On the Computational Power of Simple Dynamics

Emanuele Natale



Sapienza University of Rome December 21, 2016

Communication in Biological Systems



Flocks of birds [Ben-Shahar et al. '10]

Communication in Biological Systems

Flocks of birds [Ben-Shahar et al. '10]





Schools of fish [Sumpter et al. '08]

Communication in Biological Systems

Flocks of birds [Ben-Shahar et al. '10]





Schools of fish [Sumpter et al. '08]

Insects colonies [Franks et al. '02]



Animal communication:

- Chaotic
- Anonymous
- Parsimonious

- Uni-directional (Passive/Active)
- Noisy



 $\mathcal{PUSH}(h, \ell)$ model [Demers '88]: at each round each agent can send $a \ \ell$ -bit message to h other agents chosen independently and uniformly at random.

Uni-directional (Passive/Active)
Noisy





 $\mathcal{PUSH}(h, \ell)$ model [Demers '88]: at each round each agent can *send* $a \ \ell$ -bit message to h other agents chosen independently and uniformly at random.

Uni-directional (Passive/Active) Noisy



 $\mathcal{PUSH}(h, \ell)$ model [Demers '88]: at each round each agent can *send* $a \ \ell$ -bit message to h other agents chosen independently and uniformly at random.

Uni-directional (Passive/Active) Noisy $\ell \text{ bits}$



 $\mathcal{PULL}(h, \ell)$ model [Demers '88]: at each round each agent can observe h other agents chosen independently and uniformly at random, and shows ℓ bits to her observers. Uni-directional (Passive/Active)
Noisy





• Uni-directional (Passive/Active)

• Noisy

 $\mathcal{PULL}(h, \ell)$ model[Demers '88]: at eachround each agent canobserve h other agentschosen independently anduniformly at random, andshows ℓ bits to herobservers.

Very simple distributed algorithms: For every graph, agent and round, states are updated according to fixed rule of current state and symmetric function of states of neighbors.

Very simple distributed algorithms: For every graph, agent and round, states are updated according to fixed rule of current state and symmetric function of states of neighbors.

Examples of Dynamics

• 3-Median dynamics



Very simple distributed algorithms: For every graph, agent and round, states are updated according to fixed rule of current state and symmetric function of states of neighbors.

Examples of Dynamics

- 3-Median dynamics
- 3-Majority dynamics



Very simple distributed algorithms: For every graph, agent and round, states are updated according to fixed rule of current state and symmetric function of states of neighbors.

Examples of Dynamics

- 3-Median dynamics
- 3-Majority dynamics
- Undecided-state dynamics



Very simple distributed algorithms: For every graph, agent and round, states are updated according to fixed rule of current state and symmetric function of states of neighbors.

Examples of Dynamics

- 3-Median dynamics
- 3-Majority dynamics
- Undecided-state dynamics
- Averaging dynamics



The Power of Dynamics: Plurality Consensus

Computing the Median

3-Median dynamics [Doerr et al. '11]. Converge to $\mathcal{O}(\sqrt{n \log n})$ approximation of median of system in $\mathcal{O}(\log n)$ rounds w.h.p., even if $\mathcal{O}(\sqrt{n})$ states are arbitrarily changed at each round $(\mathcal{O}(\sqrt{n})$ -bounded adversary).

The Power of Dynamics: Plurality Consensus

Computing the Median

3-Median dynamics [Doerr et al. '11]. Converge to $\mathcal{O}(\sqrt{n \log n})$ approximation of median of system in $\mathcal{O}(\log n)$ rounds w.h.p., even if $\mathcal{O}(\sqrt{n})$ states are arbitrarily changed at each round $(\mathcal{O}(\sqrt{n})$ -bounded adversary).

Computing the Majority

3-Majority dynamics [SPAA '14, SODA '16]. If plurality has bias $\mathcal{O}(\sqrt{kn \log n})$, converges to it in $\mathcal{O}(k \log n)$ rounds w.h.p., even against $o(\sqrt{n/k})$ -bounded adversary. Without bias, converges in poly(k). h-majority converges in $\Omega(k/h^2)$.

The Power of Dynamics: Plurality Consensus

Computing the Median

3-Median dynamics [Doerr et al. '11]. Converge to $\mathcal{O}(\sqrt{n \log n})$ approximation of median of system in $\mathcal{O}(\log n)$ rounds w.h.p., even if $\mathcal{O}(\sqrt{n})$ states are arbitrarily changed at each round $(\mathcal{O}(\sqrt{n})$ -bounded adversary).

Computing the Majority

3-Majority dynamics [SPAA '14, SODA '16]. If plurality has bias $\mathcal{O}(\sqrt{kn \log n})$, converges to it in $\mathcal{O}(k \log n)$ rounds w.h.p., even against $o(\sqrt{n/k})$ -bounded adversary. Without bias, converges in poly(k). h-majority converges in $\Omega(k/h^2)$.

Undecided-State dynamics [SODA '15]. If majority/second-majority $(c_{maj}/c_{2^{nd}maj})$ is at least $1 + \epsilon$, system converges to plurality within $\tilde{\Theta}(\text{md}(\mathbf{c}))$ rounds w.h.p.

A Global Measure of Bias



Undecided-State dynamics [SODA '15]. If majority/second-majority $(c_{maj}/c_{2^{nd}maj})$ is at least $1 + \epsilon$, system converges to plurality within $\tilde{\Theta}(\mathrm{md}(\mathbf{c}))$ rounds w.h.p.











Noisy \mathcal{PULL} model: messages are randomly corrupted

Noise Matrix:

trix:

$$P := \begin{pmatrix} p_{\bullet,\bullet} & p_{\bullet,\bullet} & p_{\bullet,\bullet} \\ p_{\bullet,\bullet} & p_{\bullet,\bullet} & p_{\bullet,\bullet} \\ p_{\bullet,\bullet} & p_{\bullet,\bullet} & p_{\bullet,\bullet} \end{pmatrix}$$



Noisy \mathcal{PULL} model: messages are randomly corrupted

Noise Matrix:

$$P := \begin{pmatrix} p_{\blacktriangle, \diamond} & p_{\blacktriangle, \diamond} & p_{\diamond, \diamond} \\ p_{\diamond, \diamond} & p_{\diamond, \diamond} & p_{\diamond, \diamond} \\ p_{\diamond, \diamond} & p_{\diamond, \diamond} & p_{\diamond, \diamond} \end{pmatrix}$$



Noise affects animal communication, but animals cannot use *coding theory...*

FHK '14: Natural rules efficiently solve rumor spreading and majority consensus despite noise when $P = \begin{pmatrix} 1/2 + \varepsilon & 1/2 - \varepsilon \\ 1/2 - \varepsilon & 1/2 + \varepsilon \end{pmatrix}$



Configuration $\mathbf{c} := (\# \not \gg /n, \# \not \gg /n, \# \not \gg /n)$

 $\frac{\delta \text{-majority-biased}}{\text{configuration w.r.t.}} \qquad \# / n - \# / n > \delta$ $\# / n - \# / n > \delta$



Configuration $\mathbf{c} := (\# \not \gg /n, \# \not \gg /n, \# \not \gg /n)$

 $\begin{array}{ll} \delta \text{-majority-biased} \\ \text{configuration w.r.t.} & \# / n - \# / n > \delta \\ & \# / n - \# / n > \delta \end{array}$

 $\begin{array}{l} (\varepsilon, \delta) \text{-majority-preserving noise matrix:} \\ (\mathbf{c}P)_{\blacktriangle} - (\mathbf{c}P)_{\bigstar} > \varepsilon \delta, \quad (\mathbf{c}P)_{\bigstar} - (\mathbf{c}P)_{\bigstar} > \varepsilon \delta \end{array}$



Configuration $\mathbf{c} := (\# \not \gg /n, \# \not \gg /n, \# \not \gg /n)$

 $\begin{array}{ll} \delta \text{-majority-biased} \\ \text{configuration w.r.t.} & \# / n - \# / n > \delta \\ & \# / n - \# / n > \delta \end{array}$

 $\begin{array}{l} (\varepsilon, \delta) \text{-majority-preserving noise matrix:} \\ (\mathbf{c}P)_{\blacktriangle} - (\mathbf{c}P)_{\bigstar} > \varepsilon \delta, \quad (\mathbf{c}P)_{\bigstar} - (\mathbf{c}P)_{\bigstar} > \varepsilon \delta \end{array}$

[PODC '16]: Let S initial set of agents with k opinions. S is $\delta = \Omega(\sqrt{\log n/|S|})$ -majority-biased. $|S| = \Omega(\frac{\log n}{\epsilon^2})$. P is (ϵ, δ) -majority-preserving. Plurality consensus can be solved in $O(\frac{\log n}{\epsilon^2})$ rounds w.h.p., with $O(\log \log n + \log \frac{1}{\epsilon})$ memory per node.

The Median, the Mode and... the Mean

Dynamics can solve Consensus, Median, Majority, in robust and fault tolerant ways, but this is trivial in centralized setting.

The Median, the Mode and... the Mean

Dynamics can solve Consensus, Median, Majority, in robust and fault tolerant ways, but this is trivial in centralized setting.

Can dynamics solve a problem non-trivial in centralized setting?

Community Detection as Minimum Bisection

Minimum Bisection Problem. Input: a graph G with 2n nodes. Output: $S = \arg \min_{\substack{S \subset V \\ |S| = n}} E(S, V - S).$



[Garey, Johnson, Stockmeyer '76]: **Min-Bisection** is *NP-Complete*.

The Stochastic Block Model

Stochastic Block Model (SBM). Two "communities" of equal size V_1 and V_2 , each edge inside a community included with probability $p = \frac{a}{n}$, each edge across communities included with probability $q = \frac{b}{n} < p$.



The Stochastic Block Model

Reconstruction problem. Given graph generated by SBM, find original partition.



The Stochastic Block Model

Reconstruction problem. Given graph generated by SBM, find original partition.



- At t = 0, randomly pick value $x^{(t)} \in \{+1, -1\}$.
- Then, at each round 1. Set value $x^{(t)}$ to
 - 1. Set value $x^{(e)}$ to average of neighbors,
 - 2. Set label to **blue** if $x^{(t)} < x^{(t-1)}$, red otherwise.

- At t = 0, randomly pick value $x^{(t)} \in \{+1, -1\}$.
- Then, at each round 1. Set value $x^{(t)}$ to
 - average of neighbors,
 - 2. Set label to **blue** if $x^{(t)} < x^{(t-1)}$, red otherwise.









- At t = 0, randomly pick value $x^{(t)} \in \{+1, -1\}$.
- Then, at each round 1. Set value $x^{(t)}$ to
 - 1. Set value $x^{(*)}$ to average of neighbors,
 - 2. Set label to **blue** if $x^{(t)} < x^{(t-1)}$, red otherwise.

- At t = 0, randomly pick value $x^{(t)} \in \{+1, -1\}$.
- Then, at each round
 - 1. Set value $x^{(t)}$ to average of neighbors,
 - 2. Set label to **blue** if $x^{(t)} < x^{(t-1)}$, red otherwise.



- At t = 0, randomly pick value $x^{(t)} \in \{+1, -1\}$.
- Then, at each round
 - 1. Set value $x^{(t)}$ to average of neighbors,
 - 2. Set label to **blue** if $x^{(t)} < x^{(t-1)}$, red otherwise.









- At t = 0, randomly pick value $x^{(t)} \in \{+1, -1\}$.
- Then, at each round
 1. Set value x^(t) to average of neighbors,
 - 2. Set label to **blue** if $x^{(t)} < x^{(t-1)}$, **red** otherwise.





Al nodes at the same time:

- At t = 0, randomly pick value $x^{(t)} \in \{+1, -1\}$.
- Then, at each round 1. Set value $x^{(t)}$ to
 - average of neighbors,
 - 2. Set label to **blue** if $x^{(t)} < x^{(t-1)}$, red otherwise.

Well studied process [Shah '09]:

- Converges to (weighted) global average of initial values,
- Convergence time = mixing time of G,
- Important applications in fault-tolerant self-stabilizing consensus.



Al nodes at the same time:

- At t = 0, randomly pick value $x^{(t)} \in \{+1, -1\}$.
- Then, at each round 1. Set value $x^{(t)}$ to
 - average of neighbors,
 - 2. Set label to **blue** if $x^{(t)} < x^{(t-1)}$, red otherwise.

 $A = (\mathbb{1}_{((u,v)\in E)})_{u,v\in V}$ adjacency matrix of G

 ${\color{black} D}$ diagonal matrix of node degrees in ${\color{black} G}$

 $P = D^{-1}A$ transition matrix of random walk Well studied process [Shah '09]:

- Converges to (weighted) global average of initial values,
- Convergence time = mixing time of G,
- Important applications in fault-tolerant self-stabilizing consensus.



 $\mathbf{x}^{(t)} = P \cdot \mathbf{x}^{(t-1)} = P^t \cdot \mathbf{x}^{(0)}$











[SODA '17](Informal). $G = (V_1 \bigcup V_2, E)$ s.t. i) $\chi = \mathbf{1}_{V_1} - \mathbf{1}_{V_2}$ close to right-eigenvector of eigenvalue λ_2 of transition matrix of G, and ii) gap between λ_2 and $\lambda = \max\{\lambda_3, |\lambda_n|\}$ sufficiently large, then Averaging (approximately) identifies (V_1, V_2) .

Check Out my Thesis!



https://goo.gl/Q0LC6x

Thank you!