Average consensus with communication constraints

A meeting on Mathematical Control Theory: Controllability, Optimization, Stability (MIUR–PRIN 2006 Agrachev)

Paolo Frasca

Politecnico di Torino
Dipartimento di Matematica
http://calvino.polito.it/~frasca
joint work with R. Carli, F. Fagnani, and S. Zampieri

April 20-21, 2009

Average consensus

Given

- a graph $\mathcal{G} = (V, E)$, $V = \{1, \dots, N\}$
  - nodes are "agents";
  - edges are available communication links.
- $x_i \in \mathbb{R}, \forall i \in V$.

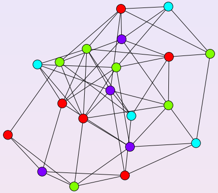

Average consensus problem is computing the average

$$x_{\text{ave}} = N^{-1} \sum_{i=1}^{N} x_i .$$

Average consensus

Given

- a graph $\mathcal{G} = (V, E)$, $V = \{1, \ldots, N\}$
  - nodes are "agents";
  - edges are available communication links.
- $x_i \in \mathbb{R}$, $\forall\, i \in V$.

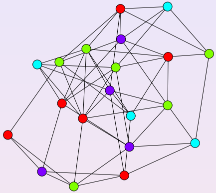

Average consensus problem is computing the average

$$x_{\mathrm{ave}} = N^{-1} \sum_{i=1}^{N} x_i .$$

The difficulty is in the communication limitations.

We want to design an iterative algorithm for that.

Consensus as rendezvous

If x_is are positions of robots,
going to consensus means going to a rendezvous point!

More general are problems of robots deployment.

Background

Average consensus is a neat example of networked problem.
More general problems [Tsi84],

- optimization;
- coverage [BCM09];
- formation;
- control.

Relations with *control with communication constraints.*

As a control problem

We want to design a linear static feedback $u(t) = Kx(t)$, such that

- K is adapted to $\mathcal{G}$, i.e. $(j, i) \notin E \implies K_{ij} = 0$
- the system

$$x(0) = x$$
$$x(t+1) = x(t) + u(t)$$

converges to average consensus, that is

$$\lim_{t \to \infty} x(t) = x_{\mathrm{ave}} \mathbf{1}.$$

Linear average consensus algorithm

Proposition

Let $P = I + K$, *and*

$$x(0) = x$$
$$x(t+1) = P\,x(t).$$

If

- $\mathcal{G}$ *is strongly connected;*
- P *is doubly stochastic;*
- $P_{ii} \geq \delta > 0$

$\Longrightarrow$ *The algorithm converges:*

$$\lim_{t\to\infty} x_i(t) = x_{\text{ave}} \quad \forall i \in \{1, \ldots, N\}.$$

Proof: based on Perron-Frobenius Theorem.

Convergence time

Some definitions

- Consensus error: $d(t) := \frac{1}{\sqrt{N}} \|y(t)\|$.
- ε-convergence time: $T_\varepsilon := \inf\{t \in \mathbb{N} \mid d(t) \leq \varepsilon\}$.
- Essential spectral radius of P:

$$\rho := \max\{|\lambda| \mid \lambda \text{ eigenvalue of } P,\ \lambda \neq 1\}.$$

Proposition

The ε-convergence time of the consensus algorithm is

$$T_\varepsilon \leq C \frac{\log \varepsilon^{-1}}{\log \rho^{-1}}.$$

Two examples of graphs

N-cycles

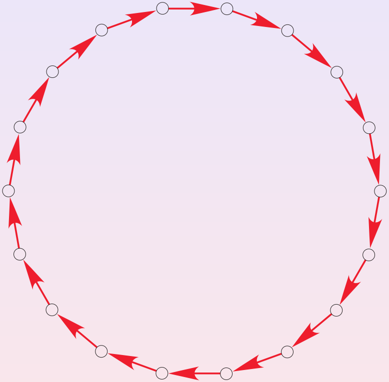

$\rho_N = 1 - \Theta(N^{-1})$ as $N \to \infty$

n-cubes (2^n nodes)

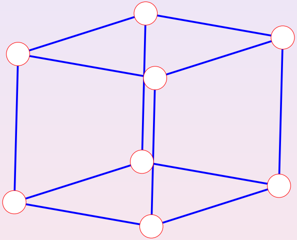

$\rho_N = 1 - \Theta(\log^{-1} N)$ as $N \to \infty$

Our contribution

The above assumed that transmission of real numbers among agents were free!

Analysis and design of novel iterative algorithms for average consensus with network models close(r) to reality: quantization, packet loss.

Examples

- At each time step, communication (with quantization) along all edges of a time-invariant network[FCFZ09, CFFZ07].
- At each time step, one randomly chosen a pair of agents communicates (with quantization) - *Gossip*[CFFZ09].
- At each time step, some randomly chosen agents broadcast to their neighbors (destructive collisions of messages).

Our contribution

The above assumed that transmission of real numbers among agents were free!

Analysis and design of novel iterative algorithms for average consensus with network models close(r) to reality: quantization, packet loss.

Examples

- At each time step, communication (with quantization) along all edges of a time-invariant network[FCFZ09, CFFZ07].
- At each time step, one randomly chosen a pair of agents communicates (with quantization) - *Gossip*[CFFZ09].
- At each time step, some randomly chosen agents broadcast to their neighbors (destructive collisions of messages).

Quantizers

- Uniform quantizers $q : \mathbb{R} \to \mathbb{Z}$
  - Deterministic quantizer

$$q_d(z) = \begin{cases} \lfloor x \rfloor & \text{if } x - \lfloor x \rfloor < 1/2 \\ \lceil x \rceil & \text{if } x - \lfloor x \rfloor \geq 1/2 \end{cases}$$

  - Probabilistic quantizer

$$q_p(x) = \begin{cases} \lfloor x \rfloor & \text{with probability} & \lceil x \rceil - x \\ \lceil x \rceil & \text{with probability} & x - \lfloor x \rfloor. \end{cases}$$

  Both can be scaled to arbitrary precision: $q^{(\varepsilon)}(x) = \varepsilon q(x/\varepsilon)$.
- Saturated quantizer. $m \in \mathbb{N}$, $q^{(m)} : \mathbb{R} \to \mathcal{S}_m$.

$$q^{(m)}(x) = \begin{cases} 1 & \text{if } x > 1 \\ -1 + \frac{2\ell - 1}{m} & \text{if } -1 + \frac{2(\ell-1)}{m} \leq x \leq -1 + \frac{2\ell}{m} \\ -1 & \text{if } x < -1 \end{cases}$$

Quantizers

- Uniform quantizers $q : \mathbb{R} \to \mathbb{Z}$
  - Deterministic quantizer

$$q_d(z) = \begin{cases} \lfloor x \rfloor & \text{if } x - \lfloor x \rfloor < 1/2 \\ \lceil x \rceil & \text{if } x - \lfloor x \rfloor \geq 1/2 \end{cases}$$

  - Probabilistic quantizer

$$q_p(x) = \begin{cases} \lfloor x \rfloor & \text{with probability} & \lceil x \rceil - x \\ \lceil x \rceil & \text{with probability} & x - \lfloor x \rfloor . \end{cases}$$

  Both can be scaled to arbitrary precision: $q^{(\varepsilon)}(x) = \varepsilon q(x/\varepsilon)$.

- Saturated quantizer. $m \in \mathbb{N}$, $q^{(m)} : \mathbb{R} \to \mathcal{S}_m$.

$$q^{(m)}(x) = \begin{cases} 1 & \text{if } x > 1 \\ -1 + \frac{2\ell - 1}{m} & \text{if } -1 + \frac{2(\ell-1)}{m} \leq x \leq -1 + \frac{2\ell}{m} \\ -1 & \text{if } x < -1 \end{cases}$$

Quantizers

- Uniform quantizers $q : \mathbb{R} \to \mathbb{Z}$
  - Deterministic quantizer

$$q_d(z) = \begin{cases} \lfloor x \rfloor & \text{if } x - \lfloor x \rfloor < 1/2 \\ \lceil x \rceil & \text{if } x - \lfloor x \rfloor \geq 1/2 \end{cases}$$

  - Probabilistic quantizer

$$q_p(x) = \begin{cases} \lfloor x \rfloor & \text{with probability} & \lceil x \rceil - x \\ \lceil x \rceil & \text{with probability} & x - \lfloor x \rfloor. \end{cases}$$

  Both can be scaled to arbitrary precision:
  $q^{(\varepsilon)}(x) = \varepsilon q(x/\varepsilon)$.
- Saturated quantizer. $m \in \mathbb{N}$, $q^{(m)} : \mathbb{R} \to \mathcal{S}_m$.

$$q^{(m)}(x) = \begin{cases} 1 & \text{if } x > 1 \\ -1 + \frac{2\ell - 1}{m} & \text{if } -1 + \frac{2(\ell-1)}{m} \leq x \leq -1 + \frac{2\ell}{m} \\ -1 & \text{if } x < -1 \end{cases}$$

Quantizers

- Uniform quantizers $q : \mathbb{R} \to \mathbb{Z}$
  - Deterministic quantizer

$$q_d(z) = \begin{cases} \lfloor x \rfloor & \text{if } x - \lfloor x \rfloor < 1/2 \\ \lceil x \rceil & \text{if } x - \lfloor x \rfloor \geq 1/2 \end{cases}$$

  - Probabilistic quantizer

$$q_p(x) = \begin{cases} \lfloor x \rfloor & \text{with probability} \quad \lceil x \rceil - x \\ \lceil x \rceil & \text{with probability} \quad x - \lfloor x \rfloor. \end{cases}$$

  Both can be scaled to arbitrary precision: $q^{(\varepsilon)}(x) = \varepsilon q(x/\varepsilon)$.

- Saturated quantizer. $m \in \mathbb{N}$, $q^{(m)} : \mathbb{R} \to \mathcal{S}_m$.

$$q^{(m)}(x) = \begin{cases} 1 & \text{if } x > 1 \\ -1 + \frac{2\ell-1}{m} & \text{if } -1 + \frac{2(\ell-1)}{m} \leq x \leq -1 + \frac{2\ell}{m} \\ -1 & \text{if } x < -1 \end{cases}$$

The update rule

We adapt the algorithm to quantized communication as

$$x(t+1) = x(t) + Kq(x(t)).$$

The update rule

We adapt the algorithm to quantized communication as

$$x(t+1) = x(t) + Kq(x(t)).$$

Results:

- It preserves the average at each time step.
- It does not converge, but we can compute as a performance index the asymptotic disagreement

$$d_\infty(P) := \sup_{x(0)} \limsup_{t\to\infty} \frac{1}{\sqrt{N}} ||y(t)||.$$

Bounded error model

We model the quantization error $x - q_d(x)$ as a bounded disturbance $\|e(t)\|_\infty \leq 1/2$. Then,

$$\begin{aligned} x_w(0) &= x(0) \\ x_w(t+1) &= Px_w(t) + (P - I)e(t). \end{aligned}$$

Let $d_\infty^w(P)$ the asymptotic disagreement for the bounded error model.

Remark: $d_\infty(P) \leq d_\infty^w(P)$

A tight worst-case bound

Theorem

Let P be symmetric. Let R be such that $0 < R < 1$ and $\mathrm{eig}(P) \subseteq B_{1-R,R}$. Then,

$$d_\infty^w(P) \leq \frac{3}{4} + \frac{1/2}{1-R} + \frac{1}{4}\log\left(\frac{1}{1-\rho}\right).$$

Theorem

Given an n-cube graph, there exist a consensus matrix P_N adapted to it, and a sequence of bounded disturbances $e(t)$, such that

$$d_\infty^w(P_N) = \frac{n}{2} = \frac{\log_2 N}{2}.$$

$\Longrightarrow$ the error increases with N: q_d not satisfactory.

Mean square analysis

With q_p we can use probability!

Mean square analysis

With q_p we can use probability!

For all $i \in V$ and $t \in \mathbb{Z}_{\geq 0}$, let $n_i(t)$ be random variables such that

- $n_i(t)$ and $n_j(s)$ uncorrelated if $i \neq j$ or $t \neq \tau$
- $n_i(t) \in [-1/2, 1/2]$
- $\mathbb{E}[n_i(t)] = 0$

Remark that $\mathbb{E}[n_i(t)^2] \leq \frac{1}{4}$.

Probabilistic model

$$\begin{cases} x_r(t+1) = Px_r(t) + (P-I)n(t), \\ x_r(0) = x(0) \end{cases}$$

Mean square analysis II

Definition: $d^r_\infty(P) := \limsup_{t\to\infty} \sqrt{\frac{1}{N}\mathbb{E}[\|y_r(t)\|^2]}.$

Theorem

If P is normal, $\implies$

$$[d^r_\infty(P)]^2 \leq \frac{1}{4N}\sum_{i=1}^{N-1}\frac{|1-\lambda_i|^2}{1-|\lambda_i|^2}.$$

Proposition

$$\frac{1}{N}\sum_{i=1}^{N-1}\frac{|1-\lambda_i|^2}{1-|\lambda_i|^2} \leq \frac{1-\delta}{\delta}.$$

It does not depend on N:

randomization has shown to be useful!

Zooming-in/zooming-out algorithm

Parameters: $m, k_{\text{in}}, k_{\text{out}}$.

$$x(0) = x_0 \qquad \text{State}$$

$$x(t+1) = x(t) + (P - I)\hat{x}(t) \qquad \forall t \in \geq 0$$

$$\hat{x}_j(0) = 0 \qquad \forall j \qquad \text{Estimate}$$

$$\hat{x}_j(t+1) = \hat{x}_j(t) + l_j(t+1) q^{(m)} \left(\frac{x_j(t+1) - \hat{x}_j(t)}{l_j(t+1)} \right) \qquad \forall j, t \geq 0$$

$$l_j(1) = l_j(0) = l_0 \qquad \forall j \qquad \text{Zooming Factor}$$

$$l_j(t+1) = \begin{cases} k_{\text{in}} l_j(t) & \text{if } |x_j(t+1) - \hat{x}_j(t)| < 1 \\ k_{\text{out}} l_j(t) & \text{if } |x_j(t+1) - \hat{x}_j(t)| \geq 1 \end{cases} \qquad \forall j, t \geq 1$$

Zooming-in/zooming-out convergence

Theorem

Assume

- $\rho < k_{\text{in}} < 1$,
- $m \geq \frac{(4+3k_{\text{in}})\sqrt{N}}{k_{\text{in}}(k_{\text{in}}-\rho)}$ *and*
- $l_0 > \frac{2(\rho+2)\|x(0)\|}{k_{\text{in}}-\frac{3\sqrt{N}}{m}}$.

Then,

$$\lim_{t\to+\infty} x_i(t) = \lim_{t\to+\infty} \hat{x}_i(t) = x_{\text{ave}}, \quad \forall i \in \{1,\ldots,N\}$$

and $T_\varepsilon \leq C \dfrac{\log \varepsilon^{-1}}{\log k_{\text{in}}^{-1}}$

ε-convergence communication cost

A given precision ε can be obtained sending enough information.
TCC_ε is the number of bits to be sent to have $\mathbb{E}[d(t)] \leq \varepsilon$.

Quantizer		Needed knowledge of $\mathcal{G}$
q_p	$\text{TCC}_\varepsilon \leq C_1 \dfrac{\log^2 \varepsilon^{-1}}{\log \rho_N^{-1}}$	None
$q^{(m)}$	$\text{TCC}_\varepsilon \leq C_2 \log N \dfrac{\log \varepsilon^{-1}}{\log \rho_N^{-1}}$	$k_{\text{in}} \geq \rho_N$

Table: Summary of results

Conclusion

- The achievable precision is “smooth” in the quantizer.
- Spectral properties of P (and of $\mathcal{G}$) matters.
- Randomization is useful.
- Encoding/decoding schemes (with memory) allow convergence.

Open problems

Main open problem is digital noisy average consensus. We have assumed that the channel is reliable in sending messages. What if not?

Open problems

Main open problem is digital noisy average consensus. We have assumed that the channel is reliable in sending messages. What if not?

Work in progress [CCFG09]: using error correcting codes one can achieve

$$\mathrm{TCC}_{\varepsilon} \leq C_3 \frac{\log^3 \varepsilon^{-1}}{\log^2 \rho_N^{-1}},$$

in spite of the errors, with no global knowledge.

Selected literature

F. Bullo, J. Cortés, and S. Martínez.
Distributed Control of Robotic Networks.
Applied Mathematics Series. Princeton University Press, 2009.
To appear. Available at http://www.coordinationbook.info.

R. Carli, G. Como, P. Frasca, and F. Garin.
Average consensus on digital noisy networks.
In *IFAC Workshop on Estimation and Control of Networked Systems*, 2009.
Submitted.

R. Carli, F. Fagnani, P. Frasca, and S. Zampieri.
Efficient quantization techniques in consensus algorithms.
In *NeCST Workshop*, Nancy, France, June 2007.

R. Carli, F. Fagnani, P. Frasca, and S. Zampieri.
Gossip consensus algorithms with quantized communication.
Automatica, 2009.
Submitted.

P. Frasca, R. Carli, F. Fagnani, and S. Zampieri.
Average consensus on networks with quantized communication.
International Journal on Robust and Nonlinear Control, 2009.
In press.

J. N. Tsitsiklis.
Problems in Decentralized Decision Making and Computation.
PhD thesis, Massachusetts Institute of Technology, November 1984.