Power-delay analysis of consensus algorithms on wireless networks with interference

S. Vanka*, V. Gupta and M. Haenggi

Department of Electrical Engineering,
University of Notre Dame,
Notre Dame, Indiana 46556, USA
E-mail: svanka@nd.edu
E-mail: vgupta2@nd.edu
E-mail: mhaenggi@nd.edu
*Corresponding author

Abstract: We study the convergence of the average consensus algorithm in wireless networks in the presence of interference. For regular lattices with periodic boundary conditions, we characterise the convergence properties of an optimal Time Division Multiple Access (TDMA) protocol that maximises the speed of convergence on these networks. We provide analytical upper and lower bounds for the convergence rate for these networks. Our results show that in an 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, increasing the transmit power to allow long-range interconnections improves the convergence rate in one-dimensional tori, while it has the opposite effect in higher dimensions.

Keywords: consensus algorithms; interference; MAC protocol; wireless networks.

Reference to this paper should be made as follows: Vanka, S., Gupta, V. and Haenggi, M. (2010) 'Power-delay analysis of consensus algorithms on wireless networks with interference', *Int. J. Systems, Control and Communications*, Vol. 2, Nos. 1/2/3, pp.256–274.

Biographical notes: Sundaram Vanka received his BTech and MTech Degrees in Electrical Engineering from the Indian Institute of Technology, Madras in 2003. After a stint in industry, he is currently working towards his PhD Degree in the Department of Electrical Engineering at the University of Notre Dame. His research interests include wireless communications and networking.

Vijay Gupta is an Assistant Professor in the Department of Electrical Engineering at the University of Notre Dame. He received his BTech Degree from the Indian Institute of Technology, Delhi and the MS and PhD Degrees from the California Institute of Technology, all in Electrical Engineering. His research interests include networked control systems and sensor networks, distributed estimation and detection, usage-based value of information and, in general, the interaction of communication, computation and control.

Martin Haenggi received his PhD in 1999 from the Swiss Federal Institute of Technology in Zurich, Switzerland (ETH). Currently, he is an Associate Professor in the Department of Electrical Engineering at the University of Notre Dame. His research interests include wireless communications and networking, in particular ad hoc and sensor networks. He serves as an Associate Editor for the *IEEE Transactions on Mobile Computing*, the *ACM Transactions on Sensor Networks*, and as a Guest Editor for the *IEEE Journal on Selected Areas in Communications.* He has published a book, four book chapters, and more than 90 research papers. He received the ETH Medal in 1999 and the CAREER award from the US National Science Foundation in 2005.

1 Introduction

Consensus in general, and average consensus in particular, has become an area of increasing research focus in recent years (e.g., see Blondel et al., 2005; Fang and Antsaklis, 2005; Jadbabaie et al., 2003; Olfati-Saber and Murray, 2004 and the references therein). Many applications, including distributed estimation (Boyd et al., 2006; Xiao and Boyd, 2003; Xiao et al., 2005), motion coordination (Ren et al., 2004) and load balancing of multiple processes (Cybenko, 1989), have been analysed in this framework.

Given n nodes each with a scalar value and a possibly time-varying interconnection graph defined on these nodes, consensus algorithms specify updating rules that every node should follow. The updated value of each node depends on the value held by itself and its neighbours at the previous time step. Initial results focused on the connectivity constraints of interconnected graphs that ensure consensus. Lately, attention has shifted to analysing the convergence with constraints imposed by communication channels between the nodes. Thus, effects such as quantisation (Nedich et al., 2007), packet erasures (Boyd et al., 2006; Hovareshti et al., 2008), additive channel noise (Huang and Manton, 2007a, 2007b), and delays (Nedich and Ozdaglar, 2007) have begun to gain attention.

Such works typically assume that the communication channels between each pair of nodes are uncoupled. However, nodes using consensus algorithms typically communicate over wireless channels that are inherently broadcast. Moreover, in wireless networks, any two nodes can communicate by spending enough energy or by lowering transmission rate. The communication topology in wireless networks thus depends on the network protocols and is, in fact, a design parameter. Higher transmission power results in a smaller graph diameter, but also to reduced network throughput due to interference. The effect of long-range interconnections on the rate of convergence of the consensus algorithm is thus not clear. We take the first steps towards analysing the effect of such communication constraints on consensus algorithms and designing the communication parameters for the consensus problem. In particular, we consider the rate of convergence of the average consensus algorithm while explicitly accounting for interference. We analyse 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 organised as follows. We begin by formulating the problem and introducing our notation. We concentrate on networks of nodes that are physically placed on a grid with periodic boundary conditions (Section 3). Some avenues for future work are presented in Section 4.

2 Problem formulation

2.1 Average consensus algorithm

In this paper, we concentrate exclusively on the average consensus algorithm, where n nodes aim to reach consensus with the final value being the average of their initial scalar values. Denote the value held by the ith node at time k as $x_i(k)$, and by $x(k)$ the n-dimensional vector obtained by stacking the values of all the nodes in a column vector.

The interconnection topology among the nodes at time k can be described by a *consensus graph* $\mathcal{G}(k)$, with an edge present between two nodes iff they can exchange information. Suppose $\mathcal{N}_i(k)$ is the neighbour set of node i at time k, and all nodes are allowed to simultaneously broadcast their states to their neighbours. Every exchange happens in a single packet transmission interval (also referred to as a *time slot* and normalised to length 1). Therefore,

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

where h is a scalar constant designed to ensure convergence of the algorithm. In this case, we consider the *iteration time* to be 1. Denoting the Laplacian of $\mathcal{G}(k)$ by $L(k)$, equation (1) can be rewritten as

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

It can be shown (see, e.g., Olfati-Saber and Murray, 2004) that under proper connectivity assumptions, provided h is small enough, average consensus is achieved.

We assume that h is fixed and $h < \frac{1}{2d_{\max}}$ where $d_{\max}$ is the maximum node degree in $\mathcal{G}(k)$ for all k. To reach average consensus, the consensus graph should be balanced. We constrain the graph to be undirected, thus satisfying this condition.

The rate of convergence of the value of the nodes is a function of the graph topology. For a time-invariant topology, it can be shown (see, e.g., Desai and Rao, 1990; Olfati-Saber and Murray, 2004; Seneta, 1981) that the convergence of the consensus protocol is geometric. The rate is 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. If L is symmetric, its SLEM can be written as its norm restricted to the subspace orthogonal to $\mathbf{1}_n \triangleq [1 \;\; 1 \;\; \dots \;\; 1]^T_{1\times n}$, where A^T denotes the transpose of a matrix A.

However, in wireless networks, a number of transmissions are necessary to set up $\mathcal{G}(k)$, since the channel must be shared by different users. Therefore if each node can receive data from say at most one neighbour at any given time, the exchange of information necessary for iteration k will require at least $1 + \max_i |\mathcal{N}_i(k)|$ transmissions. This idea is developed further in this paper.

2.2 Communication protocols

Usual treatments of the average consensus algorithm presume the existence of a consensus graph. In other words, there exists a communication channel between any two nodes connected by an edge in this graph. However, wireless channels are broadcast, and as pointed out earlier, the network topology is a *design* parameter. The broadcast nature of the channel also manifests as interference from unintended transmitters. This effect on the average consensus algorithm has not been studied previously.

We consider a situation in which the physical locations of the nodes are given. Every node then decides on the power with which it transmits. This power determines the communication radius of the node according to the relation

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

where P_0 is a normalisation constant, α is the path-loss exponent (typically $2 \leq \alpha \leq 5$), P is the transmission power and r_c is the (normalised) communication radius. All nodes at a distance smaller than r_c from the transmitter can receive the transmitted message.

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). For simplicity, we assume $r_c = r_i$, noting that the results can be generalised to other cases.

We now choose a Time Division Multiple Access (TDMA) Medium Access Control (MAC) protocol for the nodes. These protocols ensure successful communication by scheduling transmissions in time such that messages do not interfere. They demonstrate better throughput than collision-based MAC protocols, but require greater synchronisation and coordination among the nodes (see Liu and Haenggi, 2005; Xie and Haenggi, 2005).

2.3 Problem formulation

The operation of the network is divided into two phases that are repeated at every consensus iteration. Phase 1 sets up the consensus graph using multiple transmission stages. Each transmission stage consumes one time slot. In Phase 2, the nodes update their values according to equation (1). This step is assumed to be instantaneous. If Phase 1 occupies T time slots, we have

$$x(kT + T) = (I - hL)x(kT) \tag{3}$$

for the kth update. Observe that finite communication time slows down convergence.

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 characterise the convergence of the consensus algorithm for the optimal MAC protocol that sets up the consensus graph in the smallest number of time slots (thus maximising the update rate). This is studied for a regular grid of sensor nodes. A periodic boundary condition is chosen for analytical convenience.

We assume the following:

- All nodes transmit with the same power and follow a time-invariant transmission policy.
- At the time of an update, we require that the consensus graph be undirected. i.e., for any two nodes i, j in the network, $j \in \mathcal{N}_i \Leftrightarrow i \in \mathcal{N}_j$.
- We do not assume explicit routing of values through nodes.
- We assume half-duplex operation and further that collided packets cannot be decoded.

Under these assumptions, we are able to show the following results:

- We bound the rate of convergence for the optimal MAC scheduling protocol for the average consensus algorithm for tori in n dimensions.
- We show that network geometry plays a key role in identifying the optimum power allocation that maximises 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.

In the next section, we begin by studying the convergence properties of MAC protocols that maximise the speed of convergence for a given consensus graph $\mathcal{G}$.

3 Analysis of a ring and a 2D torus

3.1 The 1-D case: nodes on a ring

Consider n nodes $\{0, 1, \ldots, n-1\}$ placed uniformly on a circle of radius r centred 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 neighbours 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

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

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

An alternative interpretation of this geometry is to consider the n nodes on a regular one-dimensional torus $[0, 1]$ (hereafter called a '1-torus' or $\mathcal{T}_1(n)$), with node k assumed to be at k/n. This perspective allows us to compare the results developed here with those for higher dimensional tori.

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

$$[1 - 2mh \quad \mathbf{1}_m^T \quad 0 \quad 0 \quad \cdots \quad 0 \quad \mathbf{1}_m^T],$$

where

$$\mathbf{1}_m^T = [1 \quad 1 \quad \cdots \quad 1]_{1\times m}.$$

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. We now characterise the time necessary to form $\mathcal{G}_{1,m}$.

Figure 1 Schematic of nodes placed along a ring

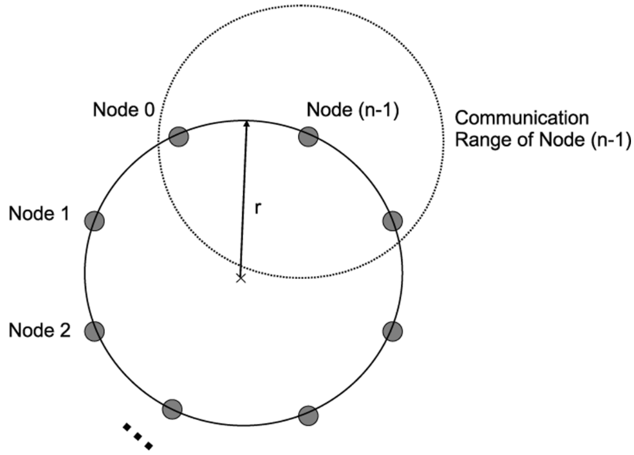

3.2 *Characterising the time for the communication phase*

We bound the length of the shortest TDMA schedule that forms $\mathcal{G}_{1,m}$. Denote its length by $T_1^*(m)$, or, more compactly, as T_1^*. Observe that for all $m_1 \leq m_2$, $\mathcal{G}_{1,m_1} \subseteq \mathcal{G}_{1,m_2}$. Therefore $\mathcal{G}_{1,m}$ 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 that a *link* is formed from node v to node u whenever the packet from v is successfully decoded at u. Since $\mathcal{G}_{1,m}$ is undirected, an edge $e \in \mathcal{G}_{1,m}$ connecting v and u is formed iff both v and u form links with each other.

Lemma 1: *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*

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

Proof: Suppose $\mathcal{G}_{1,m}$ is formed in T time slots. Each node is connected to m nearest neighbours on either side, and the node degree of $\mathcal{G}_{1,m}$ is $2m$. Suppose N_t links in $\mathcal{G}_{1,m}$ are formed in time slot t. Given that there are are $4mn$ links in the graph,

$$\sum_{t=1}^{T} N_t = 4mn. \tag{5}$$

Now suppose that $K(t)$ nodes $\{v_{i_1}, v_{i_2}, \ldots, v_{i_{K(t)}}\}$ transmit in time slot t. Let the respective power allocations be $\{P_{l_1}, P_{l_2}, \ldots, P_{l_{K(t)}}\}$.

Assume that the power allocated to node v_{i_k} allows it to broadcast its message to at most the l_k nearest neighbours on either side, or equivalently form at most $2l_k$ links. Since at most $2m$ of these links contribute to the formation of edges in $\mathcal{G}_{1,m}$,

$$N_t \leq 2 \sum_{k=1}^{K(t)} \min(m, l_k) \leq 2mK(t),$$

where the second inequality is obtained by choosing $\min(m, l_k) = m$ for all transmitting nodes. To minimise T in equation (5), we need to maximise N_t for all t.

Observe that any transmission to a neighbour beyond the m nearest neighbours does not contribute to the graph $\mathcal{G}_{1,m}$; this implies that choosing $l_k = m$ is sufficient. This argument holds for any of the $K(t)$ transmitters. Hence each of the transmitting nodes should form m links on either side, or equivalently transmit with power P_m.

Now note that the optimal protocol selects $\{v_{i_1}, v_{i_2}, \ldots, v_{i_{K(t)}}\}$ for all t while meeting this upper bound, and ensuring that T is minimised. As noted above, this minimum value of T has been called T_1^*.

Since the optimal protocol leads to each node transmitting to m-nearest neighbours whenever it communicates, for every $2m+1$ adjacent nodes, only one node can transmit in any slot to avoid interference. As a result, a TDMA schedule cannot have fewer than $2m+1$ slots. In other words, $T_1^* \geq 2m+1$.

An upper bound on T_1^* can be obtained by considering the length of a particular TDMA schedule. Consider a protocol in which each transmitter is allocated power P_m. Due to interference constraints, no two transmission intervals can overlap each other. Given that there are a total of n nodes uniformly placed on a ring, the maximum number of allowed transmitters at any time step is

$$K_{\max} = \left\lfloor \frac{n}{2m+1} \right\rfloor.$$

The transmission schedule for this MAC protocol is as follows. Consider time slot 1. Suppose some node $v \in \mathcal{V}$ transmits at P_m. We require that all nodes $(2m+1)$ nodes apart should transmit as long as the half-duplex and interference constraints are satisfied. Since the maximum number of simultaneous transmissions possible is $\lfloor \frac{n}{2m+1} \rfloor$, in $2m+1$ time slots, $(2m+1)\lfloor \frac{n}{2m+1} \rfloor$ nodes can transmit. After $2m+1$ time slots,

$$n - (2m+1)\left\lfloor \frac{n}{2m+1} \right\rfloor = \mathrm{rem}(n, 2m+1) \leq 2m$$

nodes will not have transmitted with this protocol. So this schedule forms $\mathcal{G}_{1,m}$ in $T_u = (2m+1) + \mathrm{rem}(n, 2m+1) \leq 4m+1$. Hence, we conclude $T_1^* \leq T_u \leq 4m+1$. □

We have thus bounded the length of the shortest TDMA schedule that forms the consensus graph $\mathcal{G}_{1,m}$. In other words, we have bounded the smallest time T in the update equation (3).

3.3 Bounding the rate of convergence

To characterise the fastest convergence possible for a given $\mathcal{G}_{1,m}$ we need to use the above result in conjunction with the spectral properties of $\mathcal{G}_{1,m}$ which are presented below.

Theorem 2: *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 $\epsilon(k) = x(k) - \mathbf{1}_n x_{av}$ converges geometrically to zero with the rate of decay β bounded as*

$$\rho_1^{\frac{1}{2m+1}} \leq \beta \leq \rho_1^{\frac{1}{4m+1}} \tag{6}$$

where

$$\begin{aligned} \rho_1 &= 1 - h(2m+1) + hS_1^{(m,n)} \\ S_p^{(m,n)} &= \frac{\sin\left(\frac{(2m+1)\pi p}{n}\right)}{\sin\left(\frac{\pi p}{n}\right)}, \quad p = 0, 1, \ldots n-1. \end{aligned} \tag{7}$$

Proof: The consensus graph at each update step is balanced and connected. Thus, the node values converge to the average of their initial values with the decay rate as the modulus of the second largest eigenvalue of $F_{1,m}$ (Olfati-Saber and Murray, 2004). Denote $\mathrm{e}^{-j\frac{2\pi k}{n}}$ by $W_{k,n}$. Since $F_{1,m}$ is circulant, its kth eigenvalue ρ_k is

$$\begin{aligned} \rho_k &= 1 - 2mh + h\sum_{l=1}^{m}(W_{k,n}^l + W_{-k,n}^l) \\ &= 1 - (2m+1)h + 2h\sum_{l=0}^{m}\cos\left(\frac{2\pi kl}{n}\right) \end{aligned} \tag{8}$$

which results in

$$\rho_k = 1 - (2m+1)h + hS_k^{(m,n)}, \quad k = 0, \ldots, n-1. \tag{9}$$

It is easy to see that $\rho_0 = 1$. The second largest eigenvalue is given by $\rho_1 = \rho_{n-1} < 1$, where ρ_1 was defined in equation (7). From Lemma 1, $2m+1 \leq T_1^* \leq 4m+1$, and equation (6) follows. □

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 equation (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 offset 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 *a priori* evident.

3 The effect of increasing the transmission power is most prominent at small P_m. This can again be seen from Figure 2. If $\theta = p\pi/n$ and $p \ll n$,

$$\sin\theta \approx \theta - \theta^3/3. \tag{10}$$

We use equation (10) to express the spectral gap $\mathrm{SG} \triangleq 1 - \rho_1^{\frac{1}{T}}$ when $m \ll n$ as

$$\begin{aligned}\mathrm{SG} &= 1 - \left(1 - h(2m+1) + h\frac{\sin\left(\frac{(2m+1)\pi}{n}\right)}{\sin\left(\frac{\pi}{n}\right)}\right)^{\frac{1}{T}} \\ &\approx \eta\, mh(m+1)(2m+1)T^{-1}n^{-2}\end{aligned}$$

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 $\frac{m}{n^2}$.

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

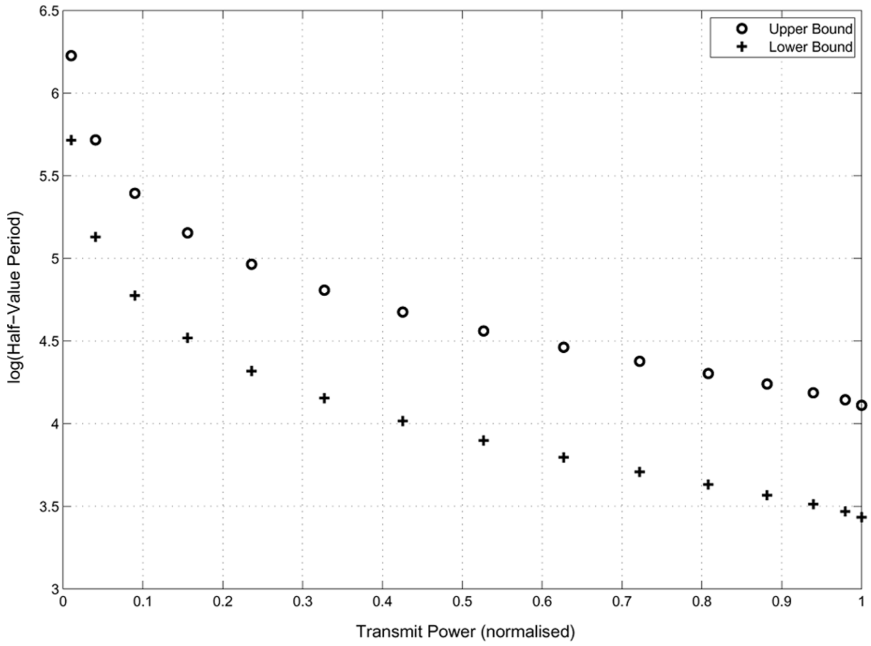

3.4 Nodes on a two-dimensional torus

We now use a similar approach to analyse consensus on two dimensional tori. A 2-torus $\mathcal{T}_2(n)$ consists of $n = l^2$ regularly spaced points on a 2-dimensional torus at $[0,1]^2$. A 2-torus with $n = 9$ nodes is depicted in Figure 3. Clearly, the 2-torus lies in the XY plane. The X and Y axes are referred to as the two *axial directions* of this torus. Choosing a node as the origin, we label each node by its position with respect to the origin node. Therefore, the node at $\left(\frac{i}{\sqrt{n}}, \frac{j}{\sqrt{n}}\right)$ is labelled (i, j).

Figure 3 The toroidal lattice $\mathcal{T}_2(9)$. Nodes in black indicate those physically placed in $[0,1]^2$. Node $(0,0)$ represents the node at the origin, with each of the black 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

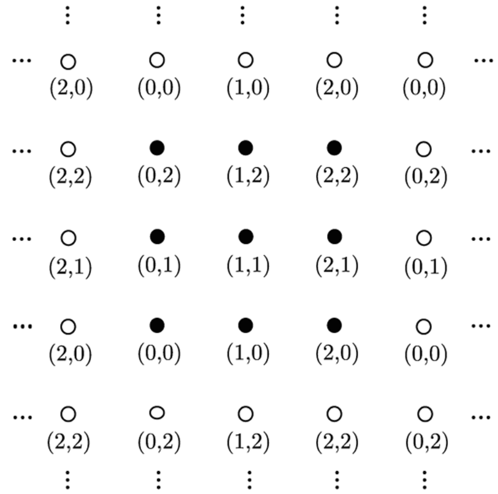

We denote the position of node (i, j) by $z_{(i,j)}$, or, more compactly, $z_{ij} \equiv \left(\frac{i}{\sqrt{n}}, \frac{j}{\sqrt{n}}\right)$. For example, in Figure 3, the position of node $(1, 1)$ is $z_{11} \triangleq \left(\frac{1}{3}, \frac{1}{3}\right)$. An alternative interpretation of a toroidal arrangement in two-dimensions is shown in Figure 4. Both these interpretations yield similar results in the limiting case of a large torus where local distances are not significantly affected by curvature. We will focus on the former interpretation in this paper.

Suppose that all nodes on a torus $\mathcal{T}_2(n)$ participate in an average consensus algorithm of the form equation (2) with a power allocation of P_m per node. As an extension to the 1-D case, we assume that with power P_m, the set of all reachable nodes will lie in a *communication sphere* of radius $\frac{m}{n^{1/2}}$ $\left(m < \lfloor \frac{l}{2} \rfloor\right)$ centred at the transmitter. We now have the desired consensus graph $\mathcal{G}_{2,m} = (\mathcal{V}, \mathcal{E}_{2,m})$ where the vertex set

$$\mathcal{V} = \{0, 1, \ldots, l-1\}^2,$$

is the set of all points in $\mathcal{T}_2(n)$. The edge set $\mathcal{E}_{2,m}$ is formed by connecting every node to all nodes on the torus that are within its communication sphere:

$$\mathcal{E}_{2,m} = \left\{ v, u \in \mathcal{V} : v \neq u, \ell_2(z_v, z_u) \leq \frac{m}{\sqrt{n}} \right\},$$

where $\ell_2(x, y)$ is the Euclidean norm between x and y.

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

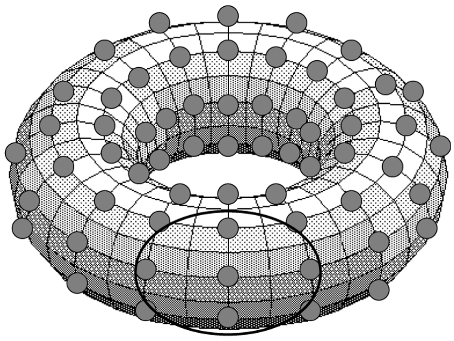

Using notation similar to that of Section 3.1, the Laplacian and the update matrix for $\mathcal{G}_{2,m}$ will be denoted by $L_{2,m}$ and $F_{2,m} \triangleq I - hL_{2,m}$, respectively. Assuming as before that each transmission occupies one time slot, we now study the convergence properties of the optimal MAC protocol that will form $\mathcal{G}_{2,m}$ in the smallest number of time slots.

We note that geometry plays a key role in scheduling transmitters in multi-dimensional networks. This is illustrated for a 2-torus with $n = 25$ nodes and $m = 1$ in Figure 5. Although identifying this schedule is non-trivial, as before we can circumvent this problem by exploiting its optimality to bound its length. Similar to Section 3.2, we denote this optimal length by T_2^*. We begin by characterising 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 the optimality of the MAC protocol and constraints imposed by our problem to bound the convergence rates for a 2-torus.

Figure 5 Effect of geometry in sphere packing, shown for $\mathcal{T}_2(25)$ and $m = 1$. The black nodes denote the transmitters (e.g., $(0,0)$). Note that the dimensionality is exploited to allow more concurrent transmissions. The diamond-shaped nodes (e.g., $(4,0)$) are covered by transmitters via their images (shown as unfilled circles)

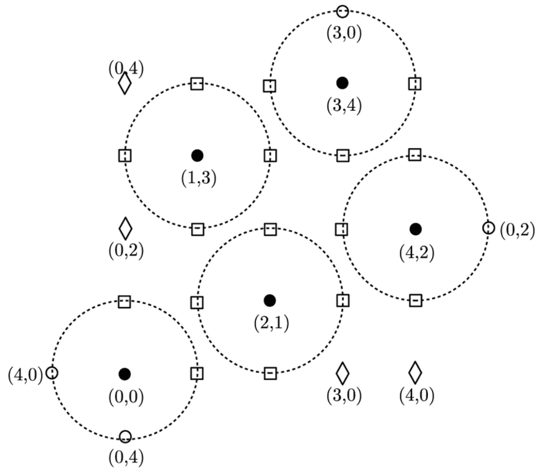

With P_m as the transmit power that enables a node to form error-free links with m neighbours in the axial directions. Given that there are $\sqrt{n}$ nodes in either of the axial directions,

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

where $\alpha \geq 2$ is the path-loss exponent. Define Π to be the elementary circulant matrix having as the first row $[0 \quad 1 \quad 0 \quad \cdots \quad 0 \quad 0]_{1\times m}$. Then the update matrix $F_{2,m}$ can be written as

$$F_{2,m} = \sum_{ij} a_{ij}\Pi^i \otimes \Pi^j$$

where $\otimes$ denotes the Kronecker product and

$$a_{ij} = \begin{cases} h & \text{if } (i,j) \neq (0,0), \\ 1 - h|\{(i,j) \in \mathbb{Z}^2 : i^2 + j^2 \leq m^2\}| & \text{if } (i,j) = (0,0). \end{cases}$$

We are now ready to bound T_2^*:

Lemma 3: *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 $1 < m < \lfloor l/2 \rfloor$, we have*

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

where $T_l = 2m^2 + 2m + 1$ 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. Let $\mathcal{B}$ denote its communication disc, i.e., a circle of radius $\frac{m}{\sqrt{n}}$ centred at $(0,0)$. From the definition of a feasible schedule, other than $(0,0)$, no other node inside $\mathcal{B}$ can transmit at this time. This means a feasible schedule can allow at most one transmission inside $\mathcal{B}$. Therefore any feasible schedule needs at least as many time slots as the number of nodes inside a sphere of radius $\frac{m}{\sqrt{n}}$ centred at $(0,0)$, in order to form $\mathcal{G}_{2,m}$. In other words, a feasible schedule has *at least* $|\mathcal{B} \cap \mathcal{T}_2(n)|$ time slots. This involves counting the number of nodes of a square grid that fall inside the circle we are considering, which can be cumbersome. For our purposes, it will suffice to find a lower bound for this value by counting only those nodes that fall in a suitably chosen square $\mathcal{U}$ circumscribed by $\mathcal{B}$.

Consider a square $\mathcal{U}$ having as its vertices, the nodes $(m,0)$, $(0,m)$, $(-m,0)$ and $(0,-m)$. Clearly $\mathcal{U} \subseteq \mathcal{B}$. Each of its diagonals contains $2m+1$ nodes. The line segments joining nodes $(-m+|k|,k)$ and $(m-|k|,k)$ for $k = -m, -m+1, \ldots, m-1, m$ are all parallel to the segment joining $(-m,0)$ and $(m,0)$. Each such segment has exactly $2m+1-2|k|$ nodes. Counting all these nodes yields the number of nodes in $\mathcal{B}$:

$$\begin{aligned}|\mathcal{U} \cap \mathcal{T}_2(n)| &= \sum_{k=-m}^{m} (2m+1) - 2|k| \\ &= 2m+1+2\sum_{k=1}^{m} 2(m-k) \\ &= 2m^2+2m+1.\end{aligned}$$

We can thus choose $T_l = 2m^2+2m+1$ as the lower bound on the length of any feasible schedule that forms $\mathcal{G}_{2,m}$. In particular, T_2^* is the length of the shortest feasible schedule. Consequently, $T_2^* \geq T_l$.

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 $p, q = 1, \ldots, \lfloor \frac{l}{2m+1} \rfloor$ 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 lies in at most one communication sphere. In each subsequent time slot, repeat this process by choosing some other node (i,j) inside the square of side $(2m+1)/\sqrt{n}$ centred 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,

$$\begin{aligned}n - \left\lfloor \frac{l}{(2m+1)} \right\rfloor^2 (2m+1)^2 = 2\left\lfloor \frac{l}{2m+1} \right\rfloor (2m+1)\,\mathrm{rem}(l,(2m+1)) \\ + (\mathrm{rem}(l,2m+1))^2\end{aligned}$$

nodes are yet to transmit. In the first term, we can schedule $\lfloor l/(2m+1) \rfloor$ nodes in each of the $2\mathrm{rem}(l,2m+1) \leq 4m$ 'rows' that require at most $4m(2m+1)$ additional time slots. Scheduling one node per time slot, all the remaining $(\mathrm{rem}(l,2m+1))^2$ 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 $T_2^* \leq 16m^2+8m+1$. □

Compared to the 1-torus, the optimal schedule for a 2-torus is bounded by two quadratic terms. As we shall see, this plays a key role in determining the effect of transmit power on convergence behaviour.

3.5 Bounding the rate of convergence

Finding the eigenvalues of $F_{2,m}$ seems intractable. This difficulty arises from the fact that the nodes that can receive data from a particular transmitter lie within discs (or spheres) centred at the transmitter. While in one-dimension such discs can cover the entire ring, such coverage is not possible in higher dimensions in general. For our purpose, we lower and upper bound such discs by squares of suitable side length that cover the entire region. The idea is illustrated in Figure 6. To this end, we begin with the following preliminary result.

Figure 6 Edges to $(2,2)$ in $\mathcal{T}_2(25)$ in $\widehat{\mathcal{G}}_{2,\tilde{m}}$, $\mathcal{G}_{2,m}$ and $\widehat{\mathcal{G}}_{2,m}$ for $m=2$. For nodes shown as squares (e.g., $(2,3)$, that lie on the innermost square), edges between $(2,2)$ and these nodes exist in all these graphs. For traingle-shaped nodes (e.g., $(4,2)$, that lie on the circle), such edges exist only in $\mathcal{G}_{2,m}$ and $\widehat{\mathcal{G}}_{2,m}$. For nodes shown as circles (e.g., $(0,0)$), such edges exist only in $\widehat{\mathcal{G}}_{2,m}$

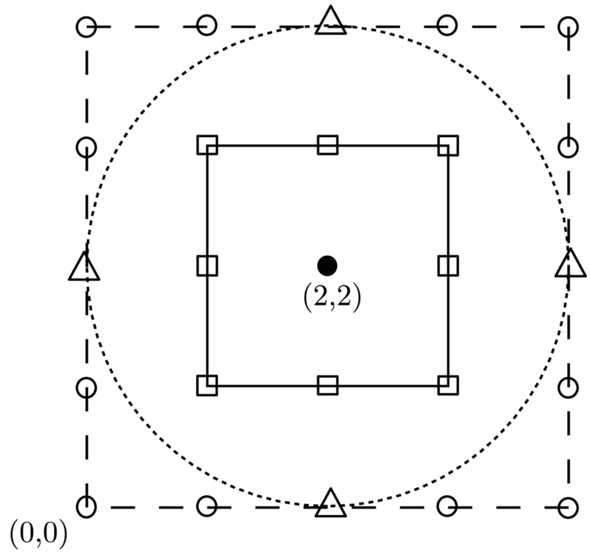

Lemma 4: *Let $\widehat{\mathcal{G}}_{2,m}$ be the consensus graph formed over $\mathcal{T}_2(n)$ by placing edges between each node (i,j) with all other nodes $(k,l) \neq (i,j)$ satisfying*

$$\ell_\infty(z_{ij}, z_{kl}) \leq \frac{m}{\sqrt{n}}, \quad 1 \leq m < \left\lfloor \frac{\sqrt{n}}{2} \right\rfloor$$

where ℓ_∞ denotes the ∞-norm. Also denote the Laplacian of $\widehat{\mathcal{G}}_{2,m}$ by $\widehat{L}_{2,m}$ and its maximum degree by $d_{\max}$ and the update matrix $\widehat{F}_{2,m} = I - h\widehat{L}_{2,m}$ for some $0 \leq h \leq \frac{1}{2d_{\max}}$. Then the eigenvalues of the $\widehat{F}_{2,m}$ are

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

where, as defined above

$$S_a^{(m,l)} = \frac{\sin\left(\frac{(2m+1)\pi a}{l}\right)}{\sin\left(\frac{\pi a}{l}\right)}, \quad a = 0, 1, \ldots l-1.$$

Proof: Let matrices A_k, $k = 0, 1$ be be $l \times l$ circulant matrices, with their first row being $[d_k \quad -\mathbf{1}_m^T \quad 0 \quad 0 \quad \cdots \quad 0 \quad -\mathbf{1}_m^T]_{1\times l}$, where $d_0 = (2m+1)^2 - 1$ and

$d_1 = -1$. Let $B_m \triangleq \mathbf{1}_m^T \otimes A_1$. It is easy to see that the Laplacian $\widehat{L}_{2,m}$ of $\widehat{\mathcal{G}}_{2,m}$ is an $n \times n$ block circulant matrix with each of its l block rows being

$$[A_0 \quad B_m \quad 0 \quad 0 \quad \cdots \quad 0 \quad B_m]_{l\times n}.$$

Exploiting the properties of block circulant matrices and the fact that A_k's are themselves circulant (and consequently share the same eigenvectors) to compute the eigenvalues of $\widehat{L}_{2,m}$:

$$\mu_{r,s} = \sum_{t=0}^{l-1} \eta_{r,t}\, \mathrm{e}^{-j\frac{2\pi st}{l}} \tag{11}$$

where $\eta_{r,t}$ is the rth eigenvalue of $A_t\ \forall r, s = 0, 1, \ldots, l-1$. Using the 1-torus result from equation (9) for $\eta_{r,t}$ and simplifying, the eigenvalues of $\widehat{F}_{2,m} = I - h\widehat{L}_{2,m}$ are

$$\lambda_{a,b} = 1 - h(2m+1)^2 + hS_a^{(m,\sqrt{n})}S_b^{(m,\sqrt{n})} \tag{12}$$

which is the desired result. $\square$

We are now in a position to bound the rate of decay for the 2-torus:

Theorem 5: *Consider a consensus algorithm of the form* (2) *on* $\mathcal{G}_{2,m}$. *If each node transmits at* P_m *for* $1 \le m < \lfloor\frac{\sqrt{n}}{2}\rfloor$, *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*

$$\lambda_1^{\frac{1}{2m(m+1)}} < \beta < \tilde{\lambda}_1^{\frac{1}{16m^2+8m+1}}$$

where $\tilde{m} \triangleq \lfloor\frac{m}{\sqrt{2}}\rfloor$ *and*

$$\lambda_1 = \left(1 - h(2m+1)^2 + h(2m+1)S_1^{(m,\sqrt{n})}\right)$$
$$\tilde{\lambda}_1 = \left(1 - h(2\tilde{m}+1)^2 + h(2\tilde{m}+1)S_1^{(\tilde{m},\sqrt{n})}\right).$$

Proof: Consider undirected graphs $\mathcal{G}_{\text{sub}}, \mathcal{G}$ and $\mathcal{G}_{\text{sup}}$ with a common vertex set $\mathcal{V}$ and edge sets $\mathcal{E}_{\text{sub}} \subseteq \mathcal{E} \subseteq \mathcal{E}_{\text{sup}}$, with the same h (refer to equation (2)) that ensures consensus for all these graphs. Call $\mathcal{G}$ the nominal graph and $\mathcal{G}_{\text{sub}}$ and $\mathcal{G}_{\text{sup}}$ as sub- and super-graphs of $\mathcal{G}$. Suppose all the graphs satisfy conditions to reach average consensus. The proof rests on two simple but important facts:

Fact 1: If iterations are performed at the same rate and if the same h is used, the consensus algorithm on a graph cannot be slower than that on its sub-graphs. This follows from the fact that the second largest eigenvalue always decreases if edges are added to a graph (cf. Mohar, 1991, Theorem 3.2).

Fact 2: Suppose the iterations on these graphs occur every T_u, T_2^* and T_l time slots respectively, with $T_u \ge T_2^* \ge T_l$. Then the graphs on which the consensus algorithm converges the fastest and slowest are still $\mathcal{G}_{\text{sup}}$ and $\mathcal{G}_{\text{sub}}$, respectively.

In what follows, we will first define the nominal consensus graph $\mathcal{G}$ and choose appropriate super- and sub-graphs $\mathcal{G}_{\text{sup}}$ and $\mathcal{G}_{\text{sub}}$. We will then use the results from Lemma 3 to obtain the values of T_l, T_2^* and T_u.

Consider the graph $\widehat{\mathcal{G}}_{2,m}$ as defined in Lemma 4. Note that

$$\ell_2(r_{ij}, r_{kl}) \le c \implies \ell_\infty(r_{ij}, r_{kl}) \le c.$$

Choosing $c = \frac{m}{\sqrt{n}}$, it is easy to see that every edge in $\mathcal{E}_{2,m}$ is present in $\widehat{\mathcal{E}}_{2,m}$. It follows that $\mathcal{G}_{2,m} \supseteq \widehat{\mathcal{G}}_{2,m}$. This is illustrated in Figure 6 for $n = 25, m = 2$.

Define $\tilde{m} = \lfloor \frac{m}{\sqrt{2}} \rfloor$ and form a graph $\widehat{\mathcal{G}}_{2,\tilde{m}} = (\mathcal{V}, \hat{\mathcal{E}}_{2,\tilde{m}})$. It is easy to see that

$$\ell_\infty(r_{ij}, r_{kl}) \le \left\lfloor \frac{\tilde{m}}{\sqrt{n}} \right\rfloor \ell_2(r_{ij}, r_{kl})$$

$$\implies \ell_2(r_{ij}, r_{kl}) \le \sqrt{2} \left\lfloor \frac{\tilde{m}}{\sqrt{n}} \right\rfloor \le \left\lfloor \frac{\sqrt{2}\tilde{m}}{\sqrt{n}} \right\rfloor \le \frac{m}{\sqrt{n}}.$$

Therefore $\widehat{\mathcal{G}}_{2,\tilde{m}} \subseteq \mathcal{G}_{2,m}$.

The differences in connectivity in the graphs $\widehat{\mathcal{G}}_{2,\tilde{m}}$, $\mathcal{G}_{2,m}$ and $\widehat{\mathcal{G}}_{2,m}$ are illustrated in Figure 6, for a torus of $n = 25$ nodes and $m = 2$. Here $\tilde{m} = \lfloor 2/\sqrt{2} \rfloor = 1$.

The consensus graphs $\widehat{\mathcal{G}}_{2,m}$ and $\mathcal{G}_{2,m}$ correspond to update matrices $\widehat{F}_{2,m}$ and $F_{2,m}$ respectively for all m. We now use the results for eigenvalues of the Laplacian $\widehat{L}_{2,m}$, derived in Lemma 4. The rate of convergence on $\widehat{\mathcal{G}}_{2,m}$ is a function of the second largest eigenvalue modulus that is obtained by setting $a = 0$ and $b = 1$ in equation (12):

$$\lambda_{0,1} = 1 - h(2m+1)^2 + h(2m+1) S_1^{(m,l)} \triangleq \lambda_1. \tag{13}$$

Making the substitution $m \leftarrow \tilde{m}$ gives

$$\tilde{\lambda}_1 \triangleq 1 - h(2\tilde{m}+1)^2 + h(2\tilde{m}+1) S_1^{(\tilde{m},l)} \tag{14}$$

which is the corresponding value for $\widehat{\mathcal{G}}_{2,\tilde{m}}$.

Consider now a consensus algorithm on $\mathcal{G}_{\text{sub}} = \widehat{\mathcal{G}}_{2,\tilde{m}}$, $\mathcal{G} = \mathcal{G}_{2,m}$ and $\mathcal{G}_{\text{sup}} = \widehat{\mathcal{G}}_{2,m}$. All the graphs are given equal edge weights h, determined by the maximum node degree in $\widehat{\mathcal{G}}_{2,m}$. Suppose an iteration on the torus forms $\mathcal{G}_{2,m}$ in T_2^* time slots. Let iterations on $\widehat{\mathcal{G}}_{2,m}$ and $\widehat{\mathcal{G}}_{2,\tilde{m}}$ occur every T_l and T_u time slots, respectively, where $T_l = 2m^2 + 2m + 1$ and $T_u = 16m^2 + 8m + 1$. From Lemma 3 we know that $T_u \ge T_2^* \ge T_l$. The values of SLEM for the corresponding to $\widehat{\mathcal{G}}_{2,m}$ and $\widehat{\mathcal{G}}_{2,\tilde{m}}$ are known. Hence the convergence rates are, respectively, λ_1^{1/T_l} and $\tilde{\lambda}_1^{1/T_u}$. The result now follows from Fact 2. □

To understand the effect of higher transmit powers on the convergence rate in 2-tori, we simplify the expressions for λ_1 and $\tilde{\lambda}_1$ in Theorem 5 using $h = \gamma/(2m+1)^2$, $0 < \gamma < 1$:

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

$$\tilde{\lambda}_1 = 1 - \gamma \frac{(2\tilde{m}+1)^2}{(2m+1)^2} + \gamma \frac{(2\tilde{m}+1)}{(2m+1)^2} \frac{\sin((2\tilde{m}+1)\pi/\sqrt{n})}{\sin(\pi/\sqrt{n})}.$$

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 the same form as ρ_1 in Theorem 2, except for the $\sqrt{n}$. For large m, the behaviour of $\tilde{\lambda}_1$ will also be similar to λ_1.

However, the two cases differ in the length of the optimal schedule, whose length was shown to be $\Theta(m)$ for the 1-torus, and $\Theta(m^2)$ for the 2-torus. Therefore, the effect of interference depends on the geometry of node placement. High transmit powers cause greater interference, thereby reducing network throughput. In scheduling MAC protocols, this effect is reflected in longer schedules. Although this is offset by the resultant long-range connections in a 1-torus, it is no longer true for higher dimensions. Figure 7 shows the lower bound for β obtained from Theorem 5 for a 2-torus with $n = 4096$ nodes arranged as a 64×64 toroidal lattice. Observe that the convergence rate *worsens* with increasing transmit power. This is surprising when compared to the 1-torus, where the result is the opposite.

Figure 7 Variation of half-value period with transmission power on a 2-torus

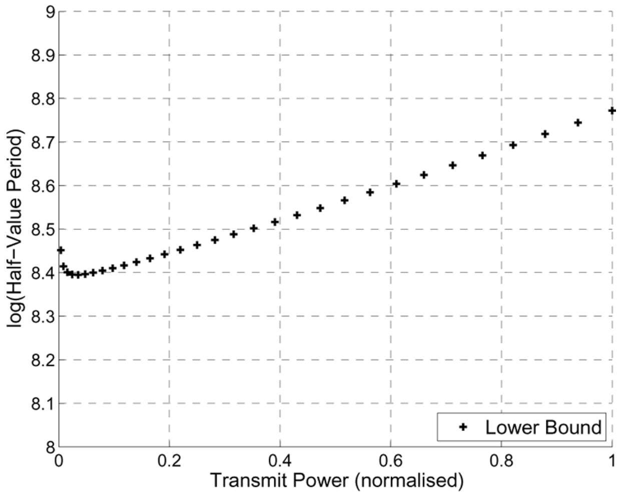

3.6 Tori in arbitrary dimensions

The results in Lemma 3 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 $|\mathcal{U} \cap \mathcal{T}_d(n)|$ where $\mathcal{U}$ is a polyhedron that can be circumscribed by the sphere of radius $m/(n^{1/d})$. Similarly, to find the upper bound one can generalise the schedule described in Lemma 1 that was used to find an upper bound. It can be shown that the optimal schedule length will be $\Theta(m^d)$.

The results in Lemma 4 can also be generalised to d-dimensions as $\lambda_1 = 1 - h(2m+1)^d + h(2m+1)^{d-1} S_1^{(m,l)}$ for the lower bound and choosing $\tilde{m} = \lfloor m/\sqrt{d} \rfloor$

to find $\tilde{\lambda}_1$ for an upper bound on the convergence rate. Thus the convergence slows down with transmit power in geometries having dimension two or more.

4 Conclusions

We introduced a framework that considers the effects of realistic communication constraints on average consensus algorithms. In particular, we analytically characterise the performance of the medium access control algorithm that maximises 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 (Carli et al., 2008). Another issue is the effect of stochastic data loss through effects due to fading and interference, using a different framework as compared to Boyd et al. (2006) and Hovareshti et al. (2008), to explicitly account for interference.

Acknowledgement

The authors thank the anonymous reviewers for their helpful suggestions. The partial support of DTRA through grant N00164-07-8510 and NSF grant 0834771 is gratefully acknowledged.

References

Blondel, V., Hendrickx, J., Olshevsky, A. and Tsitsiklis, J. (2005) 'Convergence in multiagent coordination, consensus and flocking', *Proceedings of the 44th IEEE Conference on Decision and Control*, Sevilla, Spain, pp.2996–3000.

Boyd, S., Ghosh, A., Prabhakar, B. and Shah, D. (2006) 'Randomized gossip algorithms', *IEEE Transactions on Information Theory*, Vol. 52, No. 6, pp.2508–2530.

Carli, R., Fagnani, F., Speranzon, A. and Zampieri, S. (2008) 'Communication constraints in the average consensus problem', *Automatica*, Vol. 44, No. 3, pp.671–684.

Cybenko, G. (1989) 'Dynamic load balancing for distributed memory multiprocessors', *Journal on Parallel and Distributed Computing*, Vol. 7, pp.279–301.

Desai, M.P. and Rao, V.B. (1990) '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*, New Orleans, USA, Vol. 2, pp.1211–1214.

Fang, L. and Antsaklis, P. (2005) 'On communication requirements for multi-agents consensus seeking', *Proceedings of Workshop NESC'05: University of Notre Dame, Lecture Notes in Control and Information Sciences (LNCIS)*, Springer, Vol. 331, pp.53–68.

Hovareshti, P., Gupta, V. and Baras, J.S. (2008) 'Average consensus over small world networks: a probabilistic framework', *Proceedings of the 47th IEEE Conference on Decision and Control*, Cancun, Mexico.

Huang, M. and Manton, J.H. (2007a) 'Stochastic double array analysis and convergence of consensus algorithms with noisy measurements', *Proc. American Control Conference*, New York City, USA, pp.705–710.

Huang, M. and Manton, J.H. (2007b) 'Stochastic Lyapunov analysis for consensus algorithms with noisy measurements', *Proc. American Control Conference*, New York City, USA, pp.1419–1424.

Jadbabaie, A., Lin, J. and Morse, A.S. (2003) 'Coordination of groups of mobile autonomous agents using nearest neighbour rules', *IEEE Transactions on Automatic Control*, Vol. 48, No. 6, pp.988–1001.

Liu, X. and Haenggi, M. (2005) 'Throughput analysis of fading sensor networks with regular and random topologies', *EURASIP Journal on Wireless Communications and Networking*, Vol. 4, pp.554–564, Special Issue on Wireless Sensor Networks.

Mohar, B. (1991) 'The laplacian spectrum of graphs', in Alavi, Y., Chartrand, G., Oellermann, O.R. and Schwenk, A.J. (Eds.): *Graph Theory, Combinatorics, and Applications*, Vol. 2, Wiley, pp.871–898.

Nedich, A. and Ozdaglar, A. (2007) *Convergence Rate for Consensus with Delays*, LIDS Technical Report 2774, MIT Lab. for Information and Decision Systems.

Nedich, A., Olshevsky, A., Ozdaglar, A. and Tsitsiklis, J.N. (2007) *On Distributed Averaging Algorithms and Quantization Effects*, LIDS Technical Report 2778, MIT, Lab. for Information and Decision Systems.

Olfati-Saber, R. and Murray, R.M. (2004) 'Consensus problems in networks of agents with switching topology and time-delays', *IEEE Transactions on Automatic Control*, Vol. 49, No. 9, pp.1520–1533.

Ren, W., Beard, R.W. and McLain, T.W. (2004) '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, Vol. 309, pp.171–188.

Seneta, E. (1981) *Nonnegative Matrices and Markov Chains*, 2nd ed., Springer-Verlag.

Xiao, L. and Boyd, S. (2003) 'Fast linear iterations for distributed averaging', *Proceedings of 42th IEEE Conference on Decision and Control*, Maui, Hawaii, USA, pp.4997–5002.

Xiao, L., Boyd, S. and Lall, S. (2005) 'A scheme for robust distributed sensor fusion based on average consensus', *Proceedings of International Conference on Information Processing in Sensor Networks*, Los Angeles, California, USA, pp.63–70.

Xie, M. and Haenggi, M. (2005) 'Delay performance of different MAC schemes for multihop wireless networks', *Proc. IEEE Global Communications Conference (GLOBECOM'05)*, St. Louis, Missouri, USA.