Distributed Computation over Random Wireless Networks

D. Manjunath

Bharti Centre for Communication
Department of Electrical Engineering
IIT Bombay

9 Jun 2009

Outline

Preliminaries

- Primary motivation from sensor networks where n sensors are deployed in a geographic area.
- Each sensor node makes a measurement in the area that it covers; measured value is assumed to be a 1-bit binary data. Node i has bit x_i and $\mathbf{x} = [x_1, \ldots, x_n]$.
- A special node, (sink, collector) needs to evaluate

$$f : \{0,1\}^N \to \mathcal{B}$$

where $\mathcal{B}$ is the codmain of f and assumed finite but may depend on n. Typical interest is in max (OR), min (AND), parity, histogram, and, the identity function that wants $\mathbf{x}$ at sink.

  - Symmetric functions: $f(x_1, \ldots, x_n)$ is invariant to permutations of $x_1, \ldots, x_n$.
  - Decomposable functions: $f(\mathbf{x}) = g(f(\mathbf{x}^{(1)}), \ldots, f(\mathbf{x}^{(K)}))$ where $\mathbf{x}^{(1)}, \ldots \mathbf{x}^{(K)}$ is a partition of $\mathbf{x}$.

Preliminaries

- Nodes communicate over a link;
  - In a wired channel, transmission is ‘heard’ by only one receiver.
  - In a wireless channel all ‘neighbours’ (nodes in spatial proximity) can decode the transmission, possibly with error. Further, none of the neighbours can transmit simultaneously.
  - Our interest is in wireless networks; they should exploit spatial reuse for efficiencies.
- Communication links can be assumed to be noisy or noisefree.
  - Many noise models have been considered; here we will consider the *weak noise model* of a ‘binary symmetric channel’ with error probability exactly ϵ.
- Objective: Design communication and computation algorithms to achieve the objective—obtain $f(\mathbf{x})$ at the sink.
- Performance issues: energy (number of transmissions and/or receptions), throughput and delay.

Network Models

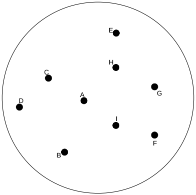

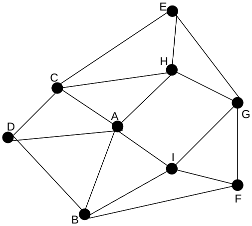

Single and multihop networks

- Single hop, collocated, or broadcast networks: Every node can decode transmissions from every other node. Network graph is fully connected.
- Multihop networks: Network nodes could be placed on a 'regular' grid or they could be randomly deployed.

Protocol

- Protocol defines a sequence of computations and transmissions.
- *Collision free protocol:* collisions avoided, sequence of transmitters may depend on history.
- *Oblivious* protocol: schedule fixed beforehand.

In a network with noisy links.

- *Oblivious* protocol guarantees no collisions.
- Valid protocol: error probability as small as required. Two variations—arbitrarily small or going to zero as $n \to \infty$.

Random, Multihop Noisefree Networks

- Nodes $1, 2, \ldots n$ are randomly distributed over $[0, 1]^2$.
- Node i is located at y_i and has data x_i.
- Define $\mathbf{x} := [x_1, \ldots, x_n]$ and $\mathbf{y} := [y_1, \ldots, y_n]$.
- x_i belongs to a finite set $\mathcal{X}$, e.g., $\mathcal{X} = \{0, 1\}$.
- We will assume y_i is uniformly distributed in $[0, 1]^2$ and x_i is arbitrary.

Random Multihop Networks: Quick Overview of Basics

- n nodes are uniformly distributed in $[0, 1]^2$.
- Boolean model for network (communication) graph: Nodes at Euclidean distance less than r_n have communicating edge, i.e., the communicating graph is a random geometric graph.
- Communication constraints: A node cannot receive and transmit simultaneously; Collisions at receiver if more than one of its neighbours are transmitting.
- Need connected network; Asymptotic connectivity threshold in a random network: $r_n = \Omega\left(\sqrt{\frac{\log n}{n}}\right)$.
- Use $r_n = c_1\left(\sqrt{\frac{\log n}{n}}\right)$ to minimise transmission energy and maximise reuse. This means localities of areas $\Theta(r_n)$ are like small colocated networks, each having $\Theta(\log n)$ nodes.

Key Background Result

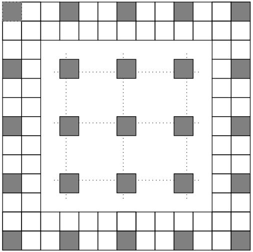

- Tessellate $[0,1]^2$ into square cells of side $\Theta\left(\sqrt{\frac{\log n}{n}}\right)$.
- The number of nodes in each square is $\Theta(\log n)$ with high probability.

Summary of Results for Noisefree Networks

[Giridhar & Kumar,2005]

- Assume block coding—each node collects N samples and $f(\cdot)$ is evaluated for N samples

	Single Hop	Random Multihop
Identity	$\Theta\left(\frac{1}{n}\right)$	$\Theta\left(\frac{1}{n}\right)$
Histogram	$\Theta\left(\frac{1}{n}\right)$	$\Theta\left(\frac{1}{\log n}\right)$
Type-Sensitive	$\Theta\left(\frac{1}{n}\right)$	$\Theta\left(\frac{1}{\log n}\right)$
Type-Threshold(min)	$\Theta\left(\frac{1}{\log n}\right)$	$\Theta\left(\frac{1}{\log \log n}\right)$

Motivating Structure-Free Networks

[Kamath & M,2008]

- Network organisation and clock synchronisation are requirements of above protocols; organisation can be expensive.
- Can we do away with knowledge of structure of the network and time synchronisation? Use random access (Aloha for MAC.
- What are the tradeoffs for differing levels of structure? To compute the max, no delay penalty!
- There will be an energy penalty.

Random, Multihop Noiseless Structure Free Networks

- Nodes know neither their absolute nor relative locations in the network
- Nodes therefore, have no idea about the n etwork topology
- We would also like that the nodes not have their clocks synchronised;
- Each node however, knows n, the total number of nodes

Computing MAX Once

Protocol

Let x_i be the data bit at Node i and $\S := \max_{1 \leq i \leq n} x_i$.

- *In each slot, Node i transmits with probability* $p = p_n$ *independently of all other transmissions in the network*
- *If Node i receives a bit successfully in slot* t*, then it sets* $X_i(t)$ *to be this bit, else it sets* $X_i(t) = 0$
- *Node i initiates* $Y_i(0) = 0$ *and updates* Y_i *using* $Y_i(t) = \max\{Y_i(t-1), X_i(t)\}$
- *If Node i transmits in slot* t*, then it would transmit* $T_i(t) = Y_i(t-1)$

- MAX will be known to all nodes if we wait long enough
- Issue: Time to obtain MAX at sink and suitable r_n and p_n.

Computing MAX Once

- An attempt probability of $p = p_n = (k_1 c_2 \log n)^{-1}$ results in a successful transmission by *some* node from a cell in any given slot with probability atleast $p_S := \frac{c_1}{c_2 e}$, which is a constant independent of n

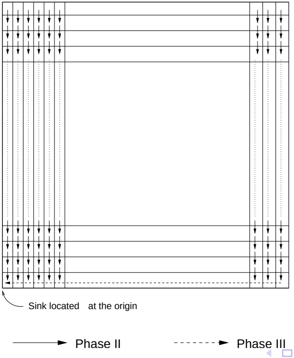

Computing MAX Once

Assume sink is located at the origin.

- All nodes in a cell have transmitted at least once in $O(\log^2 n)$ time slots w.h.p.
- Now the data has to diffuse to the cells.
- The diffusion to the bottom of square is completed in $O\left(\sqrt{\frac{n}{\log n}}\right)$ slots w.h.p.
- Now each column of cells has computed the MAX in the column and the partial result is available at some node in the bottom row.
- The result now diffuses to the sink in $O\left(\sqrt{\frac{n}{\log n}}\right)$ slots w.h.p.

Protocol One-Shot MAX computes the MAX in $O\left(\sqrt{\frac{n}{\log n}}\right)$ slots with probability atleast $(1 - \frac{k}{n^\alpha})$ for any positive constants k, α.

Pipelined Computation of MAX

- Some structure in the network can achieve higher throughput
- Nodes can compute their hop distance from the sink node; modulo 3, suffices

Protocol

Let $Z_i(r)$ be the data at Node i in round r. The sink wishes to compute $\mathcal{Z}(r) = \max_{1 \le i \le n} Z_i(r)$
Node i transmits in each slot with probability $p = p_n$ independently of all other transmissions in the network
Each round consists of τ slots.

- *Transmission: If Node i transmits in slot t of round r, then it transmits $(A_i, B_i, T_i(r,t))$ where $(A_i, B_i) \equiv h_i \mod 3$, $T_i(r,t) = \max\{Z_i(r - d + h_i), Y_i(r-1)\}, d = \frac{2}{s_n}$ is an upper bound on the hop distance of a node*

Pipelined Computation of MAX

Protocol

- *Reception: If Node i senses an idle or a collision or a successful transmission by a node with hop distance different from* $(h_i + 1)$*, then it sets* $X_i(r,t) = 0$*, else it sets* $X_i(r,t)$ *to be equal to the data bit received. It sets* $Y_i(r,0) = 0$ *at the beginning of the round and then, updates as* $Y_i(r,t) = \max\{Y_i(r,t-1), X_i(r,t)\}$*. At the end of the round, it sets* $Y_i(r) = Y_i(r,\tau)$*.*

The sink node, Node 0*, decodes the MAX as*
$\mathcal{Z}(r-d) = \max\{Z_s(r-d), Y_s(r)\}$

Pipelined Computation of MAX

- In round r, the set of nodes at hop distance $(h+1)$ relay the MAX of the data bits from round $(r-d+h+1)$, of all nodes with hop distance greater than h to some node at hop distance h.
- For successful computation of the MAX at round r, we require each node at hop distance h to have successfully transmitted atleast once in round $(r-d+h)$.
- From our analysis of Phase I of Protocol One-Shot MAX, this will happen w.h.p. if $\tau = \Theta(\log^2 n)$.

The protocol achieves a throughput of $\Omega\left(\frac{1}{\log^2 n}\right)$.

Computing the Histogram Once

- Each node generates an independent exponential random variable of mean 1 in each round.
- In each slot, each node transmits independently with probability p and listens with probability $(1-p)$.
- Identical to MAX except that the truncated, quantised random number is transmitted in each slot.

Theorem

If all the nodes execute the protocol One-Shot Histogram, then the histogram $\left(\frac{\hat{n}_0}{n}, \frac{\hat{n}_1}{n}\right)$ is available to all nodes in $O\left(n^{7/2}(\log n)^{1/2}\right)$ slots with probability atleast $(1-\frac{\pi^2}{6n})(1-\frac{k}{n^{\alpha-3}})$ and with the following accuracy: $\frac{n_b}{n}e^{-\frac{3+\pi^2}{6n}}e^{-d_{n_b}} \leq \frac{\hat{n}_b}{n} \leq \frac{n_b}{n}e^{\frac{3+\pi^2}{6n}}e^{-d_{n_b}}$ for $b = 0, 1$. ($d_n \rightarrow 0$ as $n \rightarrow \infty$)

Summary

One-shot MAX	Structure-Free	Coordinated
Time	$\Theta\left(\sqrt{\frac{n}{\log n}}\right)$	$\Theta\left(\sqrt{\frac{n}{\log n}}\right)$
Transmissions	$\Theta\left(\frac{n^{3/2}}{\log^{3/2} n}\right)$	$\Theta(n)$

Pipelined MAX	Hop Distance	Coordinated
Throughput	$\Theta\left(\frac{1}{\log^2 n}\right)$	$\Theta\left(\frac{1}{\log n}\right)$
Transmissions	$\Theta(n \log n)$	$\Theta(n)$

Approximate Computation

[Iyer, M & Sundaresan,2009]

- Connectivity requires that $r = O\left(\sqrt{\frac{logn}{n}}\right)$. This means the degree of every node is $O(\log n)$.
- Degree determines spatial reuse; And hence the throughput. Ideally, we would like a degree of $O(1)$.
- Operate in the percolation regime.
- Choose communication range $r_n(\lambda)$ to satisfy

$$nr_n^2(\lambda) = \lambda,\ \forall\ n, \tag{4.1}$$

for a particular $\lambda > \lambda_c$ where λ_c is the density at which there is percolation in the Poisson point process.

Percolating Networks

- Can reduce r_n and yet have an arbitrarily large fraction of the nodes connected in a giant component.
- Second largest component can also be precisely characterised.

Lemma

Let $r_n(\lambda)$ be a sequence satisfying $nr_n^2(\lambda) = \lambda$ with $\lambda > \lambda_c$. For every $\delta > 0$, size of largest component $L_{1,n}(\lambda)$ of $G(\mathcal{X}_n; r_n(\lambda))$ satisfies

$$\limsup_{n\to\infty} \frac{1}{\sqrt{n}} \log \Pr\left\{ \left| \frac{L_{1,n}(\lambda)}{np_\infty(\lambda)} - 1 \right| \geq \delta \right\} < 0$$

and the size of the second largest component $L_{2,n}(\lambda)$ satisfies

$$\limsup_{n\to\infty} \frac{1}{\sqrt{n}} \log \Pr\left\{ L_{2,n}(\lambda) > \delta n \right\} < 0.$$

A Graph with Bounded Degree

- The node does not have a bounded degree. But the number of nodes with large degree is small.

Lemma

For every $\delta, \varepsilon > 0$, there exists a sufficiently large λ and sufficiently large k such that the sequence of graphs $G(\mathcal{X}_n; r_n(\lambda))$ (indexed by n) satisfies the following:

(1) *The fraction of nodes in the largest component is at least $1 - \delta$, i.e., $n^{-1}L_{1,n}(\lambda) \geq (1 - \delta)$, in c.c. as $n \to \infty$;*

(2) *The fraction of nodes with degree upper bounded by k is at least $1 - \varepsilon$, i.e., $n^{-1}Z_{k,n}(\lambda) \leq \varepsilon$, in c.c. as $n \to \infty$.*

A Graph with Bounded Degree

- For nodes that have degree more than a specified constant, we can disconnect the excess nodes from the graph and rescale the transmission range to retain the giant component. We can show

Theorem

For every $\delta > 0$, there exists sufficiently large but finite λ' and k' such that the sequence of random geometric graphs $G(\mathcal{X}_n; r_n(\lambda'))$ contains a subgraph $G(V_n; r_n(\lambda'))$, where $V_n \subset \mathcal{X}_n$, with the following properties:

- *The subgraph is connected;*
- *The maximum degree of the subgraph is upper bounded by k';*
- $|V_n|/n \geq 1 - \delta$ *in c.c.*

Computation on the Graphs with Bounded Degree

Use the algorithms identical to those developed earlier on the graph with bounded degree.

Theorem

For error free networks, for any $\delta \in (0,1)$ and $\varepsilon \in (0,1)$, there is a protocol that computes the histogram with the following performances:

1. $\Pr\{e_n > 3\delta\} \leq 2\varepsilon$ *for all sufficiently large n.*
2. *The refresh rate is* $\Theta(1)$.
3. *The number of transmissions is* $\Theta(n)$ *and the total transmission energy is* $\Theta(n^{1-d/2})$. *The number of receptions is* $\Theta(n)$.
4. *The delay is* $\Theta(\sqrt{n})$

Computation on the Graphs with Bounded Degree

- The computed histogram is only probably approximately correct (PAC); the normalised histogram error is at most 3δ with probability greater than $1 - 2\varepsilon$.
- The PAC relaxation enables us to compute the PAC histogram at a constant refresh rate.
- Histogram is computed using the average function with refresh rate $\Theta(1)$. This is a $\log \log n$ over the previously best known algorithm.
- Any continuous function of the histogram can be computed in a PAC fashion at refresh rate $\Theta(1)$.
- Median and mode are type-sensitive functions but cannot be computed using this method. Computation of type-threshold functions may also result in arbitrarily large errors.

MAX in Noisy Multihop Networks

- We can construct a protocol with **no order penalty in either time or in the number of transmissions**.

Theorem (Kanoria & M,2007)

MAX(or OR) can be computed in a noisy RP network by an oblivious protocol using $\Theta(n)$ *total transmissions in a time* $\Theta\left(\sqrt{\frac{n}{\log n}}\right)$.

Parity in a Colocated Network

[Gallager,1988]

- Partition m nodes into m/k groups of k.
- Each node transmits its bit j times; each node estimates the value of the bit of each member of its group using a majority rule.
- Every node estimates the parity of its group from estimate of bits and transmits this estimate.
- Thus estimate of every group is transmitted k times.
- Each node uses majority rule to estimate parity of each group of bits (from the k estimates). And then estimate the parity of the m bits.
- k should be $O(\log m)$ and $j = O(\log k)$.
- Requires $O(m \log \log m)$ transmissions.
- Recently shown that this is indeed the optimal scheme by [Goyal, Saks & Kindler, 2006].

Histogram in a Random Planar Network

[Ying, Srikant & Dullerud,2006]

Intra cell: Each cell is a broadcast network. Adapt [Gallager,1988]

- Each node transmits its value $O(\log \log n)$ times.
- In each cell $O\left(\frac{\log n}{\log \log n}\right)$ nodes are selected to broadcast their estimates of the intra cell histogram.
- Cell centre decodes these transmissions using a majority rule to estimate the cell histrogram.

Inter Cell:

- Use block codes or repetition codes for point-to-point communication to propagate the aggregated values up the tree to the sink.

Requires $O(n \log \log n)$ bit transmissions.

Histogram in a Random Planar Network

Can we do better than $O(n \log \log n)$ bit transmissions?
No!
Let δ be a desired upper bound on the probability that $f(\cdot)$ is in error

Theorem (Dutta, Kanoria, M & Radhakrishnan,2008)

Let $R \leq n^{-\beta}$ for some $\beta > 0$. Let $\delta < \frac{1}{2}$ and $\epsilon \in (0,1)$. Then, with probability $1 - o(1)$ (over the placement of processors) every δ-error protocol on $\mathcal{N}(n,R)$ with ϵ-noise for computing the parity function $\oplus : \{0,1\}^n \to \{0,1\}$ requires $\Omega(n \log \log n)$ transmissions.

Acknowledgments

C Dutta, TIFR Mumbai
Srikanth K Iyer, IISc Bangalore
S Kamath, formerly IIT Bombay, now at UC Berkeley ’
Y Kanoria, formerly IIT Bombay, now at Stanford
R Mazumdar, Univ of Waterloo
J Radhakrsishnan, TIFR Mumbai
Rajesh Sundaresan, IISc Bangalore