(l, (2m+1)) \\ & \quad + (\text{rem}(l, 2m+1))^2 \end{aligned}$$

nodes are yet to transmit. In the first term, we can schedule $\lfloor l/(2m+1) \rfloor$ nodes in each of the $2\text{rem}(l, 2m+1) \leq 4m$ ‘‘rows’’, that require at most $4m(2m+1)$ additional time slots. Scheduling one node per time slot, all the remaining $(\text{rem}(l, 2m+1))^2$ nodes can transmit in at most $4m^2$ time slots. Therefore the schedule constructs $\mathcal{G}_{2,m}$ in $(2m+1)^2 + 4m(2m+1) + 4m^2 = 16m^2 + 8m + 1$ time slots. Thus we conclude that $T_2^* \leq 16m^2 + 8m + 1$. \square

As compared to the 1-torus, the optimal schedule for a 2-torus is bounded by two quadratic terms. This arises due to imposing interference constraints to the given geometry of node placement. As we shall see, this quadratic - rather than linear - dependence on m is the key to understanding the effect of transmit power on convergence behavior.

3.2.1 Bounding the Rate of Convergence

We now find the eigenvalues of $F_{2,m}$ by exploiting its block circulant property.

Lemma 3. *Let $\mathcal{G}_{2,m}$ be the consensus graph formed over $\mathcal{T}_2(n)$ using the Manhattan connectivity model. If $L_{2,m}$ is its Laplacian then the eigenvalues of the $F_{2,m} = I - hL_{2,m}$ are*

$$\lambda_{a,b} = 1 - 2h(2m+1) + hS_a^{(m,l)} + hS_b^{(m,l)},$$

where, as defined above $S_a^{(m,l)} = \frac{\sin(\frac{(2m+1)\pi a}{l})}{\sin(\frac{\pi a}{l})}$, $a = 0, 1, \dots, l-1$.

Proof. From (7) $F_{2,m}$ is an $n \times n$ block circulant matrix with its first block row being

$$[Q_m \ R_m \ 0 \ 0 \ \dots \ 0 \ R_m]_{l \times n},$$

where $R_m = h\mathbf{1}_m^* \otimes I$. As noted earlier, Q_m is circulant. Since the identity and the all-zero matrices are also circulant, all these matrices share the same eigenvectors. Using this in conjunction with the properties of block circulant matrices we compute the eigenvalues $\mu_{r,s}$ of $F_{2,m}$ as

$$\mu_{r,s} = \sum_{t=0}^{l-1} \eta_{r,t} e^{-j\frac{2\pi st}{l}} \quad (8)$$

where $\eta_{r,t}$ is the r^{th} eigenvalue of $Q_m \forall r, s = 0, 1, \dots, l-1$. Using the eigenvalues for the 1-torus and simplifying yields the eigenvalues of $F_{2,m} = I - hL_{2,m}$ are

$$\lambda_{a,b} = 1 - 2h(2m+1) + hS_a^{(m,\sqrt{n})} + hS_b^{(m,\sqrt{n})} \quad (9)$$

which is the desired result. \square

We are now in a position to bound the rate of decay for the case of the torus. This is presented in the form of a theorem below.

Theorem 2. Consider a consensus algorithm of the form (2) on $\mathcal{G}_{2,m}$. If each node transmits at P_m for $1 \leq m < \lfloor \frac{\sqrt{n}}{2} \rfloor$, the rate of convergence β of an optimal MAC schedule on $\mathcal{G}_{2,m}$ that drives $\delta(k) = x(k) - \mathbf{1}_n x_{av}$ to zero is bounded as

$$\lambda_1^{\frac{1}{m^2+2m+2}} < \beta < \lambda_1^{\frac{1}{16m^2+8m+1}}$$

where

$$\lambda_1 = \left(1 - h(2m+1) + h(2m+1)S_1^{(m,\sqrt{n})}\right)$$

Proof. From Lemma 3 above, the eigenvalues of the update matrix are

$$\lambda_{a,b} = 1 - 2h(2m+1) + hS_a^{(m,\sqrt{n})} + hS_b^{(m,\sqrt{n})}$$

The largest eigenvalue is obtained by maximizing both the sinc terms, i.e., by choosing $a=0, b=0$. The second largest eigenvalue is doubly degenerate, and is obtained for $a=0, b=1$ or $b=0, a=1$. Any of these choices will simplify (9) to

$$\lambda_{0,1} = \lambda_{0,1} = 1 - h(2m+1) + hS_a^{(m,\sqrt{n})} = \lambda_1.$$

From Lemma 1 we know that the length T_2^* of the optimal schedule length is bounded as $T_l \leq T_2^* \leq T_u$, where $T_l = m^2 + 2m + 2$ and $T_u = 16m^2 + 8m + 1$.

The rate of convergence of $\mathcal{G}_{2,m}$ with the optimal schedule is $\lambda_1^{1/T_2^*} \leq \lambda_1^{1/T_u}$ since $T_2^* \geq T_u$. Similarly, $\lambda_1^{1/T_l} \leq \lambda_1^{1/T_2^*}$. Hence the result follows. \square

To understand the effect of higher transmit power on the convergence rate in 2-tori, first simplify the expressions for λ_1 in Theorem 2 using $h = \gamma/(4m+1)$, $0 < \gamma < 1$:

$$\lambda_1 = 1 - \gamma \frac{(2m+1)}{(4m+1)} + \gamma \frac{\sin((2m+1)\pi/\sqrt{n})}{(4m+1)\sin(\pi/\sqrt{n})}$$

Comparing this to the 1-torus case with $h = \gamma/(2m+1)$,

$$\rho_1 = 1 - \gamma + \gamma \frac{\sin((2m+1)\pi/n)}{(2m+1)\sin(\pi/n)}.$$

Clearly λ_1 is of similar form as ρ_1 in Theorem 1, except for the \sqrt{n} .

However, there is a significant difference between the two cases in the length of the optimal schedule. This length was shown to be $\Theta(m)$ for the 1-torus, and $\Theta(m^2)$ for the 2-torus. This means that the effect of interference depends on the geometry of node placement. High transmit powers reduce available network throughput by causing more interference. In TDMA-scheduling MAC protocols, this effect is reflected in longer schedules. For a 1-torus, this is still offset by the resultant long-range connections. However, this is no longer true for higher dimensions where the convergence rate *worsens* with increasing transmit power.

3.2.2 Tori in Arbitrary Dimensions

The results in Lemma 2 can be extended to higher dimensional grids with toroidal boundary conditions. Indeed using similar arguments for power P_m the optimal schedule for a d -dimensional torus cannot be shorter than $(m-1)^d + 2md + 1$. Similarly, to find the upper bound one can generalize the schedule described in Lemma 1 that was used to find an upper bound. It can be shown that the optimal schedule will be $\Theta(m^d)$.

The results in Lemma 3 can also be generalized to d -dimensions as $\lambda_1 = 1 - h(2m+1) + hS_1^{(m,d)}$. Thus the convergence rate increases with transmit power in geometries having dimension 2 or more.

4 Hierarchical networks

Since calculation of the rates of convergence for average consensus for arbitrary graphs is not possible even without the interference constraints, we do not expect to be able to extend our results for arbitrary graphs. In this section, we consider another useful class of graphs that allow us to state analytical results. We consider a variation on the random geometric graphs by adding a backbone of dedicated (long-distance) communication nodes. Thus we consider a hierarchical network with N sensing nodes uniformly randomly placed on a torus in $[0, 1]^d$, and K^d identical regularly spaced *backbone* nodes on the torus, as shown in Figure 6 for $d = 1$.

We assume that the backbone nodes do not participate in sensing, and only communicate with each other in the network. Each backbone node has a fixed exclusive region of coverage, i.e., it (alone) collects data from all the sensing nodes within a




Fig. 6 Schematic of backbone nodes placed along a 1-dimensional torus. The i^{th} backbone node is placed at $r_i = (i - \frac{1}{2}) \frac{1}{K}$ from the origin, for $i = 1, 2, \dots, K$.

sphere of radius $r = 1/2K$. Initially, each backbone node collects and averages the data from all the sensing nodes in its region of coverage. All the backbone nodes now run an average consensus algorithm among themselves with their respective averaged data as initial values. We assume the Manhattan connectivity model. It is possible to pass on this (global) average back to all the nodes within their coverage regions in $O(1)$ steps. Therefore for analyzing the rate of convergence, it is sufficient to analyze the time taken for collecting data by the backbone nodes and the time taken to reach consensus among these nodes.

We begin by characterizing the number of sensors reporting to each backbone node.

Lemma 4. For large N and if the number of backbone nodes scales as $o(\frac{N}{\log N})$ w.h.p. the number of sensors per coverage area is $n = \frac{N\pi^{d/2}}{\Gamma(1+d/2)(2K)^d}$, where K is the number of nodes per dimension.

Proof. This can be proved by a variation of the argument used in the theory of random geometric graphs to show regularity. Consider a sphere \mathcal{S} of radius r centered at point P on the torus. Marking the points $1, 2, \dots, N$, we can associate with each point k , a random variable X_k ($1 \leq k \leq N$) be defined as:

$$X_k = \mathbf{1}_{\mathcal{S}}(k)$$

where $\mathbf{1}(\cdot)$ is the indicator function. That is, $\mathbf{1}_{\mathcal{S}}(k) = 1$ if $k \in \mathcal{S}$, and 0 otherwise. Since the sensing nodes are placed on the torus uniformly and independently of each other, $\{X_k\}$ are iid Bernoulli with success probability

$$p = \frac{\text{Vol. of sphere}}{\text{Vol. of torus}} = \frac{\pi^{d/2}}{\Gamma(1+d/2)(2K)^d}.$$

The number of sensors inside the sphere is thus a binomial random variable

$$\sum_{k=1}^N X_k \triangleq Y,$$

with $\mu \triangleq \mathbb{E}Y = Np$. We can now use the Chernoff bound:

$$\mathbb{P}(|Y - \mu| > \mu\delta) \leq 2 \exp\left(-\frac{\mu\delta^2}{2}\right).$$

For $0 < \delta < 1$, since the number of nodes $K^d = o\left(\frac{N}{\log N}\right)$ we can always choose

$$\delta = \sqrt{\frac{4 \log N}{N} \cdot \gamma K^d} = o(1), N \rightarrow \infty$$

for some $\gamma > 0$. Plugging in this value into the bound, we see that

$$\mathbb{P}(Y \notin [\mu(1 - \delta), \mu(1 + \delta)]) \leq \frac{1}{N^2 \eta \gamma} \leq \frac{1}{N^3},$$

where $\eta = \frac{\pi^{d/2}}{2^d \Gamma(1 + d/2)}$ and γ is chosen such that $\eta \gamma > 2$. So w.h.p.,

$$Y = Np(1 \pm o(1)).$$

Thus, the probability that any coverage area does not have $n(1 \pm o(1))$ nodes is

$$\begin{aligned} \mathbb{P}(\cup_k \{Y_k \notin [\mu(1 - \delta), \mu(1 + \delta)]\}) &\leq \cup_k \mathbb{P}(Y_k \notin [\mu(1 - \delta), \mu(1 + \delta)]) \\ &= K^d \frac{1}{N^3} < N \frac{1}{N^3} = \frac{1}{N^2} \end{aligned}$$

where we have used the union bound. The result now follows readily. \square

Since each of the K^d backbone nodes has n sensors in its coverage region w.h.p. when N is large, a total of nK^d sensing nodes will be covered by the backbone network. Therefore for large N , it is always possible to achieve consensus over a positive fraction

$$\kappa = \frac{nK^d}{N} = \frac{\pi^{d/2}}{2^d \Gamma(1 + d/2)}$$

of the sensing nodes, independent of N and K .

We will now study specific cases for $d = 1$ and 2 . To begin, we note that for $d = 1$,

$$\kappa_1 = \frac{\pi^{1/2}}{2\Gamma(3/2)} = 1,$$

and for $d = 2$,

$$\kappa_2 = \frac{\pi}{4\Gamma(2)} = \frac{\pi}{4} \approx 78.5\%.$$

As stated above, we assume that each backbone node collects data by polling all the sensing nodes in its coverage region. This is done in parallel over all backbone nodes, by a suitable choice of transmit power. Assuming each node transmits in one time slot, we need n time slots to initialize the consensus algorithm, where w.h.p.

$$n = \frac{\alpha N}{K^d}.$$

During the consensus phase, the network topology is identical to the (regular) ring and torus topologies discussed previously. As before, let each node transmit at fixed power P_m to reach $m \leq \lfloor K/2 \rfloor$ neighbors per direction per dimension. Using the results in Theorems 1 and 2, we obtain that for $K^d = o\left(\frac{N}{\log N}\right)$, w.h.p.,

$$\left(1 - h(2m + 1) + hS_1^{(m,K)}\right)^{\frac{1}{4m+1}} \geq \beta \geq \left(1 - h(2m + 1) + hS_1^{(m,K)}\right)^{\frac{1}{2m+1}}$$

for the ring or 1D torus and

$$\left(1 - h(2m + 1) + hS_1^{(m,K)}\right)^{\frac{1}{16m^2+8m+1}} > \beta > \left(1 - h(2m + 1) + hS_1^{(m,K)}\right)^{\frac{1}{m^2+2m+2}}$$

for the 2D torus.

5 Conclusions

We introduced a framework that considers the effects of realistic communication constraints on average consensus algorithms. In particular, we analytically characterize the performance of the medium access control algorithm that maximizes the speed of convergence. We study the effect of transmit power on convergence in the presence of interference. In interference-limited wireless networks, the geometry of node placement plays a key role in deciding the fastest converging consensus graph. While forming long-range links (using more power) always improves the convergence on ring topologies, it is not so for higher-dimensional tori.

This work could be extended to other classes of graphs, like Cayley graphs and expander graphs that have good convergence properties [3]. Another issue is the effect of stochastic data loss through effects due to fading and interference, using a different framework as compared to [2, 6], to explicitly account for interference.

Acknowledgements The partial support of DTRA grant N00164-07-8510 and NSF grant 0846631 is gratefully acknowledged.

References

1. V. Blondel, J. Hendrickx, A. Olshevsky, and J. Tsitsiklis. Convergence in multiagent coordination, consensus and flocking. *Proceedings of the 44th IEEE Conference on Decision and Control*, pages 2996–3000, 2005.
2. S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 52(6):2508–2530, 2005.

3. R. Carli, F. Fagnani, A. Speranzon and S. Zampieri, "Communication constraints in the average consensus problem," *Automatica*, 44(3):671-684, March 2008.
4. G. Cybenko. Dynamic load balancing for distributed memory multiprocessors. *Journal on parallel and distributed computing*, (7):279-301, 1989.
5. M.P. Desai, and V.B. Rao. A new eigenvalue bound for reversible Markov chains with applications to the temperature-asymptotics of simulated annealing. *Proceedings of the IEEE International Symposium on Circuits and Systems (2):1211-1214*, 1990.
6. P. Hovareshti, V. Gupta and J. S. Bars, "Average Consensus over Small World Networks: A Probabilistic Framework," IEEE Conference on Decision and Control (CDC 08) 2008, Submitted.
7. M. Huang and J.H. Manton, "Stochastic double array analysis and convergence of consensus algorithms with noisy measurements," Proc. American Control Conference, New York, pp. 705-710, July 2007.
8. M. Huang and J.H. Manton. "Stochastic Lyapunov analysis for consensus algorithms with noisy measurements," Proc. American Control Conference, New York, pp. 1419-1424, July 2007.
9. L. Fang and P. Antsaklis. On communication requirements for multi-agents consensus seeking. *Proceedings of Workshop NESC05: University of Notre Dame, Lecture Notes in Control and Information Sciences (LNCIS)*, Springer, 331:53-68, 2006.
10. A. Jadbabaie, J. Lin, and A. S. Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control*, 48(6):988-1001, 2003.
11. X. Liu and M. Haenggi. Throughput Analysis of Fading Sensor Networks with Regular and Random Topologies, *EURASIP Journal on Wireless Communications and Networking*, vol. 4, pp. 554-564, Aug. 2005. Special Issue on Wireless Sensor Networks.
12. B. Mohar. The Laplacian Spectrum of Graphs. "Graph Theory, Combinatorics, and Applications", Vol. 2, Ed. Y. Alavi, G. Chartrand, O. R. Oellermann, A. J. Schwenk, Wiley, 1991, pp. 871-898.
13. A. Nedich and A. Ozdaglar, "Convergence Rate for Consensus with Delays," LIDS Technical Report 2774, MIT, Lab. for Information and Decision Systems.
14. A. Nedich, A. Olshevsky, A. Ozdaglar, and J. N. Tsitsiklis, "On Distributed Averaging Algorithms and Quantization Effects," LIDS Technical Report 2778, MIT, Lab. for Information and Decision Systems.
15. R. Olfati Saber and R.M Murray. Consensus problems in networks of agents with switching topology and time-delays. *IEEE Transactions on Automatic Control*, 49(9):1520-1533, 2004.
16. W. Ren, R. W. Beard, and T. W. McLain. Coordination variables and consensus building in multiple vehicle systems. "*Proceedings of the Block Island Workshop on Cooperative Control, Lecture Notes in Control and Information Sciences series, Springer-Verlag*", 309:171-188, 2004.
17. E. Seneta. *Nonnegative Matrices and Markov Chains*. Springer, 2nd edition, 1981.
18. S. Vanka, V. Gupta and M. Haenggi, "Power-Delay Analysis of Consensus Algorithms on Wireless Networks with Interference," International Journal of Systems, Control and Communications (Special issue on Information Processing and Decision Making in Distributed Control), To Appear 2009.
19. L. Xiao and S. Boyd. Fast linear iterations for distributed averaging. *Proceedings of 42th IEEE Conference on Decision and Control*, pages 4997-5002, 2003.
20. L. Xiao, S. Boyd, and S. Lall. A scheme for robust distributed sensor fusion based on average consensus. *Proceedings of International Conference on Information Processing in Sensor Networks*, pages 63-70, 2005.
21. M. Xie and M. Haenggi. Delay Performance of Different MAC Schemes for Multihop Wireless Networks. *IEEE Global Communications Conference (GLOBECOM'05)*, (St. Louis, MO), Nov. 2005.