Finite Markov Information-Exchange processes

David Aldous

February 14, 2011

The voter model and coalescing MCs.

The two models considered here use the "directed" convention: in the meeting model, when agents i, j meet, choose a random direction and indicate it using an arrow $i \to j$ or $j \to i$.

Voter model. Initially each agent has a different "opinion" – agent i has opinion i. When i and j meet at time t with direction $i \to j$, then agent j adopts the current opinion of agent i.

So we can study

$$\mathcal{V}_i(t) := \text{the set of } j \text{ who have opinion } i \text{ at time } t.$$

Note that $\mathcal{V}_i(t)$ may be empty, or may be non-empty but not contain i. The number of different remaining opinions can only decrease with time.

Minor comments. (i) We can rephrase the rule as "agent i imposes his opinion on agent j".
(ii) The name is very badly chosen – people do not **vote** by changing their minds in any simple random way.
(iii) In the classical, infinite lattice, setting one traditionally took only two different initial opinions.

So $\{\mathcal{V}_i(t), i \in \textbf{Agents}\}$ is a random partition of **Agents**. A natural quantity of interest is the **consensus time**

$$T^{\text{voter}} := \min\{t : \mathcal{V}_i(t) = \textbf{Agents} \text{ for some } i\}.$$

Coalescing MC model. Initially each agent has a token – agent i has token i. At time t each agent i has a (maybe empty) collection $\mathcal{C}_i(t)$ of tokens. When i and j meet at time t with direction $i \to j$, then agent i gives his tokens to agent j; that is,

$$\mathcal{C}_j(t+) = \mathcal{C}_j(t-) \cup \mathcal{C}_i(t-), \quad \mathcal{C}_i(t+) = \emptyset.$$

Now $\{\mathcal{C}_i(t), i \in \textbf{Agents}\}$ is a random partition of **Agents**. A natural quantity of interest is the **coalescence time**

$$T^{\text{coal}} := \min\{t : \mathcal{C}_i(t) = \textbf{Agents} \text{ for some } i\}.$$

Minor comments. Regarding each non-empty cluster as a particle, each particle moves as the MC at half-speed (rates $\nu_{ij}/2$), moving independently until two particles meet and thereby coalesce. Note this factor $1/2$ in this section.

The duality relationship.

For fixed t,

$$\{\mathcal{V}_i(t), i \in \textbf{Agents}\} \stackrel{d}{=} \{\mathcal{C}_i(t), i \in \textbf{Agents}\}.$$

In particular $T^{\text{voter}} \stackrel{d}{=} T^{\text{coal}}$.
They are different as processes. For fixed i, note that $|\mathcal{V}_i(t)|$ can only change by ± 1, but $|\mathcal{C}_i(t)|$ jumps to and from 0.

In figures, time "left-to-right" gives CMC,
time "right-to-left" with reversed arrows gives VM.

Note this depends on the symmetry assumption $\nu_{ij} = \nu_{ji}$ of the meeting process.

Project. Read the abstract discussion of *duality* in Liggett (IPS sec. 2.3); put the "key identity for averaging processes" in that framework.

Schematic – the meeting model on the 8-cycle.

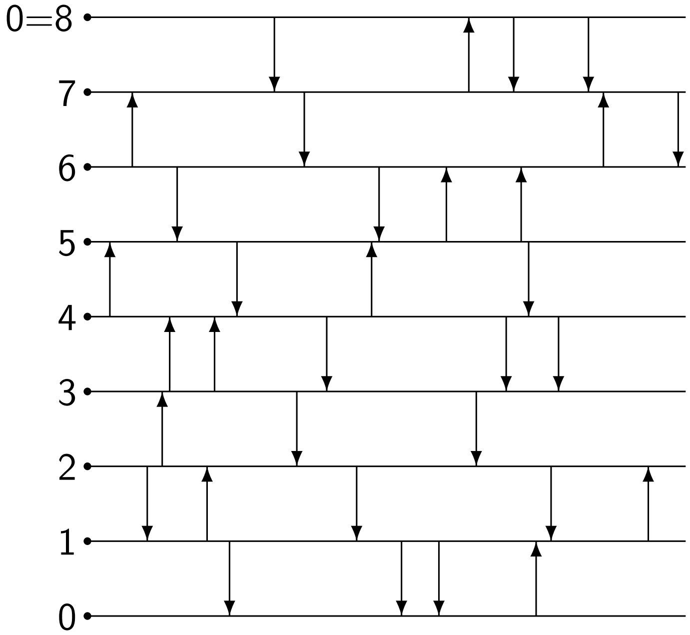

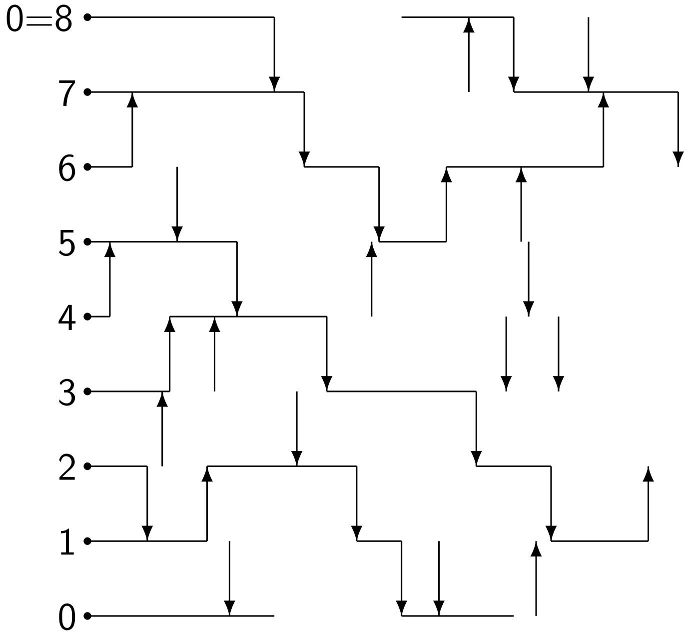

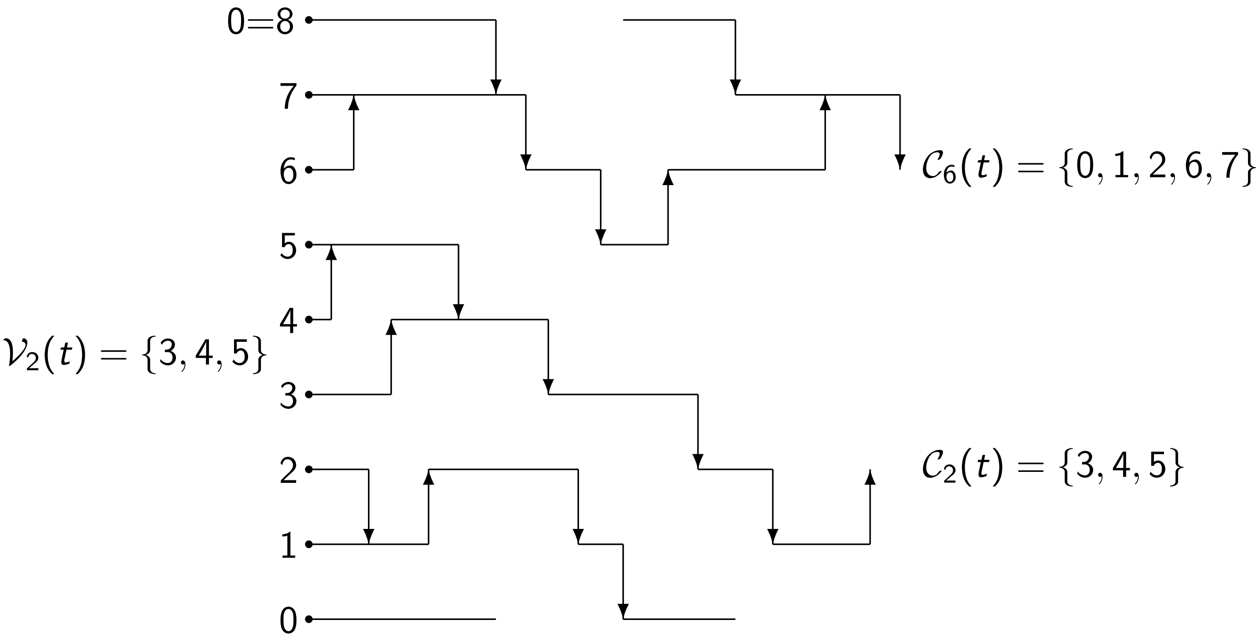

Voter model on the complete graph

There are two ways to analyze T_n^{voter} on the complete graph, both providing some bounds on other geometries.

Part of **Kingman's coalescent** is the continuous-time MC on states $\{1, 2, 3, \ldots\}$ with rates $\lambda_{k,k-1} = \binom{k}{2}, \ k \geq 2$. For that chain

$$\mathbb{E}_m T_1^{\text{hit}} = \sum_{k=2}^{m} 1/\binom{k}{2} = 2(1 - \tfrac{1}{m})$$

and in particular $\lim_{m\to\infty} \mathbb{E}_m T_1^{\text{hit}} = 2$.

In coalescing RW on the complete n-graph, the number of clusters evolves as the continuous-time MC on states $\{1, 2, 3, \ldots, n\}$ with rates $\lambda_{k,k-1} = \frac{1}{n-1}\binom{k}{2}$. So $\mathbb{E}\, T_n^{\text{coal}} = (n-1) \times 2(1 - \frac{1}{n})$ and in particular

$$\mathbb{E}\, T_n^{\text{voter}} = \mathbb{E}\, T_n^{\text{coal}} \sim 2n. \tag{1}$$

The second way is to consider the variant of the voter model with only 2 opinions, and to study the number $X(t)$ of agents with the first opinion. On the complete n-graph, $X(t)$ evolves as the continuous-time MC on states $\{0, 1, 2, \ldots, n\}$ with rates

$$\lambda_{k,k+1} = \lambda_{k,k-1} = \tfrac{k(n-k)}{2(n-1)}.$$

This process arises in classical applied probability (e.g. as the Moran model in population genetics). We want to study

$$T^{\text{hit}}_{0,n} := \min\{t : X(t) = 0 \text{ or } n\}.$$

By general birth-and-death formulas, or by comparison [board] with simple RW,

$$\mathbb{E}_k T^{\text{hit}}_{0,n} = \tfrac{2(n-1)}{n} \ \left(k(h_{n-1} - h_{k+1}) + (n-k)(h_{n-1} - h_{n-k+1})\right)$$

where $h_m := \sum_{i=1}^m 1/i$. This is maximized by $k = \lfloor n/2 \rfloor$, and

$$\max_k \mathbb{E}_k T^{\text{hit}}_{0,n} \sim (2 \log 2) \ n.$$

Now we can couple the true voter model (n different initial opinions) with the variant with only 2 opinions, initially held by k and $n-k$ agents. (Just randomly assign these two opinions, initially). From this coupling we see

$$\mathbb{P}_k(T^{\text{hit}}_{0,n} > t) \le \mathbb{P}(T^{\text{voter}}_n > t)$$

$$\mathbb{P}_k(T^{\text{hit}}_{0,n} > t) \ge \tfrac{2k(n-k-1)}{n(n-1)} \mathbb{P}(T^{\text{voter}}_n > t)$$

In particular, the latter with $k = \lfloor n/2 \rfloor$ implies

$$\mathbb{E} T^{\text{voter}}_n \le (4 \log 2 + o(1)) \ n.$$

This is weaker than the correct asymptotics (1).

Voter model on general geometry

Suppose the flow rates satisfy, for some constant κ,

$$\nu(A, A^c) := \sum_{i\in A, j\in A^c} n^{-1}\nu_{ij} \ge \kappa |A|(n-|A|)/(n-1).$$

On the complete graph this holds with $\kappa = 1$. We can repeat the analysis above – the process $X(t)$ now moves at least κ times as fast as on the complete graph, and so

$$\mathbb{E}\, T_n^{\text{voter}} \le (4\log 2 + o(1))\ n/\kappa.$$

The optimal κ is (up to a factor $(n-1)/n$) just $1/\tau_{\text{cond}}$ for the Cheeger time constant τ_{cond}, and so

$$\mathbb{E}\, T^{\text{voter}} \le (4\log 2 + o(1))\ n\ \tau_{\text{cond}}.$$

Coalescing MC on general geometry

Issues clearly related to study of the *meeting time* T^{meet} of two independent copies of the MC, a topic that arises in other contexts. Under enough symmetry (e.g. continuous-time RW on the discrete torus) the relative displacement between the two copies evolves as the same RW run at twice the speed, and study of T^{meet} reduces to study of T^{hit}.

First consider the completely general case. In terms of the associated MC define a parameter

$$\tau^* := \max_{i,j} \mathbb{E}_i T_j^{\text{hit}}.$$

The following result was conjectured long ago but only recently proved. Note that on the complete graph the mean coalescence time is asymptotically $2\times$ the mean meeting time.

Theorem (Oliveira 2010)

There exist numerical constants $C_1, C_2 < \infty$ such that, for any finite irreducible reversible MC, $\max_{i,j} \mathbb{E}_{i,j} T^{meet} \le C_1\tau^*$ and $\mathbb{E}\, T^{coal} \le C_2\tau^*$.

Proof techniques seem special, but perhaps a good "paper-talk".

To seek "$1 \pm o(1)$" limits, let us work in the meeting model setting (stationary distribution is uniform) and write τ_{meet} for mean meeting time from independent uniform starts. In a sequence of chains with $n \to \infty$, impose a condition such as the following. For each $\varepsilon > 0$

$$n^{-2}|\{(i,j) : \mathbb{E}_i T_j^{\text{hit}} \notin (1 \pm \varepsilon)\tau_{\text{meet}}\}| \to 0. \tag{2}$$

By analogy with the Kingman coalescent argument one expects some general result like

Open problem. Assuming (2), under what further conditions can we prove $\mathbb{E}\, T^{\text{coal}} \sim 2\tau_{\text{meet}}$?

This project splits into two parts.

Part 1. For fixed m, show that the mean time for m initially independent uniform walkers to coalesce should be $\sim 2(1 - \frac{1}{m})\tau_{\text{meet}}$.

Part 2. Show that for $m(n) \to \infty$ slowly, the time for the initial n walkers to coalesce into $m(n)$ clusters is $o(\tau_{\text{meet}})$.

Part 1 is essentially a consequence of known results, as follows.

From old results on mixing times (RWG section 4.3), a condition like (2) is enough to show that $\tau_{\text{mix}} = o(\tau_{\text{meet}})$. So – as a prototype use of τ_{mix} – by considering time intervals of length τ, for $\tau_{\text{mix}} \ll \tau \ll \tau_{\text{meet}}$, the events "a particular pair of walker meets in the next τ-interval" are approximately independent. This makes the "number of clusters" process behave as the Kingman coalescent.

Note. That is the hack proof. Alternatively, the explicit bound involving τ_{rel} on exponential approximation for hitting time distributions from stationarity is applicable to the meeting time of two walkers, so a more elegant way would be to find an extension of that result applicable to the case involving m walkers.

Part 2 maybe needs some different idea/assumptions.

(restate) Open problem. Assuming (2), under what further conditions can we prove $\mathbb{E} T^{\text{coal}} \sim 2\tau_{\text{meet}}$?

What is known rigorously?

Cox (1989) proves this for the torus $[0, m-1]^d$ in dimension $d \geq 2$. Here $\tau_{\text{meet}} = \tau_{\text{hit}} \sim m^d R_d$ for $d \geq 3$.

Cooper-Frieze-Radzik (2009) prove Part 1 for the random r-regular graph, where $\tau_{\text{meet}} \sim \tau_{\text{hit}} \sim \frac{r-1}{r-2} n$.
(the latter, containing other results, could be a "paper project").

Various variant models are easy to do heuristically – see e.g. Sood-Antal-Radner (2008).

In the 2-opinion case, the process $X(t)$ = number of agents with opinion 1 is a martingale. So starting with k opinion-1 agents, the chance of being absorbed in the all-1 configuration equals k/n.

One can study *biased* voter models where a agent is more likely to copy an opinion-1 neighbor. In this case the submartingale property will imply that the chance above is $> k/n$. A more challenging situation arises in the following game-theory variant, studied in Manshadi - Saberi (2011).

Symmetric prisoner's dilemma. Each agent in state C or state D. When an agent i plays an agent j
if i is C then i incurs cost $c > 0$ and j gains benefit $b > c$.
if i is D then i incurs cost 0 and j gains benefit 0.
Consider a k-regular n-vertex connected graph on agents. Take discrete time. At each time step, each vertex plays each neighbor. Represent states by

$$X_t^i = 1(\text{ agent } i \text{ in state C}).$$

So the payoff to i at time t equals

$$u_i^t = -kcX_t^i + b\sum_{j \sim i} X_t^j.$$

Agents change state as follows. Fix small $\varepsilon > 0$. At each time pick a uniform random agent. Other agents do not change state. Given we picked agent i at time t, set $X^i_{t+1} = X^J_t$, where J is a random neighbor of i chosen according to

$$\mathbb{P}(J = j) = (1 - \varepsilon)\tfrac{1}{k} + \varepsilon\theta_{i,t}(j)$$

where $\theta_{i,t}$ is the measure

$$\theta_{i,t}(j) = \tfrac{1}{k}(u^t_j + 1 - \tfrac{1}{k}\sum_{h\sim i} u^t_h)$$

which is a probability measure when we impose the condition

$$k(b + c) < 1.$$

When $\varepsilon = 0$ this is just the voter model. For $\varepsilon > 0$ we are biasing toward copying the state of a currently successful neighbor.

Theorem (Manshadi - Saberi, 2011)

Consider a connected k-regular graph with girth at least 7. Initially let a random pair of neighbors have state C and the others state D; the system then evolves according to the model above. Suppose $b/c > k^2/(k-1)$. Fix $\gamma > 0$ and set $\varepsilon = n^{-(4+\gamma)}$ and suppose n is sufficiently large. Then the probability of absorption into "all C" is $\geq \frac{2}{n} + \frac{\varepsilon}{n} f(b/c)$ for a certain strictly positive function f.