

Finite Markov Information-Exchange processes.

1. Overview

David Aldous

June 18, 2012






Analogy: **game theory** denotes a natural and useful topic, but is a terrible choice of name.

I want to delineate a topic, on which thousands of papers have been written, under unhelpful names like “interacting particle systems” or “opinion dynamics”. I’m calling this topic “FMIE processes” – offer small prize for better name!

I am **not** presenting any essentially new theorems or techniques. Instead (analogy: new version of an OS) express particular underlying viewpoint, some “nuances and themes” which suggest new questions.

What (mathematically) is a social network?

Usually formalized as a *graph*, whose vertices are individual people and where an edge indicates presence of a specified kind of relationship.



Take a finite graph with edge-weights $\nu_{ij} = \nu_{ji} > 0$.

Details: proper, undirected, connected, set $\nu_{ij} = 0$ if (i, j) not an edge. So $\mathcal{N} := (\nu_{ij})$ is a symmetric non-negative matrix.

Commentary: Vertex = individual person = agent; ν_{ij} = “strength of relationship” between person i and person j . The usual math formalization of “social network” is as the unweighted version, but in most contexts relationships are not 0/1.

We diverge from mainstream social network literature by interpreting “strength of relationship” in a specific way, as “frequency of meeting”. Being a probabilist, I mean

- Each pair i, j of agents with $\nu_{ij} > 0$ meets at the times of a rate- ν_{ij} Poisson process.

Details: independent for each pair.

Call this the **meeting model** for the given weighted graph. Not very exciting, so far

What is a FMIE process?

1. A given set of “states” for an agent (usually finite; sometimes \mathbb{Z} or \mathbb{R}).
2. Take a **meeting model** as above, specified by the weighted graph/symmetric matrix $\mathcal{N} = (\nu_{ij})$.
3. Each agent i is in some state $X_i(t)$ at time t . When two agents i, j meet at time t , they update their information according to some rule (deterministic or random). That is, the updated information $X_i(t+), X_j(t+)$ depends only on the pre-meeting information $X_i(t-), X_j(t-)$ and (perhaps) added randomness.

We distinguish between the “geometric substructure” of the meeting process and the “informational superstructure” of the update rule. Different FMIE **models** correspond to different update rules for the informational superstructure.

Write $n =$ number of agents. If there are $k < \infty$ states for an agent, then the number of **configurations** of the process is $k^n < \infty$ and the FMIE process is some finite-state (continuous-time) Markov chain. So (undergrad MC theory) qualitative time-asymptotics are determined by the strongly connected components of the transition graph of the whole process, the extreme possibilities being

- Convergence to a unique stationary distribution
- Absorption in some (random) absorbing configuration.

In most models it's easy to see which holds

. so time-asymptotics are not the issue.

FMIE = IPS? Yes, but

What probabilists call “interacting particle systems”, following Liggett’s 1985 book, are exemplified by the voter model, contact process, exclusion process, and Glauber dynamics for the Ising model, all on the infinite d -dimensional lattice, and emphasizing ideas such as phase transitions.

Our viewpoint is by analogy with the “Markov chains and mixing times” literature (e.g. Levin-Peres-Wilmer and Aldous-Fill); study how the behavior of a model depends on the underlying geometry, where “behavior” means quantitative aspects of finite-time behavior. Roughly speaking, there was sporadic work on the finite-agent setting before 2000, and a huge literature since 2000, mostly outside mathematical probability journals, emphasizing calculations/simulations for 150 variant models over the same 4 geometries:

- complete graph (= mean-field)
- d -dimensional torus
- small worlds (= torus + random long edges)
- random graph with specified degree distribution (e.g. configuration model)

(Show list of papers from 2011 course).

Here's an example from outside math probability.

<http://scnarc.rpi.edu/content/minority-rules-scientists-discover-tipping-point-spread-ideas>

One conceptual goal is to make a list of “representative” models (overlapping Liggett’s) from the “agents and information” story. My current thoughts . . .

- two “level 0” models (next slides)
- about ten “level 1” models (e.g voter model; can debate which others)

Representative in the sense that other models are recognizably variants/combinations of above.

Recall last web page: this is “just math” – NiP models not SiP models.

Background meeting model with rates ν_{ij} .

Model: Hot Potato.

There is one token. When the agent i holding the token meets another agent j , the token is passed to j .

The natural aspect to study is $Z(t)$ = the agent holding the token at time t . This $Z(t)$ is the (continuous-time) Markov chain with transition rates (ν_{ij}) .

So we have available

- general theory of finite-state MCs
- calculations in specific geometries
- In lecture 2 we'll see other FMIE models having close connection with this MC, and for which parts of "standard modern theory" of finite MCs can be used.

Digress to mention a **nuance**, as follows.

Analogy: A Google Scholar (advanced) search on
exact phrase: Galton-Watson
year 1965-1969

gets you a bunch of papers studying $Z_t =$ population size in generation t .
For years 2005-2009, half the papers talk about the Galton-Watson **tree**.

Relationship above analogous to relation between the MC and Hot
Potato; for the latter one can ask questions like

What is the expected time until agent i meets someone who has
previously held the token?

Background meeting model with rates ν_{ij} .

Model: Pandemic

Initially one agent is infected. Whenever an infected agent meets another agent, the other agent becomes infected.

- In other words, the SI epidemic with exponentially distributed infection times. Or first-passage percolation (with exponential times).
- This model is “basic” in the specific sense of fastest possible spread of information in any FMIE model.
- Like MCs, many papers doing calculations in specific geometries; see some in Lecture 3
- Unlike MCs, no “general theory”. See a conjecture in Lecture 3.
- Also see some other FMIEs (Fashionista, Precedence) built over Pandemic.

Some details/conventions

The associated MC always has uniform stationary distribution, as a consequence of $\nu_{ij} = \nu_{ji}$. We often have $\sum_{j \neq i} \nu_{ij}$ (the rate at which agent i meets other agents; loosely, i 's number of friends) is constant in i , in which case w.l.o.g.

$$\sum_{j \neq i} \nu_{ij} = 1 \quad \forall i \quad (\text{standardized rates})$$

but this is a separate regularity condition.

This condition is not realistic as “social network” – different people have different number of friends – but general inequalities under this condition will extend to cases with a bound on $\frac{\max_i \sum_j \nu_{ij}}{\min_i \sum_j \nu_{ij}}$.

Most related literature deals with unweighted graphs. In particular, a r -regular graph fits our setting via the default convention of setting rates to be $\nu_{ij} = 1/r$ for each edge.

There are several **nuances** arising from the underlying meeting process being based on a weighted graph rather than an unweighted graph. What are analogs of (in CS-algorithms settings)

1. worst-case over connected n -vertex graphs?
2. Distributed algorithms, which envisages a processor at each vertex which knows identity of neighbors. For a meeting model analog what do we want to assume each agent i knows? All values ν_{ij} for all other agents j , or just observed data on meetings over $[0, t_0]$?

Interesting issues for future work

Digress to mention technical result where “weighted graphs” is (mathematically) right setting.

In the exclusion (**interchange**) process there are n distinguishable tokens, one at each agent. When two agents meet they exchange tokens. The motion of an individual token is as the associated MC, so has some spectral gap $\lambda_{RW}(G) > 0$. The whole interchange process is itself a reversible MC, so has a spectral gap which satisfies (easy)

$\lambda_{IP}(G) \leq \lambda_{RW}(G)$. Longstanding conjecture “always equal”, proved by Caputo - Liggett - Rochthammer (2009).

FMIE encompasses many meeting models and many update rules, so we can't expect any remarkable general results. But there are 6 “general principles” – all kinda obvious once you say them – which seem worth saying.

One we've seen already:

(repeat of previous slide)

General principle 1

Write $n =$ number of agents. If there are $k < \infty$ states for an agent, then the number of **configurations** of the process is $k^n < \infty$ and the FMIE process is some finite-state (continuous-time) Markov chain. So (undergrad MC theory) qualitative time-asymptotics are determined by the strongly connected components of the transition graph of the whole process, the extreme possibilities being

- Convergence to a unique stationary distribution
- Absorption in some (random) absorbing configuration.

In most models it's easy to see which holds

. so time-asymptotics are not the issue.

General principle 2

The FMIE setup encourages you to think about coupling different processes.

Discussion. The modern prevalence of coupling techniques in studying Markovian processes actually dates from the 1970s introduction of IPS. Typically we think of “clever” couplings of two copies of the same process from different starts.

Of course it’s “obvious” that we can and do also couple different processes. But consider our two basic models – MC and epidemics. Does any paper actually consider coupling them?

In Hot Potato can define

$S_1(t)$ = agents who have held token before time t

$S_2(t)$ = agents who, before time t , met agent who had previously held token

$S_k(t)$ = “ k ’th hand process” = agents who, at some $T < t$, meet agent in $S_{k-1}(T)$

Now $\cup_k S_k(t)$ is the “infected” set in Pandemic.

General principle 3

For many simple models there is some specific aspect which is invariant in the sense of depending only on n , not on the geometry.

Here we often assume

$$\sum_{j \neq i} \nu_{ij} = 1 \quad \forall i \quad (\text{standardized rates})$$

For the MC we have

$$\mathbb{E}_i(\text{return time to } i) = n$$

Curiously, Pandemic seems an exception to this principle

[Other 3 general principles in lecture 2 – or think for yourself . . .]

Digression: what do graph theorists know about weighted graphs that might be useful?

Most familiar to some of us is

- Electric networks and reversible MCs and spanning trees
- but there are classical topics such as
- max-flow = min-cut.

What else?

Just for fun: An inequality says (unweighted case) “your friends have more friends than you do (on average)”. The “weighted” version says “your friends party more often than you do (on average)”.

Write $f(i) := \sum_j \nu_{ij}$

Pick (I, J) via:

$$I \text{ uniform on agents, } \mathbb{P}(J = j | I = i) = \nu_{ij}/f(i).$$

Then $\mathbb{E}f(J) \geq \mathbb{E}f(I)$.

[graph-structured size-biasing]

Digression: finite weighted graphs \leftrightarrow finite metric spaces.

A weighted graph is specified by a matrix ν_{ij}

A metric space is specified by a matrix $d(i, j)$.

There are many general ways to use one structure to define the other structure, with the intuition: large ν_{ij} implies small $d(i, j)$. Some are related to our FMIE world. I'll give a few examples – is there a survey?

Given a metric $d(i, j)$:

1. (boring!) Define $\nu_{ij} := f(d(i, j))$ for decreasing f . Get weighted version of complete graph.
2. First define some unweighted graph, then define weights as **traffic flow**:

$\nu_{ij} = \mathbb{P}(\text{edge } (i, j) \text{ in shortest route between uniform random start and end})$

where edge-lengths given by the metric.

Two ways to get the unweighted graph.

(a) Proximity graphs. Simplest version: create edge (i, j) if $\nexists k$ s.t.

$\max(d(j, k), d(k, j)) < d(i, j)$.

(b) First create random graph $\mathbb{P}((i, j) \text{ is edge}) = f(d(i, j))$ for decreasing f . Then use traffic flow, then take \mathbb{E} .

3. Exercise.

Given a weighted graph ν_{ij} :

1. (boring!) Assign length $1/\nu_{ij}$ to edge (i, j) , then define $d(k, \ell) =$ length of shortest route from k to ℓ . (Graph distance).
2. (a) (well-known “interesting”). Regard as electrical network with $\nu_{ij} = 1/(\text{resistance of edge } (i, j))$, define

$$d(i, j) = \text{effective resistance between } i \text{ and } j.$$

With standardized rates, this is equivalent (up to scaling) to mean commute time between i and j for the associated MC.

(b) But many other ways to use the associated MC, for instance as integrated variation distance

$$d(i, j) := \int_0^\infty \|P_i(X_t \in \cdot) - P_j(X_t \in \cdot)\|_{VD} dt$$

or its L^2 analog.

3. Use the Pandemic (FPP) times T_{ij}^{epi} to define $d(i, j) := \mathbb{E} T_{ij}^{\text{epi}}$.
4. or use other FMIE processes.

Reminder of some topics (on board)

- VD mixing time and cut-off
- Relaxation time = $1/(\text{spectral gap})$
- Configuration model for graph with prescribed degree distribution.