Plurality Dynamics in Distributed Synchronous Systems

Emanuele Natale

PhD student, Sapienza Università di Roma

April 4, 2014

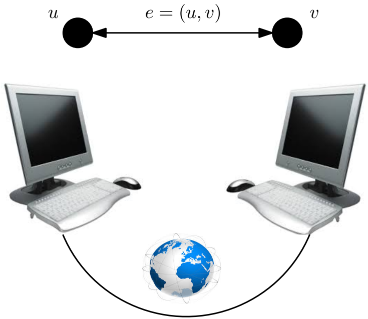

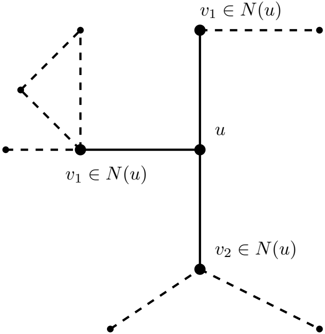

Given a graph G

- Nodes are *computational entities*
- Edges are *possible communications*
- The nodes' knowledge is only *local*

We are interested in **dynamics**:
protocols that apply simple rules using very little memory
$\implies$ Nodes have low computational power (small memory)
$\implies$ interactions are very limited

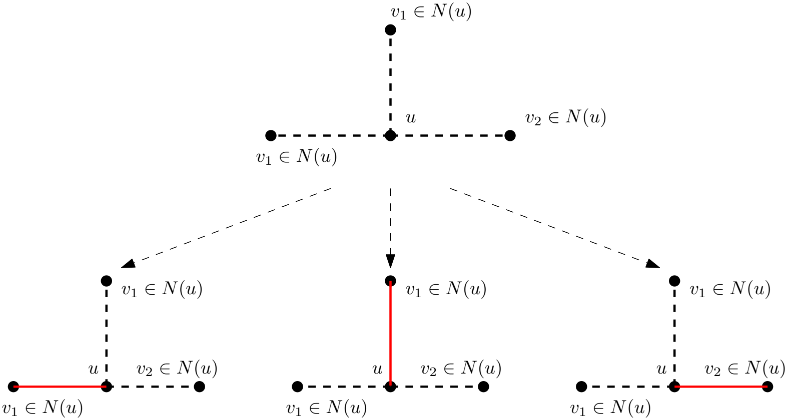

Gossip model: one communication at a time

With whom does a node communicate?

- deterministic choice $\implies$ deterministic protocol
- choose a neighbor *uniformly at random* $\implies$ randomized protocol

With whom does a node communicate?

- deterministic choice $\implies$ deterministic protocol
- choose a neighbor *uniformly at random* $\implies$ randomized protocol

Randomized protocols:

- Fault tolerant! (Probabilistic output)
- Motivated by the modeled phenomenon

$\implies$ **random-uniform Gossip model**

When does a node communicate?

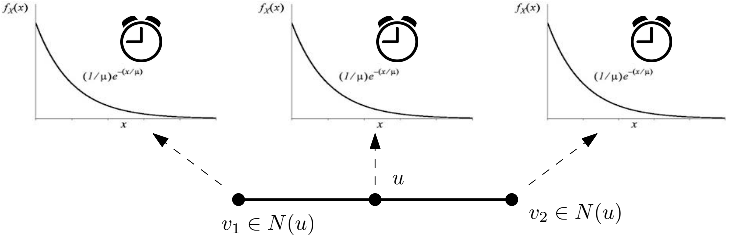

- Poisson clocks on nodes
- Poisson clocks on edges
- Population protocols

No simultaneous communications $\Longrightarrow$ asynchronous models

Simultaneous communications are typical: all nodes communicate in parallel
$\Longrightarrow$ *synchronous Gossip model*

But... synchronous vs asynchronous... why should I care?

Simultaneous communications are typical: all nodes communicate in parallel

$\Longrightarrow$ *synchronous Gossip model*

But. . . synchronous vs asynchronous. . . why should I care?
If we know each node's *state*. . . we have a **markovian process**!

$$\mathbf{P}\left(\underset{G_{t+1}}{\triangle} \;\middle|\; \underset{G_t}{\triangle} \,,\, \underset{G_{t-1}}{\triangle} \right) = \mathbf{P}\left(\underset{G_{t+1}}{\triangle} \;\middle|\; \underset{G_t}{\triangle} \right)$$

Suppose that each node's state is either 0 or 1.

- Asynchronous model $\Longrightarrow$ birth&death chain
- Synchronous model $\Longrightarrow$ . . .

The Majority Consensus Problem

A very natural problem: each node has some opinion (say some $i \in \{1, \dots, k\}$) and every node would know what is the most frequent one (**plurality**).

We would like to solve the problem with a dynamic such that

- uses little memory: if the size of the network double, it requires only an additional constant memory $\rightarrow O(\log n)$ bits
- is very fast: if the size of the network double, it requires only an additional constant time $\rightarrow O(\log n)$ steps
- is very "reliable": the probability of a correct output is $1 - \frac{1}{n^{const}}$, and we say that "the protocol converge *with high probability*" (w.h.p.)

The Voter Model

Suppose that the opinion are binary: each node has either 0 or 1.

The Voter Model. During each step, each node adopts the opinion of a neighbor picked at random.

The Voter Model

Suppose that the opinion are binary: each node has either 0 or 1.

The Voter Model. During each step, each node adopts the opinion of a neighbor picked at random.

Reversing time, each opinion make a random walk!

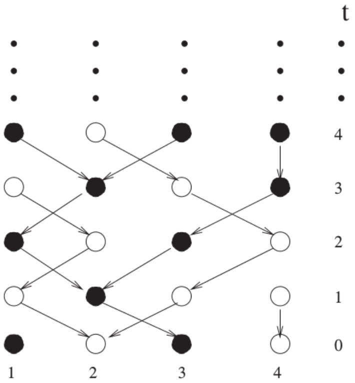

What is the probability of converging to the right plurality opinion? With a nice martingale argument...

Theorem

For a connected non-bipartite graph, the probability that the uniform process is absorbed into the all-0 states is

$$\sum_{\substack{u \text{ started} \\ \text{with op. } 0}} \frac{deg(u)}{2|E|}$$

Y. Hassin and D. Peleg. *Distributed probabilistic polling and applications to proportionate agreement.* Information and Computation 2001.

What is the convergence time?

Theorem

Let G be a connected graph with n nodes, m edges, average node degree d, and maximum degree $\Delta = O\left(m^{1-\epsilon}\right)$, for an arbitrary constant $\epsilon > 0$. Let $C(n)$ be the expected coalescence time for a system of n particles making a lazy random walks on G, where originally one particle starts at each node. Then

$$C(n) = \frac{n}{\nu\left(1-\lambda_2\right)}$$

where $\nu = \frac{\sum_u (deg(u))^2}{d^2 n}$. Thus by the equivalence between coalescence and voting, the expected time $\mathbb{E}\left[C_v\right]$ to complete voting on G has the same upper bound.

C. Cooper, R. Elsässer, H. Ono, and T. Radzik. *Coalescing random walks and voting on graphs*. PODC 2012

The median rule

The voter model is not efficient...
What if we look at more than one neighbor (e.g. by deciding every two or three steps...)?

The median rule

The voter model is not efficient...
What if we look at more than one neighbor (e.g. by deciding every two or three steps...)?

Let us consider the binary-opinion case on the complete graph. Suppose that each node look at two nodes choosen uniformly at random and adopts the more frequent opinion among theirs and its own.

Theorem

Provided that the initial difference between 0s and 1s is greater than $c\sqrt{n \log n}$ (where c is a sufficiently large constant), then with high probability within $= (\log n)$ time steps all nodes adopt the plurality opinion.

B. Doerr, L. A. Goldberg, L. Minder, T. Sauerwald, and C. Scheideler. *Stabilizing consensus with the power of two choices.* SPAA 2011.

The majority rule

The median rule is a good dynamic in the binary case on the complete graph, but with *many* initial opinions the median is not the plurality value...

What if we look at three random neighbors (on the complete graph) and adopt the *majority* opinion among their three opinions (breaking ties at random)?

The majority rule

The median rule is a good dynamic in the binary case on the complete graph, but with *many* initial opinions the median is not the plurality value...

What if we look at three random neighbors (on the complete graph) and adopt the *majority* opinion among their three opinions (breaking ties at random)?

Theorem

Let $3 \leq \lambda < \sqrt{n}$. Consider any initial assignment of k opinions such that the plurality opinion belongs to at least $\frac{n}{\lambda}$ nodes and surpasses any other one by at least $22\sqrt{\lambda n \log n}$ nodes, then w.h.p. the 3-majority protocol converges to the plurality in $O(\lambda \log n)$ time.

L. Becchetti, A. Clementi, E. Natale, F. Pasquale, R. Silvestri, and L. Trevisan. *Simple dynamics for majority consensus.* SPAA 2014

What if we look at h random neighbors (on the complete graph) and adopt the *majority* opinion among their h opinions (breaking ties at random)?

We can show the improvement is limited in the following way.

Theorem

If the plurality opinion is initially hold by no more than $\frac{3}{2} \cdot \frac{n}{k}$ where k is the number of initial opinions, then w.h.p. the h-majority takes at least $\Omega\left(\frac{k}{h^2}\right)$ time steps to converge.

The extra-state dynamic

Let us consider this dynamic:

- If a node "see" a node with a different opinion, becomes "undecided";
- If an undecided node see any opinion, adopts that one;
- In all other cases, it does nothing.

Theorem

In the asynchronous random-uniform Gossip model, on the complete graph and with binary opinion, the extra-state dynamic converge w.h.p. to the plurality in time $O(n \log n)$ *even when the number of 1s and 0s differ only by* $\omega\left(\sqrt{n \log n}\right)$.

D. Angluin et al. *A simple population protocol for fast robust approximate majority.* Distributed Computing 2008 A. Babaee et al. *Distributed multivalued consensus.*The Computer Journal 2013

E. Perron et al. *Using three states for binary consensus on complete graphs.* INFOCOM 2009

In the synchronous random-uniform Gossip model, on the complete graph we have proved the following.

Theorem

If the plurality opinion is greater than $\omega\left(\sqrt{n \log^3 n}\right)$, *the number of opinions does not exceed* $O\left(\sqrt[3]{\frac{n}{\log n}}\right)$, *and the plurality surpasses every other opinion by a factor* $(1+\alpha)$ *where* α *is an arbitrarily small positive constant, then with high probability the extra-state protocol converge to the plurality in time* $O\left(R \cdot \log n\right)$ *where*

$$R = \sum_{i \text{ opinion}} \frac{|\{\text{nodes who start with } i\}|^2}{|\{\text{nodes who start with the plurality}\}|^2}.$$

Quoting a reviewer of our work (SPAA 2014):
"*This paper analyzes an interesting open problem. These dynamics questions are often very difficult to handle, and we need more ideas in this field. [. . .] In fact, it is interesting because known techniques suffice. It turns out that there is enough structure in this problem, so that if you carefully exploit it, everything goes through. That in itself is slightly surprising.*"