

Degree Fluctuations and the Convergence Time of Consensus Algorithms

Alex Olshevsky

John N. Tsitsiklis

Abstract—We consider a consensus algorithm in which every node in a time-varying undirected connected graph assigns equal weight to each of its neighbors. Under the assumption that the degree of any given node is constant in time, we show that the algorithm achieves consensus within a given accuracy ϵ on n nodes in time $O(n^3 \ln(n/\epsilon))$. Because there is a direct relation between consensus algorithms in time-varying environments and inhomogeneous random walks, our result also translates into a general statement on such random walks. Moreover, we give simple proofs that the convergence time becomes exponentially large in the number of nodes n under slight relaxations of the above assumptions. We prove that exponential convergence time is possible for consensus algorithms on fixed directed graphs, and we use an example of Cao, Spielman, and Morse to give a simple argument that the same is possible if the constant degrees assumption is even slightly relaxed.

I. INTRODUCTION

Consensus algorithms are a class of iterative update schemes that are commonly used as building blocks for the design of distributed control laws. Their main advantage is robustness in the presence of time varying environments and unexpected communication link failures. Consensus algorithms have attracted significant interest in a variety of contexts such as distributed optimization [18],[17] or coverage control [12], and many other contexts involving networks in which central control is absent and communication capabilities are time-varying.

While the convergence properties of consensus algorithms in time-varying environments are well understood, much less is known about the corresponding convergence times. An inspection of the classical convergence proofs ([3], [13]) leads to convergence time upper bounds that grow exponentially with the number of nodes. It is then natural to look for conditions under which the convergence time only grows polynomially, and this is the subject of this paper.

In our main result, we show that a consensus algorithm in which every node assigns equal weight to each of its neighbors in an undirected, connected graph (where the graph can be time-varying) has polynomial convergence time if the degree of any given node is constant in time. Because there is a direct relation between consensus algorithms in time-varying environments and nonhomogeneous random walks, our result also translates into a general statement on such random walks.

A. Model, notation, and background

In this subsection, we define our notation, the model of interest, and some background on consensus algorithms.

Given a directed graph G , we will use $N_i(G)$ to denote the set $\{j \mid (i, j) \text{ is an edge}\}$ of direct successors of node i in G , and $d_i(G)$ to denote the cardinality of $N_i(G)$. Given a sequence of directed graphs $G(0), G(1), \dots, G(k-1)$, we will use the simpler notation $N_i(t), d_i(t)$ in place of $N_i(G(t)), d_i(G(t))$, and we will make a similar simplification for other variables of interest.

We are interested in analyzing a consensus algorithm in which a node assigns equal weight to each one of its neighbors. We consider n nodes and assume that at each discrete time t , node i stores a real number $x_i(t)$. We let $x(t) = (x_1(t), \dots, x_n(t))$. For any given sequence of directed graphs $G(0), G(1), G(2), \dots$, and any initial vector $x(0)$, the algorithm is described by the update equation

$$x_i(t+1) = \frac{1}{d_i(t)} \sum_{j \in N_i(t)} x_j(t), \quad i = 1, \dots, n, \quad (1)$$

which can also be written in the form

$$x(t+1) = A(t)x(t),$$

for a suitably defined sequence of matrices $A(0), A(1), \dots, A(t-1)$. The graphs $G(t)$, which appear in the above update rule through $d_i(t)$ and $N_i(t)$, correspond to information flow among the agents; the edge (i, j) is present in $G(t)$ if and only if agent i uses the value $x_j(t)$ of agent j in its update at time t . To reflect the fact that every agent always has access to its own information, we assume that every graph $G(t)$ contains all the self-loops (i, i) . Note that we have $[A(t)]_{ij} > 0$ if and only if (i, j) is an edge in $G(t)$.

It is well known ([18], [13]) that, subject to some natural conditions on the graph sequence, every component of $x(t)$ converges to a common value. In this paper, we focus on the convergence rate of this process in some natural settings. To quantify the progress of the algorithm towards consensus, we will use the function

$$S(x) = \max_i x_i - \min_i x_i.$$

For any $\epsilon > 0$, we will say that a sequence of graphs $G(0), G(1), \dots, G(k-1)$ (alternatively, a sequence of matrices $A(0), A(1), \dots, A(k-1)$) results in ϵ -consensus if $S(x(k)) \leq \epsilon S(x(0))$ for all initial vectors $x(0)$.

We will focus on graph sequences in which every graph $G(t)$ is bidirectional, meaning that if (i, j) is an edge in

Research partially supported by the NSF under grant CMMI-0856063.

Alex Olshevsky is with the Department of Mechanical and Aerospace Engineering, Princeton University, Email: aolshevs@princeton.edu

John N. Tsitsiklis is with the Laboratory for Information and Decision Systems, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology. Email: jnt@mit.edu.

$G(t)$, then so is (j, i) . In practice, graphs that capture information flows are often bidirectional. For example, $G(t)$ is bidirectional if: (i) $G(t)$ contains all the edges between agents that are physically within some distance of each other; (ii) $G(t)$ contains all the edges between agents that have line-of-sight views of each other; (iii) $G(t)$ contains the edges corresponding to pairs of agents that can send messages to each other using a protocol that relies on acknowledgements.

It is an immediate consequence of existing convergence proofs ([3], [13]) that any sequence of $Cn^n \ln(1/\epsilon)$ connected bidirectional graphs, with self-loops at every node, results in ϵ -consensus. Here, C is a constant that does not depend on the problem parameters n and ϵ . We are interested in simple conditions under which the undesirable $O(n^n)$ scaling becomes polynomial in n .

B. Our results

Our contributions are twofold. First, in Section II, we prove our main result.

Theorem 1. *Consider a sequence $G(0), G(1), \dots, G(k-1)$ of connected bidirectional graphs, with self-loops at each node. Suppose, furthermore, that the degree of each node stays constant in time, i.e.,*

$$d_i(t) = d_i(t'), \quad \forall i, t, t'.$$

If the length k of the graph sequence is at least

$$n^3 \ln \frac{2n^2}{\epsilon},$$

then ϵ -consensus is achieved.

To put Theorem 1 in perspective, we note that polynomial convergence times were only known for the cases where:

- (a) the graphs $G(t)$ are the same at each time t , bidirectional, connected, with all self-loops present [14]; or
- (b) in addition to some natural connectivity assumptions, the underlying iteration matrices are doubly stochastic, which, for the consensus algorithms considered here amounts to an assumption that each graph $G(t)$ is regular [16].

Theorem 1 can be viewed as a generalization of the above two results.

In Section III, we give an interpretation of our results in terms of Markov chains. Theorem 1 can be interpreted as providing a sufficient condition for a random walk on a time-varying graph to forget its initial distribution in polynomial time.

In Section IV, we show through examples that relaxing the assumptions of Theorem 1 even slightly can lead to a convergence time which is exponential in n . Specifically, we show the following.

- (i) If we do not require the graphs $G(t)$ to be bidirectional, exponential convergence time is possible, even if the graphs $G(t)$ do not change with time;
- (ii) If we replace the assumption that each $d_i(t)$ is independent of t with the weaker assumption that the

sorted degree sequence (say, in non-increasing order) is independent of t (thus allowing nodes to “swap” degrees), exponential convergence time is possible. While this fact was known (although unpublished) [5], our contribution is to provide a simple proof. The presentation of our counterexamples capitalizes on the correspondence with Markov chains discussed in Section III.

In summary: for connected bidirectional graphs with self-loops, unchanging degrees is a sufficient condition for polynomial time convergence, but relaxing it even slightly by either allowing the nodes to “swap” degrees or by losing link symmetry leads to the possibility of exponential convergence time.

C. Previous work

There is considerable and growing literature on the convergence time of consensus algorithms. We only mention papers that are closest to our own work, omitting references to the literature on various aspects of consensus convergence times that we do not address here, such as topology design, performance in geometric random graphs, etc.

Worst-case upper bounds on the convergence times of consensus algorithms have been established in [8], [6], [7], [1], [2], [11], [10], [9]. The papers [8], [6], [7] considered a setting slightly more general than ours, and established exponential upper bounds. The papers [1], [2] addressed the convergence times of consensus algorithms in terms of spanning trees that capture the information flow between the nodes. It was observed that in several cases this approach produces tight estimates of the convergence times. The works [11], [10], [9] take a geometric approach, and consider the convergence time in a somewhat different model, involving interactions between geographic nearest neighbors. These papers find that the convergence time is quite high (either singly exponential or iterated exponential, depending on the model). The original papers [18], [13] also considered the effect of delays on convergence; some more recent work on this subject may be found in [7] and [4].

Our work differs from these papers in that our convergence time bounds are polynomial in n . To the best of our knowledge, polynomial bounds on the particular consensus algorithm considered in this paper had been derived earlier only in [14] and [16]. Our work encompasses a much wider class of situations than [14], which required the graphs $G(t)$ to be constant in time. Moreover, our results may be viewed as complementary to those in [16], which proved a polynomial convergence time bound for averaging algorithms, involving doubly stochastic matrices.

II. PROOF OF THEOREM 1

As in the statement of Theorem 1, we assume that we are given a sequence of bidirectional connected graphs $G(0), G(1), \dots$, with self-loops at each node, and such that $d_i(t)$ is the same in each $G(t)$. We will thus drop the parameter t and refer to the degree of node i simply as d_i . Observe that $d_i > 0$ due to the connectivity of the graphs $G(t)$.

We will use \mathcal{G} to refer to the class of bidirectional connected graphs with self-loops at every node such that the degree of node i is d_i . We let D be the $n \times n$ diagonal matrix whose i th diagonal entry is d_i . We will use $E'(G)$ to denote the edge set of a graph G . We will sometimes find it convenient to use the notation $E(G)$ to refer to the set of *unordered pairs* (i, j) such that the *ordered pairs* (i, j) and (j, i) belong to $E'(G)$.

Definition 1. We define the inner product $\langle \cdot, \cdot \rangle_d$ by

$$\langle x, y \rangle_d = \sum_{i=1}^n d_i x_i y_i.$$

Note that because $d_i > 0$ for all i , $\langle x, y \rangle_d$ is a valid inner product.

Definition 2. Given a directed graph G , we define the update matrix $A(G)$ by

$$[A(G)]_{ij} = \begin{cases} 1/d_i(G), & \text{if } j \in N_i(G), \\ 0, & \text{otherwise.} \end{cases}$$

We use $A(t)$ as a shorthand for $A(G(t))$, so that Eq. (1) can be written as

$$x(t+1) = A(t)x(t). \quad (2)$$

Conversely, given an update matrix A of the above form, we will use $G(A)$ to denote the graph G whose update matrix is A . We use $N_i(A)$ as shorthand for $N_i(G(A))$; the quantities $d_i(A)$, $E(A)$, and $E'(A)$ are defined similarly. Finally, we use \mathcal{A} to denote the set of update matrices $A(G)$ associated with graphs $G \in \mathcal{G}$.

Note that DA is the adjacency matrix associated with the graph corresponding to A . Given that we restrict to bidirectional graphs, DA is symmetric for every $A \in \mathcal{A}$.

Lemma 1. For any $A \in \mathcal{A}$, we have $\langle x, Ay \rangle_d = \sum_{(i,j) \in E'(A)} x_i y_j$.

Proof: We have

$$\langle x, Ay \rangle_d = x^T D A y = \sum_{i=1}^n \sum_{j \in N_i(A)} x_i y_j = \sum_{(i,j) \in E'(A)} x_i y_j$$

■

Lemma 2. Each $A \in \mathcal{A}$ is self-adjoint with respect to the inner product $\langle \cdot, \cdot \rangle_d$.

Proof: Using the fact that DA is symmetric and D is diagonal, we have

$$\langle x, Ay \rangle_d = x^T D A y = x^T (DA)^T y = x^T A^T D y = \langle Ax, y \rangle_d.$$

■

Lemma 2 is the reason for introducing the inner product $\langle \cdot, \cdot \rangle_d$. The fact that matrices in the set \mathcal{A} are self-adjoint plays a central role in the analysis of the algorithm (2). One of its consequences is that matrices in \mathcal{A} have real eigenvalues. We use the notation $\lambda_i(A)$ to denote the i th largest eigenvalue of a matrix $A \in \mathcal{A}$. Note that every $A \in \mathcal{A}$ is a stochastic matrix, and therefore $\lambda_1(A) = 1$.

Next, we identify a weighted average that is preserved by the iteration $x(t+1) = A(t)x(t)$. For any x , we let

$$\bar{x} = \frac{\langle x, \mathbf{1} \rangle_d}{\langle \mathbf{1}, \mathbf{1} \rangle_d}.$$

We observe that for any $A \in \mathcal{A}$,

$$\langle x, \mathbf{1} \rangle_d = x^T D \mathbf{1} = x^T D A \mathbf{1} = x^T A^T D \mathbf{1} = \langle Ax, \mathbf{1} \rangle_d,$$

where the second equality used the fact that A is a stochastic matrix. Therefore,

$$\overline{Ax} = \bar{x}, \quad \forall A \in \mathcal{A}.$$

With these preliminaries in place, we now proceed to the main part of our analysis, which is based on the Lyapunov function

$$V(x) = \langle x - \bar{x}\mathbf{1}, x - \bar{x}\mathbf{1} \rangle_d = \sum_{i=1}^n d_i (x_i - \bar{x})^2.$$

The next lemma quantifies the decrease of $V(\cdot)$ when a vector x is multiplied by some matrix $A \in \mathcal{A}$.

Lemma 3. For any $A \in \mathcal{A}$ and any vector x , we have

$$V(Ax) \leq \lambda_2(A^2)V(x).$$

Proof: Fix some $A \in \mathcal{A}$ and some vector x . Since $\overline{Ax} = \bar{x}$, the lemma asserts that

$$\langle Ax - \bar{x}\mathbf{1}, Ax - \bar{x}\mathbf{1} \rangle_d \leq \lambda_2(A^2) \langle x - \bar{x}\mathbf{1}, x - \bar{x}\mathbf{1} \rangle_d.$$

Let $y = x - \bar{x}\mathbf{1}$, and note that $\langle y, \mathbf{1} \rangle_d = 0$. It therefore suffices to show that

$$\langle Ay, Ay \rangle_d \leq \lambda_2(A^2) \langle y, y \rangle_d,$$

for all y with $\langle y, \mathbf{1} \rangle_d = 0$; or, equivalently, that

$$\frac{\langle y, y \rangle_d - \langle Ay, Ay \rangle_d}{\langle y, y \rangle_d} \geq 1 - \lambda_2^2(A), \quad \text{if } \langle y, \mathbf{1} \rangle_d = 0 \text{ and } y \neq 0. \quad (3)$$

To establish Eq. (3), we note that

$$\begin{aligned} \min_{\langle y, \mathbf{1} \rangle_d = 0} \frac{\langle y, y \rangle_d - \langle Ay, Ay \rangle_d}{\langle y, y \rangle_d} &= \min_{\langle y, \mathbf{1} \rangle_d = 0} \frac{\langle y, y \rangle_d - \langle y, A^2 y \rangle_d}{\langle y, y \rangle_d} \\ &= \min_{\langle y, \mathbf{1} \rangle_d = 0} \frac{\langle y, (I - A^2)y \rangle_d}{\langle y, y \rangle_d}, \end{aligned}$$

where in the minima above we only consider nonzero vectors y ; note that these minima are attained — it suffices to consider vectors y on the unit ball, a compact set. Now, observe that $I - A^2$ is also self-adjoint under the inner product $\langle \cdot, \cdot \rangle_d$. Moreover, since A (and, therefore, A^2 as well) is stochastic, the smallest eigenvalue of $I - A^2$ is 0, with an associated eigenvector of $\mathbf{1}$. Consequently, by the Courant-Fischer variational characterization of eigenvalues, the expression on the right is the second smallest eigenvalue of $I - A^2$:

$$\min_{\langle y, \mathbf{1} \rangle_d = 0} \frac{\langle y, y \rangle_d - \langle Ay, Ay \rangle_d}{\langle y, y \rangle_d} = \lambda_{n-1}(I - A^2) = 1 - \lambda_2(A^2),$$

which concludes the proof. \blacksquare

Thus, to bound how much $V(x)$ decreases at each step, it suffices to obtain an upper bound on $\lambda_2(A^2)$, for matrices $A \in \mathcal{A}$. This can be done using the next lemma, which is the main result of [14]. We include a short proof for completeness.

Lemma 4. *Let $A \in \mathcal{A}$, and let ℓ be the diameter of the graph $G(A)$. Then,*

$$\sqrt{\lambda_2(A^2)} = \max\{|\lambda_n(A)|, \lambda_2(A)\} \leq 1 - \frac{1}{nd_{\max}\ell},$$

where d_{\max} is the largest of the degrees d_i .

Proof: The first inequality follows because the eigenvalues of A^2 are the squares of the eigenvalues of A . For the second inequality, using again the Courant-Fisher characterization, and some easy algebra, we have

$$\begin{aligned} \lambda_2(A) &= \max_{\substack{\langle x, \mathbf{1} \rangle_d = 0 \\ \langle x, x \rangle_d = 1}} \langle x, Ax \rangle_d = \max_{\substack{\langle x, \mathbf{1} \rangle_d = 0 \\ \langle x, x \rangle_d = 1}} \sum_{(i,j) \in E'(A)} x_i x_j \\ &= 1 - \min_{\substack{\langle x, \mathbf{1} \rangle_d = 0 \\ \langle x, x \rangle_d = 1}} \sum_{(i,j) \in E(A)} (x_i - x_j)^2. \end{aligned}$$

Thus, it suffices to show that

$$\min_{\substack{\langle x, \mathbf{1} \rangle_d = 0 \\ \langle x, x \rangle_d = 1}} \sum_{(i,j) \in E(A)} (x_i - x_j)^2 \geq \frac{1}{nd_{\max}\ell}$$

Towards this purpose, we carry out a variation of an argument first used in [14]. Fix some x that satisfies $\langle x, \mathbf{1} \rangle_d = 0$ and $\langle x, x \rangle_d = 1$. Without loss of generality, we assume that (i) node 1 has the largest value of $d_i x_i^2$, (ii) node k has the smallest value of x_i , and (iii) the shortest path from node 1 to k is $(1, 2), (2, 3), \dots, (k-1, k)$. The condition $\langle x, x \rangle_d = 1$ implies that $d_1 x_1^2 \geq 1/n$, and consequently that $x_1 \geq 1/\sqrt{nd_{\max}}$; the requirement $\langle x, \mathbf{1} \rangle_d = 0$ implies that $x_k < 0$. Thus, $x_1 - x_k \geq 1/\sqrt{nd_{\max}}$, which we write as

$$(x_1 - x_2) + (x_2 - x_3) + \dots + (x_{k-1} - x_k) \geq \frac{1}{\sqrt{nd_{\max}}}.$$

Applying the Cauchy-Schwarz inequality, we get

$$k \sum_{i=1}^{k-1} (x_i - x_{i+1})^2 \geq \frac{1}{nd_{\max}}.$$

We then use the fact that $k \leq \ell$, to obtain the claimed bound on $\lambda_2(A)$.

As for λ_n , we observe that the diagonal entries of A are at least $1/n$ and the row sums are 1. The Gershgorin circle theorem immediately gives $\lambda_n \geq -1 + 1/n$, which is stronger than the bound we have claimed. \blacksquare

We can now complete the proof of Theorem 1. Lemma 3 describes the decrease in the variance $V(x(t))$ in terms of $\lambda_2(A^2(t))$, and Lemma 4 gives us a way to upper bound the latter quantity.

Proof of Theorem 1: Using Lemmas 3 and 4, and the bounds $d_{\max} \leq n$, $\ell \leq n$, we see that for every $A \in \mathcal{A}$, and

every x , we have

$$\begin{aligned} V(Ax) &\leq \lambda_2(A^2)V(x) = \max\{|\lambda_n(A)|^2, \lambda_2^2(A)\}V(x) \\ &\leq \left(1 - \frac{1}{n^3}\right)^2 V(x). \end{aligned}$$

Because the definition of ϵ -consensus is in terms of $S(x)$ rather than $V(x)$, we need to relate these two quantities. On the one hand, for every x , we have

$$V(x) = \sum_{i=1}^n d_i (x_i - \bar{x})^2 \leq n \sum_{i=1}^n (x_i - \bar{x})^2 \leq n^2 S^2(x).$$

On the other hand, for every x , we have

$$V(x) \geq \max_i (x_i - \bar{x})^2 \geq \frac{1}{4} (\max_i x_i - \min_i x_i)^2 = \frac{1}{4} S^2(x).$$

Suppose that $t \geq n^3 \ln(2n^2/\epsilon)$. Then,

$$\begin{aligned} S(x(t)) &\leq \sqrt{4V(x(t))} \leq 2 \left(1 - \frac{1}{n^3}\right)^{n^3 \ln(2n^2/\epsilon)} V(x(0)) \\ &\leq 2n^2 \left(1 - \frac{1}{n^3}\right)^{n^3 \ln(2n^2/\epsilon)} S(x(0)) \\ &\leq 2n^2 e^{-\ln(2n^2/\epsilon)} S(x(0)) = \epsilon S(x(0)). \end{aligned}$$

(We have used here the inequality $(1 - c/n)^n \leq e^{-c}$, for $c > 0$.) \blacksquare

III. MARKOV CHAIN INTERPRETATION

In this section, we give an alternative interpretation of the convergence time of a consensus algorithm in terms of inhomogeneous Markov chains. In the next section, we will use this interpretation to give some examples of graph sequences that do not satisfy Theorem 1 and which have exponentially large convergence times.

We consider an inhomogeneous Markov chain whose transition probability matrix at time k is $A(k)$. We fix some time t and define

$$P = A(0)A(1) \dots A(t-1).$$

This is the associated t -step transition probability matrix: the ij -th entry of P , denoted by p_{ij} , is the probability that the state at time t is j , given that the initial state is i . Let \mathbf{p}_i be the vector whose k th component is p_{ik} ; thus \mathbf{p}_i^T is the i th row of P .

We address a question which is generic in the study of Markov chains, namely, whether the chain eventually “forgets” its initial state, i.e., whether for all i, j , $\mathbf{p}_i - \mathbf{p}_j$ converges to zero as t increases, and if so, at what rate. We will say that the sequence of matrices $A(0), A(1), \dots, A(t-1)$ is ϵ -forgetful if for all i, j , we have

$$\frac{1}{2} \sum_k |p_{ik} - p_{jk}| \leq \epsilon.$$

The above quantity, $\frac{1}{2} \max_{i,j} \|\mathbf{p}_i - \mathbf{p}_j\|_1$ is known as the *coefficient of ergodicity* of the matrix P , and appears often in the study of consensus algorithms (see, for example, [8]).

The matrix P can also be interpreted in terms of consensus updates: an initial vector x is multiplied by the matrices $A(t-1)$

1), $A(t-2), \dots, A(0)$, to produce the vector Px ; note that the different matrices are now applied in the reverse order. The result that follows relates the times to achieve ϵ -consensus or ϵ -forgetfulness, and is essentially the same as Proposition 4.5 of [15].

Proposition 2. *The sequence of matrices $A(0), A(1), \dots, A(t-1)$ is ϵ -forgetful if and only if the sequence of matrices $A(t-1), A(t-2), \dots, A(0)$ results in ϵ -consensus (i.e., $S(Px) \leq \epsilon S(x)$, for every vector x .)*

Proof: Suppose that the matrix sequence $A(0), A(1), \dots, A(t-1)$ is ϵ -forgetful, i.e., that $\frac{1}{2} \sum_k |p_{ik} - p_{jk}| \leq \epsilon$, for all i and j . Given a vector x , let $c = (\max_k x_k + \min_k x_k)/2$. Note that $\|x - c\mathbf{1}\|_\infty = (\max_k x_k - \min_k x_k)/2 = S(x)/2$. We then have

$$\begin{aligned} |[Px]_i - [Px]_j| &= \left| \sum_k (p_{ik} - p_{jk})x_k \right| \\ &= \left| \sum_k (p_{ik} - p_{jk})(x_k - c) \right| \\ &\leq \|\mathbf{p}_i - \mathbf{p}_j\|_1 \cdot \|x - c\mathbf{1}\|_\infty \leq 2\epsilon S(x)/2 = \epsilon S(x). \end{aligned}$$

Since this is true for every i and j , we obtain $S(Px) \leq \epsilon S(x)$, and the sequence $A(t-1), A(t-2), \dots, A(0)$ results in ϵ -consensus.

Conversely, suppose that the sequence of matrices $A(t-1), A(t-2), \dots, A(0)$ results in ϵ -consensus. Fix some i and j . Let x be a vector whose k th component is $1/2$ if $p_{ik} \geq p_{jk}$ and $-1/2$ otherwise. Note that $S(x) = 1$. We have

$$\frac{1}{2} \|\mathbf{p}_i - \mathbf{p}_j\|_1 = (\mathbf{p}_i^T - \mathbf{p}_j^T)x = [Px]_i - [Px]_j \leq \epsilon S(x) = \epsilon,$$

where the last inequality made use of the ϵ -consensus assumption. Thus, the sequence of matrices $A(0), A(1), \dots, A(t-1)$ is ϵ -forgetful. ■

We will use Proposition 2 for the special case of Markov chains that are random walks. Given a directed graph $G(t)$, we let, as before, $N_i(t)$ be the set of nodes j for which the edge (i, j) is present. In the random walk associated with this graph, if the state at time t is i , the state at time $t+1$ is chosen to be one of the elements of $N_i(t)$, with equal probability. We let $A(t)$ be the associated transition probability matrix. We will say that a sequence of graphs is ϵ -forgetful whenever the corresponding sequence of transition probability matrices is ϵ -forgetful. Proposition 2 allows us to reinterpret Theorem 1 as follows: random walks on time-varying bidirectional connected graphs with self-loops forget their initial distribution in a polynomial number of steps.

Proposition 2 also has a corollary which we will use later. It is based on the following observation: concatenating two sequence of graphs, each of which achieves ϵ -consensus, results in a sequence which achieves ϵ^2 -consensus.

Corollary 3. *Suppose that a sequence of graphs is ϵ -forgetful. Then, concatenating this sequence with itself k times results in a sequence which is ϵ^k -forgetful.*

IV. SOME COUNTEREXAMPLES

A. *The bidirectionality requirement in Theorem 1 cannot be relaxed*




Fig. 1. The graph used in Proposition 4.

In this subsection, we show that it is impossible to drop the assumption that the graph is bidirectional, even when the graph does not change with time.

Proposition 4. *Let G be the graph shown in Figure 1. Consider the graph sequence consisting of G , repeated t times. For this graph sequence to result in $(1/8)$ -consensus, we must have $t \geq 2^{n/2}/6$.*

Proof: Suppose that this graph sequence of length t results in ϵ -consensus. By Proposition 2, it is $(1/8)$ -forgetful. In particular, $|p_{11} - p_{n1}| \leq 1/4$ (recall that p_{ij} stands for the t -step transition probability from i to j).

For $i \leq n/2$, let T_i be the first time that a random walk that starts at state i visits the bottom part of the graph, and let δ_i be the probability that T_i is less than or equal to t . Note that conditional on starting at 1 and never transitioning to the bottom half, the probability of being at state 1 at time t is $1/2$; thus, $p_{11} \geq \frac{1}{2}(1 - \delta_1)$. Furthermore, $p_{n1} \leq \delta_1$ by symmetry. Therefore,

$$\frac{1}{4} \geq |p_{11} - p_{n1}| \geq \frac{1}{2}(1 - \delta_1) - \delta_1 = \frac{1}{2} - \frac{3}{2}\delta_1,$$

which yields $\delta_1 \geq 1/6$.

Using a straightforward coupling argument, we have $\delta_i \geq \delta_1 \geq 1/6$, for $i = 2, \dots, n/2$. By viewing periods of length t as a single attempt to get to the bottom half of the graph, with each attempt having probability at least $1/6$ to succeed, we conclude that $E[T_1] \leq 6t$.

On the other hand, T_1 is just the first time until the walk moves to the right $n/2$ consecutive times. Each time that the random walk returns to state 1, we have a new trial with probability of success equal to $2^{-n/2}$. Thus, $2^{n/2} \leq E[T_1] \leq 6t$, which yields the desired result. ■

Remark: It is not hard to see that if we add a self-loop to each node in Figure 1, the result holds with minor modifications.




Fig. 2. The top-left figure shows graph $G(0)$; top-right shows $G(1)$; bottom-left shows $G((n/2) - 2)$; bottom-right shows $G((n/2) - 1)$. As these figures illustrate, $G(t + 1)$ is obtained by applying a circular shift to each half of $G(t)$. All edges are bidirectional, and every node has a self-loop which is not shown. For aesthetic reasons, instead of labeling the nodes as $1, \dots, n$, we label them with $1, \dots, n/2$ and $1', \dots, (n/2)'$.

B. The unchanging degrees condition in Theorem 1 cannot be relaxed

In this subsection, we show that it is impossible to relax the condition of unchanging degrees in Theorem 1. In particular, if we only impose the slightly weaker condition that the sorted degree sequence (the non-increasing list of node degrees) does not change with time, the time to achieve ϵ -consensus can grow exponentially with n . This is an unpublished result of Cao, Spielman, and Morse [5]; we provide here a simple proof.

Proposition 5. *Let n be even and let t be an integer multiple of $n/2$. Consider the graph sequence of length $t = kn/2$, consisting of periodic repetitions of the reversal of the length- $n/2$ sequence described in Figure 2. For this graph sequence to result in $(1/4)$ -consensus, we must have $t \geq 2^{(n/2)}/8$.*

Proof: Suppose that this graph sequence of length t results in $(1/4)$ -consensus. Then Proposition 2 implies that the sequence of length $kn/2$ consisting of periodic repetitions of the length $n/2$ sequence described in Figure 2 is $(1/4)$ -forgetful. Let p_{ij} be the associated t -step transition probabilities.

Let T be the time that it takes for a random walk that starts at state $n/2$ to cross into the right-hand side part of the graph. Let δ be the probability that T is less than or equal to t . Let $R = \{1' \dots (n/2)'\}$. We have $\sum_{j' \in R} p_{(n/2), j'} \leq \delta$ and, using symmetry, $\sum_{j' \in R} p_{(n/2)', j'} \geq 1 - \delta$. Using the fact that the graph sequence is $(1/4)$ -forgetful in the first inequality below, we have

$$\begin{aligned} \frac{1}{2} &\geq \sum_{j' \in R} |p_{(n/2)', j'} - p_{(n/2), j'}| \\ &\geq \sum_{j' \in R} p_{(n/2)', j'} - \sum_{j' \in R} p_{(n/2), j'} \\ &\geq (1 - \delta) - \delta = 1 - 2\delta, \end{aligned}$$

which yields $\delta \geq 1/4$. Arguing as in the proof of Proposition 4 (node $n/2$ is the least favorable “non-central” starting state in the left-hade side of the graph), we obtain $E[T] \leq 4t$.

Note that for a walk that starts at state $n/2$ to cross into the right-hand side part of the graph, it must first take a self-loop $(n/2) - 1$ consecutive times. Consequently, $2^{(n/2)-1} \leq$

$E[T] \leq 4t$, which yields the desired result. \blacksquare

V. CONCLUSIONS

The main contribution of this paper is Theorem 1, which shows that consensus algorithms converge in polynomial time for a large class of graph sequences. We also gave simple proofs showing that this finding is fragile, and even slight relaxations of the hypotheses cause the conclusion to fail.

A similar result is available for the case of doubly stochastic update matrices $A(t)$, and our result can be viewed as complementary [16]. Interestingly, both results rely on a suitable quadratic Lyapunov function as well as on the fact that all update matrices share a common left eigenvector.

REFERENCES

- [1] D. Angeli, P.-A. Bliman, “Tight estimates for convergence of some non-stationary consensus algorithms,” *Systems and Control Letters*, vol. 57, no. 12, pp. 996-1004, 2008.
- [2] D. Angeli, P.-A. Bliman, “Convergence speed of unsteady distributed consensus: decay estimate along the settling spanning-trees,” *SIAM Journal on Control and Optimization*, vol. 48, no. 1, pp. 1-32, 2009.
- [3] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*, Prentice Hall, 1989.
- [4] P.-A. Bliman, A. Nedic, A. Ozdaglar, “Rate of convergence for consensus with delays,” *Proceedings of the 47th IEEE Conference on Decision and Control*, Cancun, Mexico, Dec. 2008.
- [5] Ming Cao, personal communication.
- [6] M. Cao, A. S. Morse, B. D. O. Anderson, “Reaching a consensus in a dynamically changing environment: a graphical approach,” *SIAM Journal on Control and Optimization*, vol. 47, no. 2, pp. 575-600, 2008.
- [7] M. Cao, A. S. Morse, B. D. O. Anderson, “Reaching a consensus in a dynamically changing environment: convergence rates, measurement delays, and asynchronous events,” *SIAM Journal on Control and Optimization*, vol. 47, no. 2, pp. 575-600, 2008.
- [8] M. Cao, D. Spielman, A. S. Morse, “A lower bound on convergence of a distributed network consensus algorithm,” *Proceedings of the 44th IEEE Conference on Decision and Control and European Control Conference*, Madrid, Spain, Dec. 2005.
- [9] B. Chazelle, “Analytical tools for natural algorithms,” *Proceedings of the First Symposium on Innovations in Computer Science*, Beijing, China, Jan. 2010.
- [10] B. Chazelle, “The convergence of bird flocking,” *Proceedings of the 26th Annual Symposium on Computational Geometry*, Snowbird, USA, Jun. 2010.
- [11] B. Chazelle, “Natural algorithms,” *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, New York, USA, Jan. 2009.
- [12] C. Gao, J. Cortes, F. Bullo, “Notes on averaging over acyclic graphs and discrete coverage control,” *Automatica*, vol. 44, no. 8, pp. 2120-2127, 2008.
- [13] A. Jadbabaie, J. Lin, and A. S. Morse, “Coordination of groups of mobile autonomous agents using nearest neighbor rules,” *IEEE Transactions on Automatic Control*, vol. 48, no. 3, pp. 988-1001, 2003.
- [14] H. J. Landau, A. M. Odlyzko, “Bounds for eigenvalues of certain stochastic matrices,” *Linear Algebra and Its Applications*, vol. 38, pp. 5-15, 1981.
- [15] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov Chains and Mixing Times*, American Mathematical Society, 2008.
- [16] A. Nedic, A. Olshevsky, A. Ozdaglar, and J. N. Tsitsiklis, “On distributed averaging algorithms and quantization effects,” *IEEE Transactions on Automatic Control*, vol. 54, no. 11, pp. 2506-2517, 2009.
- [17] A. Nedic, A. Ozdaglar, “Distributed subgradient methods for multi-agent optimization,” *IEEE Transactions on Automatic Control*, vol. 54, no. 1, pp. 48-61, 2009.
- [18] J. N. Tsitsiklis, D. P. Bertsekas, and M. Athans, “Distributed asynchronous deterministic and stochastic gradient optimization algorithms,” *IEEE Transactions on Automatic Control*, vol. 31, no. 9, 1986, pp. 803-812.