6.207/14.15: Networks
Lectures 7: Information Spread, Distributed Computation

Outline

Information spread.

Conductance determines how long.

Distributed computation.

Equals information spread.

Markov chain convergence.

Spectral gap and *conductance*.

References:

Shah, Chapter 3 (3.1-3.2), Chapter 5 (5.1-5.2)

Information spread

Network graph G over $N = \{1, \ldots, n\}$ nodes, edges E

- Given information at one of the nodes, spread it to *all* nodes
- By "Gossiping"
- How long does it take?

Gossip dynamics:

- At each time, each node $i \in N$ does the following:
- if node i does not have information, nothing to spread or gossip
- else if it does have information, it sends it to one of it's neighbors
  - let $P_{ij} = \mathbb{P}(i \text{ sends information to } j)$
  - by definition, $\sum_{j \in N} P_{ij} = 1$, and
  - $P_{ij} = 0$ if j is not neighbor of i
- Example: *uniform* gossip
  - $P_{ij} = 1/k_i$ for all $(i, j) \in E$

Information spread

Why study Gossip dynamics

This is how socially information spreads

More generally, this is how "contact" driven network effect spreads

This is how large scale distributed computer systems are built

e.g. Cassandra, an Apache Open Source Distributed DataBase used by some of the largest organizations including Netflix, etc.

Key question

How long does it take for all nodes to receive information?

How does it depend on the graph structure, P?

Information spread

Let us consider few examples:

Ring graph

Star graph

Complete graph

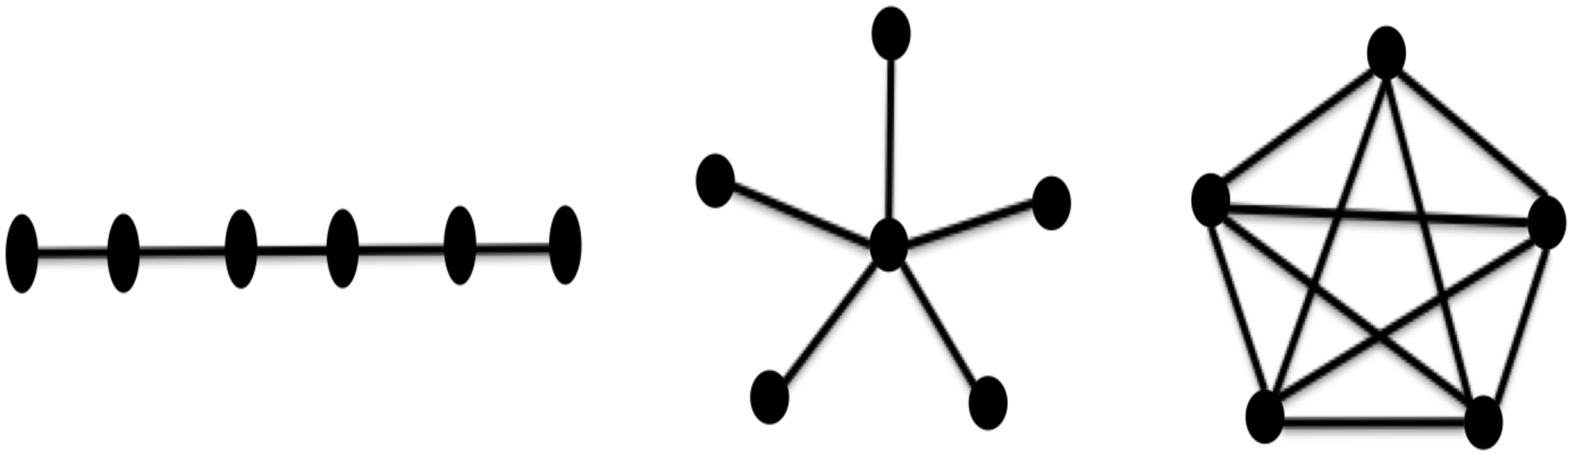

Information spread

Key question

How long does it take for all nodes to receive information?
How does it depend on the graph structure, P?

A crisp answer

$$\mathrm{T}_{spr} \sim \frac{\log n}{\Phi(P)}$$

where $\Phi(P)$ is the *conductance* of P (and hence graph)

Conductance

Conductance of $P = [P_{ij}]$ is defined as

$$\Phi(P) = \min_{S \subset N : |S| \leq n/2} \frac{\sum_{i \in S, j \in S^c} P_{ij}}{|S|} \tag{1}$$

Examples: uniform gossip

Ring graph: $\Phi \sim 1/n$

Star graph: $\Phi \sim 1/n$

Complete graph: $\Phi \sim 1$

Conductance and Information spread

Let us consider how information spreads

Let $S(k) \subset N$ be nodes that possess information at time k

Initially, $S(1) = \{i\}$ for some $i \in N$

Consider $j \notin S(k)$. Then $j \in S(k+1)$

some $\ell \in S(k)$ contacts it

– Therefore,

$$\begin{aligned}
\mathbb{P}(j \in S(k+1) | j \notin S(k)) &= 1 - \mathbb{P}(j \notin S(k+1) | j \notin S(k)) \\
&= 1 - \mathbb{P}(\cap_{\ell \in S(k)} \ \ell \text{ does not contact } j) \\
&= 1 - \prod_{\ell \in S(k)} \mathbb{P}(\ell \text{ does not contact } j) \\
&= 1 - \prod_{\ell \in S(k)} (1 - P_{\ell j}) \\
&\approx \sum_{\ell \in S(k)} P_{\ell j}
\end{aligned}$$

Conductance and Information spread

In summary:

$$\mathbb{E}[|S(k+1)| \quad |S(k)||S(k)] \approx \sum_{j \notin S(k)} \sum_{\ell \in S(k)} P_{\ell j}.$$

Therefore:

$$\begin{aligned}
\mathbb{E}[|S(k+1)||S(k)] &\approx |S(k)| + \sum_{j \notin S(k)} \sum_{\ell \in S(k)} P_{\ell j} \\
&\approx |S(k)| \Big(1 + \frac{\sum_{j \notin S(k)} \sum_{\ell \in S(k)} P_{\ell j}}{|S(k)|}\Big) \\
&\geq |S(k)| \Big(1 + \min_{S \subset N} \frac{\sum_{j \notin S} \sum_{\ell \in S} P_{\ell j}}{|S|}\Big)
\end{aligned}$$

Conductance and Information spread

Continuing:

$$\mathbb{E}[|S(k+1)||S(k)] \geq |S(k)|\Big(1 + \min_{S \subset N} \frac{\sum_{j \notin S} \sum_{\ell \in S} P_{\ell j}}{|S|}\Big)$$

If $S(k) \leq n/2$, then we can restrict for $S \subset N$ s.t. $|S| \leq n/2$

Using definition of conductance

$$\begin{aligned} \mathbb{E}[|S(k+1)||S(k)] &\geq |S(k)|\Big(1 + \Phi(P)\Big) \\ &\geq |S(1)|\Big(1 + \Phi(P)\Big)^k \\ &\approx \exp\Big(k\Phi(P)\Big). \end{aligned}$$

Therefore, approximately it takes $\frac{\log n}{\Phi(P)}$ steps to reach $n/2$ nodes

Conductance and Information spread

Spreading time

Invoking symmetries to go from $n/2$ to n nodes

Using a little sophisticated probabilistic analysis

And some, it can be concluded that

$$\mathrm{T}_{spr} \sim \frac{\log n}{\Phi(P)}$$

where $\Phi(P)$ is the *conductance* of P (and hence graph)

In general, it is the *information bottleneck*:
the information spread takes $1/\Phi(P)$

Distributed computation

Generic question

Given network G over nodes N with edges E

Each node $i \in N$ has information x_i

Compute a global function:

$$f(x_1, \ldots, x_n)$$

by communicating along the network links

processing *local* information at each of the node continually

while keeping *limited* local state at each node

Distributed computation

The simplest possible example

(estimate) number of nodes in the network at each node locally

there is no globally agreed unique names for each node

only local communications are allowed while keeping local state small

Well, we can play a game to understand this

Let's figure out how many students are there in this class?

Know your neighbors

A distributed algorithm

Every node generates a random number

Node $i \in N$ draws random variable R_i as per an Exponential distribution of mean 1
(in Python: import random; random.expovariate(1))

Compute global minimum, $R^\star = \min_{i \in N} R_i$

Using *Gossip* mechanism

Repeat the above for L times

$R^\star_\ell, \; 1 \leq \ell \leq L$ be global minimum computed during these L trials

Estimate of number of neighbors

$\hat{n} = \frac{L}{\sum_{\ell=1}^{L} R^\star_\ell}$

Exponential distribution

Exponential distribution with parameter $\lambda > 0$

- X be random variable with this distribution: for any $t \in \mathbb{R}$,

$$\mathbb{P}\big(X > t\big) = \exp\big(-\lambda t\big).$$

- Minimum of exponential random variables
  - Let X_i, $i \in N$ be independent random variables
  - Distribution of X_i is Exponential with parameter λ_i, $i \in N$
  - $X^* = \min_{i \in N} X_i$

$$\begin{aligned}\mathbb{P}\big(X^* > t\big) &= \mathbb{P}\big(\cap_{i \in N} X_i > t\big)\\ &= \prod_{i \in N} \mathbb{P}\big(X_i > t\big)\\ &= \prod_{i \in N} \exp\big(-\lambda_i t\big)\\ &= \exp\Big(-\big(\sum_i \lambda_i\big)t\Big).\end{aligned}$$

Exponential distribution

- Exponential distribution with parameter $\lambda > 0$
  - X be random variable with this distribution: for any $t \in \mathbb{R}$,

$$\mathbb{P}\big(X > t\big) = \exp\big(-\lambda t\big).$$

- Minimum of exponential random variables
  - $X^* = \min_{i \in N} X_i$ has exponential distribution with parameter $\sum_{i \in N} \lambda_i$
- Mean of exponential variable X with parameter $\lambda > 0$

$$\begin{aligned}\mathbb{E}[X] &= \int_0^\infty \mathbb{P}(X > t)dt \\ &= \int_0^\infty \exp\big(-\lambda t\big)dt \\ &= \frac{1}{\lambda}\Big[\exp\big(-\lambda t\big)\Big]_\infty^0 \\ &= \frac{1}{\lambda}.\end{aligned}$$

Exponential distribution

Back to counting nodes

- Node i's random number has exponential distribution of parameter 1
- All nodes computed minimum of these numbers
- Hence minimum had exponential distribution with parameter n
- That is, mean of the minimum is $1/n$
- Averaging over multiple trials gives a good estimation of $1/n$

Adding up numbers

- Node i has a number x_i
- Node i draws random variable per exponential distribution of parameter x_i
- Then minimum would have exponential distribution with parameter $\sum_i x_i$
- Subsequently, algorithm is computing estimation of $\sum_i x_i$

Computing minimum

Gossip algorithm

- Node $i \in N$ has value R_i and goal is to compute $R^\star = \min_i R_i$
- Node $i \in N$ keeps an estimate of global minimum, say $\hat{R}_i^\star$
- Initially, $\hat{R}_i^\star = R_i$ for all $i \in N$
- Whenever node j contacts i
  - Node j sends $\hat{R}_j^\star$ to i
  - Node i updates $\hat{R}_i^\star = \min\left(\hat{R}_j^\star, \hat{R}_i^\star\right)$

How long does it take for everyone to know minimum?

- Suppose $R_1 = R^\star$.
- Then the spread of minimum obeys exactly same dynamics as
  - Spreading information starting with node!
- That is, *information spread* $=$ *minimum computation*!

Distribution computation = Information spread

Distribution computation

Given network G over nodes N with edges E

Each node $i \in N$ has information x_i

Goal is to compute a global function

$$f(x_1, \ldots, x_n) = \sum_i f_i(x_i)$$

Gossip algorithm

P be gossip probability matrix over G

Computing $f(\cdot)$ via multiple minimum computation

For $(1 \pm \varepsilon)f(\cdot)$ estimation, need $1/\varepsilon^2$ computations

$$T_{dist-comp} \sim \frac{1}{\varepsilon^2} T_{min}$$

And time to compute minimum, T_{min} is information spread

$$T_{min} = T_{spr} \sim \frac{\log n}{\Phi(P)}$$

Markov chain

Properties of a Markov chain

- P be probability transition matrix
- Assume it be irreducible and aperiodic
- By Perron-Frobenius theorem
  - The largest eigenvalue $\lambda_1 = 1$
  - And $|\lambda_2| < \lambda_1 = 1$
- Define gap $g(P) = 1 - |\lambda_2|$

For simplicity, assume $P = P^T$

- By definition, for any P, $P\mathbf{1} = \mathbf{1}$
- Since $P = P^T$, we have $P^T\mathbf{1} = \mathbf{1}$
- That is, $\frac{1}{n}\mathbf{1}$ is stationary distribution

Markov chain

Dynamics

Let $p(k)$ be probability distribution at time k

$$p(k+1) = P^T p(k)$$

Let $\mathbf{1}, s_2, \ldots, s_n$ be eigenvectors of P^T

with associated eigenvalues $1, \lambda_2, \ldots, \lambda_n$

$0 \leq |\lambda_n| \leq \cdots \leq |\lambda_2| < 1$

Then, as argued for linear dynamics, we have

$$p(k) = c_1 \mathbf{1} + \sum_{i=2}^{n} \lambda_i^k c_k s_k$$

with some constants $c_1, \ldots, c_n$

Markov chain

Dynamics (continuing)

Therefore:

$$\|p(k) - c_1 \mathbf{1}\| \leq \sum_{i=2}^{n} |\lambda_i|^k |c_i| \|s_i\|$$

$$\leq (n-1) C |\lambda_2|^k$$

where $C = \max_{i=2}^{n} |c_i| \|s_i\|$

Subsequently

$$k \geq \frac{\log n + \log C + \log \frac{1}{\varepsilon}}{\log \frac{1}{|\lambda_2|}} \quad \Rightarrow \quad \|p(k) - c_1 \mathbf{1}\| \leq \varepsilon.$$

Markov chain

Convergence

The ε-convergence time scales as

$$T_{conv}(\varepsilon) \sim \frac{\log n + \log \frac{1}{\varepsilon}}{\log \frac{1}{|\lambda_2|}}.$$

Using $\log(1-x) \approx -x$ for $x \in (0, 1)$, we get

$$\log \frac{1}{|\lambda_2|} = -\log|\lambda_2| = -\log(1-(1-|\lambda_2|))$$
$$= -\log(1-g(P)) \approx g(P)$$

That is, the ε-convergence time scales as

$$T_{conv}(\varepsilon) \sim \frac{\log n + \log \frac{1}{\varepsilon}}{g(P)}$$

Markov chain

(Spectral) gap and conductance

Markov chain can not converge faster than information spread

And information spreads in time $1/\Phi(P)$

That is (ignoring constants)

$$\frac{1}{\Phi(P)} \leq \frac{1}{g(P)} \quad \Leftrightarrow \quad g(P) \leq \Phi(P)$$

A remarkable fact known as Cheeger's inequality:

$$\frac{1}{2}\Phi(P)^2 \leq g(P) \leq 2\Phi(P).$$

MIT OpenCourseWare
https://ocw.mit.edu

14.15J/6.207J Networks
Spring 2018

For information about citing these materials or our Terms of Use, visit: https://ocw.mit.edu/terms.