Weighted Gossip: Distributed Averaging Using Non-Doubly Stochastic Matrices

Florence Bénézit, ENS-INRIA, France
Vincent Blondel, UCL, Belgium
Patrick Thiran, EPFL, Switzerland
John Tsitsiklis, MIT, USA
Martin Vetterli, EPFL, Switzerland

***Abstract*—This paper presents a general class of gossip-based averaging algorithms, which are inspired from Uniform Gossip [1]. While Uniform Gossip works synchronously on complete graphs, weighted gossip algorithms allow asynchronous rounds and converge on any connected, directed or undirected graph. Unlike most previous gossip algorithms [2]–[6], Weighted Gossip admits stochastic update matrices which need not be doubly stochastic. Double-stochasticity being very restrictive in a distributed setting [7], this novel degree of freedom is essential and it opens the perspective of designing a large number of new gossip-based algorithms. To give an example, we present one of these algorithms, which we call One-Way Averaging. It is based on random geographic routing, just like Path Averaging [5], except that routes are one way instead of round trip. Hence in this example, getting rid of double stochasticity allows us to add robustness to Path Averaging.**

I. Introduction

Gossip algorithms were recently developed to solve the distributed average consensus problem [1]–[6]. Every node i in a network holds a value x_i and wants to learn the average x_{ave} of all the values in the network in a distributed way. Most gossip algorithms were designed for wireless sensor networks, which are usually modeled as random geometric graphs and sometimes as lattices. Ideally a distributed averaging algorithm should be efficient in terms of energy and delay without requiring too much knowledge about the network topology at each node, nor sophisticated coordination between nodes.

The simplest gossip algorithm is Pairwise Gossip, where random pairs of connected nodes iteratively and locally average their values until convergence to the global average [2]. Pairwise local averaging is an easy task, which does not require global knowledge nor global coordination, thus Pairwise Gossip fulfills the requirements of our distributed problem. However, the convergence speed of Pairwise Gossip suffers from the locality of the updates, and it was shown that averaging random geographic routes instead of local neighborhoods is an order-optimal communication scheme to run gossip. Let n be the number of nodes in the network. On random geometric graphs, Pairwise Gossip requires $\Theta(n^2)$ messages whereas Path Averaging requires only $\Theta(n \log n)$ messages under some conditions [5].

The previous algorithm gained efficiency at the price of more complex coordination. At every round of Path Averaging, a random node wakes up and generates a random route. Values are aggregated along the route and the destination node computes the average of the values collected along the route. Then the destination node sends the average back through the same route so that all the nodes in the route can update their values to the average. Path Averaging is efficient in terms of energy consumption, but it demands some long distance coordination to make sure that all the values in the route were updated correctly. Routing information back and forth might as well introduce delay issues, because a node that is engaged in a route needs to wait for the update to come back before it can proceed to another round. Furthermore, in a mobile network, or in a highly dynamic network, routing the information back on the same route might even not succeed.

This work started with the goal of designing a unidirectional gossip algorithm fulfilling the following requirements:

- Keep a geographic routing communication scheme because it is highly diffusive,
- Avoid routing back data: instead of long distance agreements, only agreements between neighbors are allowed,
- Route crossing is possible at any time, without introducing errors in the algorithm.

As we were designing One-Way Averaging, we happened to prove the correctness of a broad set of gossip-based algorithms, which we present in this paper along with One-Way Averaging. These algorithms can be asynchronous and they use stochastic diffusion matrices which are not necessarily doubly stochastic, as announced by the title of the paper.

In Section II, we give some background on gossip algorithms, and we explain why Uniform Gossip is a key algorithm to get inspired from when building a unidirectional gossip algorithm. In Section III, we present Weighted Gossip, an asynchronous generalization of Uniform Gossip, which was already suggested in [1] but had remained unnamed. We show in Section IV that weighted gossip algorithms converge to x_{ave}, which is a novel result to the best of our knowledge. In Section V, we describe in detail One-Way Averaging and we show on simulations that the good diffusivity of geographic routes in Path Averaging persists in One-Way Averaging. Computing the speed of convergence of weighted gossip algorithms remains open and is part of future work.

II. Background on Gossip Algorithms

The values to be averaged are gathered in a vector $\boldsymbol{x}(0)$ and at any iteration t, the current estimates of the average x_{ave} are gathered in $\boldsymbol{x}(t)$. Gossip algorithms update estimates linearly. At any iteration t, there is a matrix $\boldsymbol{W}(t)$ such that:

$$\boldsymbol{x}(t)^T = \boldsymbol{x}(t-1)^T \boldsymbol{W}(t).$$

In gossip algorithms that converge to average consensus, $\boldsymbol{W}(t)$ is doubly stochastic: $\boldsymbol{W}(t)\mathbf{1} = \mathbf{1}$ ensures that the global average is conserved, and $\mathbf{1}^T\boldsymbol{W}(t) = \mathbf{1}^T$ guarantees stable consensus. To perform averaging on a one way route, $\boldsymbol{W}(t)$ should be upper triangular (up to a node index permutation). But the only matrix that is both doubly stochastic and upper triangular matrix is the identity matrix. Thus, unidirectional averaging requires to drop double stochasticity.

Uniform Gossip solves this issue in the following way. Instead of updating one vector $\boldsymbol{x}(t)$ of variables, it updates a vector $\boldsymbol{s}(t)$ of *sums*, and a vector $\boldsymbol{\omega}(t)$ of *weights*. Uniform Gossip initializes $\boldsymbol{s}(0) = \boldsymbol{x}(0)$ and $\boldsymbol{\omega}(0) = \mathbf{1}$. At any time, the vector of estimates is $\boldsymbol{x}(t) = \boldsymbol{s}(t)/\boldsymbol{\omega}(t)$, where the division is performed elementwise. The updates are computed with stochastic *diffusion* matrices $\{\boldsymbol{D}(t)\}_{t>0}$:

$$\boldsymbol{s}(t)^T = \boldsymbol{s}(t-1)^T\boldsymbol{D}(t), \tag{1}$$
$$\boldsymbol{\omega}(t)^T = \boldsymbol{\omega}(t-1)^T\boldsymbol{D}(t). \tag{2}$$

Kempe et al. [1] prove that the algorithm converges to a consensus on x_{ave} ($\lim_t \boldsymbol{x}(t) = x_{ave}\mathbf{1}$) in the special case where for any node i, $\boldsymbol{D}_{ii}(t) = 1/2$ and $\boldsymbol{D}_{ij}(t) = 1/2$ for one node j chosen i.i.d. uniformly at random. As a key remark, note that here $\boldsymbol{D}(t)$ is *not* doubly stochastic. The algorithm is synchronous and it works on complete graphs without routing, and on other graphs with routing. We show in this paper that the idea works with many more sequences of matrices $\{\boldsymbol{D}(t)\}_{t>0}$ than just the one used in Uniform Gossip.

III. Weighted Gossip

We call Weighted Gossip the class of gossip-based algorithms following the sum and weight structure of Uniform Gossip described above (Eq. (1) and (2)). A weighted gossip algorithm is entirely characterized by the distribution of its diffusion matrices $\{\boldsymbol{D}(t)\}_{t>0}$. Let $\boldsymbol{P}(s,t) := \mathbf{D}(s)\mathbf{D}(s+1)\ldots\mathbf{D}(t)$ and let $\mathbf{P}(t) := \boldsymbol{P}(1,t)$. Then

$$\boldsymbol{s}(t)^T = \boldsymbol{x}(0)^T\boldsymbol{P}(t), \tag{3}$$
$$\boldsymbol{\omega}(t)^T = \mathbf{1}^T\boldsymbol{P}(t). \tag{4}$$

If a weighted gossip algorithm is asynchronous, then, $\boldsymbol{D}_{ii}(t) = 1$ and $\boldsymbol{D}_{ij,j\neq i}(t) = 0$ for the nodes i that do not contribute to iteration t. If $\boldsymbol{D}_{ij}(t) \neq 0$, then node i sends $(\boldsymbol{D}_{ij}(t)\boldsymbol{s}_i(t-1), \boldsymbol{D}_{ij}(t)\boldsymbol{\omega}_i(t-1))$ to node j, which adds the received data to its own sum $\boldsymbol{s}_j(t-1)$ and weight $\boldsymbol{\omega}_j(t-1)$. At any iteration t, the estimate at node i is $\boldsymbol{x}_i(t) = \boldsymbol{s}_i(t)/\boldsymbol{\omega}_i(t)$.

Because $\mathbf{1}^T\boldsymbol{D}(t) \neq \mathbf{1}^T$, sums and weights do not reach a consensus. However, because $\boldsymbol{D}(t)\mathbf{1} = \mathbf{1}$, sums and weights are conserved: at any iteration t,

$$\sum_{i=1}^{n}\boldsymbol{s}_i(t) = \sum_{i=1}^{n}\boldsymbol{x}_i(0) = nx_{ave}, \tag{5}$$
$$\sum_{i=1}^{n}\boldsymbol{\omega}_i(t) = n. \tag{6}$$

This implies that Weighted Gossip is a class of non-biased estimators for the average (even though $\sum_{i=1}^{n}\boldsymbol{x}_i(t)$ is not conserved through time!):

Theorem 3.1 (Non-biased estimator): If the estimates $\boldsymbol{x}(t) = \boldsymbol{s}(t)/\boldsymbol{\omega}(t)$ converge to a consensus, then the consensus value is the average x_{ave}.

Proof: Let c be the consensus value. For any $\epsilon > 0$, there is an iteration t_0 after which, for any node i, $|\boldsymbol{x}_i(t) - c| < \epsilon$. Then, for any $t > t_0$, $|\boldsymbol{s}_i(t) - c\boldsymbol{\omega}_i(t)| < \epsilon\boldsymbol{\omega}_i(t)$ (weights are always positive). Hence, summing over i,

$$\left|\sum_i(\boldsymbol{s}_i(t) - c\boldsymbol{\omega}_i(t))\right| \leq \sum_i|\boldsymbol{s}_i(t) - c\boldsymbol{\omega}_i(t)| < \epsilon\sum_i\boldsymbol{\omega}_i(t).$$

Using Eq. (5), (6), the previous equation can be written as $|nx_{ave} - nc| < n\epsilon$, which is equivalent to $|x_{ave} - c| < \epsilon$. Hence $c = x_{ave}$. ■

In the next section, we show that, although sums and weights do not reach a consensus, the estimates $\{x_i(t)\}_{1\leq i\leq n}$ converge to a consensus under some conditions.

IV. Convergence

In this section we prove that Weighted Gossip succeeds in other cases than just Uniform Gossip.

Assumption 1: $\{\boldsymbol{D}(t)\}_{t>0}$ is a stationary and ergodic sequence of stochastic matrices with positive diagonals, and $\mathbb{E}[\boldsymbol{D}]$ is irreducible.

Irreducibility means that the graph formed by edges (i,j) such that $\mathbb{P}[\mathbf{D}_{ij} > 0] > 0$ is connected, which requires the connectivity of the network. Note that i.i.d. sequences are stationary and ergodic. Stationarity implies that $\mathbb{E}[\boldsymbol{D}]$ does not depend on t. Positive diagonals means that each node should always keep part of its sum and weight: $\forall i, t, \mathbf{D}_{ii}(t) > 0$.

Theorem 4.1 (Main Theorem): Under Assumption 1, Weighted Gossip using $\{\boldsymbol{D}(t)\}_{t>0}$ converges to a consensus with probability 1, i.e. $\lim_{t\to\infty}\boldsymbol{x}(t) = x_{ave}\mathbf{1}$.

To prove Th. 4.1, we will start by upper bounding the error $\|\mathbf{x}(t) - x_{ave}\mathbf{1}\|_\infty$ with a non-increasing function $f(t)$ (Lemma 4.1): let $\eta_{ji}(t) = \mathbf{P}_{ji}(t) - \sum_{j=1}^{n}\mathbf{P}_{ji}(t)/n = \mathbf{P}_{ji}(t) - \boldsymbol{\omega}_i(t)/n$, then f is defined as $f(t) = \max_{1\leq i\leq n} f_i(t)$, where $f_i(t) = \sum_{j=1}^{n}|\eta_{ji}(t)|/\boldsymbol{\omega}_i(t)$. Then, we will prove that $f(t)$ vanishes to 0 by showing that $\eta_{ji}(t)$ vanishes to 0 (weak ergodicity argument of Lemma 4.3) and that $\boldsymbol{\omega}_i(t)$ is bounded away from 0 infinitely often (Lemma 4.4).

Lemma 4.1: If $\{\mathbf{D}(t)\}_{t>0}$ is a sequence of stochastic matrices, then the function $f(t)$ is non increasing. Furthermore,

$$\|\mathbf{x}(t) - x_{ave}\mathbf{1}\|_\infty \leq \|\mathbf{x}(0)\|_\infty f(t). \tag{7}$$

Proof: By Eq. (3), for any node i,

$$\begin{aligned}
|x_i(t) - x_{ave}| &= \left|\frac{\sum_{j=1}^{n}\mathbf{P}_{ji}(t)x_j(0)}{\boldsymbol{\omega}_i(t)} - x_{ave}\right| \\
&= \left|\frac{\sum_{j=1}^{n}(\boldsymbol{\omega}_i(t)/n + \eta_{ji}(t))x_j(0)}{\boldsymbol{\omega}_i(t)} - x_{ave}\right| \\
&= \left|\frac{\sum_{j=1}^{n}\eta_{ji}(t)x_j(0)}{\boldsymbol{\omega}_i(t)}\right| \\
&\leq \|\mathbf{x}(0)\|_\infty\frac{\sum_{j=1}^{n}|\eta_{ji}(t)|}{\boldsymbol{\omega}_i(t)} \\
&= \|\mathbf{x}(0)\|_\infty f_i(t),
\end{aligned}$$

which proves Eq (7). Next, we need to prove that $f(t)$ is a non-increasing function. For any node i, by Eq. (1) and (2),

$$\begin{aligned} f_i(t) &= \sum_{j=1}^{n} \frac{|\eta_{ji}(t)|}{\boldsymbol{\omega}_i(t)} = \sum_{j=1}^{n} \frac{\left|\sum_{k=1}^{n} \eta_{jk}(t-1)\mathbf{D}_{ki}(t)\right|}{\sum_{k=1}^{n} \boldsymbol{\omega}_k(t-1)\mathbf{D}_{ki}(t)} \\ &\leq \sum_{j=1}^{n} \frac{\sum_{k=1}^{n} |\eta_{jk}(t-1)|\,\mathbf{D}_{ki}(t)}{\sum_{k=1}^{n} \boldsymbol{\omega}_k(t-1)\mathbf{D}_{ki}(t)} \\ &= \frac{\sum_{k=1}^{n}\sum_{j=1}^{n} |\eta_{jk}(t-1)|\,\mathbf{D}_{ki}(t)}{\sum_{k=1}^{n} \boldsymbol{\omega}_k(t-1)\mathbf{D}_{ki}(t)} \\ &\leq \max_k \frac{\sum_{j=1}^{n} |\eta_{jk}(t-1)|\,\mathbf{D}_{ki}(t)}{\boldsymbol{\omega}_k(t-1)\mathbf{D}_{ki}(t)} \qquad (8) \\ &= \max_k \frac{\sum_{j=1}^{n} |\eta_{jk}(t-1)|}{\boldsymbol{\omega}_k(t-1)} \\ &= \max_k f_k(t-1) = f(t-1), \end{aligned}$$

which implies that $f(t) \leq f(t-1)$. Eq. (8) comes from the following equality: for any $\{a_k\}_{1\leq k\leq n} \geq 0, \{b_k\}_{1\leq k\leq n} > 0$,

$$\frac{\sum_{k=1}^{n} a_k}{\sum_{k=1}^{n} b_k} = \sum_{k=1}^{n} \frac{b_k}{\sum_{j=1}^{n} b_j}\frac{a_k}{b_k} \leq \max_k \frac{a_k}{b_k}.$$

■

The following lemma is useful to prove Lemmas 4.3 and 4.4.

Lemma 4.2: Under Assumption 1, there is a deterministic time T and a constant c such that

$$\mathbb{P}[\mathbf{D}(1)\mathbf{D}(2)\ldots\mathbf{D}(T) > c] > 0,$$

where $\boldsymbol{A} > c$ means that every entry of $\boldsymbol{A}$ is larger than c.

Proof: The proof of this lemma can be found in [8]. For the case where $\{\boldsymbol{D}(t)\}_{t>0}$ is i.i.d., a simpler proof can be found in [9]. Note that the theorems proven in [8] and [9] are slightly different than our lemma because the authors multiply matrices on the left, whereas we multiply them on the right. However the multiplication side does not change the proof.

For completeness and simplicity, we give the proof in the i.i.d. case. $\mathbb{E}[\boldsymbol{D}]$ being irreducible and having a positive diagonal, it is primitive as well: there is an $m > 0$ such that $\mathbb{E}[\mathbf{D}]^m > 0$ (elementwise). $\{\mathbf{D}(t)\}_{t\geq 1}$ is i.i.d., hence $\mathbb{E}[\mathbf{D}(1)\mathbf{D}(2)\ldots\mathbf{D}(m)] = \mathbb{E}[\mathbf{D}]^m > 0$, and $\mathbb{P}\left[(\mathbf{D}(1)\mathbf{D}(2)\ldots\mathbf{D}(m))_{ij} > 0\right] > 0$ for any entry (i,j). For any time t, the diagonal coefficients of $\mathbf{D}(t)$ are non-zero, thus, if the $(i,j)^{th}$ entry of $\mathbf{P}(k,k+m-1) = \mathbf{D}(k)\mathbf{D}(k+1)\ldots\mathbf{D}(k+m-1)$ is positive, then $\mathbf{P}_{ij}(t) > 0$ for all $t \geq k+m-1$. Now take $T = n(n-1)m$. The probability that $\mathbf{P}(T) > 0$ is larger than or equal to the joint probability that $\mathbf{P}_{12}(1,m) > 0$, $\mathbf{P}_{13}(m+1,2m) > 0$, …, $\mathbf{P}_{n,n-1}(T-m+1,T) > 0$. By independence of $\{\mathbf{D}(t)\}_{t\geq 1}$,

$$\begin{aligned} \mathbb{P}[\mathbf{P}(T) > 0] &\geq \mathbb{P}[\mathbf{P}_{1,2}(1,m) > 0]\mathbb{P}[\mathbf{P}_{1,3}(m+1,2m) > 0] \\ &\quad \ldots\ldots\mathbb{P}[\mathbf{P}_{n,n-1}(T-m+1,T) > 0] > 0. \end{aligned}$$

Therefore, there is a $c > 0$ such that $\mathbb{P}[\mathbf{D}(1)\mathbf{D}(2)\ldots\mathbf{D}(T) > c] > 0$. ■

Lemma 4.3 (Weak ergodicity): Under Assumption 1, $\{\mathbf{D}(t)\}_{t\geq 1}$ is weakly ergodic.

Weak ergodicity means that when t grows, $\mathbf{P}(t)$ tends to have identical rows, which may vary with t. It is weaker than strong ergodicity, where $\mathbf{P}(t)$ tends to a matrix $\mathbf{1}\boldsymbol{\pi}^T$, where $\boldsymbol{\pi}$ does not vary with t. Interestingly, simple computations show that if $\mathbf{P}(t)$ has identical rows, then consensus is reached. All we need to know in this paper is that weak ergodicity implies that

$$\lim_{t\to\infty} \max_{i,j} \sum_{k=1}^{n} |\mathbf{P}_{ik}(t) - \mathbf{P}_{jk}(t)| = 0,$$

and we suggest [10] for further reading about weak ergodicity.

Proof: Let $\mathbf{Q}$ be a stochastic matrix. The Dobrushin coefficient $\delta(\mathbf{Q})$ of matrix $\mathbf{Q}$ is defined as:

$$\delta(\mathbf{Q}) = \frac{1}{2}\max_{ij} \sum_{k=1}^{n} |\mathbf{Q}_{ik} - \mathbf{Q}_{jk}|\,.$$

One can show [10] that $0 \leq \delta(\mathbf{Q}) \leq 1$, and that for any stochastic matrices $\mathbf{Q}_1$ and $\mathbf{Q}_2$,

$$\delta(\mathbf{Q}_1\mathbf{Q}_2) \leq \delta(\mathbf{Q}_1)\delta(\mathbf{Q}_2). \qquad (9)$$

Another useful fact is that for any stochastic matrix $\mathbf{Q}$

$$1 - \delta(\mathbf{Q}) \geq \max_j \min_i \mathbf{Q}_{ij} \geq \min_{i,j} \mathbf{Q}_{ij}. \qquad (10)$$

A block criterion for weak ergodicity [10] is based on Eq. (9): $\{\mathbf{D}(t)\}_{t\geq 1}$ is weakly ergodic if and only if there is a strictly increasing sequence of integers $\{k_s\}_{s\geq 1}$ such that

$$\sum_{s=1}^{\infty} (1 - \delta(\mathbf{P}(k_s+1, k_{s+1}))) = \infty. \qquad (11)$$

We use this criterion with $k_s = sT$, where T was defined in Lemma 4.2.

A joint consequence of Lemma 4.2 and of Birkhoff's ergodic theorem [11], [8] (in the i.i.d. case, one can use the strong law of large numbers instead) is that the event $\{\mathbf{D}(k_s+1)\mathbf{D}(k_s+2)\ \ldots\mathbf{D}(k_{s+1}) > c\}$ happens infinitely often with probability 1. Hence, using Eq. (10), the event $\{1 - \delta\left(\mathbf{P}(k_s+1, k_{s+1})\right) > c\}$ happens infinitely often with probability 1. We can thus conclude that the block criterion (11) holds with probability 1 and that $\{\mathbf{D}(t)\}_{t\geq 1}$ is weakly ergodic. ■

The next lemma shows that, although weights can become arbitrarily small, they are uniformly large enough infinitely often.

Lemma 4.4: Under Assumption 1, there is a constant α such that, for any time t, with probability 1, there is a time $t_1 \geq t$ at which $\min_i \omega_i(t_1) \geq \alpha$.

Proof: As mentioned in the proof of Lemma 4.3, the event $\{\mathbf{D}(k_s+1)\mathbf{D}(k_s+2)\ \ldots\mathbf{D}(k_{s+1}) > c\}$, where $k_s = sT$, happens infinitely often with probability 1. Let t_1 be the first time larger than t such that $\mathbf{D}(t_1-T+1)\mathbf{D}(t_1-T+2)\ldots\mathbf{D}(t_1) > c$. Then the weights at time t_1 satisfy

$$\begin{aligned} \boldsymbol{\omega}(t_1)^T &= \boldsymbol{\omega}(t_1-T)^T\mathbf{D}(t_1-T+1)\ldots\mathbf{D}(t_1) \\ &> c\boldsymbol{\omega}(t_1-T)^T\mathbf{1}\mathbf{1}^T, \end{aligned}$$

because weights are always positive. Now, because the sum of weights is equal to n, $\boldsymbol{\omega}(t_1 - T)^T \mathbf{1} = n$. Hence $\boldsymbol{\omega}(t_1)^T > cn\mathbf{1}^T$. Taking $\alpha = cn$ concludes the proof. ■

To prove Theorem 4.1, it remains to show that $f(t)$ converges to 0.

Proof: (Theorem 4.1) For any $\varepsilon > 0$, according to Lemma 4.3, there is a time t_0 such that for any $t \geq t_0$,

$$\max_{i,j} \sum_{k=1}^{n} |\mathbf{P}_{ik}(t) - \mathbf{P}_{jk}(t)| < \varepsilon.$$

As a consequence $|\mathbf{P}_{ik}(t) - \mathbf{P}_{jk}(t)| < \varepsilon$ for any i, j, k. Hence $|\eta_{jk}(t)| < \varepsilon$ as well. Indeed,

$$\begin{aligned} |\eta_{jk}(t)| &= \left| \mathbf{P}_{jk}(t) - \sum_{i=1}^{n} \frac{\mathbf{P}_{ik}(t)}{n} \right| = \left| \sum_{i=1}^{n} \frac{\mathbf{P}_{jk}(t) - \mathbf{P}_{ik}(t)}{n} \right| \\ &\leq \sum_{i=1}^{n} \frac{|\mathbf{P}_{jk}(t) - \mathbf{P}_{ik}(t)|}{n} < \sum_{i=1}^{n} \frac{\varepsilon}{n} = \varepsilon. \end{aligned}$$

Therefore, for any $t \geq t_0$ and any $1 \leq i \leq n$,

$$f_i(t) < \frac{n\varepsilon}{\boldsymbol{\omega}_i(t)},$$

and therefore

$$f(t) < \frac{n\varepsilon}{\min_i \boldsymbol{\omega}_i(t)}.$$

Using Lemma 4.4, there is a constant α such that, with probability 1, there is a time $t_1 \geq t_0$ at which $\min_i \boldsymbol{\omega}_i(t_1) \geq \alpha$. Then, for any ε', it suffices to take $\varepsilon = \alpha\varepsilon'/n$ to conclude that there is a time t_1 with probability 1 such that $f(t_1) < \varepsilon'$. Since f is non increasing (Lemma 4.1), for all time $t \geq t_1$, $f(t) < \varepsilon'$; in other words $f(t)$ converges to 0. Using (7) concludes the proof. ■

Remark: A similar convergence result can be proved without Assumption 1 (stationarity and ergodicity of the matrices $\mathbf{D}(t)$), in a setting where the matrices are chosen in a perhaps adversarial manner. One needs only some minimal connectivity assumptions, which then guarantee that there exists a finite number T such that, for all t, all entries of $\mathbf{D}(t+1)\cdots\mathbf{D}(t+T)$ are bounded below by a positive constant c (see, e.g., Lemma 5.2.1 in [12]).

V. One-Way Averaging

In this section, we describe in detail a novel weighted gossip algorithm, which we call One-Way Averaging.

A. Assumptions and Notations

Assume that the network is a random geometric graph on a convex area $\mathcal{A}$, with a connection radius $r(n)$ large enough to enable geographic routing [3]. For every node i, let $\mathcal{Z}_i$ be a distribution of points outside of the area $\mathcal{A}$, and let $\mathcal{H}_i$ be a distribution of integers larger than 2. Each node has an independent local exponential random clock of rate λ, and initiates an iteration when it rings. Equivalently, time is counted in terms of a global and virtual exponential clock of rate $n\lambda$. Each time the global clock rings, a node wakes up independently and uniformly at random. In the analysis, t indicates how many times the global clock rang. A detailed analysis of this time model can be found in [2].

B. Description of One-Way Averaging

Each node i initializes its sum $\boldsymbol{s}_i(0) = \boldsymbol{x}_i(0)$ and its weight $\boldsymbol{\omega}_i(0) = 1$. For any iteration $t > 0$, let i be the node whose clock rings. Node i draws a target Z according to distribution $\mathcal{Z}_i$ and a number $H \geq 2$ of hops according to distribution $\mathcal{H}_i$. Node i chooses uniformly at random a neighbor which is closer to the target Z than itself. If there is no such neighbor then the iteration terminates. If such a neighbor j exists, then node i divides its sum $\boldsymbol{s}_i(t-1)$ and its weight $\boldsymbol{\omega}_i(t-1)$ by H and sends $(\boldsymbol{s}_i(t-1), \boldsymbol{\omega}_i(t-1)) * (H-1)/H$ to node j. It also sends the remaining number $H-1$ of hops and the target Z. Node j adds the received sum and weight to its sum $\boldsymbol{s}_j(t-1)$ and its weight $\boldsymbol{\omega}_j(t-1)$. Then it performs the same operation as node i towards a node that is closer to the target, except that it divides its new sum and weight by $H-1$ instead of H (formally, $H \leftarrow H-1$). Messages are greedily sent towards the target, H being decremented at each hop. The iteration ends when $H = 1$ or when a node does not have any neighbor to forward a message to. At any time, the estimate of any node is the ratio between its sum and its weight.

C. Diffusion Matrices

Suppose that at round t, a whole route of H nodes is generated. Then, after re-indexing nodes starting with the nodes in the route, the diffusion matrix $\mathbf{D}(t)$ can be written as:

$$\begin{pmatrix} 1/H & 1/H & \dots & 1/H & 1/H & \mathbf{0} \\ 0 & 1/(H-1) & \dots & 1/(H-1) & 1/(H-1) & \mathbf{0} \\ 0 & 0 & \ddots & \vdots & \vdots & \mathbf{0} \\ 0 & 0 & 0 & 1/2 & 1/2 & \mathbf{0} \\ 0 & 0 & 0 & 0 & 1 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{Id} \end{pmatrix},$$

where $\mathbf{Id}$ denotes the identity matrix. If the route stops early and has for example only 3 nodes while $H = 4$, then, after re-indexing the nodes, $\mathbf{D}(t)$ can be written as:

$$\begin{pmatrix} 1/4 & 1/4 & 1/2 & \mathbf{0} \\ 0 & 1/3 & 2/3 & \mathbf{0} \\ 0 & 0 & 1 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{Id} \end{pmatrix}.$$

Note that $\mathbf{D}(t)$ is indeed stochastic for all t. It is upper-triangular as well: One-Way Averaging does not require to route information backwards along the path. Furthermore, $\{\mathbf{D}(t)\}_{t>0}$ verifies Assumption 1. First, $\{\mathbf{D}(t)\}_{t>0}$ is an i.i.d. sequence. Second, $\{\mathbf{D}(t)\}_{t>0}$ have positive diagonals. Third, if the network is connected and if the routes generated by distributions $\{\mathcal{Z}_i\}_{1\leq i\leq n}$ and $\{\mathcal{H}_i\}_{1\leq i\leq n}$ connect the network, then $\mathbb{E}[\mathbf{D}]$ is irreducible. Therefore, One-Way Averaging is a successful distributed averaging algorithm. Finally, routes can cross each other without corrupting the algorithm (the resulting diffusion matrices are still stochastic).

D. Simulation

One-Way Averaging and Path Averaging were run (Matlab) on random geometric graphs on the unit square, using the same routes for a fair comparison. At each iteration t, the number $H(t)$ of hops was generated with $\mathcal{H}$ uniform in $[\lceil 1/\sqrt{2}r(n)\rceil, \lceil \sqrt{2}/r(n)\rceil]$ and the target $Z(t)$ was drawn in the following way: let I be the coordinates of the woken node, and let U be a point drawn uniformly at random in the unit square, then

$$Z(t) = I + 3\frac{U - I}{\|U - I\|_2}.$$

Let $C(t_1, t_2)$ be the message cost of a given algorithm from iteration t_1 to iteration t_2. For One-Way Averaging, $C(t_1, t_2) = \sum_{t=t_1}^{t_2} R(t)$, where $R(t) \leq H(t)$ is the effective route length at iteration t. Because Path Averaging routes information back and forth, the cost of one iteration is taken to be equal to twice the route length: $C(t_1, t_2) = 2\sum_{t=t_1}^{t_2} R(t)$. Let $\boldsymbol{\epsilon}(t) = \mathbf{x}(t) - x_{ave}\mathbf{1}$. The *empirical consensus cost* is defined as:

$$C^{emp}(t_1, t_2) = \frac{C(t_1, t_2)}{\log \|\boldsymbol{\epsilon}(t_1)\| - \log \|\boldsymbol{\epsilon}(t_2)\|},$$

so that

$$\|\boldsymbol{\epsilon}(t_2)\| = \|\boldsymbol{\epsilon}(t_1)\| \exp\left(-\frac{C(t_1, t_2)}{C^{emp}(t_1, t_2)}\right).$$

In Fig. 1, we display the empirical consensus cost of both algorithms, with $t_1 = 750$ and t_2 growing linearly with n. We can see that One-Way Averaging performs better than Path Averaging on this example. Although One-Way Averaging converges slower in terms of iterations, spending twice as few messages per iteration is sufficient here to outperform Path Averaging.

The speed of convergence depends on the network but also on $\{\mathcal{Z}_i\}_{1\leq i\leq n}$ and $\{\mathcal{H}_i\}_{1\leq i\leq n}$, which we have not optimized. It would be interesting in further work to compute the speed of convergence of Weighted Gossip, and to derive optimal distributions $\{\mathcal{Z}_i\}_{1\leq i\leq n}$ and $\{\mathcal{H}_i\}_{1\leq i\leq n}$ for a given network using One-Way Averaging. As a conclusion, One-Way Averaging seems to have the same diffusive qualities as Path Averaging while being more robust at the same time.

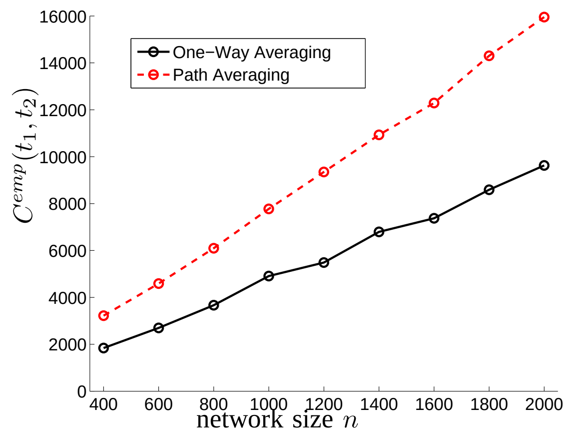

Fig. 1. Comparison of the consensus cost for One-Way Averaging and Path Averaging in random geometric graphs of increasing sizes n. The connection radius scales as $r(n) = \sqrt{6 \log n/n}$. Display of $C^{emp}(t_1, t_2)$ averaged over 15 graphs and 4 simulation runs per graph.

Acknowledgements

The work presented in this paper was supported (in part) by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under grant number 5005-67322, and by the NSF under grant ECCS-0701623.

VI. Conclusion

We proved that weighted gossip algorithms converge to average consensus with probability 1 in a very general setting, i.e. in connected networks, with stationary and ergodic iterations, and with a simple stability condition (positive diagonals). We believe that dropping double stochasticity opens great opportunities in designing new distributed averaging algorithms that are more robust and adapted to the specificities of each network. One-Way Averaging for example is more robust than Path Averaging, and it surprisingly consumes fewer messages on simulations. Also, double stochasticity is difficult to enforce in a distributed manner in directed graphs using unidirectional communications. With Weighted Gossip, one could easily build averaging algorithms for directed networks that are reliable enough not to require acknowledgements.

The next step of this work is to compute analytically the speed of convergence of Weighted Gossip. In classical Gossip, double stochasticity would greatly simplify derivations, but this feature disappears in Weighted Gossip, which makes the problem more difficult.

References

[1] D. Kempe, A. Dobra, and J. Gehrke, “Gossip-based computation of aggregate information,” in *FOCS*, vol. 44. IEEE, 2003, pp. 482–491.

[2] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, “Gossip algorithms : Design, analysis and applications,” in *IEEE, INFOCOM*, 2005.

[3] A. G. Dimakis, A. D. Sarwate, and M. J. Wainwright, “Geographic gossip: efficient aggregation for sensor networks,” in *ACM/IEEE Symposium on Information Processing in Sensor Networks*, 2006.

[4] B. Nazer, A. Dimakis, and M. Gastpar, “Local interference can accelerate gossip algorithms,” in *Allerton Conference on Communication, Control, and Computing*, 2008, pp. 591–598.

[5] F. Bénézit, A. Dimakis, P. Thiran, and M. Vetterli, “Order-optimal consensus through randomized path averaging,” *submitted for publication*.

[6] A. Nedic, A. Olshevsky, A. Ozdaglar, and J. Tsitsiklis, “On distributed averaging algorithms and quantization effects,” in *IEEE Conference on Decision and Control*, 2008, pp. 4825–4830.

[7] B. Gharesifard and J. Cortés, “When does a digraph admit a doubly stochastic adjacency matrix?” in *submitted to the American Control Conference, 2010*.

[8] A. Tahbaz-Salehi and A. Jadbabaie, “Consensus over ergodic stationary graph processes,” *IEEE Transactions on Automatic Control*, 2009.

[9] ——, “Necessary and sufficient conditions for consensus over random independent and identically distributed switching graphs,” in *IEEE Conference on Decision and Control*, 2007, pp. 4209–4214.

[10] P. Brémaud, *Markov Chains. Gibbs Fields, Monte Carlo Simulation, and Queues*. Springer, 1999.

[11] R. Durrett, *Probability: theory and examples*. Duxbury Press Belmont, CA, 1996.

[12] J. Tsitsiklis, “Problems in decentralized decision making and computation,” Ph.D. dissertation, M. I. T., Dept. of Electrical Engineering and Computer Science, 1984.