Distributed Averaging in Dynamic Networks

Shreevatsa Rajagopalan and Devavrat Shah

Abstract—**The question of computing average of numbers present at nodes in a network in a distributed manner using gossip or message-passing algorithms has been of great recent interest across disciplines—algorithms, control and robotics, estimation, social networks, etc. It has served as a non-trivial, representative model for an important class of questions arising in these disciplines and thus guiding intellectual progress over the past few decades. In most of these applications, there is inherent dynamics present, such as changes in the network topology in terms of communication links, changes in the values of numbers present at nodes, and nodes joining or leaving. The effect of dynamics in terms of communication links on the design and analysis of algorithms for averaging is reasonably well understood, e.g., J. Tsitsiklis, "Problems in decentralized decision making and computation," Ph.D. dissertation, MIT, 1984; D. Bertsekas *et al.*, "Convergence theories of distributed iterative processes: A survey," MIT, 1983; A. Jadbabaie *et al.*, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Trans. on Automatic Control*, vol. 48, no. 6, pp. 988–1001, 2003; S. Boyd *et al.*, "Gossip algorithms: Design, analysis and applications," *Proc. IEEE INFOCOM*, 2005, pp. 1653–1664. However, little is known about the effect of other forms of dynamics. In this paper, we study the effect of such types of dynamics in the context of maintaining average in the network. Specifically, we design *dynamics-aware* message-passing or gossip algorithm that maintains good estimate of average in presence of continuous change in numbers at nodes. Clearly, in presence of such dynamics the best one can hope for is a tradeoff between the accuracy of each node's estimate of the average at each time instant and the rate of dynamics. For our algorithm, we characterize this tradeoff and establish it to be *near* optimal. The dependence of the accuracy of the algorithm on the rate of dynamics as well as on the underlying graph structure is quantified.**

Index Terms—**Averaging, consensus, dynamic, gossip.**

I. INTRODUCTION

DYNAMICS is inherent to any networked control and processing system. Typically, the observed or sensed state of the system is summarized in form of parameters that change over time. To achieve desired functionality of the system, it is necessary to continually or periodically evaluate or compute some function of these parameters. For instance, in a communication network such as the Internet, the problem of *routing* is to decide routes based on network parameters such as the load on the links, the connectivity of the network, etc.; popular Border Gateway Protocol (BGP) for routing does precisely the same [5], i.e., to update routes continually based on measured link loads, connectivity, etc. More generally, to share network resources in a generic constrained data network modeled as a queueing network, a theoretically well-accepted algorithm, known as the maximum weight or pressure [6], makes decision at each time instant based on the network state summarized via queue-sizes. In summary, the generic problem of interest is of the following form: the system has external input which is observed, and an appropriate algorithm is devised by a "controller" that often performs some form of optimization. Subsequently "control" or "action" is fed back into the system, and the goal is to keep the system "stable"—close to some "desired" state. In a networked setup, this task has to be performed in a decentralized manner by means of simple algorithm. This is because of the need of scalability as well as engineering constraints. This has led to the study of simple, distributed or gossip or message-passing algorithms across disciplines in the recent years.

Most of the known results implicitly or explicitly assume a certain *time-scale separation* in the system dynamics: the feedback is assumed to be *instantaneous* or equivalently, the algorithm is assumed to be executed at a *much faster* timescale than the timescale at which changes happen in the system state. Thus, while algorithm designed with such a timescale separation assumption works well when the system is static or slowly changing, it may not work properly when the system is dynamic, which is the case in practice. For example, this is the primary conceptual reason why routing protocols like the BGP suffer from instability: BGP updates routing tables to reflect changes in the network, but the network state changes at the same timescale and subsequently, this results in the instability of BGP, known as *route flapping*—a major limitation of existing Internet architecture [7].

Therefore, it would be ideal to have algorithms that can deal with the dynamics. It is to be expected that robust algorithms come at a cost in terms of performance—algorithm becomes resistant to changes in the system by trading off exactness in performance. That is, it is reasonable to expect a trade-off between the accuracy of the algorithm and the degree of system dynamics that it can handle. Therefore, the goal is to design robust algorithms that are "dynamics-aware." An ideal algorithm would work across a range of dynamics, adapting itself according to what is feasible. And it would achieve as good a tradeoff between dynamics and performance accuracy as possible. Now roughly speaking, there are three categories of dynamics: highly dynamic scenario, moderately dynamic scenario, and static or slowly changing scenario. Goal is to investigate performance of

averaging algorithms in this entire range of scenarios and understand the interplay between dynamics, network structure, and performance.

A. Related Work

Tsitsiklis [1] initiated the inquiry of interplay between network dynamics and performance of decentralized algorithm in the context of networked control. Specifically, in [1] the question of reaching *consensus* among a collection agents is studied. The algorithm proposed performs linear iterative averaging. The effect of network dynamics in terms of link connectivity on performance of algorithm was studied. This seemingly simple task of (weighted) averaging is used as a "subroutine" for performing estimation as well as solving a class of optimization problems in a distributed manner. An interested reader is referred to book by Bertsekas and Tsitsiklis [8]. In the above-mentioned work as well as in the more recent works (see Jadbabaie, Lin, and Morse [3], Blondel *et al.* [9]), the dynamics in terms of the network link connectivity or topology was addressed in a bounded adversarial setup. In a nutshell, the above results established that asymptotically, the algorithm will find the correct average at all nodes (or consensus will be reached) if and only if the network, formed by links that are present infinitely often over time, is connected; the rate of convergence depends on the "how often" each link is present. Similar concerns and many more applications have brought this question to *life* in recent years—for example, work by Boyd *et al.* [4] on gossip algorithms considers effect of such topological dynamics in a randomized model inspired by peer-to-peer networks. All of these models consider dynamics in terms of the topology, but *not* in terms of 1) the values at the nodes, or 2) node arrival and departure. In this paper, we shall focus on the effect of dynamics of type 1) and 2) on the performance of algorithm.

Before we proceed towards the content of this paper, it is worth noting similar aspects of dynamics captured in online and streaming algorithms as well as stochastic networks. Now online and streaming algorithms can be viewed as setup in which the state of the system is gradually revealed. However, it is different from the setup where the system state is continually changing with time. Generically, such algorithms are studied in centralized setup and hence the network structure does not play any role in determining performance. In the context of stochastic queueing networks, dynamics is inherent. Here, the primary performance goal is the stochastic stability. This is an asymptotic and it usually does not quantify the precise tradeoffs between performance and dynamics that can be achieved by designing dynamics-aware algorithms.

We take note of recent survey by Dimakis *et al.* [10] on the state of art of consensus and gossip algorithm, a recent monograph of Gossip algorithms by Shah [11] and work on dynamic linear estimation in sensor networks by Kar and Moura [12].

B. Our Contributions

In this paper, we address the question of designing dynamics aware algorithm for the problem of estimating the sum (or equivalently average) in a network, where the values of nodes are changing continually with time. This problem is chosen because it is the simplest nontrivial problem. As noted above,

it is also a well-studied problem and of great interest in control theory, computer science, and networks, both independently and as well as a subroutine for many applications such as linear estimation, consensus, projections for dimensionality reduction, maintaining "sketches," computing gradients in convex optimization, etc. An interested reader is referred to a recent monography by Shah [11].

In order to study this problem, design algorithms for it, and analyze their performance, we introduce a natural model for the dynamics in terms of changes in the values or numbers present at nodes in the network. These changes may involve either high dynamics or moderate dynamics. We model the former as *multiplicative changes* in the values, and the latter as *additive changes*. As the main result, we design algorithms for both of these scenarios that are dynamics-aware and analyze the tradeoff it induces in terms of accuracy of estimation and the rate of dynamics.

For the high or fast dynamics, modeled as "multiplicative changes," we design an algorithm using extremal properties of the exponential distribution, relation between exponential random variables of different rates and a novel distributed algorithm to maintain minimum of numbers in a network in dynamic setup. This algorithm can be viewed as a (nontrivial) generalization of the algorithm by Mosk-Aoyama and Shah [13] to compute summation in a static network. Specifically, the algorithm of [13] does not extend to dynamic setup readily for two reasons: 1) the algorithm of [13] requires estimation of minimum of numbers in the network and the obvious scheme used does not work in dynamic setup (see Section I-C for details); 2) the natural extension of algorithm of [13] would require drawing new random numbers every time nodes change their values. This will lead to a situation where some form of "time-stamp" would be required associated with each random number and destroying elegance (as well as distributed property) of the algorithm. Our algorithm is presented in Section III and its near optimality properties are stated in Theorems 1 and 2.

For the moderate or slow dynamics, modeled as "additive changes," we study the property of the known linear iterative averaging algorithm. We find that the error in the estimation in such a scenario is bounded in terms of the spectral gap of the communication matrix. The precise result is stated in Theorem 3.

In both cases, the accuracy of the algorithm depends on the "rate of dynamics" and the graph topology. For multiplicative changes, the network diameter and for additive changes, the spectral gap of the "averaging matrix" affects the accuracy.

C. Dynamic Minimum

As mentioned earlier, an important hurdle that one needs to overcome to adapt algorithm by [13] in the presence of multiplicative changes, is to estimate minimum of these dynamic numbers in the network in a distributed manner. While this question of "dynamically computing minimum" is not the main result of this paper, it is an interesting byproduct contains the essence of the challenges encountered in designing algorithms in presence of dynamics. Therefore, we shall describe the problem and challenges involved here. Appropriate algorithm is described in Section III-A.

Now, the problem. Given a network with connectivity graph $G = (V, E)$, value $Y_v(t)$ at node $v \in V$ that changes with time which is indexed by $t \in \{0, 1, 2, \ldots\}$. We wish to maintain $Y_{\min}(t) = \min_{v \in V} Y_v(t)$ at each node using simple, distributed algorithm. That is, an algorithm that can only maintain limited data structure (does not scale with network size $|V|$), communicate with its neighbors as per G and preferably does not utilized global network structure.

Now when $Y_v(t) = Y_v$ for all t, i.e., values at nodes does not change, an obvious algorithm for this is as follows: each node $v \in V$ maintains estimate $\tilde{Y}_v$ with initially $\tilde{Y}_v = Y_v$; at each subsequent time step, it updates them as $\tilde{Y}_v = \min(\tilde{Y}_v, \min_{u \in N(v)} \tilde{Y}_u)$, where $N(v) = \{u \in V : (u, v) \in E\}$ denotes the set of neighbors of v. If G is connected, then within diameter many steps, each node learns the precise value of the minimum.

However, this simple algorithm does not work when the values are changing. The reason is that once the values $\tilde{Y}_v(t)$ all become small (e.g., if all $Y_v(t)$ remain 0 for some time), then, even if the $Y_v(t)$ increase unboundedly, the small estimates $\tilde{Y}_v$ will continue to remain small! In essence, this approach forever "remembers" the lowest value ever attained. Conceptually, the algorithm needs to "forget" the history quickly enough. At the same time, since the algorithm is distributed, it cannot "forget" it too quickly or else it may not even be good in the static case! In other words, an appropriate algorithm must learn *discounting* of old information, while not entirely discarding it. This is a general problem for online algorithms dealing with dynamic systems. For instance, consider the updating rule (used by, e.g., Kalman filtering) which maintains a model $y(t)$, and, on obtaining new information $n(t)$ at time t, updates $y(t + 1) = \alpha y(t) + (1 - \alpha) n(t)$ for some α. In this case, we have $y(t + 1) = \sum_{k \geq 0} \alpha^k (1 - \alpha) n(t - k)$. By doing so, the effects of older $n(s)$, $s < t$, are dampened; for instance the contributions of $\{n(s) \text{ for } s \leq t - 1/(1 - \alpha)\}$ can effectively be ignored. Of course, in our setup finding such proper α requires knowledge of entire network structure and such property α may not even exist in our setup.

Indeed, there is an elegant (and different) way to resolve this question as explained in Section III-A. The algorithm presented maintains an estimate of the minimum to within an accuracy that depends on the diameter of the network and it is (order) optimal. In general, devising similar function-dependent discounting procedures for other functions would be an interesting direction of future research.

II. MODEL AND RESULTS

This section describes setup and problem statement followed by main results.

A. Setup

We have a network with an underlying connectivity graph $G = (V, E)$ that has $n = |V|$ nodes. We shall assume that network graph G is connected and has D as its diameter. At each node v, there are variables $X_v(t)$ taking non-negative real values that depend on time $t \geq 0$ with initially $t = 0$ and $X_v(0) = 1$ for all $v \in V$. We shall consider deterministic

communication with synchronous time: time is discrete with $t \in \{0, 1, 2, \ldots, \}$ and communication is synchronous. At each time t, each node can communicate with all of its neighbors, and exchange one number[1] per unit time. That is, if $(u, v) \in E$ then nodes u and v can send a number to each other. The values of the nodes change over time and at most once each discrete time. Specifically, values at nodes change in a bounded manner either additively or multiplicatively: let $\delta > 0$ be fixed and given.

1) Additive change: for any $v \in V$, $|X_v(t + 1) - X_v(t) \leq \delta|$ for any $t \geq 0$.
2) Multiplicative change: for any $v \in V$, $e^{-\delta} \leq (X(t + 1))/(X(t)) \leq e^{\delta}$ for any $t \geq 0$.

It is worth noting here that the above model requires all changes to be uniformly bounded. As the reader will notice, the algorithm for multiplicative change even uses the knowledge of this bound explicitly to obtain desired result. We strongly believe that these worst case bounds can be replaced by stochastic bounds on changes in value (both in model and the algorithm). Extending results for more generic dynamic setup remains important direction going forward.

B. Problem Statement

The goal is to estimate at each node v, using a distributed (message-passing) algorithm, a certain function of these values $X_u(t)$, namely the sum $\sum_{u \in V} X_u(t)$ (or equivalently when n is known, the average), and maintain these estimates as the values change with time. Ideally, one wishes to minimize the error in estimation over all times.

C. Results

The results are described separately for multiplicative changes and additive changes.

1) Multiplicative Changes:

Theorem 1: For any given $p \in (0, 1)$ and $\epsilon \in (0, 0.35)$, there exists a randomized algorithm (described in Section III) that maintains estimates $\tilde{X}_v(t)$ at each $v \in V$, so that under the deterministic communication model and for any $m = \lceil (3 \ln(2/p))/(\epsilon^2) \rceil$

$$(1 - \epsilon) e^{-m(D+1)\delta} \leq \frac{\tilde{X}_v(t)}{\sum_v X_v(t)} \leq (1 + \epsilon) e^{m(D+1)\delta}$$

for all $t \geq mD$, with probability at least $1 - p$.

Theorem 1 suggests the tradeoff between dynamics rate δ and accuracy which depends on diameter D. Next, we state a lower bound on the accuracy of estimation achievable by *any* randomized (or deterministic) algorithm.

Theorem 2: Consider class of distributed algorithms as described in Section II-A in which each node v can obtain information only from its neighbors $\{u \in V : (u, v) \in E\}$ and the information communicated by a node u to v only depends on the state of u. Suppose there exists such an algorithm that maintains estimates $\tilde{X}_v(t)$ for $v \in V$ such that $e^{-\Delta} \leq (\tilde{X}_v(t))/(\sum_v X_v(t)) \leq e^{\Delta}$ for all $v \in V$ with probability at least 3/4, then $\Delta \geq D\delta$.

---

[1] In practice the bit-rate of communication is bounded and a number can be exchanged only up to some accuracy. We shall ignore this issue here.

Fig. 1. Ring graph on $n = 7$ nodes. Each node communicates with its two neighbors.

Fig. 2. Complete graph.

Ignoring the ϵ and p, note that the exponents in Theorems 1 and 2 match, up to a factor of m. This shows that the exponent $D\delta$ is inherent effect of dynamics on the performance of algorithm.

2) Additive Changes: The algorithm that we study is the well studied, known linear iterative algorithm. It utilizes a doubly stochastic (averaging) matrix $A = [A_{uv}] \in \mathbb{R}_+^{n \times n}$ that is graph G conformant, i.e., $A_{uv} = 0$ if $(u, v) \notin E$, and irreducible. The following result states the accuracy of estimation of the algorithm in presence of additive changes. This result is elementary and should be known in literature. We state it here for completeness.

Theorem 3: The linear iterative algorithm utilizing matrix A described in Section IV maintains estimate $\hat{X}_v(t)$ at each node $v \in V$ so that for $t \geq 0$

$$\|\hat{X}(t) - X_{\text{ave}}(t)\mathbf{1}\| \leq \frac{\delta\sqrt{n}}{1 - \lambda}$$

where $\hat{X}(t) = [\hat{X}_v(t)]$ is the vector of estimates, $X_{\text{ave}}(t) = (1/n)\sum_v X_v(t)$ is the actual average at time t and $\lambda = \lambda(A)$ defined as

$$\lambda(A) = \sup_{x \in \mathbb{R}^n : \sum_v x_v = 0} \frac{\|Ax\|}{\|x\|}.$$

D. Examples: Network Graph Models

For concreteness, we apply the results to several graph models and see how the performance scales under multiplicative model.

1) Line/Ring Graphs: A line graph (see Fig. 1) consists of n nodes, say $1, 2, \ldots, n$, with an edge between i and $i + 1$ for $1 \leq i \leq n - 1$. A ring graph is similar, with, in addition to the edges in the line graph, an edge between n and 1. For a line graph or ring graph with n nodes, its diameter is $D = \Theta(n)$. Therefore, in Theorems 1 and 2, for the deterministic communication model, we have, with probability at least $1 - p$ at all times t

$$e^{-\Theta(n\delta)} \leq \frac{\tilde{X}_v(t)}{\sum_v X_v(t)} \leq e^{\Theta(n\delta)}$$

ignoring the constants that depend on p and ϵ.

2) Grid Graphs: A grid graph, natural extension of line graph to two dimensions, consists of n nodes, say $1, 2, \ldots, n$, placed on an $\sqrt{n} \times \sqrt{n}$ grid with each node connected to four of its

neighbor with exception of nodes on boundary which may be connected to three or two neighbors. The diameter for such a graph is $D = \Theta(\sqrt{n})$. Therefore, in Theorems 1 and 2, we have that with probability at least $1 - p$ at all times t

$$e^{-\Theta(\sqrt{n}\delta)} \leq \frac{\tilde{X}_v(t)}{\sum_v X_v(t)} \leq e^{\Theta(\sqrt{n}\delta)}$$

ignoring the constants that depend on p and ϵ.

3) Geometric Random Graphs: A geometric random graph is the probabilistic analogue of grid graph. Here n nodes are placed independently and uniformly at random in a two-dimensional square of unit area. Two nodes are connected to each other if they are within distance $r(n)$ of each other. For the connectivity of the graph with high probability, it is well understood that $r(n) \geq r^*(n)$ is required. Here this critical connectivity radius $r^*(n) = \Theta(\sqrt{\log n / n})$. For graph obtained with $r(n) \geq r^*(n)$, with high probability the diameter is $\Theta(1/r(n))$ and equals $\Theta(\sqrt{n/\log n})$ when $r(n) = \Theta(r^*(n))$. In this particular case, in Theorems 1 and 2, we have that with probability at least $1 - p$ at all times t

$$e^{-\Theta(\delta\sqrt{n/\log n})} \leq \frac{\tilde{X}_v(t)}{\sum_v X_v(t)} \leq e^{\Theta(\delta\sqrt{n/\log n})}$$

ignoring the constants that depend on p and ϵ.

4) Complete Graphs: A complete graph (see Fig. 2) on n nodes has an edge between each pair of nodes; thus its diameter is 1. In this case, the theorems give good results, but a straightforward algorithm would be even simpler: since each node communicates with *all nodes* at each time t, it has all the information that was in the network at time $t - 1$, and estimation is therefore trivial, to within a difference of δ or factor of e^δ for additive and multiplicative changes, respectively. Note that this error exactly matches the lower bound of $D\delta = \delta$ from Theorem 2:

$$e^{-\delta} \leq \frac{\tilde{X}_v(t)}{\sum_v X_v(t)} \leq e^{\delta}$$

5) Star Graphs: A star graph (see Fig. 3) is a graph whose edges are exactly all those between one fixed node and every other node. Compared to complete graphs which have $\binom{n}{2}$ edges, star graphs are very sparse, having only $n - 1$ edges. However, since their diameter is small $(D = 2)$, the behavior of the algorithms on the star graph is qualitatively the same as for complete graphs (with deterministic communication).

Fig. 3. Star graph.

This shows that the diameter is the predominant feature of the topology which affects the performance.

6) Expander Graphs: Expander graphs are graphs that are sparse but have high connectivity. In expander graphs, the diameter is $O(\log n)$ [14]. Since $e^{\Theta(\log n)} = n^{\Theta(1)}$, the estimates can be made to satisfy

$$(1 - \epsilon)n^{-\Theta(m\delta)} \leq \frac{\tilde{X}_v(t)}{\sum_v X_v(t)} \leq (1 + \epsilon)n^{\Theta(m\delta)}.$$

It should be noted that *small-world networks* [15], [16] though not necessarily expander graphs, have a diameter of $O(\text{polylog}(n))$.

III. MULTIPLICATIVE CHANGES

This section describes the algorithm for maintaining estimate of sum at each node in the presence of multiplicative changes in nodes.

Our main algorithm, as an important subroutine uses an algorithm for maintaining *minimum* (rather than the sum) of values in a network in presence of dynamics. This is described in Section III-A. In a way, this answers question posed in Section I-C.

Next, in Section III-B, we show how the algorithm of Section III-A can be utilized to estimate the sum with values changing multiplicatively. We shall present two algorithms. The first algorithm described in Section III-B1) merely uses the maximum of the values as an estimate for the sum. The second algorithm described in Section III-B2 involves transforming the problem of summation into the problem of minimum computation by means of randomization. These two algorithms collectively establish Theorem 1.

Finally, in Section III-C, we establish the lower bound, using adversarial arguments, on the performance achievable by *any* algorithm. This establishes Theorem 2.

A. Estimating the Minimum

In this section, we state an algorithm for maintaining estimates of the minimum of values in the network: let $Y_v(t)$ be non-negative value at node $v \in V$ so that for any $v \in V$ and $t \geq 0$

$$e^{-\delta} \leq \frac{Y_v(t+1)}{Y_v(t)} \leq e^{\delta}.$$

For this section, consider this as the problem of interest. However, as mentioned earlier, in Section III-B2, we shall transform the problem of estimating $\sum X_v(t)$ into the problem of computing minimum of values $Y_v(t)$, where $Y_v(t)$ will be a (random) function of the actual value $X_v(t)$.

Let $N_k(v)$ denote the set of nodes at distance (with respect to the shortest path metric in G) $\leq k$ from v. Let $N(v) = N_1(v)$ denote the set of neighbors of v including v itself, and let $d(u, v)$ denote the (shortest-path) distance between u and v. Clearly, diameter of G can be defined as

$$D = \max_{u,v \in V} d(u, v).$$

As discussed in Section I-C, if $\delta = 0$ or the case when $Y_v(t) = Y_v$ for all t, essentially a trivial message-passing algorithm can find the correct minimum in time D: each node maintains an estimate $\hat{Y}_v(t)$ of the actual minimum $\min_v Y_v$, and, on communicating with all its neighbors and receiving their estimates $\hat{Y}_v(t-1)$, updates its estimate simply as the minimum of all known values: $\hat{Y}_v(t) = \min_{u \in N(v)} \hat{Y}_u(t-1)$. While this algorithm works for $\delta = 0$, it can be arbitrarily bad for $\delta > 0$ as explained in Section I-C. We present our algorithm for any $\delta \geq 0$.

1) Algorithm: Each node v maintains an estimate $\tilde{Y}_v(t)$ of $\min_u Y_u(t)$. At the beginning of each time t, each node v exchanges its current estimate $\tilde{Y}_v(t-1)$, from time $t-1$, with all its neighbors, and similarly receives all their estimates. Then, the new estimate $\tilde{Y}_v(t)$ is computed as

$$\tilde{Y}_v(t) = \min \left\{ Y_v(t), \min_{u \in N(v)} \tilde{Y}_u(t-1)e^{\delta} \right\}. \tag{1}$$

In other words, the new estimate of node v, at time t, is the minimum of its own new value, and the estimates of all its neighbors at time $t-1$, scaled up by e^{δ}. This last factor is the essential component of the discounting. We shall assume that initially, $t = 0$ and $\tilde{Y}_v(0) = Y_v(0)$. And we shall use the definition $\tilde{Y}_v(s) = \infty$ for $s < 0$.

Theorem 4: Under the above setup, for any $t \geq D$

$$\left(\min_{u \in V} Y_u(t) \right) \leq \tilde{Y}_v(t) \leq e^{2D\delta} \left(\min_{u \in V} Y_u(t) \right). \tag{2}$$

Proof: To start with, we claim the following identity: for any $v \in V$ and $t \geq 0$

$$\tilde{Y}_v(t) = \min_{k \geq 0} \min_{u \in N_k(v)} \{ Y_u(t-k)e^{k\delta} \} \tag{3}$$

$$= \min_{u \in V} \min_{k \geq d(u,v)} \{ Y_u(t-k)e^{k\delta} \}. \tag{4}$$

In the above, we use definition $Y_u(s) = \infty$ for $s < 0$. To establish this identity, observe that the two expressions (3) and (4) are equal—only the order of the mins is reversed; they both evaluate the minimum over the same set of pairs (u, k). Thus, it is enough to prove the first equality.

The equality follows by recursive application of the update rule (1). To that end, application of the rule once and twice respectively, we obtain

$$\tilde{Y}_v(t) = \min\left\{Y_v(t), \min_{u \in N(v)} \tilde{Y}_u(t-1)e^\delta\right\}$$
$$= \min\left\{Y_v(t), \min_{u \in N(v)} Y_u(t-1)e^\delta,\right.$$
$$\left.\min_{w \in N_2(v)} \tilde{Y}_w(t-2)e^{2\delta}\right\}.$$

More generally, for any $m \geq 1$, it follows that

$$\tilde{Y}_v(t) = \min\left\{\min_{k=0}^{m} \min_{u \in N_k(v)}\{Y_u(t-k)e^{k\delta}\},\right.$$
$$\left.\min_{w \in N_{m+1}(v)} \tilde{Y}_w(t-m-1)e^{(m+1)\delta}\right\}.$$

By definition, we assume $\tilde{Y}_w(t) = \infty$ for $t < 0$. Therefore, by taking $m \to \infty$ and with notation $Y_u(s) = \infty$ for any $u \in V$ and $s < 0$, we obtain

$$\tilde{Y}_v(t) = \min_{k \geq 0} \min_{u \in N_k(v)}\{Y_u(t-k)e^{k\delta}\}.$$

This establishes identity (3) [and hence (4)] as desired.

Now, we shall utilize (4) for $\tilde{Y}_v(t)$ to establish (2) for any given $v \in V$. To this end, for a given $v \in V$ using (4) it follows that

$$\tilde{Y}_v(t) = \min_{u \in V}\left\{Y_u(t-d(u,v))e^{d(u,v)\delta}\right\}. \qquad (5)$$

This follows immediately from (4) and the fact that the value $Y_u(\cdot)$ increases (re. decreases) by multiplicative factor e^δ (re. $e^{-\delta}$) in unit time.

Now we are ready to establish (2). Observe that due to bound on multiplicative change, $Y_u(t) \leq Y_u(t-d(u,v))e^{d(u,v)\delta}$. Therefore, from (5), it follows that for any $t \geq 0$

$$(\min_{u \in V} Y_u(t)) \leq \tilde{Y}_v(t). \qquad (6)$$

Similarly, due to bound on multiplicative change, it follows that $Y_u(t-d(u,v)) \leq Y_u(t)e^{d(u,v)\delta}$ for $t \geq d(u,v)$. Therefore, from (5) it follows that for $t \geq D$

$$\tilde{Y}_v(t) \leq \min_{u \in V}\left\{Y_u(t)e^{2d(u,v)\delta}\right\}$$
$$\leq \min_{u \in V}\left\{e^{2D\delta}Y_u(t)\right\}$$
$$\leq e^{2D\delta}(\min_{u \in V} Y_u(t)). \qquad (7)$$

From (6) and (7), the desired result (2) follows. ∎

Note that the error is one-sided, but by taking the estimate to be $\tilde{Y}_v(t)e^{D\delta}$ instead, the guarantee within a factor $e^{D\delta}$ can be provided. In fact, it can be shown using argument similar to that used in Section III-C, that this is the best possible bound that any algorithm can guarantee.

Although we proved these results about estimating the minimum for multiplicative changes, since that is the form in which we will use this algorithm, it is worth noting that a similar

bound, simpler in form, holds for estimating the minimum under additive changes. This requires using an update rule analogous to (1).

Theorem 5 (Additive Version): Suppose value at node $v \in V$ at time t be $Y_v(t)$. For all $t \geq 0$ and $v \in V$, let

$$|Y_v(t+1) - Y_v(t)| \leq \delta.$$

Let the estimation at node v at time $t \geq 0$, denoted by $\tilde{Y}_v(t)$ be updated as

$$\tilde{Y}_v(t) = \min\left\{Y_v(t), \min_{u \in N(v)} \tilde{Y}_u(t-1) + \delta\right\}, \qquad (8)$$

with $\tilde{Y}_v(s) = \infty$ for all $s < 0$. Then, for all $t \geq D$

$$(\min_{u \in V} Y_u(t)) \leq \tilde{Y}_v(t) \leq (\min_{u \in V} Y_u(t)) + 2D\delta. \qquad (9)$$

Proof: Define $Z_v(t) = \exp(Y_v(t))$ and $\tilde{Z}_v(t) = \exp(\tilde{Y}_v(t))$. Observe that then $Z_v(\cdot)$ and $\tilde{Z}_v(\cdot)$ are similar to the setup of multiplicative change. Therefore, by Theorem 4, it follows that

$$(\min_{u \in V} Z_u(t)) \leq \tilde{Z}_v(t) \leq e^{2D\delta}(\min_{u \in V} Z_u(t)).$$

This is equivalent to (9). ∎

It is easy to see that similar algorithm can be devised to estimate the maximum instead with similar guarantees, a fact which we shall use in the next section.

B. Estimating the Sum

In this section, we describe our algorithms for maintaining estimates of the sum, under multiplicative changes. We shall start with a simpler algorithm, that maintains maximum as an estimate of the summation. This will be followed by a more involved algorithm that builds on the algorithm of Mosk-Aoyama and Shah [13].

1) A Simple Estimator Using Maximum: An estimate for the sum, that would be within a factor of n, is the *maximum*. This is because the sum of a set of n positive numbers is larger than the maximum of the numbers, but no more than n times the maximum. Therefore, $\sqrt{n}$ times the maximum provides an estimator of summation that is within factor $1/\sqrt{n}$ and $\sqrt{n}$ of the summation. In Section III-A we described an algorithm to maintain estimation of minimum within multiplicative factor $e^{D\delta}$. A similar algorithm (with $\min$ replaced by $\max$ and e^δ by $e^{-\delta}$ in (1)) for maximum provides its estimation within factor $e^{D\delta}$ as well. Therefore, effectively we have an estimation of summation within factor $\sqrt{n}e^{D\delta}$. Next, we shall describe a more involved algorithm for which the estimation will be within factor $e^{\Theta(D\delta)}$ of summation. Now, when δ is large enough, i.e., $\sqrt{n} \ll e^{D\delta}$, both the naive maximum-based estimation and the more involved algorithm perform similarly. However, when $e^{D\delta} \ll \sqrt{n}$ the more involved algorithm dominates the naive maximum based estimator. In general, the randomized algorithm described next has essentially the best possible performance with error within factor $e^{\Theta(D\delta)}$.

2) Randomized Estimator: We now turn to our algorithm for estimating the sum of values in a network, using a combination

of the algorithm for estimating the minimum found in the previous section, and an appropriate transformation on exponentially distributed random variables.

Recall that $X_v(t)$ denote the non-negative values at nodes $v \in V$ at time t. Without loss of generality, let us assume that $X_v(0) = 1$ for all $v \in V$. We shall assume that for any $t \geq 0$ and $v \in V$

$$e^{-\delta} \leq \frac{X_v(t+1)}{X_v(t)} \leq e^{\delta}.$$

We shall associate an auxiliary variable $Y_v(t)$ to each node as follows: $Y_v(t)$ is a random function of $X_v(t)$, distributed as an exponential random variable with parameter $X_v(t)$, denoted $Y_v(t) \sim \mathrm{Exp}(X_v(t))$. That is,

$$\mathbb{P}(Y_v(t) \geq \zeta) = e^{-X_v(t)\zeta}, \quad \text{for any } \zeta \geq 0.$$

As we shall see, $Y_v(t)$ will also change dynamically so that for any $v \in V$ and $t \geq 0$

$$e^{-\delta} \leq \frac{Y_v(t+1)}{Y_v(t)} \leq e^{\delta}.$$

Let $\tilde{Y}_v(t)$ be estimation of $\min_{u \in V} Y_u(t)$ node v as per algorithm described in Section III-A using the multiplicative factor e^{δ}. Using $\tilde{Y}_v(t)$, we shall obtain estimate $\tilde{X}_v(t)$ of the sum $\sum_{u \in V} X_u(t)$. Next, we describe key intuition behind such an algorithm followed by precise algorithm.

The basic idea behind the algorithm, which enables us to transform the problem of computing the sum to that of the minimum, is the following elementary but unique extremal property of the exponential distribution: *The minimum of a set of exponential random variables is an exponential random variable whose rate is the sum of their rates*. Therefore, if we had random variables $Y_v(t) \sim \mathrm{Exp}(X_v(t))$, then their minimum is distributed $\sim \mathrm{Exp}(\sum_{u \in V} X_u(t))$. Therefore, the inverse of mean (or average) of minimum $\min_{u \in V} Y_u(t)$ is $\sum_{u \in V} X_u(t)$. Therefore, by computing several independent samples of minimum $\min_{u \in V} Y_u(t)$, we can obtain reasonable estimate of $\sum_{u \in V} X_u(t)$.

The above is precisely the idea (and algorithm) used by [13] for static setup. However, its adaptation to dynamic scenario is not immediate. Specifically, we need to make sure that $Y_v(t) \sim \mathrm{Exp}(X_v(t))$ for all t and $v \in V$ while ensuring that $Y_v(t+1)/Y_v(t)$ is within factor $e^{-\delta}$ and e^{δ} with probability 1 for all t. This is where another property of exponential distribution comes to our rescue. Specifically, consider the following rule for generating $Y_v(t)$ for all $t \geq 0$ and $v \in V$:

$$\text{at} \quad t = 0, \quad Y_v(0) \sim \mathrm{Exp}(X_v(0)),$$
$$\text{for} \quad t > 0, \quad Y_v(t+1) = \frac{X_v(t)}{X_v(t+1)} Y_v(t). \quad (10)$$

Thus, the variable $Y_v(0)$ is randomly generated once at $t = 0$, and for all subsequent t, $Y_v(t)$ is a deterministic function of $Y_v(0)$ and changes in $X_v(\cdot)$. It is easy to check that this achieves $Y_v(t) \sim \mathrm{Exp}(X_v(t))$. Further, we have

$$e^{-\delta} \leq \frac{Y_v(t)}{Y_v(t+1)} \leq e^{\delta}.$$

Next, we shall describe the precise algorithm.

a) Algorithm: Given $p \in (0,1)$ and $\epsilon \in (0, 0.35)$, let $m = \lceil (3\ln(2/p))/(\epsilon^2) \rceil$. For each $v \in V$ and $t \geq 0$, define $Y_{v,i}(t)$ with $1 \leq i \leq m$ as follows:
- for $t = 0$, generate $Y_{v,i}(0)$ independently of everything else and as per exponential distribution with parameter $X_v(0)$;
- for $t \geq 1$, update $Y_{v,i}(t+1) = X_v(t)Y_{v,i}(t)/X_v(t+1)$.

As discussed above, $Y_{v,i}(t) \sim \mathrm{Exp}(X_v(t))$ for all (v, i) and for all t. The m different indices i can be thought of as m independent copies of the algorithm running in parallel. The update rule involves the different i's "taking turns" in round-robin fashion: at each time t, if $t \equiv i \bmod m$, then the nodes exchange their i-values and update them as

$$\tilde{Y}_{v,i}(t) = \min \left\{ Y_{v,i}(t), \min_{u \in N(v)} \tilde{Y}_{u,i}(t-1)e^{m\delta} \right\}. \quad (11)$$

Define

$$\tilde{Y}_v(t) = \frac{1}{m} \sum_{i=1}^{m} \tilde{Y}_{v,i}(t). \quad (12)$$

Then, the estimate of $\sum_{u \in V} X_u(t)$ at node v is given by

$$\tilde{X}_v(t) = \frac{e^{m(D+1)\delta}}{\tilde{Y}_v(t)}. \quad (13)$$

Next, we establish the bound claimed in Theorem 1 for this algorithm.

Proof of Theorem 1: As $\tilde{Y}_{v,i}(t)$ changes only at $t \equiv i \bmod m$, we have $\tilde{Y}_{v,i}(t-1) = \tilde{Y}_{v,i}(t-2) = \cdots = \tilde{Y}_{v,i}(t-m)$. In addition, $e^{-m\delta} \leq (Y_{v,i}(t))/(Y_{v,i}(t-m)) \leq e^{m\delta}$, so focusing on the index i and times $t \equiv i \bmod m$, (11) is simply the update rule (1) for the values $Y_{v,i}(t)$ and associated estimates $\tilde{Y}_{v,i}(t)$:

$$\tilde{Y}_{v,i}(t) = \min \left\{ Y_{v,i}(t), \min_{u \in N(v)} \tilde{Y}_{u,i}(t-m)e^{m\delta} \right\}.$$

Therefore, from Theorem 4 of Section III-A, we have that for $t \equiv i \bmod m$ and for each i

$$(\min_{u \in V} Y_{u,i}(t)) \leq \tilde{Y}_{vi}(t) \leq e^{2Dm\delta}(\min_{u \in V} Y_{u,i}(t)). \quad (14)$$

Therefore, for any $t \geq mD$ and each i

$$e^{-m\delta}(\min_{u \in V} Y_{u,i}(t)) \leq \tilde{Y}_{vi}(t) \leq e^{(2D+1)m\delta}(\min_{u \in V} Y_{u,i}(t)). \quad (15)$$

For ease of notation, let $Z_i = \min_{u \in V} Y_{u,i}(t)$ with $1 \leq i \leq m$ and $Z = (1/m)\sum_{i=1}^{m} Z_i$. Summing up the inequalities in (15) over i, and using (12), we can write

$$e^{-m\delta}Z \leq \tilde{Y}_v(t) \leq Ze^{(2D+1)m\delta}.$$

The Z_is are indepenent and identically distributed (i.i.d.) random variables, each exponentially distributed with rate $\lambda = \sum_{v \in V} X_v(t)$, which is the sum we want to estimate.

By large deviation estimation for exponential distribution, it follows that

$$\mathbb{P}\left(Z \geq \frac{c}{\lambda}\right) \leq \exp(-m(c - 1 - \ln c)) \quad \text{for } c > 1,$$

$$\mathbb{P}\left(Z \leq \frac{c}{\lambda}\right) \leq \exp(-m(c - 1 - \ln c)) \quad \text{for } c < 1.$$

Therefore, it can be verified that for any $\epsilon \in (0, 0.35)$

$$\mathbb{P}\left(\frac{1}{\lambda(1 - \epsilon)} \leq Z \leq \frac{1}{\lambda(1 + \epsilon)}\right) \leq \exp(-m\epsilon^2/3).$$

Thus, for any $\epsilon \in (0, 0.35)$ and $p \in (0, 1)$, if $m = \lceil (3\ln(2/p))/(\epsilon^2) \rceil$, then with probability at least $1 - p$

$$
\begin{aligned}
\frac{e^{-m\delta}}{\lambda(1 + \epsilon)} &\leq e^{-m\delta} Z \\
&\leq \tilde{Y}_v(t) \\
&\leq e^{(2D+1)m\delta} Z \\
&\leq \frac{e^{(2D+1)m\delta}}{\lambda(1 - \epsilon)}.
\end{aligned}
$$

Finally, since the estimate of node v is $\tilde{X}_v(t) = (e^{m(D+1)\delta})/(Y_v(t))$, it follows that with probability at least $1 - p$ and for $t \geq mD$

$$(1 + \epsilon)e^{m(D+1)\delta} \geq \frac{\tilde{X}_v(t)}{\lambda} \geq (1 - \epsilon)e^{-(1+D)m\delta}.$$

This completes the proof of Theorem 1.   ∎

C. Lower Bound

This section establishes lower bound on the accuracy for any algorithm operating under multiplicative changes. Specifically, we shall prove Theorem 2.

To that end, suppose we have an algorithm that, for any values $X_v(t)$ with the promise that

$$e^{-\delta} \leq \frac{X_v(t + 1)}{X_v(t)} \leq e^{\delta}$$

can maintain estimates $\tilde{X}_v(t)$ with the guarantee that for any given $t \geq 0$ and all v

$$e^{-\Delta} \leq \frac{\tilde{X}_v(t)}{\sum_u X_u(t)} \leq e^{\Delta}$$

with probability at least 3/4.

We observe that the estimate $\tilde{X}_v(t)$ can depend only on $X_v(t)$, $X_u(t - 1), u \in N(v)$, and $\tilde{X}_u(t - 1), u \in N(v)$, and not on $X_u(t)$. (Formally, $\tilde{X}_v(t)$ is independent of $X_u(t)$ when conditioned on $X_u(t - 1)$ and $\tilde{X}_u(t - 1)$, for $u \neq v$.) More generally, $\tilde{X}_v(t)$ can depend on $X_u(t - k)$ or $\tilde{X}_u(t - k)$ only if $k \geq d(u, v)$, since any information at node u takes time at least $d(u, v)$ to reach node v.

Let u, v be any two nodes, with $d(u, v) = k$. Consider the following scenario. For all times t up to a certain time (say some large enough time t_0), let $X_w(t) = 1$ for all $w \neq u$, and $X_u(t) = M$ where M is some sufficiently large constant. Now consider the following two cases for $t > t_0$.

1) Values at all other nodes remain the same and $X_u(t)$ increases by e^{δ} at each $t > t_0$.

2) Values at all other nodes remain the same and $X_u(t)$ decreases by factor e^{δ}, i.e., multiplied by $e^{-\delta}$ at each $t > t_0$.

In either case, at time $t = t_0 + k$, the information that is made available at node v (including that about node u) is exactly the same. Therefore, the (randomized) decision taken by the algorithm is the same (distributionally) in either case. Since algorithm can predict the correct value within factor e^{Δ} with probability at least 3/4, it must be that $e^{2\Delta}$ is at least the ratio of the summation of values at node u in the above two cases. Now in the first case, the summation could be as large as $Me^{k\delta} + (n - 1)$ and in the second case summation could be as small as $Me^{-k\delta} + (n - 1)$. Therefore,

$$e^{2\Delta} \geq \frac{Me^{k\delta} + (n - 1)}{Me^{-k\delta} + (n - 1)}.$$

Applying this argument for pair (v, u) so that $d(u, v) = D$ and taking M arbitrary large, we obtain that

$$\Delta \geq D\delta.$$

This establishes Theorem 2.

IV. ADDITIVE CHANGES

Here we consider the scenario with additive changes. As mentioned in Section II-C2, we shall study the known linear iterative algorithm and establish Theorem 3.

To that end, let $X_v(t)$ be value at node $v \in V$ at time $t \geq 0$. Let $\Delta(t) = X(t + 1) - X(t)$ denote the vector of changes in the values at nodes from time t to $t + 1$. By the definition of additive change model, for all $t \geq 0$ and $v \in V$

$$|\Delta_v(t)| = |X_v(t + 1) - X_v(t)| \leq \delta.$$

The aim here is to estimate average $X_{\text{ave}}(t)$ where

$$X_{\text{ave}}(t) = \frac{1}{n} \sum_{i=1}^{n} X_i(t).$$

Note that this is equivalent to computing summation assuming n is known. In what follows, we describe the known algorithm maintains estimate $\hat{X}_v(t)$ of $X_{\text{ave}}(t)$ at each node $v \in V$ for all $t \geq 0$. The algorithm utilizes a doubly stochastic, graph G conformant and irreducible matrix $A = [A_{uv}] \in \mathbb{R}_+^{n \times n}$. That is, A satisfies the following properties:

1) $A_{uv} = 0$ if $(u, v) \notin E$;
2) $\sum_v A_{uv} = 1$ for all $u \in V$;
3) $\sum_u A_{uv} = 1$ for all $v \in V$;
4) directed graph $G(A) = (V, E(A))$ is connected where a directed edge $(u, v) \in E(A)$ iff $A_{uv} > 0$.

A. Algorithm

In word, the algorithm is simple: at each time, each node communicates with its neighbor and exchanges their current estimates. Subsequently, each node updates its own estimate as weighted summation of the estimates received from its neighbors and the change in its value. Specifically, for any node $v \in V$, its estimate $\hat{X}_v(t + 1)$ is updated as

$$\hat{X}_v(t + 1) = \sum_u A_{vu} \hat{X}_u(t) + \Delta_v(t). \tag{16}$$

Equivalently,

$$\hat{X}(t+1) = A\hat{X}(t) + \boldsymbol{\Delta}(t). \tag{17}$$

This is immediate extension of the known linear iterative averaging algorithm (for static case) whose update rule is given by

$$\hat{X}(t+1) = A\hat{X}(t).$$

In literature, it is well known that when values at nodes do not change (static case) and A satisfies properties listed above, then $\hat{X}(t) \to X_{\mathrm{ave}}$, where X_{ave} is the average of the node values. Here, node values change and $X_{\mathrm{ave}}(t)$ over time. We shall establish that $\hat{X}(t)$ remains "close to" $X_{\mathrm{ave}}(t)\mathbf{1}$.

B. Example of A

There are many ways to design such doubly stochastic, graph conformant and irreducible matrices A. A simple choice, that requires only a known upper bound d on the max-vertex degree, i.e., $d \geq \max_v d_v$, is derived from the Metropolis–Hastings rule [17]. As per this, define A as

$$A_{uv} = \begin{cases} 0, & \text{if } v \text{ is not a neighbor of } u \\ \frac{1}{d}, & \text{if } v \neq u \text{ is a neighbor of } u \\ 1 - \frac{d_u}{d}, & \text{if } v = u. \end{cases}$$

Since row u or column u has exactly d_u entries equal to $1/d$, and one (diagonal) entry equal to $1 - d_u/d$, it follows that A is doubly stochastic. It is graph G conformant by definition and since $G(A) = G$, if G is connected (which we assume here) then so is $G(A)$.

C. Analysis

Here, we establish Theorem 3. Define $Y(t) = \hat{X}(t) - X_{\mathrm{ave}}(t)\mathbf{1}$, the error vector in estimation at time t. Initially, we assume that $X_v(0) = 1$ for all v and $\hat{X}_v(0) = X_v(0) = 1$. Therefore, $Y(0) = \mathbf{0}$. Also define

$$\boldsymbol{\Delta}_{\mathrm{ave}}(t) = \frac{1}{n} \sum_v \boldsymbol{\Delta}_v(t).$$

Observe that, since A is doubly stochastic, $A(t)\mathbf{1} = \mathbf{1}$. Therefore, using (17) it follows that

$$\begin{aligned} Y(t+1) &= \hat{X}(t+1) - X_{\mathrm{ave}}(t+1)\mathbf{1} \\ &= A(t)\hat{X}(t) + \boldsymbol{\Delta}(t) - (X_{\mathrm{ave}}(t) + \boldsymbol{\Delta}_{\mathrm{ave}}(t))\mathbf{1} \\ &= A(t)Y(t) + e(t). \end{aligned} \tag{18}$$

where $e(t) = \boldsymbol{\Delta}(t) - \boldsymbol{\Delta}_{\mathrm{ave}}(t)\mathbf{1}$. The goal is to bound $\|Y(\cdot)\|$ using identity (18). To that end, consider $e = e(t)$ for any t. Under additive change model $|e_v| \leq 2\delta$ for all $v \in V$ and $\sum_v e_v = 0$. This means that

$$\|e\|^2 = \sum_v e_v^2 \leq \sum_v 4\delta^2 = 4n\delta^2. \tag{19}$$

Using this, we shall bound $\|Y(\cdot)\|$. By recursive application of (18), it follows that

$$Y(t+1) = \sum_{k=0}^{t} A^k e(t-k) + Y(0) = \sum_{k=0}^{t} A^k e(t-k).$$

Therefore, using triangle's inequality

$$\begin{aligned} \|Y(t+1)\| &\leq \sum_{k=0}^{t} \|A^k e(t-k)\| \\ &\overset{(a)}{\leq} \sum_{k=0}^{t} \lambda^k 2\delta\sqrt{n} \\ &\leq \frac{2\delta\sqrt{n}}{1-\lambda} \end{aligned}$$

where recall that $\lambda = \lambda(A)$ is defined as

$$\lambda(A) = \sup_{x \in \mathbb{R}^n : \sum_v x_v = 0} \frac{\|Ax\|}{\|x\|}.$$

To obtain (a), we have used the fact that if $x \in \mathbb{R}^n$ such that $\sum_v x_v = 0$, then $\sum_v z_v = 0$ where $z = Ax$ since A is doubly stochastic. This completes the proof of Theorem 3.

V. CONCLUSION

In this paper, we presented "dynamics-aware" distributed algorithms for estimating the sum or average of values in a network, under dynamics. The algorithm described exhibit near optimal tradeoff between accuracy of estimation and rate of dynamics. Specifically, the error in accuracy depends on network topology. Going forward, developing such dynamics-aware algorithm for generic distributed function computation remains an important research challenge.

REFERENCES

[1] J. Tsitsiklis, "Problems in decentralized decision making and computation," Ph.D. dissertation, Dept. of Elect. Eng. and Comput. Sci., Mass. Inst. of Technol., Cambridge, MA, 1984.

[2] D. Bertsekas, J. Tsitsiklis, and M. Athans, "Convergence theories of distributed iterative processes: A survey," Lab. for Inf. and Dec. Syst., Mass. Inst. of Technol., Cambridge, MA, 1983.

[3] A. Jadbabaie, J. Lin, and A. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Transa. Autom. Control*, vol. 48, no. 6, pp. 988–1001, Jun. 2003.

[4] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Gossip algorithms: Design, analysis and applications," in *Proc. IEEE INFOCOM*, Mar. 2005, vol. 3, pp. 1653–1664.

[5] Y. Rekhter *et al.*, RFC 1771: A Border Gateway Protocol 4 (BGP-4), Mar. 1995.

[6] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. Autom. Control*, vol. 37, no. 12, pp. 1936–1948, Dec. 1992.

[7] T. Griffin and G. Wilfong, "An analysis of BGP convergence properties," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 29, no. 4, pp. 277–288, 1999.

[8] D. Bertsekas and J. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Upper Saddle River, NJ: Prentice-Hall, 1989.

[9] V. Blondel, J. Hendrickx, A. Olshevsky, and J. Tsitsiklis, "Convergence in multiagent coordination, consensus, and flocking," in *Proc. IEEE Conf. Decision Control*, 2005, vol. 44, p. 2996.

[10] A. Dimakis, S. Kar, J. Moura, M. Rabbat, and A. Scaglione, "Gossip algorithms for distributed signal processing," *Proc. IEEE*, vol. 98, no. 11, pp. 1847–1864, Nov. 2010.

[11] D. Shah, "Gossip algorithms," in *Foundations and Trends in Networking*. Delft, The Netherlands: Now Publishers, Jun. 2009 [Online]. Available: http://web.mit.edu/devavrat/www/Gossipbook.pdf

[12] S. Kar and J. F. M. Moura, "Distributed linear parameter estimation in sensor networks: Convergence properties," in *Proc. 42nd IEEE Asilomar Conf. Signals, Syst., Comput.*, pp. 1347–1351.

[13] D. Mosk-Aoyama and D. Shah, "Computing separable functions via gossip," in *Proc. 25th Annu. ACM Symp. Principles Distrib. Comput.*, 2006, p. 122, ACM.

[14] S. Hoory, N. Linial, and A. Wigderson, "Expander graphs and their applications," *Bull.-Ameri. Math. Soc.*, vol. 43, no. 4, p. 439, 2006.

[15] D. Watts and S. Strogatz, "Collective dynamics of 'small-world' networks," *Nature*, vol. 393, no. 6684, pp. 440–442, 1998.

[16] J. Kleinberg, "The small-world phenomenon: An algorithm perspective," in *Proc. 32nd Annu. ACM Symp. Theory of Comput.*, New York, 2000, pp. 163–170, ACM.

[17] S. Chib and E. Greenberg, "Understanding the metropolis-hastings algorithm," *Amer. Statist.*, vol. 49, no. 4, pp. 327–335, 1995.

Shreevatsa Rajagopalan received the B.S. degree in computer science and mathematics from Chennai Institute of Mathematics, Chennai, India, in 2007 and the M.S. degree in operations research from Massachusetts Institute of Technology, Cambridge, in 2010.

Devavrat Shah received the B.Tech. degree in computer science and engineering from the IIT Bombay. Mumbai, India, in 1999 and the Ph.D. degree in computer science from Stanford University, Stanford, CA, in 2004.

He is currently the Jamieson Career Development Associate Professor with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology (MIT), Cambridge. He is a member of the Laboratory for Information and Decision Systems (LIDS) and Operations Research Center (ORC), MIT. His research focus is on theory of large complex networks which includes network algorithms, stochastic networks, network information theory, and large-scale statistical inference. He is currently an Associate Editor of *Operations Research*.

Prof. Shah was co-awarded the best paper awards at the IEEE IN-FOCOM'04, ACM SIGMETRICS/Performance'06; and supervised best student paper awards at Neural Information Processing Systems '08 and ACM SIGMETRICS/Performance '09. He received the 2005 George B. Dantzig Best Dissertation Award from INFORMS. He received the ACM SIGMETRICS Rising Star Award 2008 for his work on network scheduling algorithms. He was recently awarded the 2010 Erlang Prize from INFORMS which is given to a young researcher for outstanding contributions to applied probability.