

November 16, 2007
Prepared by Adam Guetz

Diffusion, gossip, and protocol design

Suppose you have a distributed system with many components. For example, an airline, a sensor network, etc. How to best disseminate information in this network without overloading it?

One way: build a binary tree, send information to neighbors.

Problem: this is not very robust, is vulnerable to disconnection of edges. Other topologies are also possible, such as expander graphs, regular graphs, etc., but these may have problems as well.

Three main types of randomized rumor spreading algorithms have been proposed:

- **Push based methods:** If you know the rumor, call a random person, inform them if they haven't heard it. exponential growth until about $n/2$ nodes are informed ($\log(n)$ rounds, n total messages) . Assume that there are cn nodes informed. The probability that an uninformed node receives the rumor in this round is $1 - (1 - 1/n)^{cn} > 1 - 1/e$. ($O(\log n)$ rounds, $O(n \log n)$ calls. So push does well for the first half of the nodes, but then does worse for the second half
- **Pull based methods:** If you don't know the rumor, call a random person and ask them. Takes $O(\log n)$ rounds to inform the first half. Let u_t be the number of people who are not informed.

$$E(u_{t+1}/n) = (u_t/n)^2$$

After $\log \log n$ rounds, everybody is informed.

- **Push and Pull based methods :** It seems reasonable to combine the two methods, but does it achieve better results? Karp et al. (2000) were able to show that it does.

Theorem 11.1 (Karp et al. 2000) *The push-pull method terminates after $\log_3 n + O(\log \log n)$ rounds and $O(n \log \log n)$ messages.*

Proof:

Let s_t be the number of informed nodes and let $u_t = n - s_t$. We split the process into four phases, ordered by the number of infected nodes.

Phase 1 (start): $1 \leq s_t \leq \log^4 n$

The probability that a message is pushed to an informed node is $\frac{\text{polylog}n}{n}$. So with high probability, phase 1 ends after $O(\log \log n)$ rounds.

Phase 2 (exponential growth): $\log^4 n \leq s_t \leq n/\log n$

Let m_t be the number of messages sent at time t .

Then $E(m_t) = 2s_t$, because each informed node calls one player and is called by one player on average. Applying a Chernoff bound shows that is tight within $o(1/\log(n))$ w.h.p. Some of the messages are *wasted*, but the probability of wasting a message can be bounded by

$$s_{t-1}/n + m/n \leq (3 + o(1/\log n))/\log n.$$

Therefore,

$$S_{t+1} \geq S_t(3 - O(1/\log n)).$$

The number of rounds in this phase is $\leq \log_3 n + O(\log \log n)$.

Phase 3 (quadratic shrinking): $n/\log n \leq s_t \leq n - \sqrt{n} \log^4 n$

Even if we only take into account the pull transmissions, we obtain

$$E\left(\frac{u_{t+1}}{n}\right) \leq \left(\frac{u_t}{n}\right)^2.$$

Applying a Chernoff bound gives

$$u_{t+1} \leq \frac{u_t^2}{n} - O\left(\frac{1}{\log n}\right).$$

This round takes $O(\log \log n)$ rounds.

Phase 4 (finish): $u_t \leq \sqrt{n} \log^4 n$

Each uninformed person has at least probability

$$1 - \frac{\log^4(n)}{\sqrt{n}}$$

to receive the message by a pull transmission. Therefore, in a constant number of rounds, phase 4 terminates. ■

Average temperature in a sensor network

Suppose we wish to compute the average temperature of a region throughout which sensor nodes have been placed. The following procedure is described and analyzed in Boyd et al. 2005:

At each timestep t , a node chosen uniformly at random contacts one of its neighbors with probability proportional to edge weight, and each node replaces its temperature value with the average of the two previous temperatures. Let P be the stochastic matrix of edge weights. Let $x(t)$ be the vector of temperature values at each node at timestep t .

The *averaging time* $T_{ave}(\epsilon, P)$ is defined as

$$T_{ave}(\epsilon, P) = \inf \left\{ t : \Pr \left(\frac{\|x(t) - x_{avg}\|}{\|x(0)\|} \geq \epsilon \right) \leq \epsilon \right\},$$

i.e. the number of timesteps before the total deviation of node values from the average is bounded by ϵ .

Denote

$$W = I - \frac{1}{2n}D + \frac{P + P^T}{2n},$$

where D is the diagonal matrix with entries $D_i = \sum_{j=1}^n [P_{ij} + P_{ji}]$. Then the following holds:

Theorem 11.2 (Boyd et al. 2005)

$$T_{ave}(\epsilon, P) = O \left(\frac{\log \epsilon^{-1}}{\log \lambda_2(W)^{-1}} \right)$$

When P is symmetric, this is closely related to the mixing time of the random walk defined by P .

See Boyd et al. 2005 for proof and more details.