Information Theoretic Bounds for Distributed Computation

Munther A. Dahleh, MIT

In collaboration with

Ola Ayaso and Devavrat Shah

Frontiers in Distributed Communication, Sensing and Control August, 2008.

Motivations and Challenges

- Optimal algorithms for distributed computation/control in networks:
  - Wireless and Sensor Networks,
  - Social networks,
  - Example: Consensus and Belief Propagation with bit constraints.

- Computation is built on top of an unreliable transport layer
  - Nodes communicate with Neighbors at ‘random’,
  - Communication is Asynchronous,
  - Neighbors may broadcast a message without confirmation,
  - Nodes have Limited Storage, Limited Computation Power.

- Communication/Computation-Algorithm-independent lower bounds:
  - Capture tradeoffs: computation performance and network constraints.
  - Assert optimality of computation performance:

    Given an algorithm, if computation time matches the lower bound, then it is optimal.

Summary of Results

- Information Theoretic formulations and techniques for distributed computation.
  - Algorithm-independent lower bounds on computation time of distributed algorithms as a function of communication constraints.
- For a particular scenario,
  - Lower bounds on algorithm run-time as a function of network topology.
  - An algorithm and upper bound on its run-time.
  - Both bounds scale reciprocally with “conductance.”

Lower bound is tight in capturing effect of network topology via conductance.
Our algorithm’s run-time is optimal in its dependence on conductance.

Comparison to Other Work

- Information Theory Literature
  - Two nodes computing a function
    - Ahlswede/Csiszar(81), Orlitsky/Roche(01).
  - Reliable communication in networks (Network Information Theory).
  - Control over communication networks
    - Martins/Dahleh, Martins/Dahleh/Doyle.
- Communication Literature
  - Algorithms for computation in networks
    - Gallager(88), Ying/Srikant/Dullerud(06), Luo(03),
    - Boyd/Ghosh/Prabhakar/Shah(06),
  - Consensus literature.
  - Most upper/lower bounds on algorithm-performance are algorithm-specific.

This work draws on formulations and tools of Network Information Theory and applies them in the context of distributed computation.

Formulation: Computation in Networks

- Source: Joint PDF $(X_1(0), \ldots, X_n(0))$
- Each node has:
  - Partial Information, $X_i(0)$,
  - Communication Algorithm,
  - Unlimited Computation Resources.
- Goal of nodes:
  - Computation, $K_i = f_i\left(X_1(0), \ldots X_n(0)\right)$,
  - Performance, $E\left(X_i(T) - K_i\right)^2 \leq 2^{-\alpha}$.
- Communication Network
  - Fixed Topology,
  - Noisy, discrete Channels.
- Objective: Impact of communication constraints on time, T, until performance is guaranteed.

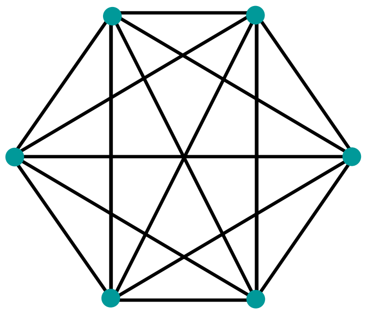

$T?$

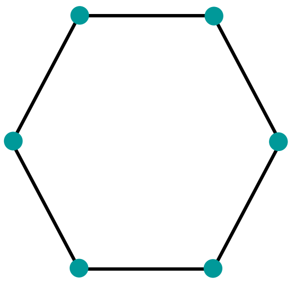

$T?$

Information Theory to Know

- Information Theory gives algorithm-independent lower bound.

- $h(X)$ is the differential entropy: Uncertainty or Randomness.
  - Conditioning reduces entropy,
    $$h(X|Y) \leq h(X).$$
  - The normal distribution maximizes entropy over all distributions with the same variance, and.
    $$h(X) \leq \frac{1}{2} \log 2\pi e Var(X).$$

- $I(X;Y)$ is the mutual information: Degree of dependence.
  - Related to differential entropy
    $$I(X;Y) = h(X) - h(X|Y).$$
  - Processing doesn't increase information in Y about X.
    $$I(X;Y) \geq I(X; f(Y)).$$

X and Y are scalar continuous RVs.

Lower Bound for Computation in Arbitrary Network

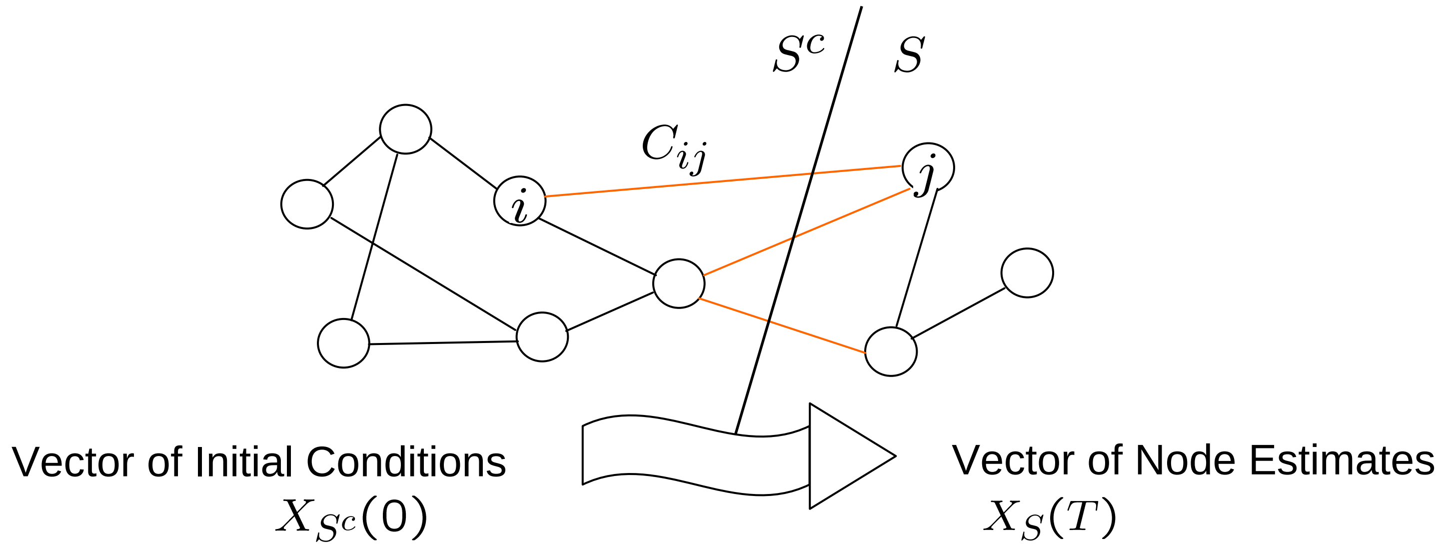

- K_S is the vector of random variables that nodes in S compute.

$$
\begin{aligned}
T\left(\sum_{i\in S^c}\sum_{j\in S}\mathbf{C}_{ij}\right) &\geq I(X_S(T); X_{S^c}(0)|X_S(0)) \\
&\geq I(X_S(T); K_S|X_S(0)) \\
&\geq \underbrace{h(K_S|X_S(0))}_{\text{Uncertainty}} - \frac{1}{2}\log\left(\prod_{j\in S} 2\pi e\,\underbrace{E(X_j(T)-K_j)^2}_{\text{Accuracy}}\right)
\end{aligned}
$$

Lower bound depends on uncertainty in function K and desired accuracy.

Simple Example: Two Nodes Computing Sum

- $K = X_1(0) + X_2(0)$
- $X_1(0) \sim N(0, v_1)$
  $X_2(0) \sim N(0, v_2)$
  Independent

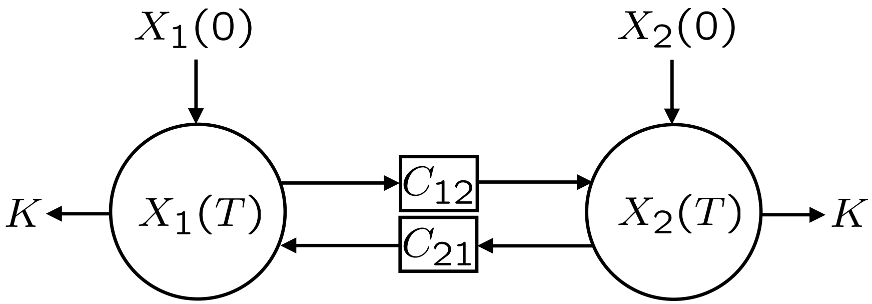

- $h(K|X_1(0)) = \gamma + \frac{1}{2}\log v_2$
  $h(K|X_2(0)) = \gamma + \frac{1}{2}\log v_1$

- If $E\,(X_1(T) - K)^2 \leq 2^{-\alpha}$ and $E\,(X_2(T) - K)^2 \leq 2^{-\alpha}$

then, $$T \geq \max\left\{\frac{1}{2C_{21}}\left(\log v_2 + \alpha\right), \frac{1}{2C_{12}}\left(\log v_1 + \alpha\right)\right\}$$

More information must be communicated about initial condition with larger variance.

$\gamma = \frac{1}{2}\log 2\pi e$

General Sum: Algorithm-Independent Lower Bound on T

- Each node, i :
  - Has an initial value $X_i(0) \sim U[1,\ B+1]$ IID
  - Seeks to compute $K = \sum_{j=1}^{n} X_j(0)$
  - Maintains $X_i(T)$: Estimate of K at time T.
- Communication Network:
  - Graph, G: Connected.
  - Channels: Independent, Symmetric, Erasure, $C_{ij} = p_{ij}$.

Any algorithm that guarantees for all nodes

$$P\left(|X_i(T) - k| \le \varepsilon k \,\middle|\, K = k\right) > 1 - \delta$$

must have

$$T \ge \frac{1}{2\Phi(G)} \log \frac{1}{\varepsilon^2 B + \left(\frac{1}{B}\right)^{\frac{1}{n}} + \kappa\delta}$$

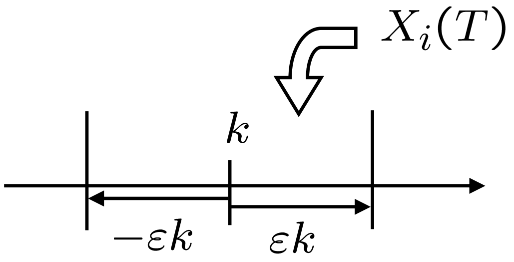

Conductance: Definition and Computation

- Definition resembles Graph-Theoretic definition

$$\Phi(G) = \min_{\substack{S \subset V \\ |S| \leq \frac{n}{2}}} \frac{\sum_{i \in V \setminus S} \sum_{j \in S} C_{ij}}{|S|}$$

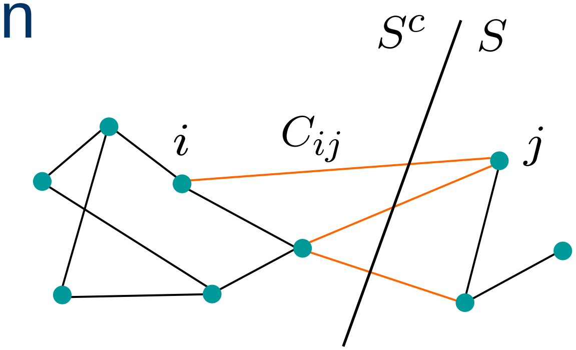

- Computation for two topologies with n nodes
  - Ring Graph: Severe Topology Constraints
    - Links weighted equally $C = \frac{1}{4}$

$$\Phi(G) = \min_{|S| \leq \frac{n}{2}} \frac{1/2}{|S|} = \frac{1/2}{n/2} = \frac{1}{n}$$

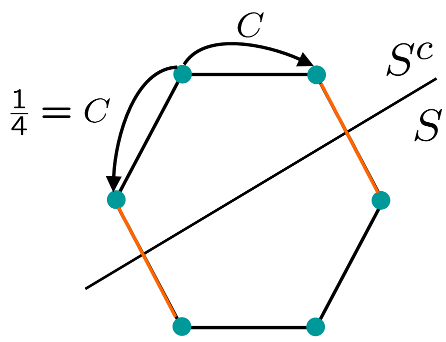

  - Complete Graph: No Topology Constraints
    - Links weighted equally $C = \frac{1}{n}$

$$\Phi(G) = \min_{|S| \leq \frac{n}{2}} \frac{n - |S|}{n} = \frac{n - \frac{n}{2}}{n} = \frac{1}{2}$$

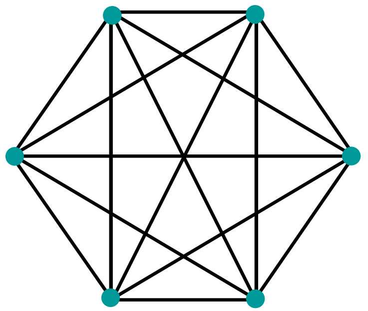

The more severe network constraints, the smaller conductance.

Lower Bound is Tight!

- n Node Network
- Each node, i :
  - Has an initial value $x_i(0) \in [1,\ B+1]$
  - Needs to compute $k = \sum_{j=1}^{n} x_j(0)$
  - Maintains $x_i(T)$: Estimate of K at time T.
- Communication Network:
  - Graph, G: Connected.
  - Channels: Symmetric, Erasure, $C_{ij} = p_{ij}$.

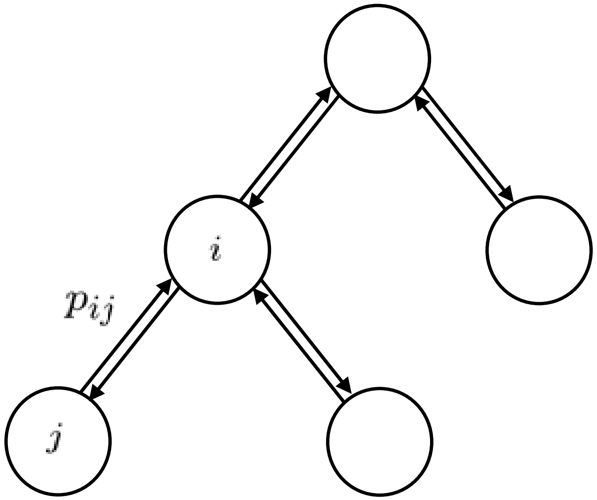

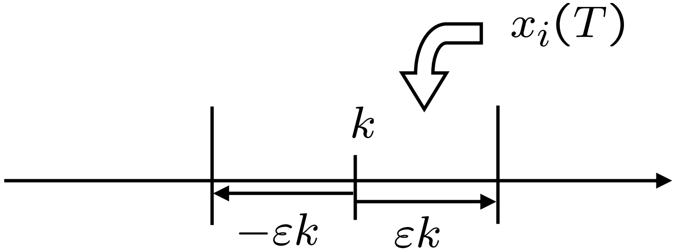

our algorithm guarantees for all nodes,

$$P\left(|x_i(T) - k| \le \varepsilon k\right) > 1 - \frac{1}{n^2}$$

in

$$T = O\left(\frac{\varepsilon^{-2}\log^3 n}{\Phi(G)}\right)$$

Conductance: Implication on Run-Time

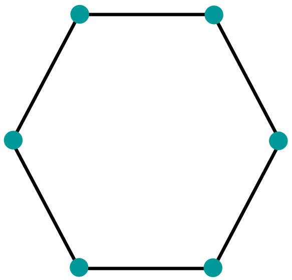

- For Ring with n nodes, conductance scales like $1/n$

$$\Phi(G) \approx \frac{1}{n}$$

- Known: algorithms performing linear computations converge in run-time T that scales like

$$\frac{1}{\Phi^2(G)} \approx n^2$$

We provide an algorithm with run-time T that scales like

$$\frac{1}{\Phi(G)} \approx n$$

Our algorithm-independent lower bound: No matter what algorithm is used, cannot have run-time T scaling better than this.

Examples

- Internet: *Preferential connectivity model*
  - Network is built incrementally
  - A new node joins existing nodes with probability proportional to their degree
  - Conductance: $\Phi(G) = \Omega(1)$
  - Run-Time: $T = O\left(\varepsilon^{-2}(\log^3 n)\right)$
- Wireless Networks: Gupta-Kumar Model
  - Two nodes are connected if they are within a disc of radius r
  - Critical radius: $r = \Theta(\sqrt{\log n / n})$
  - Conductance: $\Phi(G) = \Theta(\sqrt{1/n})$
  - Run-Time: $T = O\left(\varepsilon^{-2}\sqrt{n}(\log^3 n)\right)$
- Social Networks: Connection is proportional to the distance
  - Run-Time: $T = O\left(\varepsilon^{-2}(\log^3 n)\right)$

Algorithm Run-Time depends on Conductance

- Our algorithm depends on information spreading as a subroutine.
- Time until information spreads in a network is inversely related to conductance.

$$T_{spr} = O\left(\frac{\log n + \log \delta^{-1}}{\Phi(G)}\right)$$

- More severe network constraints imply
  - Smaller conductance and longer time until information spreads.

Run-time of our algorithm is inversely related to conductance because it depends on information spreading.

An Algorithm for Nodes Communicating Real-Valued Messages

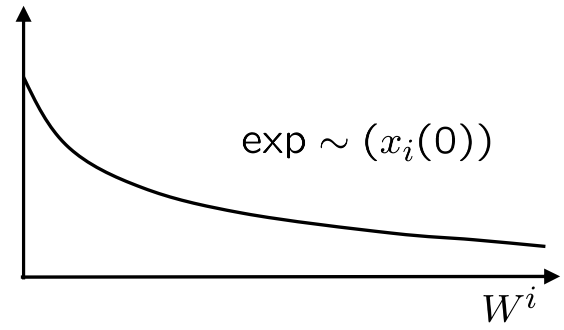

- Algorithm and run-time based on Mosk-Aoyama/Shah
- Node i has $x_i(0)$, seeks $\sum_{j=1}^{n} x_j(0)$
  - Generates W^i : Sample from exp ~ $\mu_i = \frac{1}{x_i(0)}$
  - Communicates with neighbors until learns

$$W^* = \min(W^1, W^2, \ldots, W^n)$$

(Information Spreading)

⇒ Exponentially distributed with mean $\frac{1}{\sum_{j=1}^{n} x_j(0)}$

  - Repeat r times; Obtain $W_1^*, W_2^*, \ldots, W_r^*$
  - Compute $\frac{1}{r}\sum_{i=1}^{r} W_i^*$ ⇒ For large r, approaches $\frac{1}{\sum_{j=1}^{n} x_j(0)}$
  - Set $x_i(r) = \frac{r}{\sum_{i=1}^{r} W_i^*}$ ⇒ For large r, $P\left(|x_i(r) - k| \leq \varepsilon k\right) > 1 - \delta$

Calculation of Algorithm Running Time

- Information Spreading:
  Time until all nodes have the minimum $W^* = \min(W^1, W^2, \ldots, W^n)$

$$T_{spr} = O\left(\frac{\log n + \log \delta^{-1}}{\Phi(P)}\right) \qquad \Phi(P) = \min_{\substack{S \subset V \\ |S| \le \frac{n}{2}}} \frac{\sum_{i \in V \setminus S} \sum_{j \in S} p_{ij}}{|S|}$$

- Algorithm Running Time: Until all nodes have $W_1^*, W_2^*, \ldots, W_r^*$

$$T = O(rT_{spr})$$

  - For all nodes, $P\left(|x_i(r) - k| \le \varepsilon k\right) > 1 - \delta$ when,

    $$r \ge \varepsilon^{-2}(1 + \log \delta^{-1})$$

  - So, $T = O\left(\frac{\varepsilon^{-2}(\log n + \log \delta^{-1})}{\Phi(P)}\right)$

Run-time of Mosk-Aoyama/Shah algorithm is inversely related to conductance because it depends on information spreading.

An Algorithm for Nodes Communicating Bits

- In our formulation, nodes communicate via erasure channels
  - With probability p_{ij} : Node i sends M bits to j
    $1 - p_{ij}$: Node j gets nothing 'erasure'
  - $C_{ij} = p_{ij}$

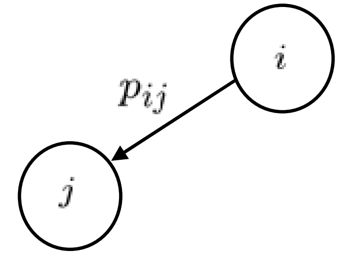

- Use same algorithm, except,
  - Truncate exponentials at large enough m ,
  - Quantize samples and communicate bits.

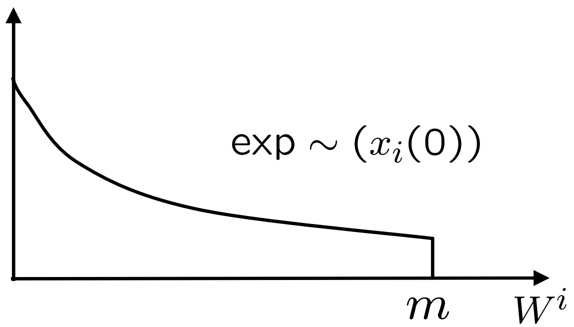

- For all nodes, $P\left(|x_i(T) - k| \leq \varepsilon k\right) > 1 - \delta$, when,

$$T = O\left(\frac{\varepsilon^{-2}(\log n + \log \delta^{-1})}{\Phi(G)} \log n\right)$$

Communicating bits slows algorithm down by $\log n$

Summary

- Information Theoretic tools in context of distributed computation.
    - Captured information needed to reduce variance of estimates.
    - Assumed unlimited computation resources at nodes.
  - Advantages of the technique
    - Inequalities used are simple.
    - No assumptions about the process governing estimates.
    - No assumptions about communication algorithm.

- Tight bound for run-time when nodes compute sum via erasure channels.
  - Our algorithm's run-time scales reciprocally with conductance.
  - Algorithm-independent lower bound scales reciprocally with conductance.

Lower bound is tight in capturing effect of network topology via conductance.
Our algorithm's run-time is optimal in its dependence on conductance.

Future Work

- Limitation on Computation
  - Nodes compute different parts of the functions
  - Belief Propagation

- The role of Mobility on distributed decision
  - Tradeoffs between mobility and communication

Networked systems-Quad Chart

Understand the impact of a network on computation/control

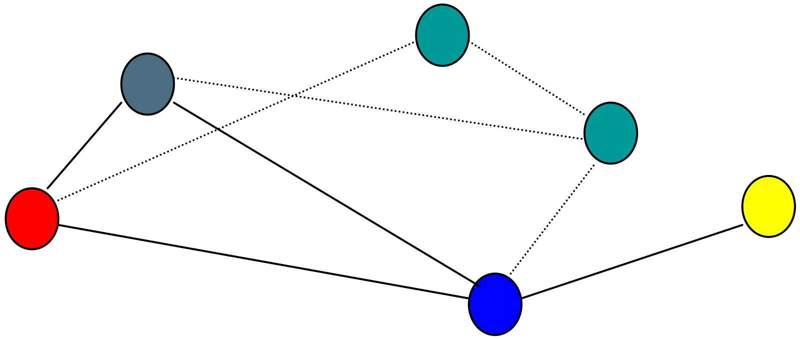

- Long term: Characterize fundamental limitations and capabilities of distributed decision systems
- Short term: Impact of topology of an unreliable transport layer on run-time for computation

Results

- Information Theoretic formulations and techniques for distributed computation
- Algorithm-independent lower bounds on run-time for computation
- A specific scenario for which the bound is tight.

Funding

- Partially supported by AFOSR grant # FA9550-08-1-0085
- People partially supported
  - Dahleh (PI)
  - Ayaso (Student, Post Doc)
  - Como (Post Doc)
  - Itani (student)