Average consensus on networks with quantized communication

Paolo Frasca

paolo.frasca@polito.it

joint work with *Ruggero Carli*, *Fabio Fagnani*, *Thomas Taylor*, and *Sandro Zampieri*

Università di Padova, Arizona State University and Politecnico di Torino.

Average consensus over networks

Several agents in a network have to communicate in order to achieve an agreement about the average of their states.

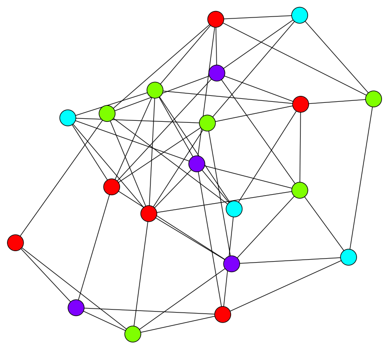

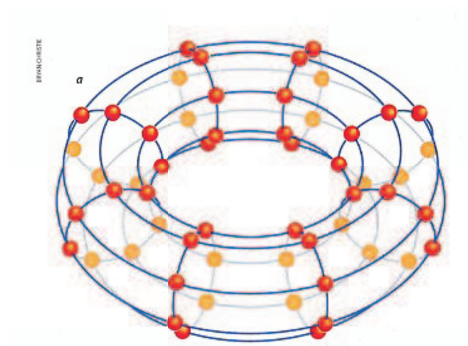

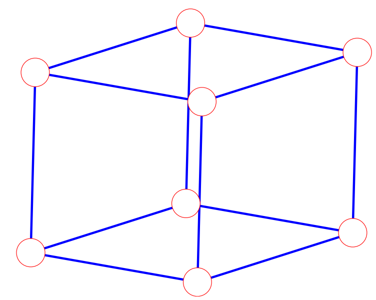

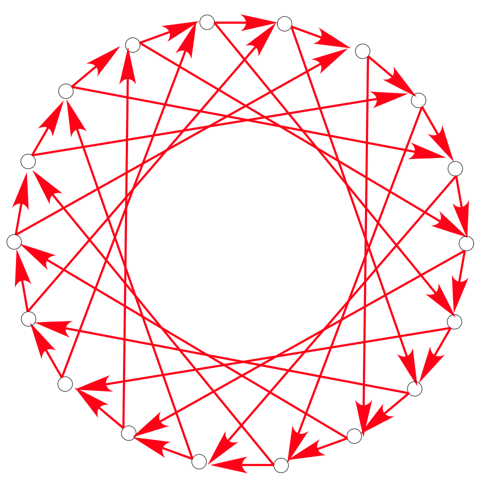

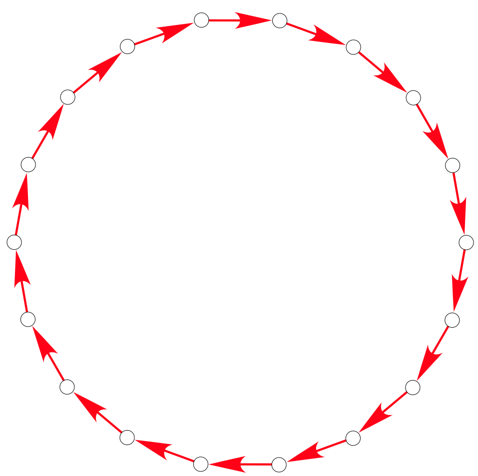

Average consensus problem

Linear dynamical system on $\mathbb{R}^N$

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

componentwise $\quad x_i(t+1) = \sum_{i=1}^{N} P_{ij}\,x_j(t), \quad i = 1 \dots N$

The matrix P has to depend on the communication network:
if j does *not* communicate with i, then $P_{ij} = 0$

Goal: design K so that all agents tend to share the same state

$$\lim_{t \to +\infty} x_i(t) = \frac{1}{N} \sum_{j=1}^{N} x_j(0) \quad \forall\, i.$$

Applications & motivations

In distributed control and information theory.

- Data fusion in sensor networks
- Coordination and rendezvous of UAV and robots
- Load balancing between processors

Our problem

We suppose that the agents can exchange information through a

- time-invariant
- strongly connected
- digital

communication network.

Information has to be symbolic, i.e. quantized:

Normalized to 1 uniform quantization: $q[x(t)] = \text{ round } [x(t)]$.

$\Rightarrow$ *the classical consensus can not be reached.*

See the seminal works by A. Kashyap, T. Basar, R. Srikant (2006), and by L. Xiao, S. Boyd and S. J. Kim (2007).

Perfect links & naive quantization

- If links are not digital, it is known that we can choose P such that
  - $P_{ii} > 0 \quad \forall i$
  - P is non negative
  - P is doubly stochastic, i.e. $\sum_{j=1}^{N} P_{ij} = \sum_{i=1}^{N} P_{ij} = 1.$
- With digital links, the *naive* approach

$$x_i(t+1) = P_{ii}\, x_i(t) + \sum_{j \neq i} P_{ij}\, q[x_j(t)],$$

fails because this non linear map

- does not drive the agents to equal states,
- does not preserve their average.

Our proposal

Take a slightly different evolution map:

$$x_i(t+1) = x_i(t) - (1 - P_{ii})q[x_i(t)] + \sum_{j \neq i} P_{ij}\, q[x_j(t)],$$

$$x(t+1) = x(t) + (P - I)\, q[x(t)],$$

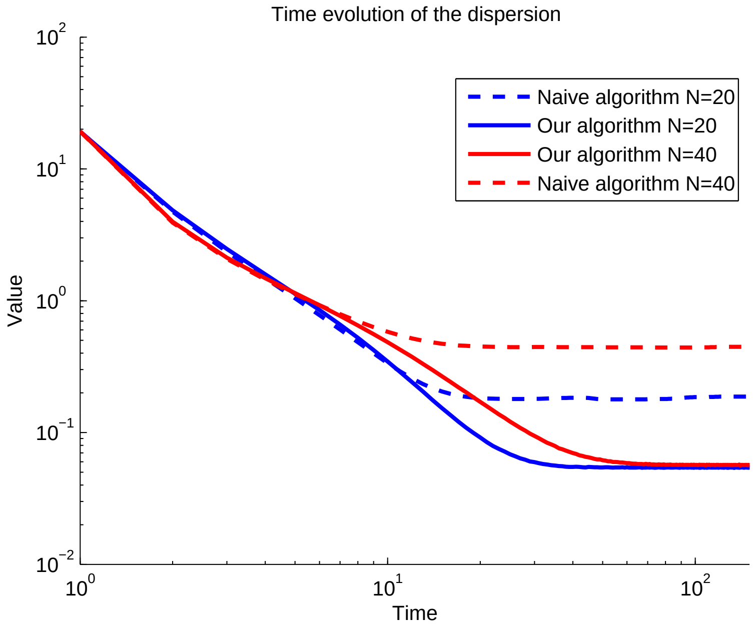

This map preserves the average of states and drives them nearer to consensus.

Example: random geometric graph.

Three approaches to the analysis

Remarks:

- we are interested in the asymptotic of average consensus disagreement

$$\Delta(t) := x(t) - \frac{1}{N}\sum_{i=1}^{N} x_i(0)\mathbb{1}.$$

- The quantization errors $e(t) := x(t) - q[x(t)]$ are bounded

Three approaches can be useful:

- Working on the actual quantized map. $d_\infty(P) := \limsup_{t\to\infty} \frac{1}{\sqrt{N}}||\Delta(t)||_2.$
- Worst case analysis. $d_\infty^w := \lim_{t\to\infty} \sup_{||e||_\infty \le 1/2} \frac{1}{N}||\Delta(t)||_2$
- Probabilistic analysis: consider e as a random variable.

$$d_\infty^r(P) := \sqrt{\lim_{t\to\infty} \frac{1}{N}\mathbb{E}[||\Delta(t)||]^2}$$

Exact analysis

+ we study the actual system
- very difficult: results only in special cases.

- Complete graphs with uniform weights, $P = \frac{1}{N}\mathbb{1}\mathbb{1}^T$.
  $\Rightarrow d_\infty(P) \leq 1$

- Directed circuits with uniform weights, $k = 1/2$.
  Thanks to a symbolic dynamics underlying the system, $\Rightarrow$ $d_\infty(P) \leq 1/2$

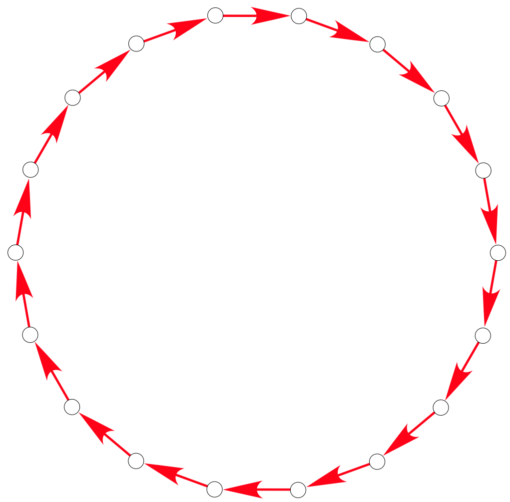

$$x_i(t+1) = (1-k)x_i(t) + kx_i(t)$$

Worst case

- + it's easier
- + it gives upper bounds on the actual system, since $d_\infty(P) \leq d_\infty^w(P)$
- - results are often very conservative.

We find two main bounds

- $d_\infty^w(P) \leq \dfrac{C_P}{1-\rho_{ess}(P)}$
- if P is symmetric $\Rightarrow$ $d_\infty^w(P) \leq \frac{1}{2}\sum_{s=0}^{\infty} \rho(P^s(I-P))$ ρ spectral radius

Worst case II

Consequences on the dependence on N.

- If there is a uniform lower bound G on the spectral gaps, the performance does not worsen in N.

- If P is symmetric Cayley matrix with bounded degree, (its spectral gap is infinitesimal in N) $\Rightarrow \quad d_\infty(P_N) = O(\log N)$.
- This bound is tight for the hypercube graphs: $d_\infty^w(P) = \frac{\log_2 N}{2}$.

Is this divergence intrinsical in the system?

Probabilistic I

+ it's easier
- there is little *a priori* justification, since the original system is not random
- it gives no upper bound
+ results are near to typical simulated results (*a posteriori* justification).

e is supposed to be a random variable acting as an additive noise, having

- zero mean,
- variance σ^2.

Then $d_\infty^r(P) = \sqrt{\frac{1}{N}\sigma^2 \sum_{i=1}^{N-1} \frac{|1-\lambda_i|^2}{1-|\lambda_i|^2}}$, λ_i are the eigenvalues of P.

Probabilistic II

In general, if $P_{ii} \geq \varepsilon \, \forall \, i$, $\Rightarrow$ $d^r_\infty(P) \leq \frac{1-\varepsilon}{\varepsilon}$.

This applies to sequences of Cayley graphs.

Examples:

Hypercube $d^r_\infty(P) = \sqrt{\frac{N-1}{N}}\sigma$

Directed circuit $d^r_\infty(P) = \sqrt{\frac{N-1}{N}\frac{k}{1-k}\sigma^2}$

Undirected circuit $\lim_{N\to\infty} d^r_\infty(P) = \left(\frac{1}{\sqrt{1-2k}} - 1\right)^{1/2} \sigma$

+ In most cases, $d^r_\infty(P)$ can be bounded uniformly in N
  $\Longrightarrow$ good scalability

- $d^r_\infty(P)$ can not be bounded uniformly on other parameters (e.g, k in directed circuit)

Simulations: hypercube graph

The method scales very well in N.

- No logarithmic divergence.
- Well compatible with the probabilistic (uniform noise) result.

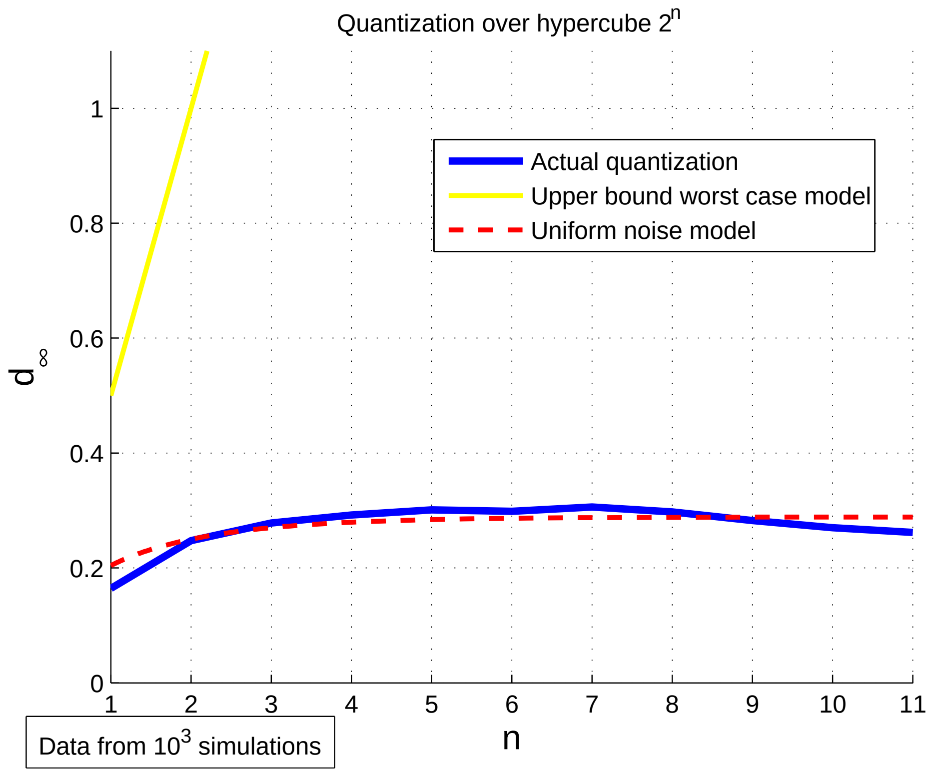

Remark that the x-axis is logarithmic.

Parameters tuning

In the directed circuit, the performance depends on the parameter k.

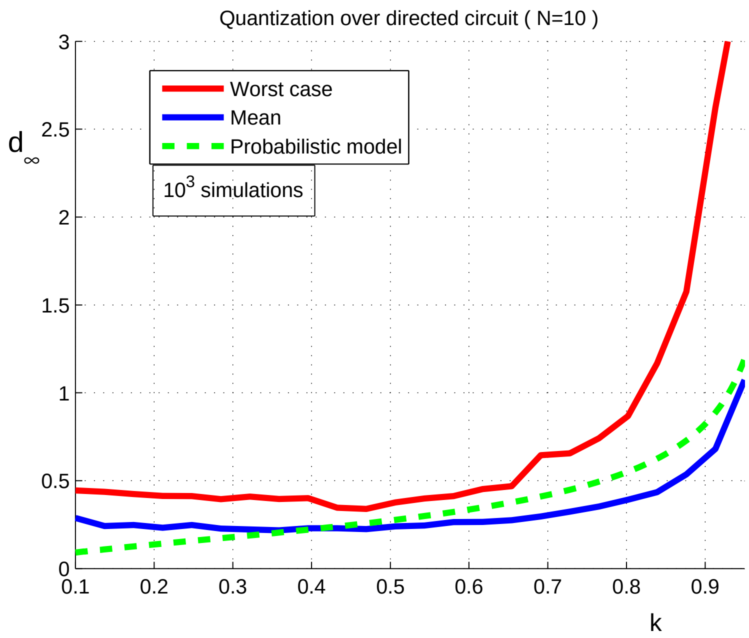

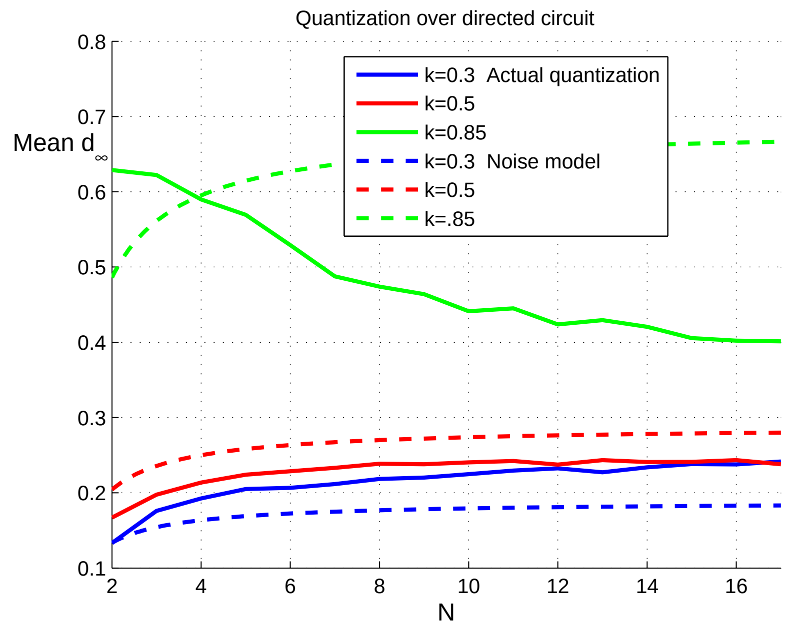

- No logarithmic divergence in N
- qualitatively compatible with probabilistic results:
  - in k
  - in N for small k

Discussion and open problems

The algorithm works very well.

- it preserves the average of initial conditions,
- it drives the agents near to the consensus in typical cases (e.g. uniform weights),

But:

- which is the right theoretical approach?
- are there trade-offs between asymptotical vicinity to consensus and speed of convergence?
- are there better algorithms? E.g., we know that using uniform quantization in an encoder-decoder scheme with memory (Zooming in-zooming out), average consensus is reachable.