Majority Consensus by Local Polling

Moez Draief

Imperial College London

Lund Workshop on Dynamics and Control in Networks
October 2014

Aggregating through Majority

1785, Marquis de Condorcet's weak law of large numbers

- in a large population of voters, and each one independently votes correctly with probability $\alpha > 1/2$
- as population size grows, probability that the outcome of a majority vote is correct converges to one

Information is efficiently aggregated

Aggregating in a network

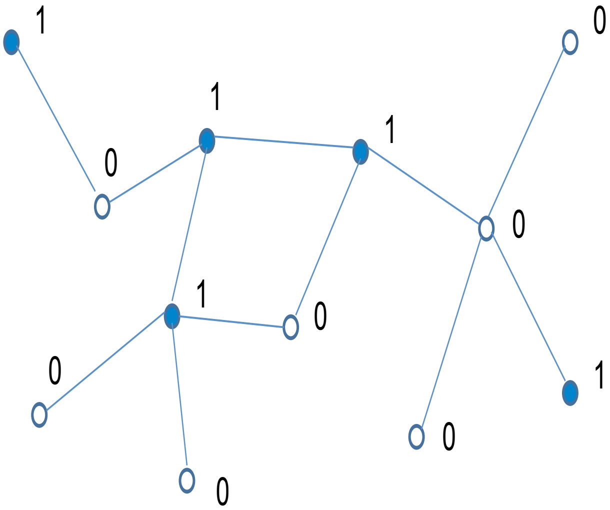

Binary majority consensus

Desired outcome and metrics

- Nodes end with opinion held by majority of nodes
- Node can probe neighbours and update opinion accordingly using little (constant) memory
- Probability of error (convergence to incorrect consensus)
- Time to convergence

Applications

- Occurrence of a given event in cooperative decision making
- Voting in distributed systems
- Routine to solve more elaborate distributed decision making instances

Local Majority Dynamics

$G = (V, E)$ simple connected graph on $|V| = n$ vertices

Each vertex either red (**1**) or blue (**0**).
Initial proportion of blues is $\alpha \in (1/2, 1)$

GOAL: Local algorithm for inferring the majority state.

- Does the graph settle into one colour?
- If so, how does the graph structure and the initial distribution affect which colour wins?
- How long does it take?

Outline

- Distributed consensus [known results]
- Interval consensus [Draief, Vojnovic '12]
- Local polling [Abdullah, Draief '14]

Continuous-time Interaction Model

- Connected undirected graph $G = (V, E)$, $|V| = n$
- αn nodes hold 0 and $(1-\alpha)n$ nodes hold 1, $\alpha \in (1/2, 1)$
- Nodes i and j interact at rate $q_{ij} = q_{ji}$, $q_{ij} \neq 0$ iff $(i, j) \in E$

Markov chain

- $(X_t)_{t \geq 0}$ continuous-time Markov chain with rate matrix Q, $q_{ii} = -\sum_{i \neq j} q_{ij}$
- $(\pi_i)_{i \in V}$ stationary distribution is uniform on V. Mixing time:

$$\left|\mathbb{P}_j(X_t = i) - 1/n\right| = O\left(e^{-\lambda_2(Q)t}\right)$$

where $\lambda_2(Q) = \inf\{\sum_{i,j} q_{ij}(x_i - x_j)^2/2, ||x|| = 1,\ x^T\mathbf{1} = 0\}$

Performance of voter model

Node i contacts j at rate q_{ij} and i updates to j's state

Theorem [Hassin-Peleg '01]

- The number of nodes in state 1 is a martingale.
- Probability of reaching (wrong) consensus at 1 is $1 - \alpha$.
- Time to convergence of voter model $O(n/(\lambda_2(Q)))$.

Time to convergence

Complete graph

- Each edge has rate $1/(n-1)$. Number of agents with opinion 1 evolves as Birth-Death process

$$\lambda_{k,k+1} = \lambda_{k,k-1} = \frac{k(n-k)}{n-1}.$$

- Time to convergence $= O(n)$

General graphs

- Conductance $\eta(Q) = \inf_{A \subset V} \frac{\sum_{i \in A, j \in A^c} q_{ij}}{|A||A^c|/n}$
- Markov chain tracking the number of nodes in state 0 evolves at least $\eta(Q)$ times as fast as on the complete graph, since

$$\sum_{i \in A,\, j \in A^c} q_{ij} \geq \eta(Q) \underbrace{\frac{|A||A^c|}{n}}_{\text{complete graph}}$$

- Time to convergence $O(n/\eta(Q))$,

D. Aldous, "Interacting particle systems as stochastic social dynamics" Bernoulli 19(4), 1122-1149, 2013.

Time to convergence

Cheeger's inequality

- Conductance: $\eta(Q) = \inf_{A \subset V} \frac{\sum_{i \in A, j \in A^c} q_{ij}}{|A||A^c|/n}$
- Spectral Gap:
  $\lambda_2(Q) = \inf\{\sum_{i,j} q_{ij}(x_i - x_i)^2/2, ||x|| = 1,\ x^T \mathbf{1} = 0\}$

$$2\lambda_2(Q) \leq \eta(Q) \,.$$

- Time to convergence of voter model $O(n/(\lambda_2(Q)))$.

Let S of size k be the subset realising the inf in $\eta(Q)$ and let x such that $x_i = -\sqrt{\frac{n-k}{kn}}$, $i \in S$ and $x_i = \sqrt{\frac{k}{(n-k)n}}$, $i \in S^c$.

Distributed averaging

At each interaction of (i, j) occurring at rate q_{ij}

$$x_i(t) = x_j(t) = \frac{x_i(t-) + x_j(t-)}{2} .$$

Theorem [Boyd et al '06, Aldous '12]

- Algorithm converges to the average value, using $O(Poly(\log(n))$ memory per node
- Time to convergence to up $O(1/n)$ error of the average is

$$\Theta(\log(n)/\lambda_2(Q)) ,$$

Distributed averaging: Proof

Let $R(t) = ||x(t)||^2$. When an i, j interaction takes place $R(t)$ reduces by $(x_i - x_j)^2/2$.

$$\begin{aligned}\mathbb{E}(dR(t) \mid x(t) = x) &= \sum_{i,j} q_{ij} \left(2 \left(\frac{x_i + x_j}{2} \right)^2 - (x_i^2 + x_j^2) \right) \\ &= -\sum_{i,j} q_{ij} \frac{(x_i - x_j)^2}{2} dt \\ \text{(Assume that } \textstyle\sum_i x_i(0) = 0) &\leq -\lambda_2(Q) ||x||^2 dt\end{aligned}$$

In particular

$$\mathbb{E}||x(t)||^2 \leq ||x(0)||^2 e^{-\lambda_2(Q)t}$$

Small memory

Could we use less memory and still guarantee small error?

Theorem: Impossibility

- Connected undirected graph $G = (V, E)$, $|V| = n$,
- αn nodes in 0 and $(1-\alpha)n$ nodes in 1, $\alpha \in (1/2, 1)$, $2\alpha - 1$ is the *voting margin*.

No 1-bit distributed algorithm can solve the majority consensus problem.

Land, Belew, "No perfect two-state cellular automata for density classification exists", PRL 74, 5148-5150, 1995

Ternary Consensus

- αn nodes hold 0 and $(1-\alpha)n$ nodes hold 1,
- Additional state e for undecided nodes, $q_{i,j} = 1/n$, $\forall i, j$

Theorem [PVV '09]

Probability of reaching wrong consensus **1**. For n large,

$$P_{error} = (1 + o(1))2^{-D(\alpha||\frac{1}{2})n}$$

where $D(\alpha||\frac{1}{2})$ is the Kullback-Leibler divergence. T time to convergence, $\mathbb{E}(T) = (1 + o(1))\log n$.

- Results (seem to) hold for expander but fail for the line.
- Generalises beyond binary consensus [Babaee, Draief '14]

Averaging-like updates: States $0 < e_0 < e_1 < 1$.
Rules: Swaps + Annihilation

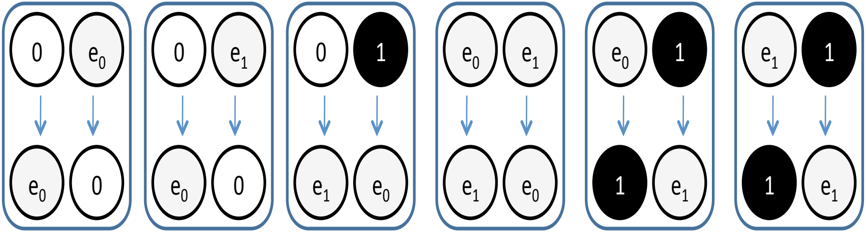

Kashyap, Basar, Srikant, “Quantized consensus” Automatica, 1192-1203, 2007

Bénézit, Thiran, Vetterli, Interval consensus: From quantized gossip to voting, ICASSP 2009

Mean-field analysis (Complete graph)

Let $q_{ij} = \frac{1}{n-1}$, $i \neq j$ and $\mathbf{X}(t) = (|S_0(t)|, |S_{e_0}(t)|, |S_{e_1}(t)|, |S_1(t)|)$ is a Markov process with the following transition rates

$$\rightarrow \begin{cases} (|S_0(t)|-1, |S_{e_0}(t)|+1, |S_{e_1}(t)|+1, |S_1(t)|-1) & : \frac{|S_0(t)||S_1(t)|}{n-1} \\ (|S_0(t)|, |S_{e_0}(t)|-1, |S_{e_1}(t)|+1, |S_1(t)|) & : \frac{|S_{e_0}(t)||S_1(t)|}{n-1} \\ (|S_0(t)|, |S_{e_0}(t)|+1, |S_{e_1}(t)|-1, |S_1(t)|) & : \frac{|S_0(t)||S_{e_1}(t)|}{n-1} \end{cases}.$$

By Kurtz's theorem, $\mathbf{X}(t)/n$ converges to $(s_0(t), s_{e_0}(t), s_{e_1}(t), s_1(t))$

$$\begin{aligned} s_0'(t) &= -s_1(t)s_0(t) \\ s_1'(t) &= -s_0(t)s_1(t) \\ s_{e_1}'(t) &= s_1(t)(1 - s_1(t)) - (s_0(t) + s_1(t))s_{e_1}(t) \end{aligned}$$

with $s_{e_0}(t) = 1 - s_0(t) - s_{e_1}(t) - s_1(t)$, $t \geq 0$.

Proposition [Draief, Vojnovic '10]

For large t,

$$s_{e_1}(t) \sim (2\alpha - 1)\frac{1-\alpha}{\alpha} t e^{-(2\alpha-1)t}$$
$$s_1(t) \sim (2\alpha - 1)\frac{1-\alpha}{\alpha} e^{-(2\alpha-1)t}.$$

In particular, $t^1_{n,\alpha}$ and $t^{e_1}_{n,\alpha}$ times nodes in 1 and e_1 to disappear

$$t^1_{n,\alpha} = \frac{1}{2\alpha - 1}\log(n) + O(1)$$
$$t^{e_1}_{n,\alpha} = \frac{1}{2\alpha - 1}[\log(n) + \log(\log(n))] + O(1).$$

Minority states

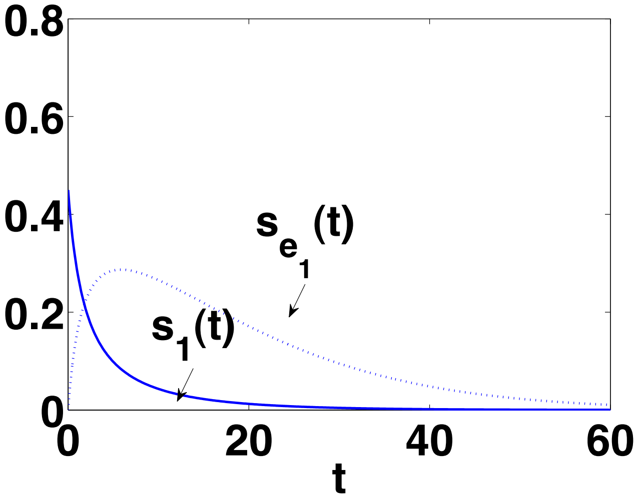

General bound

Theorem [Draief, Vojnovic '12]

Let T be the time until there are only nodes in state 0 and e_0.

$$\mathbb{E}(T) = O\left(\log n / \delta(Q, \alpha)\right)$$

where $\delta(Q, \alpha) = \min_{S \subset V, |S| = (2\alpha - 1)n} \min_{\lambda \in Spec(Q_S)} |\lambda|$

$$Q_S = \left[\begin{array}{c|c} \text{diag}(q_{ii},\ i \in S) & \mathbf{0} \\ \hline (q_{ij})_{i \in S^c,\, j \in S} & (q_{ij})_{i,j \in S^c} \end{array}\right]$$

Master Equation

First phase: $Z_i(t)$ ($A_i(t)$) indicator that i in state 0 (1) at t

$$(Z,A) \to \begin{cases} (Z-e_i, A-e_j) & : \quad q_{i,j} Z_i A_j \\ (Z-e_i+e_j, A) & : \quad q_{i,j} Z_i (1-A_j-Z_j) \\ (Z, A-e_i+e_j) & : \quad q_{i,j} A_i (1-A_j-Z_j) \end{cases}$$

Second phase: $B_i(t)$ indicator that node i is in state e_1 at t

$$(Z,B) \to \begin{cases} (Z-e_i+e_j, B-e_j) & : \quad q_{i,j} Z_i B_j \\ (Z-e_i+e_j, B) & : \quad q_{i,j} Z_i (1-B_j-Z_j) \\ (Z, B-e_i+e_j) & : \quad q_{i,j} B_i (1-B_j-Z_j) \end{cases}$$

(random) Piecewise-linear dynamical system

$$\frac{d}{dt}\mathbb{E}(Y_i(t)) = -\left(\sum_{l\in V} q_{i,l}\right)\mathbb{E}(Y_i(t)) + \sum_{j\in V} q_{i,j}\mathbb{E}\left(Y_j(t)(1 - Z_i(t))\right).$$

Dynamics reduces to $Y(t) = (Y_i(t))_{i\in V}$,

$$\frac{d}{dt}\mathbb{E}_k(Y(t)) = Q_{S_k}\mathbb{E}_k(Y(t)),$$

for $t \in [t_k, t_{k+1})$ during which $\{S_0(t) = S_k\}$ and Q_{S_k} is given by

$$Q_S(i,j) = \begin{cases} -\sum_{l\in V} q_{i,l}, & i = j \\ q_{i,j}, & i \notin S, j \neq i \\ 0, & i \in S, j \neq i. \end{cases}$$

Proposition

$$\mathbb{E}(Y(t)) = \mathbb{E}\left[e^{\lambda(t)} Y(0)\right]$$

where $\lambda(t) = Q_{S_k}(t - t_k) + \sum_{l=0}^{k-1} Q_{S_l}(t_{l+1} - t_l)$.

Lemma

For any finite graph G, there exists $\delta(Q, \alpha) > 0$ such that, for any non-empty subset of vertices S with $|S| \in [(2\alpha - 1)n, \alpha n]$, if λ is an eigenvalue of the matrix Q_S, then

$$\lambda \leq -\delta(G, \alpha) < 0.$$

Proof: Spectrum of Q_S

$$Q_S = \left[\begin{array}{c|c} \text{diag}(q_{ii},\ i \in S) & \mathbf{0} \\ \hline (q_{ij})_{i\in S^c,\, j\in S} & (q_{ij})_{i,j\in S^c} \end{array}\right]$$

- First $\left(q_{ii} = -\sum_{l\neq i} q_{i,l}\right)$, $i \in S$ are eigenvalues of Q_S
- The remaining eigenvalues correspond to eigenvectors $\underline{x} = (\underbrace{0,\ldots,0}_{S}, \underbrace{x}_{S^c})^T$. Let $W \subset S^c$, for $i \in W$, $x_i \neq 0$

$$\begin{aligned} -\lambda &= \underline{x}^T Q_S \underline{x} \\ &= \sum_{i\in W}\sum_{j\in S} q_{i,j}x_i^2 + \sum_{i\in W, j\in S^c\setminus W} q_{i,j}x_i^2 + \frac{1}{2}\sum_{i,j\in W} q_{i,j}(x_i - x_j)^2 \end{aligned}$$

Proof

Note that

$$\mathbb{E}(Y(t)) = \mathbb{E}\left[e^{\lambda(t)} Y(0)\right]$$

where $\lambda(t) = Q_{S_k}(t - t_k) + \sum_{l=0}^{k-1} Q_{S_l}(t_{l+1} - t_l)$
By Jensen's and matrix norm inequalities,

$$||\mathbb{E}(Y(t))||_2 \leq \mathbb{E}\left[||e^{Q_{S_k}(t-t_k)}|| \prod_{l=0}^{k-1} ||e^{Q_{S_l}(t_{l+1}-t_l)}|| \, ||Y(0)||_2\right] \leq \sqrt{n} e^{-\delta(G,\alpha)t}$$

Therefore, by Cauchy-Schwartz, we have

$$\mathbb{P}(Y(t) \neq \mathbf{0}) \leq \sum_{i \in V} \mathbb{E}(Y_i(t)) \leq n \, e^{-\delta(G,\alpha)t}$$

We conclude since $\mathbb{E}(T_0) = \int_0^{\infty} \mathbb{P}(Y(t) \neq \mathbf{0}) dt$.

Complete graph

Corollary

An application of the theorem to complete graph $q_{i,j} = \frac{1}{n-1}$ for all $i \neq j$, yields

$$\mathbb{E}(T) \leq 2\frac{1}{2\alpha - 1}\log(n).$$

Exact asymptotics

A direct analysis of the dynamics of the 1st phase

$$\mathbb{E}(T_1) = \frac{n-1}{|S_0| - |S_1|}\left(H_{|S_1|} + H_{|S_0|-|S_1|} - H_{|S_0|}\right)$$

where $H_k = \sum_{i=1}^{k} \frac{1}{i}$

Various initial conditions

- $|S_0| - |S_n| = (2\alpha - 1)n$, α a constant larger than $1/2$

$$\mathbb{E}(T_1) = \frac{1}{2\alpha - 1}\log(n) + O(1).$$

- If $|S_0| = |S_1|$

$$\mathbb{E}(T_1) = \frac{\pi^2}{6}n(1 + o(1)).$$

- $\mu_n = (|S_0| - |S_1|)/n$ is strictly positive but small ($o(1)$),

$$\mathbb{E}(T_1) = \frac{1}{\mu_n}\log(n\mu_n) + O(1).$$

Complete Graph: Theory v. Simulation

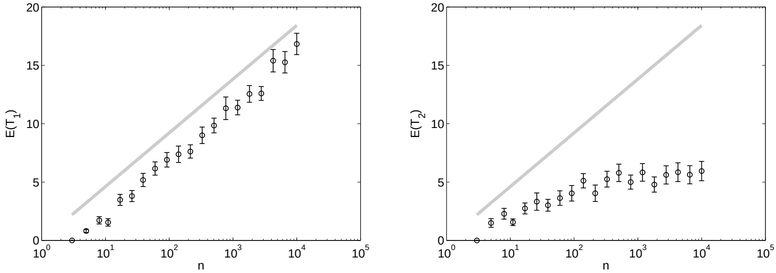

Star

- **Star Network:** $q_{1,i} = q_{i,1} = \frac{1}{n-1}$, $i \neq 1$ and $q_{i,j} = 0$, $i, j \neq 1$. $\mathbb{E}(T_i) \leq \frac{1}{2\alpha-1} n \log(n)$. Using, direct calculation

$$\mathbb{E}(T_1) = \frac{1}{(2\alpha-1)(3-2\alpha)} n \log(n) + O(n)$$

- **ER-graph:** $q_{i,j} = \frac{1}{np_n} X_{i,j}$ $X_{i,j}$ i.i.d. Bernoulli r.v. with mean $c\frac{\log(n)}{n}$, $c > \frac{2}{2\alpha-1}$, for h^{-1} the inverse of $h(x) = x\log(x) + 1 - x$,

$$\mathbb{E}(T_i) \leq \frac{1}{(2\alpha-1)h^{-1}\left(\frac{2}{c(2\alpha-1)}\right)} \log(n) + O(1)$$

- **Path**: $\mathbb{E}(T_i) \leq \frac{16(1-\alpha)^2}{\pi^2} n^2 \log(n) + O(1)$
- **Ring**: $\mathbb{E}(T_i) \leq \frac{4(1-\alpha)^2}{\pi^2} n^2 \log(n) + O(1)$.

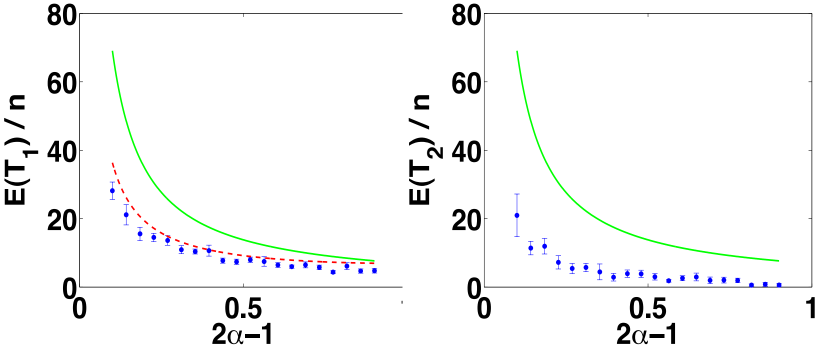
80
60
40
20
0
E(T_1) / n
0
0.5
2α−1
80
60
40
20
0
E(T_2) / n
0
0.5
1
2α−1

ER-graph

- **Star Network:** $q_{1,i} = q_{i,1} = \frac{1}{n-1}$, $i \neq 1$ and $q_{i,j} = 0$, $i, j \neq 1$. $\mathbb{E}(T_i) \leq \frac{1}{2\alpha - 1} n \log(n)$. Using, direct calculation

$$\mathbb{E}(T_1) = \frac{1}{(2\alpha - 1)(3 - 2\alpha)} n \log(n) + O(n)$$

- **ER-graph:** $q_{i,j} = \frac{1}{np_n} X_{i,j}$ $X_{i,j}$ i.i.d. Bernoulli r.v. with mean $p_n = c\frac{\log(n)}{n}$, $c > \frac{2}{2\alpha - 1}$, for h^{-1} the inverse of $h(x) = x\log(x) + 1 - x$,

$$\mathbb{E}(T_i) \leq \frac{1}{(2\alpha - 1) h^{-1}\left(\frac{2}{c(2\alpha - 1)}\right)} \log(n) + O(1)$$

- **Path**: $\mathbb{E}(T_i) \leq \frac{16(1-\alpha)^2}{\pi^2} n^2 \log(n) + O(1)$
- **Ring**: $\mathbb{E}(T_i) \leq \frac{4(1-\alpha)^2}{\pi^2} n^2 \log(n) + O(1)$.

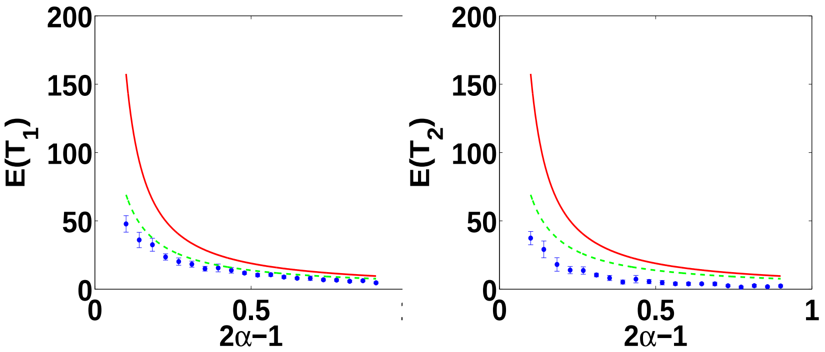
200
150
100
50
0
E(T1)
0
0.5
2α−1
200
150
100
50
0
E(T2)
0
0.5
1
2α−1

Path and Ring

- **Star Network:** $q_{1,i} = q_{i,1} = \frac{1}{n-1}$, $i \neq 1$ and $q_{i,j} = 0$, $i, j \neq 1$. $\mathbb{E}(T_i) \leq \frac{1}{2\alpha-1} n \log(n)$. Using, direct calculation

$$\mathbb{E}(T_1) = \frac{1}{(2\alpha - 1)(3 - 2\alpha)} n \log(n) + O(n)$$

- **ER-graph:** $q_{i,j} = \frac{1}{np_n} X_{i,j}$ $X_{i,j}$ i.i.d. Bernoulli r.v. with mean $c\frac{\log(n)}{n}$, $c > \frac{2}{2\alpha-1}$, for h^{-1} the inverse of $h(x) = x\log(x) + 1 - x$,

$$\mathbb{E}(T_i) \leq \frac{1}{(2\alpha - 1)h^{-1}\left(\frac{2}{c(2\alpha-1)}\right)} \log(n) + O(1)$$

- **Path**: $\mathbb{E}(T_i) \leq \frac{16(1-\alpha)^2}{\pi^2} n^2 \log(n) + O(1)$
- **Ring**: $\mathbb{E}(T_i) \leq \frac{4(1-\alpha)^2}{\pi^2} n^2 \log(n) + O(1)$.

Summary

- Upper bound on the expected convergence time for a number of distributed for solving Majority consensus
- Bounds based on the location of the spectral gap of rate matrix (generalised-cut: quick for expander graphs).
- For binary consensus, expected convergence time critically depends on the voting margin
- Application to particular network topologies: complete graphs, stars, ER graph, paths, cycles.

(Discrete-time $k-$choice local majority protocol

At $t = 0$, each vertex of G is blue **independently with constant probability** $\alpha \in (1/2, 1)$.

Local Majority

We then run $\mathcal{MP}^k$ on G. Choose k odd ($k \geq 5$ in what follows).

- At each time t, each vertex v polls k neighbours **uar**, and assumes majority colour
- If v doesn't have k neighbours, poll all, or all minus one

What is the probability that there will be a red consensus?

How long does it take to reach consensus?

Graphs of a given degree sequence

Let $V = [n]$

$\mathcal{G}_n(\mathbf{d})$: the set of connected simple graphs with degree sequence $\mathbf{d} = (d_1, d_2, \ldots, d_n)$, where d_i is the degree of vertex $i \in V$.

Need some restrictions on degree sequence to make it graphical, e.g., $\sum_i d_i$ is even

Nice degree sequences

Let $V_j = \{i \in V : d_i = j\}$, $= \frac{1}{n} \sum_{i=1}^{n} d_i$ be the average degree, $0 < \kappa \leq 1$, $0 < c < 1/8$ constants, and let $\gamma = (\sqrt{\log n}/)^{1/3}$. A degree sequence **d** is **nice** if it satisfies

(i) Average degree $= o(\sqrt{\log n})$.

(ii) Minimum degree $\delta \geq 3$.

(iii) Let $d \geq 5$ be such that $|V_d| = \kappa n + o(n)$. We call d the **effective minimum degree.**

(iv) Number of little vertices $\sum_{j=\delta}^{d-1} |V_j| = O(n^{\frac{1}{11}})$; a vertex i is **little** if $d_i \leq d - 1$.

(v) Maximum degree $\Delta = O(n^{\frac{1}{11}})$.

(vi) Upper tail size $\sum_{j=\gamma}^{\Delta} n_j = O(\Delta)$.

The effective minimum degree

(iii) Let $d \geq 5$ be such that $|V_d| = \kappa n + o(n)$. We call d the **effective minimum degree.**

Need not be a constant, can have $d \to \infty$ as $n \to \infty$

Not necessarily the minimum degree (though it can be)

Can have "little" vertices with smaller degree, as long as not too many of them:

(iv) Number of little vertices $\sum_{j=\delta}^{d-1} |V_j| = O(n^{\frac{1}{11}})$; a vertex i is **little** if $d_i \leq d - 1$.

Examples of nice degree sequences

- Any d-regular graph with $d \geq 5$ and $d = o(\sqrt{\log n})$
- 'Bi-regular' graph where half the vertices are degree $d \geq 5$ and half of degree $\Delta = o(\sqrt{\log n})$.
- Truncated power-law

Results: informal statement

Suppose G is typical with effective min degree d. If we run $\mathcal{MP}^k$ then

Upper bound

If $d/k = O(1)$ and α is 'not too close' to $1/2$, then **whp**, correct consensus is reached within $(A \log_k d) \log_k \log_k n$ steps

$(A \leq 5$ and $A \to 1$ if $k \to \infty)$

Lower bound

Any algorithm where a vertex keeps its colour if same as all neighbours, will take at least $\log_d \log_d n$ steps to reach correct consensus, **whp**

Bias condition

"α is not too close to $1/2$" means

$$\left[\left(1+\frac{1}{\sqrt{k}}\right)2\right]^{\frac{2}{k-2}}\alpha(1-\alpha) < 1/4$$

Since $\alpha \neq 1/2 \Rightarrow \alpha(1-\alpha) < 1/4$, so inefficiency is in $\left[\left(1+\frac{1}{\sqrt{k}}\right)2\right]^{\frac{2}{k-2}}$

$k = 5$ needs $1 - \alpha < 0.143$

$k = 20$ needs $1 - \alpha < 0.350$

$k = 100$ needs $1 - \alpha < 0.437$

Compare with other works

E. Mossel, J. Neeman, O. Tamuz ('14) Study local majority on d-regular λ-expanders. Show sufficient bias implies certain correct consensus.

better bias condition but only regular graphs, no timing information, full polling only

Compare with other works

E. Mossel, J. Neeman, O. Tamuz ('14) Study local majority on d-regular λ-expanders. Show sufficient bias implies certain correct consensus.

better bias condition but only regular graphs, no timing information, full polling only

Y. Kanoria and A. Montanari ('10) Study local majority on d-regular infinite tree. Give bias conditions for convergence to majority

+better bias condition, -only infinite regular graph

Compare with other works

E. Mossel, J. Neeman, O. Tamuz ('14) Study local majority on d-regular λ-expanders. Show sufficient bias implies certain correct consensus.

better bias condition but only regular graphs, no timing information, full polling only

Y. Kanoria and A. Montanari ('10) Study local majority on d-regular infinite tree. Give bias conditions for convergence to majority

+better bias condition, -only infinite regular graph

J. Cruise and A. Ganesh ('10) Study (m,d)-generalisation of local majority on complete graphs with unit rate exponential on each vertex. Give exponential decay error probability and $O(\log n)$ timing

+stronger error probability, -only complete graph

Typical graphs

Typical graphs: For a nice degree sequence **d**, the space $\mathcal{G}_n(\mathbf{d})$ is the set of nice graphs

We do not analyse for the whole space, only for those graphs called **typical**

Informally, G is typical if it is nice and:

- most vertices are locally tree-like
- little vertices and very high-degree vertices, should they exist, are far from each other and small cycles

Let $\mathcal{G}'_n(\mathbf{d}) \subset \mathcal{G}_n(\mathbf{d})$ be the typical graphs, then $|\mathcal{G}'_n(\mathbf{d})|/|\mathcal{G}_n(\mathbf{d})| \to 1$ as $n \to \infty$

Modified Majority

Let $\mathcal{T} = G[v, c \log_k \log_k n]$.
At $t+1$, each $x \in V$ randomly picks a $x(k)$-subset of neighbours $N_x(t+1)$

- $x \notin \mathcal{T}$ then x becomes at $t+1$ the majority colour of the vertices in $N_x(t+1)$.

$$X_{t+1}^{\mathcal{MMP}^k(v,s)}(x) = \mathbf{1}_{\left\{\left(\sum_{y \in N_x(t+1)} X_t^{\mathcal{MP}^k}(y)\right) > x(k)/2\right\}}.$$

- non-leaf $x \in \mathcal{T}$ and $\mathrm{Par}(x)$ the parent of x in $\mathcal{T}$. At $t+1$, x becomes the majority colour of the vertices in $N_x(t+1)$, with the added assumption that $\mathrm{Par}(x)$ was red at time t.

$$X_{t+1}^{\mathcal{MMP}^k(v,s)}(x) = \mathbf{1}_{\left\{\left(\sum_{y \in N_x(t+1) \setminus \{\mathrm{Par}(x)\}} X_t^{\mathcal{MMP}^k(v,s)}(y)\right) > x(k)/2\right\}}.$$

Modified majority protocol

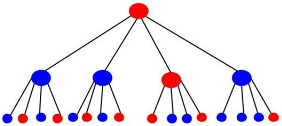

For a vertex v, let $X_v(t)$ be the indicator v is red at time t under $\mathcal{MPP}^k$. Let $k = 2r + 1$.

- At time $t = 0$, for each level 2 (i.e., leaf) vertex v, $\mathbb{P}(X_v(0) = 0) = p_0 = 1 - \alpha$
- At time $t = 1$, for each level 1 vertex v $\mathbb{P}(X_v(1) = 0) = p_1 = \mathbb{P}(\text{Bin}(2r, p_0) \geq r)$
- At time $t = 2$, for each level 0 vertex v (i.e., the root) $\mathbb{P}(X_v(2) = 0) = p_2 = \mathbb{P}(\text{Bin}(2r, p_1) \geq r)$

Modified majority protocol

If height of the tree is H, then given p_t, at $t+1$, for v at distance $H-t-1$ from root,

$$\mathbb{P}(X_v(t+1)=0) = p_{t+1} = \mathbb{P}(\mathrm{Bin}(2r, p_t) \geq r)$$

and we get a rapidly decaying sequence $p_0 > p_1 > \ldots > p_t$ with $p_0 = \alpha \gg p_t$ when t large

When $t = \Omega(\log \log n)$, p_t is very small and we conclude by union bound over all n vertices

The root will have the correct colour.

Now we are left to deal with vertices not locally tree-like...

Theorem: Erdös-Renyi graphs

Let $p = \frac{c \log n}{n}$ where $c > 2 + \epsilon$ for some constant $\epsilon > 0$, $k \geq 5$ and $\nu = \lfloor \frac{k-1}{2} \rfloor$. Run $\mathcal{MP}^k$ on $G \in \mathcal{G}(n, p)$.
Let $A = \frac{1+\varepsilon}{\log_k(k-1) - \log_k 2}$ where $\varepsilon > 0$ is a small constant. Subject to condition

$$\left[\left(1 + \frac{1}{\sqrt{2\nu}}\right) 2\right]^{\frac{1}{\nu-1}} 4\alpha(1-\alpha) < 1$$

by time $A \log_k \log_k n$, $\mathcal{MP}^k$ will have reached consensus on the initial majority **whp**.

Summary

- Asymptotic correct and efficient consensus using local polling. What happens for other values of k? [Cooper-Elsasser-Radzik'14]
- Analysis for a sparse family of graphs and dense E-R graphs.
- Still lot of ongoing interest...

Related work

- Kanoria, Montanari, Majority dynamics on trees and the dynamic cavity method, Annals Applied Prob., 2011
- Jung, Kim, Vojnovic, Distributed Ranking in Networks with Limited Memory and Communication, IEEE ISIT 2012
- Mossel, Neeman and Tamuz, Majority Dynamics and Aggregation of Information in Social Networks, Autonomous Agents and Multi-Agent Systems, 2014
- Mertzios, Nikoletseas, Raptopoulos, Spirakis, Determining Majority in Networks with Local Interactions and very Small Local Memory, ICALP 2014
- Becchetti, Clementi, Natale, Pasquale, Silvestri, Trevisan, Simple dynamics for plurality consensus. SPAA 2014.