PROGRESSIVE QUANTIZATION IN DISTRIBUTED AVERAGE CONSENSUS

Dorina Thanou[†]*, Effrosyni Kokiopoulou*[‡] *and Pascal Frossard*[†]

[†]Signal Processing Laboratory (LTS4), Ecole Polytechnique Fédérale de Lausanne (EPFL)
[‡] Seminar for Applied Mathematics (SAM), ETH Zurich
Email:{dorina.thanou, pascal.frossard}@epfl.ch

ABSTRACT

We consider the problem of distributed average consensus in a sensor network where sensors exchange quantized information with their neighbors. In particular, we exploit the increasing correlation between the exchanged values throughout the iterations of the consensus algorithm in order to design a novel quantization scheme, particularly efficient at low bit rates. We implement a low complexity, uniform quantizer in each sensor, where refined quantization is achieved by progressively reducing the quantization intervals with the convergence of the consensus algorithm. We propose a recurrence relation for computing the quantization parameters that depend on the network topology and the communication rate. Finally, simulation results demonstrate the effectiveness of the progressive quantization scheme that leads to the consensus solution even at low communication rate.

Index Terms— Distributed average consensus, sensor networks, progressive quantization.

1. INTRODUCTION

Distributed consensus algorithms have attracted a lot of research interest due to their applications in wireless network systems. They are mainly used in ad-hoc sensor networks in order to compute the global average of sensor data in a distributed fashion, using only local inter-sensor communication. Some of their most important applications include distributed coordination and synchronization in multi-agent systems, distributed estimation, distributed classification and distributed control problems.

While in theory convergence to the global average is mostly dependent on the sensor network topology, the performance of distributed average consensus algorithms is largely connected to the power or communication constraints and limited precision operations in practical systems. In general, the information exchanged by the network nodes has to be quantized prior to transmission due to limited communication bandwidth and limited computational power. However, this quantization process induces some quantization noise that is accumulated throughout the iterative consensus algorithm and affects its convergence, leading to significant performance degradation [1].

A few works have been proposed recently to address the problem of quantization in distributed average consensus. In particular, it was shown in [1] that if the quantization noise is modeled as white and additive with fixed variance then consensus cannot be achieved. The authors in [2] propose a probabilistic quantization scheme that reaches consensus almost surely to a random variable whose expected value is equal to the desired average. Moreover, Kashyap et al. [3] designed an average consensus algorithm with the additional constraint that the states of the agents are integers while the authors in [4] introduced a predictive coding scheme that exploits the temporal correlation among successive iterations. On the other hand, modifications of the classical consensus algorithm have been proposed in [5, 6] where the average is preserved at each iteration with good convergence properties. In general, all the above mentioned algorithms either maintain the average value in the network but cannot reach a consensus effectively, or converge to a random variable that is not always the target average value. More recently, quantization strategies have been proposed in [7] and [8] that maintain the average of the initial states and at the same time converge asymptotically to the true average value. Although these last two solutions perform quite well at high bit rates, the convergence rate appears to be slow when the quantization is coarse. In addition, the stability of both quantization schemes depends on the choice of globally defined parameters that do not seem easy to determine a priori.

In this paper, we address the problem of average consensus with quantized communication and we overcome the limitations of the above algorithms by designing a novel progressive quantization scheme that limits the quantization noise and, contrary to the existing works, leads to faster convergence to the average value even at low bit rates. Motivated by the observation that the correlation between the values communicated by the nodes increases with the consensus iterations, we propose to progressively reduce the range of the quantizer in order to refine the information exchanged in the network. The proposed quantization scheme is consistent and of reduced complexity since at each iteration all the nodes implement the same quantization scheme, with the same parameters. We describe a method for computing offline the parameters of the quantizer, which depend on the network topology and the communication constraints. We illustrate the performance of the proposed scheme through simulations, which confirm that consensus to the true average is achieved even in the case where the information is hardly quantized.

2. PROGRESSIVE QUANTIZER FOR DISTRIBUTED AVERAGE CONSENSUS

We consider a sensor network topology that is modeled as a weighted, undirected graph $\mathcal{G} = (V, \mathcal{E})$, where $V \in \{1, \ldots, m\}$ represents the sensor nodes and $m = |V|$ denotes the number of nodes. An edge denoted by an unordered pair $\{i, j\} \in \mathcal{E}$, represents a link between two sensor nodes i and j that communicate with each other. Moreover, a positive weight $W(i, j) > 0$ is assigned to each edge if $\{i, j\} \in \mathcal{E}$. The set of neighbors for node i is finally denoted as $\mathcal{N}_i = \{j|\{i, j\} \in \mathcal{E}\}$.

[‡] Current address: Google Research, Zurich. Email: kokiopou@google.com

978-1-4673-0046-9/12/$26.00 ©2012 IEEE

The node states over the network at time t can be expressed as a vector $z_t = [z_t(1), \dots, z_t(m)]^T$, where $z_t(i)$ represents a real scalar assigned to node i at time t. The distributed average consensus problem consists in computing iteratively at every node the average $\mu = \frac{1}{m}\sum_{i=0}^{m} z_0(i)$, where $z_0(i)$ is the initial state at sensor i. Under ideal conditions, consensus can be achieved by linear iterations of the form $z_{t+1} = W z_t$, where the symmetric weight matrix W satisfies the conditions that are required to achieve asymptotic average consensus [9], expressed as

$$\mathbf{1}^T W = \mathbf{1}^T,\ W\mathbf{1} = \mathbf{1},\ \rho(W - \mathbf{1}\mathbf{1}^T/m) < 1, \tag{1}$$

with $\rho(\cdot)$ the spectral radius of the matrix and $\mathbf{1}$ the vector of ones.

When the communication rate is limited, the value $z_t(i)$ of a sensor node i at each step t is quantized prior to its transmission to neighbor nodes. The quantized value $\hat{z}_t(i)$ can be written as

$$\hat{z}_t(i) = z_t(i) + \epsilon_t(i), \tag{2}$$

where $\epsilon_t(i)$ models the additive quantization noise of sensor i at iteration t. We assume that the initial sensors' states lie in a finite interval, with $z_0^{(\min)}, z_0^{(\max)}$ the minimum and the maximum values of the interval respectively. In the case of a n-bit uniform quantizer, the quantized value can be written as $\hat{z}_t(i) = \left\lfloor \frac{z_t(i) - z_0^{(\min)}}{\Delta} \right\rfloor \cdot \Delta + \frac{\Delta}{2} + z_0^{(\min)}$. The parameter $\Delta = S/2^n$ is the quantization step-size while the size of the interval, i.e., the quantization range, is $S = z_0^{(\max)} - z_0^{(\min)}$.

In the presence of quantization noise, we assume the following linear iterations that preserve the average of the initial states [5]

$$z_{t+1} = z_t + (W - I)\hat{z}_t, \tag{3}$$

where I is the identity matrix. An analytical expression of Eq.(3) shows that the quantization error propagates through the iterations of the consensus algorithm. More specifically, the states $z_{t+1}, \hat{z}_t$ are expressed as

$$\begin{aligned} \hat{z}_t &= W^t z_0 + \sum_{s=0}^{t-1} W^s (W - I)\epsilon_{t-s-1} + \epsilon_t \\ z_{t+1} &= W^{t+1} z_0 + \sum_{s=0}^{t} W^s (W - I)\epsilon_{t-s}. \end{aligned} \tag{4}$$

We note that the effect of the accumulated quantization noise becomes particularly significant at low bit rate. However, as the number of iterations increases, the correlation between the sensors' states increases and the transmitted values fall into an interval of decreasing size. Hence, quantization in the range S results in a waste of bits or in limited precision that prevents the algorithm to converge to the true average value. We therefore propose a new progressive quantization algorithm that adapts the quantization step-size as the number of linear iterations increases. We keep a simple uniform quantizer with a fixed number of bits, but we reduce the quantization range so that quantization becomes finer along the iterations.

In particular, we denote as $S_t(i)$ the quantization range of the proposed scheme in node i at time t. This range is adapted in each sensor as the iterations proceed and decreases over time. The quantizer is further centered around the previously quantized value of the consensus algorithm $\hat{z}_{t-1}(i)$. In more details, the sensor i encodes its state $z_{t+1}(i)$ by using a quantization range that is defined as $[\hat{z}_t(i) - S_{t+1}(i)/2, \hat{z}_t(i) + S_{t+1}(i)/2]$, where $S_{t+1}(i) > 0$. The data is uniformly quantized in this reduced range, which leads to a step-size $\Delta_{t+1} = \frac{S_{t+1}(i)}{2^n}$ that decreases over time. The values falling out of the quantization range are clipped and coded to the nearest quantizer value. The size of the quantization range however stays identical for all the sensors, independently of their previous state and their position in the network (i.e., $S_{t+1}(i) = S_{t+1},\ \forall i = 1, \dots, m$). The latter simplifies the design of the quantizer as the only parameter to be determined is the size of the quantization range S_{t+1} at each iteration. Since each neighbor node $j \in \mathcal{N}_i$ knows the value $\hat{z}_t(i)$ received at the previous iteration, it is able to perform inverse quantization and to compute correctly the value $\hat{z}_{t+1}(i)$. We call the proposed quantization scheme Progressive Quantizer.

3. DESIGN OF THE PARAMETERS OF THE PROGRESSIVE QUANTIZER

In this section, we propose a constructive method to compute the size of the quantization range S_{t+1} a priori, based on the properties of the network topology and the communication constraints. For effective quantization, the quantization range should be chosen such that the values computed in the nodes fall with high probability within the quantization range. Thus, for each sensor i, the value $z_{t+1}(i)$ should fall within the range $[\hat{z}_t(i) - S_{t+1}/2, \hat{z}_t(i) + S_{t+1}/2]$. If this is not the case, it is mapped to the representative value of the closest interval. Hence, we need to compute S_{t+1} such that the absolute difference between two consecutive sensor states is, with high probability, upper-bounded by $S_{t+1}/2$. One way to estimate S_{t+1} is to bound the mean square difference between two successive values such that

$$E[\|z_{t+1} - \hat{z}_t\|^2] \leq m \left(\frac{S_{t+1}}{2} \right)^2. \tag{5}$$

In this work, $\|\cdot\|$ denotes the L2 norm. Since the quantization range should be positive, without loss of generality, we pose first $S_{t+1} = 2 \cdot e^{-\beta_{t+1}}$. Hence determining the size of the quantization range becomes equivalent to computing β_{t+1}. In the sequel, we derive first an upper-bound of $E[\|z_{t+1} - \hat{z}_t\|^2]$ that depends on the previous $\{\beta_1, ..., \beta_t\}$ values. In order to derive the upper-bound we have exploited the properties of the matrix W and we have modeled the quantization noise samples $\epsilon_t(i)$ in (2) as (spatially and temporally) independent random variables that are uniformly distributed with zero mean and variance $\Delta_t^2/12$ [10]. In particular, the upper-bound of $E[\|z_{t+1} - \hat{z}_t\|^2]$ follows from the following Proposition whose technical details are provided in [11], due to lack of space.

Proposition 1 *Let $\hat{z}_t$ and z_{t+1} defined as in (4). Let also $\lambda_2 := \rho(W - \frac{\mathbf{1}\mathbf{1}^T}{m})$ and $\lambda_{\min}$ be the smallest algebraically eigenvalue of W. Then, it holds that*

$$\begin{aligned} E[\|z_{t+1} - \hat{z}_t\|^2] \leq{}& \|z_0\|^2 \lambda_2^{2t} (1 - \lambda_{\min})^2 \\ &+ (1 - \lambda_{\min})^2 \sum_{s=0}^{t-1} \|W^s (W - I)\|^2 m \frac{S_{t-s-1}^2}{2^{2n} \cdot 12} \\ &+ (2 - \lambda_{\min})^2 m \frac{S_t^2}{2^{2n} \cdot 12}. \end{aligned} \tag{6}$$

Eqs. (5) and (6) along with the fact that $\|z_0\|^2 \leq m\|z_0\|_\infty^2$ and

$S_{t+1} = 2 \cdot e^{-\beta_{t+1}}$, imply that

$$\begin{aligned} e^{-2\beta_{t+1}} &= \|z_0\|_\infty^2 \lambda_2^{2t} (1-\lambda_{\min})^2 \\ &\quad + (1-\lambda_{\min})^2 \sum_{s=0}^{t-1} \|W^s(W-I)\|^2 \frac{e^{-2\beta_{t-s-1}}}{2^{2n}\cdot 3} \\ &\quad + (2-\lambda_{\min})^2 \frac{e^{-2\beta_t}}{2^{2n}\cdot 3}, \quad t \geq 1. \end{aligned} \tag{7}$$

The formula above leads to a recursive computation of β_{t+1} at each time step $t+1$ of the consensus algorithm. As boundary conditions for the recursion, we first compute β_0 using the initial range of the quantizer i.e., $2e^{-\beta_0} = \Delta_0$ and next compute β_1 from a simplified version of (7) where the intermediate term from the right hand side is dropped. Finally, we note that the values of the quantization range depend on the convergence rate λ_2 of the average consensus algorithm in the absence of quantization noise, on the absolute maximum value of the initial data $\|z_0\|_\infty^2 = \max\{|z_0^{(\min)}|, |z_0^{(\max)}|\}$, on the network topology W and on the number of quantization bits n.

4. SIMULATION RESULTS

In this section we provide simulation results that verify the effectiveness of the Progressive Quantizer. We consider a network of 40 sensors (i.e., $m = 40$) following the random geographic graph model, i.e., the sensors are uniformly random distributed over the unit square $[0,1] \times [0,1]$. We assume that two neighbor sensors are connected if their Euclidean distance is less than $r = \sqrt{(\log m)/m}$, which ensures connectivity with high probability [12]. We assume static network topologies, which implies that the edge set does not change over the iterations. As an illustration, we consider the Metropolis weight matrice [9] defined as:

$$W[i,j] = \begin{cases} \frac{1}{1+\max\{d(i),d(j)\}}, & \text{if } \{i,j\} \in \mathcal{E} \\ 1 - \sum_{(i,k)\in\mathcal{E}} W[i,k], & \text{if } i = j \\ 0, & \text{otherwise,} \end{cases}$$

where $d(i)$ denotes the degree of the i^{th} sensor.

4.1. Evolution of the quantization range

We first validate the decreasing behavior of the quantization range $S_t = 2 \cdot e^{-\beta_t}$ over the iterations. For implementation issues, we fix a parameter $\delta = 10^{-16}$. At iteration t, if the quantization range S_t becomes smaller than δ, we quantize with the range computed at the previous iteration i.e., we set $S_t = S_{t-1}$. We compute recursively the values $\{\beta_1, ..., \beta_{t+1}\}$ from Eq.(7) for a random network topology with initial states in the range $[0,1]$ and communication rates of $n = [2,4,6]$ bits. We observe in Fig.1(a) that the value of β_t appears to increase linearly with the number of iterations, which implies that Eq.(7) leads to a size of the quantization range that decreases exponentially over time. Finally, we notice that the slope of the function β_t is independent of the bitrate, while the y-intercept value depends on the number of quantization bits.

4.2. Average consensus performance of the Progressive Quantizer

We first compare the proposed quantization scheme (ProgQ) with a baseline uniform quantizer with a constant range $S = 1$ (UnifQ). Fig. 1(b) illustrates the average consensus performance corresponding to the absolute error $\|z_t - \mu\mathbf{1}\|_2$ versus the number of iterations for $n = [2,4,6]$ bits. In order to obtain statistically meaningful results we average the error over 200 random realizations of the network topology with random initial values. Observe that the performance of the proposed quantization scheme is very satisfactory even at a very low bit rate (2 bits) within a few iterations. In particular, the error $\|z_t - \mu\mathbf{1}\|_2$ shows a decreasing behavior over the iterations, which means that the quantizer does not saturate. It rather follows the evolution of the average consensus algorithm in the noiseless case. On the other hand, the performance of the uniform quantizer with a constant range saturates even at high bit rate.

In addition, we compare the proposed Progressive Quantizer with (a) the adaptive quantizer (AdaptQ) [7] and (b) the zoom in-zoom out uniform encoder (ZoomQ) [8]. In particular, the scheme proposed in [7] is based on the Delta modulation with variable step-size. The step-size is adapted throughout the iterations based on the previously sent bits and a constant K. However, the scheme is quite sensitive to the value of K and the performance can deteriorate for non-carefully chosen values. On the other hand, the differential encoding scheme proposed in [8] uses a uniform quantizer and the transmitted value is the quantized difference of the current value from the previous estimate, scaled by a factor f that is adapted over time. This factor is similar to the step-size of [7] and it grows or decreases depending on the difference between the new state z_{t+1} and the previously quantized state $\hat{z}_t$. The decrease or the increase depends on the constants k_{in} and k_{out} respectively and the way that these constants have to be determined seems to be an open question. In our experiments, for the first scheme (AdaptQ) we choose $K = 1.2$ as defined in [7], while for the second scheme (ZoomQ) we choose the parameters $k_{in} = 0.5, k_{out} = 2$ and the scaling factor $f_0 = 0.5$ as defined in [8]. Fig 2 illustrates the obtained results. Notice that our scheme outperforms both AdaptQ [7] and ZoomQ [8]. AdaptQ appears to saturate especially for a small number of bits. The performance of ZoomQ seems to be quite good for 4 and 6 bits, but it suffers at low bit rate.

Our scheme bear some resemblance with these two schemes in the sense that we also propose to adapt a scaling function with the difference that the scaling function in our case has a very specific definition that consists of the sensors' dynamic range. Moreover, we impose that range to consistently decrease, which is intuitively supported by the increasing correlation of the sensors' states throughout the iterations. The latter is done by taking into consideration both the available number of bits and the converging behavior of the average consensus algorithm. Finally, the a priori estimation of the quantization range based on the recursive Eq.(7) reduces the complexity of the quantizer in comparison to the previously mentioned schemes where the step-size is adapted online, independently for each sensor.

5. CONCLUSIONS

In this paper, we have proposed a novel quantization scheme for solving the average consensus problem when sensors exchange quantized state information. In particular, our scheme is based on progressively reducing the range of a uniform quantizer with constant bit rate. The proposed quantizer exploits efficiently the available number of bits without increasing the communication complexity over the network. Simulation results show the effectiveness of our scheme that outperforms the existing solutions and leads to convergence to the average value even at low bit rates.

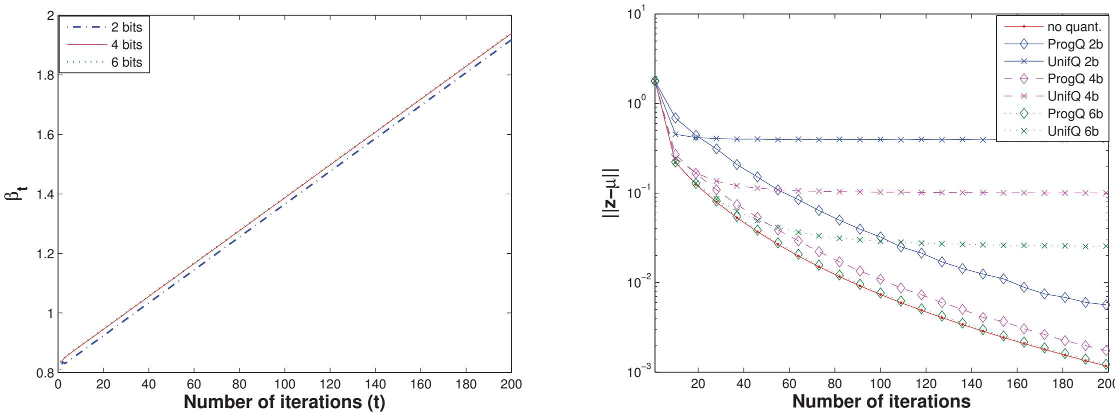

(a) (b)

Fig. 1. (a) Evolution of the β_t values over the iterations. (b) Average consensus performance of the proposed quantization scheme (ProgQ) vs uniform quantizer with a constant range (UnifQ) for 2, 4 and 6 bits.

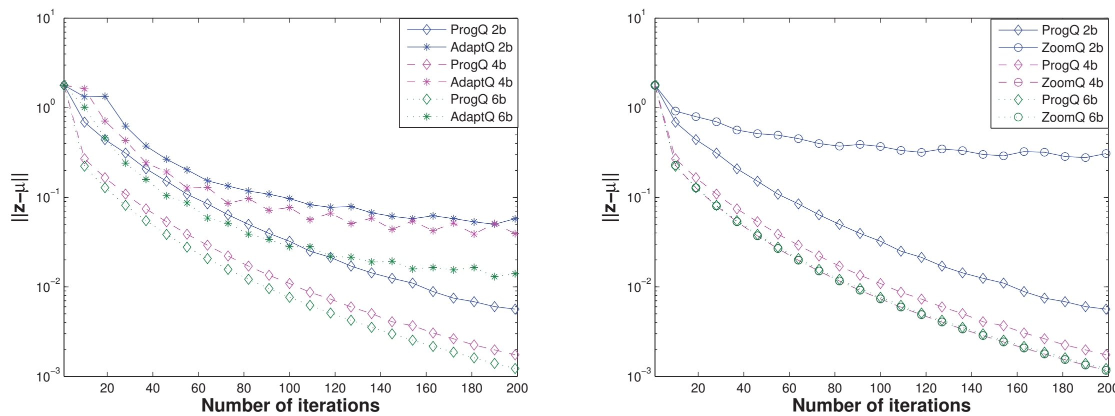

(a) (b)

Fig. 2. Average consensus performance of the proposed quantization scheme (ProgQ) vs (a) adaptive quantizer (AdaptQ) [7] and (b) zoom-in, zoom-out uniform quantizer (ZoomQ) [8] for 2, 4 and 6 bits.

6. REFERENCES

[1] L. Xiao, S. Boyd, and S-J. Kim, "Distributed average consensus with least-mean-square deviation," *Journal of Parallel and Distributed Computing*, vol. 67, no. 1, pp. 34–46, 2007.

[2] T. Aysal, M.Coates, and M. Rabbat, "Distributed Average Consensus with Dithering Quantization ," *IEEE Trans. Signal Process.*, vol. 56, no. 10, Oct. 2008.

[3] T. Basar A. Kashyap and R. Srikant, "Quantized consensus," *Automatica*, vol. 43, no. 7, pp. 1192–1203, Apr. 2007.

[4] M. Yildiz and A. Scaglione, "Coding with side information for rate-constrained consensus," *IEEE Trans. Signal Process.*, vol. 56, no. 8, Aug. 2008.

[5] P. Frasca, R. Carli, F. Fagnani, and S. Zampieri, "Average consensus on networks with quantized communication," *Inter. Journal of Robust and Nonlinear Control*, vol. 19, no. 16, pp. 1787–1816, Nov. 2009.

[6] J. Fang and H. Li, "Distributed consensus with quantized data via sequence averaging ," *IEEE Trans. Signal Process.*, vol. 58, no. 2, Feb. 2010.

[7] J. Fang and H. Li, "An adaptive quantization scheme for distributed consensus," in *Proc. IEEE Int. Conf. Acoust., Speech and Signal Processing*, Taipei, Taiwan, Apr. 2009.

[8] R. Carli, F. Bullo, and S. Zampieri, "Quantized average consensus via dynamic coding/decoding schemes," *Inter. Journal of Robust and Nonlinear Control*, vol. 20, no. 2, pp. 156–175, 2010.

[9] L. Xiao and S. Boyd, "Fast linear iterations for distributed averaging," *System and Control Letters*, vol. 53, pp. 65–78, Feb. 2004.

[10] A. Sripad and D. Snyder, "A necessary and sufficient condition for quantization errors to be uniform and white," *IEEE Trans. Acoust., Speech and Signal Processing*, vol. 25, no. 5, pp. 442–448, Oct. 1977.

[11] D. Thanou, E. Kokiopoulou, and P. Frossard, "Progressive quantization in distributed average consensus," *Technical Report, available at http://arxiv.org/pdf/1105.1074*, May 2011.

[12] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. on Information Theory*, vol. 46, no. 2, pp. 388–404, March 2000.