POLYNOMIAL FILTER DESIGN FOR QUANTIZED CONSENSUS

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

[†]Signal Processing Laboratory (LTS4), Swiss Federal Institute of Technology, Lausanne (EPFL)
[*]Multimedia Communications and Networking Laboratory, Ewha Womans University, Seoul, Korea
[‡] Seminar for Applied Mathematics (SAM), Swiss Federal Institute of Technology, Zürich (ETHZ)
Email:{dorina.thanou, pascal.frossard}@epfl.ch, hyunggon.park@ewha.ac.kr, effrosyni.kokiopoulou@sam.math.ethz.ch

ABSTRACT

We consider the problem of distributed average consensus where sensors exchange quantized data with their neighbors. We deploy a polynomial filtering approach in the network nodes in order to accelerate the convergence of the consensus problem. The quantization of the values computed by the sensors however imposes a careful design of the polynomial filter. We first study the impact of the quantization noise in the performance of accelerated consensus based on polynomial filtering. It occurs that the performance is clearly penalized by the quantization noise, whose impact directly depends on the filter coefficients. We then formulate a convex optimization problem for determining the coefficients of a polynomial filter, which is able to control the quantization noise while accelerating the convergence rate. The simulation results show that the proposed solution is robust to quantization noise while assuring a high convergence speed to the average value in the network.

Index Terms— Distributed averaging, distributed consensus, polynomial filtering, uniform quantization.

1. INTRODUCTION

Distributed average consensus (DAC) algorithms are becoming increasingly popular and are attractive for 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 agreement and synchronization problems, distributed coordination of mobile autonomous agents and distribution data fusion in sensor networks (e.g., [1, 2]).

Without any communication rate restriction and considering that the data are sent over a reliable channel, the convergence rate of the distributed average consensus problem is accelerated significantly. Apart from the classical approach based on linear iterations (successive multiplications of the network weight matrix with the vector of initial sensor values [3], [4]) a more efficient method tries to accelerate the convergence rate by using previous estimates [5]. This approach can be achieved by applying a matrix polynomial on the weight matrix, in order to shape its spectrum by minimizing its second largest eigenvalue. Polynomial filtering has been shown to outperform the simple iterative method in terms of convergence speed and robustness to dynamic topologies. Both methods allow every node state to converge to the average of the initial values after some iterations.

Note that these approaches have been designed and optimized based on the assumption that there is no error in the information exchanged among nodes. However, in practice, this assumption is generally infeasible due to several constraints such as limited communication bandwidth, unreliable communication channels, limited computation power, etc. The information exchanged by the nodes has therefore to be *quantized* in order to reduce the communication overheads. As a result, however, this incurs quantization noise that is accumulated during the iterations. Therefore, existing consensus acceleration solutions only provide limited performance in the presence of quantization noise.

While there exists a substantial body of work that discusses average consensus problems with quantized communication (e.g., [6–13]), they all assume that consensus is achieved by linear iterations. Thus, these prior works do not explore the impact of the quantization noise on the acceleration methods. In this paper, we propose an algorithm that enables the polynomial filtering acceleration method to become robust to quantization noise, while ensuring fast convergence. We analytically investigate the impact of the quantization noise on the performance of this approach and show that both the convergence rate and the accuracy of the average consensus depend on the filter coefficients. We show that the coefficients can be efficiently obtained by solving a convex optimization problem. Finally, we study the tradeoff between the convergence rate and the accuracy of the average consensus, which provides a guideline for the polynomial filter design.

The rest of the paper is organized as follows. In Section 2, we briefly review the distributed average consensus problem based on the polynomial filtering methodology. We propose a noise robust polynomial filter design in Section 3. Several simulation results are presented in Section 4, and the conclusions are drawn in Section 5.

2. PRELIMINARIES

A sensor network topology is modeled as an undirected graph $G = (V, \mathrm{E})$, where $V \in \{1, \ldots, m\}$ represents the sensor nodes and $m = |V|$ denotes the number of nodes. An edge that represents a link between two sensor nodes i and j is denoted by

This work was supported by the Swiss National Science Foundation grants 200021-118230. This work was done while H. Park was with EPFL.

ISSN 2076-1465

an unordered pair $\{i, j\} \in \mathrm{E}$, which can be established if sensors i and j communicate with each other. We denote the set of neighbors for node i as $N_i = \{j | \{i, j\} \in \mathrm{E}\}$.

By denoting with $z_t(i)$ a real scalar assigned to node i at time t, the node states (i.e., node values) over the network at t can be expressed as a vector $z_t = [z_t(1), \ldots, z_t(m)]^T$. Correspondingly, the initial node state is z_0. Then, the distributed average consensus problem consists of computing iteratively at every node the average $\mu = \frac{1}{m}\sum_{i=0}^{m} z_0(i)$. In order to compute μ, we consider distributed linear iterations at each sensor of the following form:

$$z_{t+1}(i) = W[i,i]z_t(i) + \sum_{j \in N_i} W[i, j]z_t(j), \qquad (1)$$

where $W[i, j]$ represents the weight associated with edge $\{i, j\}$. The weight matrix W can be specified by the topology of the network graph. In this paper, we assume that W satisfies the conditions that are required to achieve asymptotic average consensus [3], expressed as

$$1^T W = 1^T,\ W1 = 1,\ \rho(W - 11^T/m) < 1,$$

where $\rho(\cdot)$ denotes the spectral radius of the matrix. It is known that a smaller value of the second largest eigenvalue $\lambda_2(W)$ of W leads to faster convergence [3]. Moreover, the convergence rate can be accelerated by applying a polynomial filter $p_k(\cdot)$ of degree k every $k+1$ steps [5]. Specifically, given W, a polynomial filter p_k is applied to W, leading to

$$p_k(W) = \sum_{l=0}^{k} a_l W^l = a_0 I + a_1 W + \ldots + a_k W^k.$$

This is equivalent to applying the iterative method to $p_k(W)$ and in practice, it implies a periodic update of the current sensor's values while exploiting the memory of the sensors:

$$\begin{aligned} z_{t+k+1} = p_k(W)z_t &= a_0 z_t + a_1 W z_t + \ldots + a_k W^k z_t \\ &= a_0 z_t + a_1 z_{t+1} + \ldots + a_k z_{t+k}. \end{aligned} \qquad (2)$$

The polynomial filter update is followed by the distributed linear iteration of Eq.(1). It is known that the eigenvalues of $p_k(W)$ are simply the polynomial filtered eigenvalues of W, i.e., $p_k(\lambda_i(W))$. Thus the application of the polynomial filter on the spectrum of W can impact the magnitude of $\lambda_2(W)$ which mainly drives the convergence rate. As a result, the convergence rate can be significantly influenced by the polynomial filter design (i.e., selection of the polynomial coefficients). Two alternative techniques for computing the coefficients a_l of the filter p_k are discussed in [5]. The first approach, Newton's interpolation polynomial, is based on Hermite interpolation and its objective is to dampen the smallest eigenvalues of W by imposing smoothness constraints of p_k at the endpoints of the interval of the spectrum of W. The second technique is based on solving a semi-definite program (SDP) for computing the optimal coefficients. The optimization problem can be formulated as:

$$\begin{aligned} \mathbf{a} = \arg\min \eta & \\ \text{subject to } \rho\left(\sum_{l=0}^{k} a_l W^l - 11^T/m\right) &\le \eta \\ \left(\sum_{l=0}^{k} a_l W^l\right) 1 = 1, & \end{aligned} \qquad (3)$$

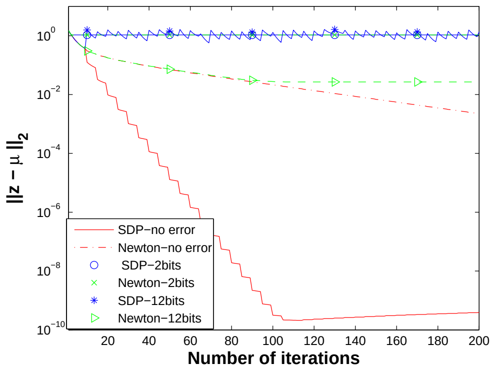

Fig. 1. Average consensus performance of polynomial filter based approaches [5] (no error, quantization with 2 and 12 bits).

This approach outperforms the Newton's interpolating polynomial approach in terms of convergence rate.

The above mentioned filters have been originally designed without considering quantization effects. If quantization noise is introduced, they may provide a limited performance as we show in the following example. We consider that the sensors quantize their state by using a uniform quantizer with 2 and 12 bits respectively. For more details about the simulation settings see Section 4.1. Fig. 1 shows the absolute error $\|z_t - \mu 1\|_2$ over t iterations when SDP-based polynomial filter and Newton's interpolating polynomial filter are deployed. The red curves follow [5] and correspond to the performance of the SDP and the Newton method in the noiseless case. As expected, if quantization noise is introduced in the distributed average consensus, the existing filter based solutions provide only a limited performance. Interestingly, polynomial filtering with Newton's polynomial seems to perform better than the one with SDP polynomial. This result confirms that the optimal coefficients computed with Problem (3) were designed for ideal conditions, without taking into consideration rate constraints. More importantly, both solutions do not improve the consensus performance even when refined quantizers are deployed. This can be attributed to the fact that by applying the polynomial filter, the introduced noise is amplified, resulting in significant performance degradation. Therefore, we can conclude that the polynomial filters need to be redesigned by explicitly considering the impact of the noise on their average consensus performance. This is discussed in the next section.

3. FILTER DESIGN FOR QUANTIZED CONSENSUS

We assume that sensors exchange *quantized* information, which enables them to reduce the communication overhead. Thus, additive quantization errors are introduced in the exchanged information. In particular, we assume that $z(i) \in \mathbb{R}$ lies in a finite interval of size S and is quantized by a q-bit uniform quantizer before it is transmitted to the neighbor sensors. The q-bit uniform quantizer output of a scalar z can be expressed as

$$Q(z) = \left\lfloor \frac{z - z_{min}}{\Delta} \right\rfloor \times \Delta + \frac{\Delta}{2} + z_{min},$$

where Δ is the quantizer step-size and z_{min} is the minimum dynamic range. Note that the quantizer step size and the range S are linked by the relation $S = 2^q \Delta$.

Before state information is exchanged, the value $\tilde{z}_t(i)$ of a sensor node i at each step t is quantized, such that

$$\hat{z}_t(i) = \tilde{z}_t(i) + \varepsilon_t(i), \tag{4}$$

where $\varepsilon_t(i)$ is the incurred quantization error in step t, $\tilde{z}_t(i)$ is the current state of sensor i (before quantization) and $\hat{z}_t(i)$ is the quantized value that the sensor i will send to its neighbors. We set as initial condition $\tilde{z}_0(i) = z_0(i)$ which means that $\hat{z}_0(i) = \tilde{z}_0(i) + \varepsilon_0(i) = z_0(i) + \varepsilon_0(i)$. Then, each node updates its state as a linear combination of its own quantized state as well as the quantized states of its neighbors based on the recursive update in (1):

$$\tilde{z}_{t+1} = W \cdot (\tilde{z}_t + \varepsilon_t), \quad t \geq 0$$

where $\varepsilon_t = [\varepsilon_t(1), \varepsilon_t(2), ..., \varepsilon_t(m)]^T$. After k iterations, the state $\tilde{z}_{t+k}$ can be correspondingly expressed as

$$\tilde{z}_{t+k} = W^k \tilde{z}_t + \sum_{l=0}^{k-1} W^{k-l} \varepsilon_{t+l}.$$

After polynomial filtering (see also Eq. (2)) the resulting $\tilde{z}_{t+k+1}$ can be expressed as

$$\begin{aligned} \tilde{z}_{t+k+1} &= a_0 \tilde{z}_t + \sum_{l=1}^{k} a_l \tilde{z}_{t+l} \\ &= \sum_{l=0}^{k} a_l W^l \tilde{z}_t + \sum_{l=1}^{k} \left[\sum_{j=0}^{l-1} a_l W^{l-j} \varepsilon_{j+t} \right] \\ &= \sum_{l=0}^{k} a_l W^l \tilde{z}_t + \sum_{l=0}^{k-1} \left[\sum_{j=1}^{k-l} a_{l+j} W^j \right] \varepsilon_{l+t}. \end{aligned} \tag{5}$$

The obtained values from the above equation are quantized and sent to the neighboring sensors. We observe that the quantization error due to the previous k steps, is accumulated over the iterations and it is represented by the second term of the above equation. As shown in Eq. (5), the accumulated quantization noise of each sensor significantly depends on the filter coefficients. Specifically, the output of the polynomial filter is a linear combination of the quantization noise vectors $(\varepsilon_t, \ldots, \varepsilon_{t+k})$ introduced during the previous iterations, the weight matrix W and the filter coefficients. We observe also that each quantization noise vector is multiplied by a matrix polynomial of W. Note that the weight matrix depends on the topology and it is fixed. Also, the quantization errors depend on the number of the available bits (rate constraints). Hence the only factor that we can modify in order to reduce the impact of the noise in the consensus performance is the value of the coefficients. Therefore, given W and a polynomial degree k, the polynomial coefficients should be determined such that they can lead to a fast convergence while at the same time reducing the accumulated error.

In order to limit the effects of the quantization noise, the filter coefficients need to be computed such that they minimize $\sum_{j=1}^{k-l} a_{l+j} W^j$ in Eq.(5) for $0 \leq l \leq k-1$. This can be achieved by minimizing the L2-norm of the above matrix polynomials which leads to diminishing the effect on the quantization noise vectors. Moreover, in order to accelerate the convergence rate and assure convergence, the spectral radius $\rho(\sum_{l=0}^{k} a_l W^l - \frac{1}{m} 11^T)$ also needs to be minimized subject to the constraint $(\sum_{l=0}^{k} a_l W^l) 1 = 1$ [5]. Putting all the above facts together, the filter coefficients $\mathbf{a} = (a_0, \ldots, a_k) \in \mathbb{R}^{k+1}$ can be determined as

$$\begin{aligned} &\mathbf{a} = \arg\min \eta \\ &\quad \text{subject to } \rho \left(\sum_{l=0}^{k} a_l W^l - 11^T/m \right) \leq \eta \\ &\quad \left\| \sum_{j=1}^{k-l} a_{l+j} W^j \right\|_2 \leq v \cdot \eta, \; 0 \leq l \leq k-1 \\ &\quad \left(\sum_{l=0}^{k} a_l W^l \right) 1 = 1, \end{aligned} \tag{6}$$

where $\eta \in \mathbb{R}$ is an auxiliary variable. The new optimization problem consists of the constraints in Problem (3) and one additional constraint that controls the allowable amount of quantization noise. Since the spectral radius is a convex function of the polynomial coefficients [5], Problem (6) is also a convex optimization problem. Thus, a global optimal solution can be efficiently obtained. Notice that a parameter v is introduced in the inequality constraints of Problem (6), which determines the tolerable noise level. A lower spectral radius (i.e., faster convergence speed) can be achieved by allowing a higher tolerable noise level (i.e., a larger v), as a larger value of v leads to an enlarged constraint set in Problem (6). Since our primary goal of filter design is to accelerate the convergence speed, we assume that $v > 1$. However, the impact of quantization noise on the average consensus performance increases as v increases. This tradeoff between convergence and robustness to noise will be quantified and investigated in the next section.

4. SIMULATION RESULTS

In this section, we quantify the performance of the proposed noise robust filter design for distributed average consensus and we compare it with that of the already existing methods. The performance is estimated in terms of the convergence speed and the accuracy of the achieved consensus value.

4.1. Simulation setup

We consider a network that consists of 40 sensors (i.e., $m = 40$) uniformly 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 the connectivity radius $r = \sqrt{(\log m)/m}$. Each sensor performs 200 iterations. As an illustration, we consider the maximum-degree weight matrix,

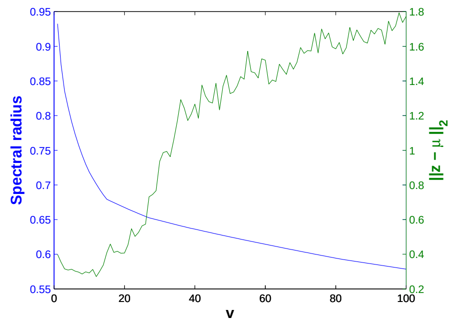

Fig. 2. Impact of v on the spectral radius and the consensus accuracy.

defined as

$$W[i,j] = \begin{cases} 1/m, & \text{if } \{i,j\} \in \mathrm{E} \\ 1-d(i)/m, & \text{if } i=j \\ 0, & \text{otherwise} \end{cases}$$

where $d(i)$ denotes the degree of the i^{th} sensor. We assume static network topologies. This implies that the edge set does not change over the iterations so the matrix W is considered fixed. We deploy a polynomial filter of degree $k=4$ and we investigate the average performance based on 400 random realizations of the sensor network and random initial measurements. For more details about the chosen degree of the polynomial see [5]. In the simulations, the quantization step-size is chosen to be $\Delta = 1/2^q$ and the exchanged data are quantized uniformly with $q = 2, 6, 12$ bits.

4.2. Performance of Noise Robust Polynomial Filters

In this section, we quantify the performance of the noise robust polynomial filter proposed in Section 3. As an illustration, we assume that $k=4$ in Eq.(6). The filter coefficients are determined by solving the convex Optimization Problem (6) in MATLAB using the SeDuMi solver [1].

We first investigate the impact of the parameter v introduced in Problem (6) on the convergence speed and the accuracy of the average consensus algorithm. Fig.2 confirms the tradeoff between these two objectives. The accuracy is measured by computing the absolute error $\|z_t - \mu 1\|_2$ over the iteration $t=200$, with varying v, for a $q=6$ bits quantization. As v increases, the spectral radius decreases (i.e., the convergence is faster). Note that the improvement rate of the spectral radius is significant in the range of $1 < v \leq 15$, while it becomes smaller for $v \geq 16$. However, by increasing v, we loose much in terms of accuracy since we tolerate a high noise level. As a result, v should be determined by considering this tradeoff.

Fig.3 presents the average consensus performance achieved based on the proposed noise robust polynomial filter (denoted by robust-SDP) which is designed for different values of $v = \{3, 10, 50\}$. For comparison purpose, we also present the results achieved based on Newton's polynomial methodology and the simple iterative method.

[1]Publically available at: http://sedumi.mcmaster.ca/

We compare first the performance of the two polynomial filtering methodologies. We observe that for a small number of bits (2 bits) the performance of both methods is extremely deteriorated. However, if we increase the number of the available bits, the proposed noise robust polynomial filter provides a generally improved performance in comparison to Newton's interpolating polynomial filter. These simulation results confirm that the proposed filter is adaptively designed by considering both the convergence speed as well as the accuracy of the average consensus. As we increase the amount of the tolerable noise level ($v=50$) the performance of the robust-SDP tends to deteriorate since, for a large value of the variable v, the additional constraint in the Problem (6) does not have any impact on the constraint set. Thus the solutions that we obtain by solving the Optimization Problem (6) are close to the one obtained by solving the initial Optimization Problem (3) without the noise constraints. By comparing the performance of the two filtering methodologies to the results shown in Fig.1, it becomes clear that the proposed polynomial filter outperforms both original SDP and Newton's interpolating polynomial filter approaches.

For the sake of completeness we compare the robust-SDP method with the simple iterative method. When v is relatively small (i.e., $v=3$) and when the number of the available bits is limited, the performance of the noise robust SDP is similar to that of the simple iterative method. This is consistent with our discussion in Section 3 since a very small amount of noise is allowed at the cost of a low convergence speed. Thus the convergence rate of the robust SDP is close to the one achieved by the simple iterative method as the accelerating effect of polynomial filtering is penalized. As we increase the amount of noise that we tolerate ($v=10$), the convergence speed is significantly improved and the robust SDP clearly outperforms the iterative method. If we keep increasing v, ($v=50$) the impact of the quantization noise on the performance becomes significant and the achieved performance tend to approach the one obtained by solving the Optimization Problem (3). This means that during the first iterations, the sensors will tend to converge faster to some value (due to the lower value of the spectral radius) but the difference of the achieved value from the true average is quite high and it becomes higher each time that we apply the polynomial filter. So even though we would expect a gain in the convergence speed, in practice the increased noise effect does not permit to observe significant differences in terms of convergence rate compared to the performance achieved when $v=10$.

We conclude that even in the case of quantized communication the proposed filtering methodology outperforms the classical iterative algorithm even though the gain in the convergence rate is not as high as in the case when communication is performed under ideal conditions [5]. The above results confirm that the value of v should be determined by considering the tradeoff between the convergence speed and the accuracy of the average consensus.

5. CONCLUSIONS

In this paper, we have investigated the performance of the polynomial filtering methodology for average consensus when sensors exchange quantized state information. We

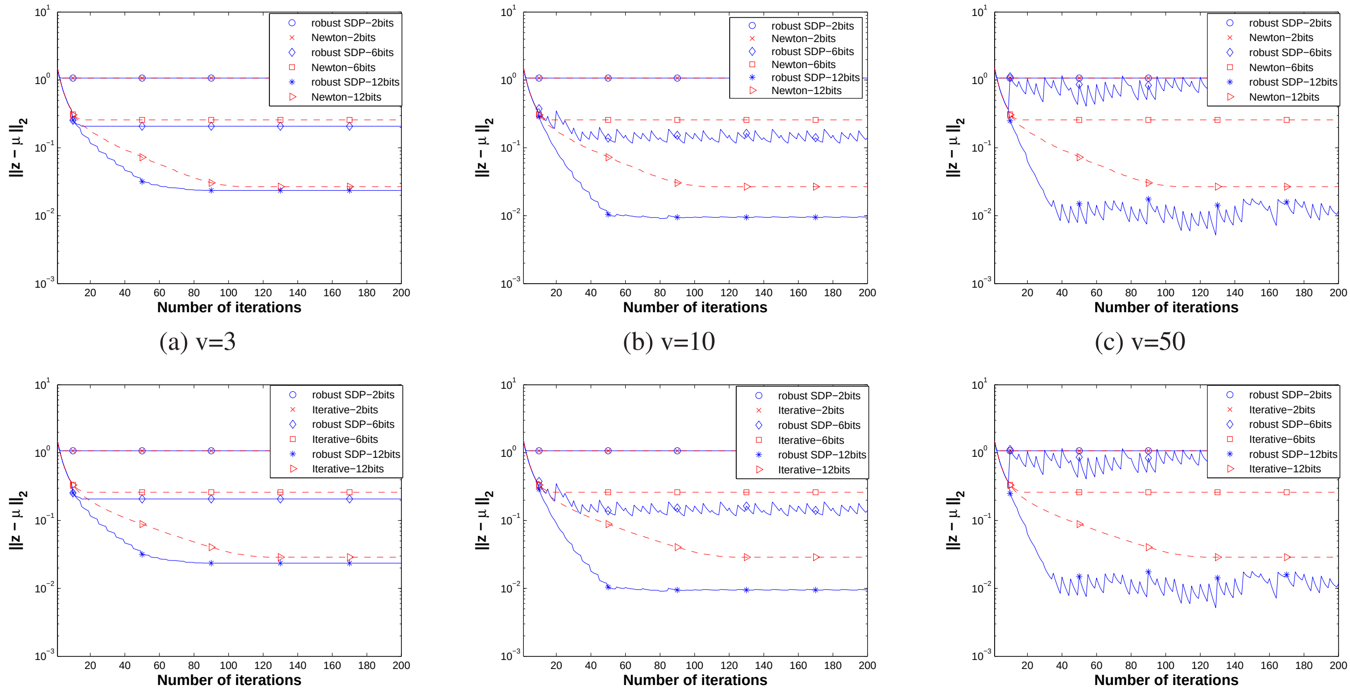

(a) v=3 (b) v=10 (c) v=50

(d) v=3 (e) v=10 (f) v=50

Fig. 3. Average consensus performance of robust SDP vs Newton's polynomial (a),(b),(c) and robust SDP vs the simple iterative method (d),(e),(f) when the data are uniformly quantized with 2, 6 and 12 bits.

show that, under the presence of quantization noise, the existing methods for designing the optimal polynomial achieve a limited performance. We propose a noise robust approach for computing the polynomial that is based on a tradeoff between minimizing the quantization noise and maintaining a high convergence speed. We conclude that even in the case of quantized communication, the polynomial filtering methodology with a proper design of the polynomial coefficients, achieves a better performance than the simple iterative method. Simulation results show the effectiveness of the proposed methodology which seems to outperform both the simple iterative method and the Newton polynomial filtering methodology. Future research work should define the optimal balance between the two contradicting factors and investigate its relation to the weight matrix of the network. Moreover, a challenging problem would be to adapt the polynomial filtering methodology to the proposed iterative scheme in [7].

6. REFERENCES

[1] N.A. Lynch, *Distributed Algorithms*, San Francisco, CA:Morgan Kaufmann Publishers, 2006.

[2] W. Ren and R.W. Beard, "Consensus seeking in multiagent systems under dynamically changing interaction topologies," *IEEE Trans. Autom. Control*, vol. 50, no. 5, pp. 655–661, 2005.

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

[4] L. Xiao, S. Boyd, and S. Lall, "Distributed average consensus with time-varying metropolis weights," *Automatica*, Jun. 2006.

[5] E. Kokiopoulou and P. Frossard, "Polynomial filtering for fast convergence in distributed consensus," *IEEE Trans. Signal Process.*, vol. 57, no. 1, pp. 342–354, Jan. 2009.

[6] S.-J Kim L. Xiao, S. Boyd, "Distributed average consensus with least-mean-square deviation," *International Journal of Robust and Nonlinear Control*, vol. 67, no. 1, pp. 33–46, 2007.

[7] F. Fagnani P. Frasca, R. Carli and S. Zampieri, "Average consensus on networks with quantized communication," *Journal of Parallel and Distributed Computing*, vol. 19, no. 16, pp. 1787–1816, Nov. 2008.

[8] A. Pages-Zamora S. S. Pereira, "Distributed consensus in wireless sensor networks with quantized information exchange ," in *IEEE 9th Workshop on Signal Processing Advances in wireless Communications, SPAWC 2008*, July 2008, pp. 241–245.

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

[10] T. Basar A. Kashyap and R. Srikant, "Quantized consensus," *Automatica*, vol. 43, pp. 1192–1203, Apr. 2007.

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

[12] S. Kar and J. M. F. Moura, "Distributed consensus algorithms in sensor networks: quantized data," *IEEE Trans. Signal Process.*, 2007.

[13] J. Fang and H. Li, "An adaptive quantization scheme for distributed consensus," *Proc. IEEE Int. Conf. Acoust., Speech and Signal Processing*, pp. 2777–2780, Apr. 2009.