Weierstrass Institute for
Applied Analysis and Stochastics

Leibniz
Association

A large-deviations principle for all the components in a sparse inhomogeneous Erdős-Rényi graph

Wolfgang König
TU Berlin and WIAS

based on joint work (in progress) with Luisa Andreis (Florence), Heide Langhammer and Robert Patterson (WIAS)

The MARCUS-LUSHNIKOV model is a non-spatial (i.e., a mean-field) coagulation model [MARCUS 1968], [GILLESPIE 1972], [LUSHNIKOV 1978]:

Continuous-time Markov process of vectors of particle masses at time $t \in [0, \infty)$:

$$M_1^{(N)}(t) \geq M_2^{(N)}(t) \geq M_3^{(N)}(t) \geq \cdots \geq M_{n(t)}^{(N)}(t) \geq 1, \qquad \sum_{i=1}^{n(t)} M_i^{(N)}(t) = N.$$

We start with $M_i^{(N)}(0) = 1$ for any i.

Coagulation mechanism:

Particles with masses m and $\widetilde{m}$ coagulate after an exponential random time with parameter $K_N(m, \widetilde{m})$ (the coagulation kernel) independently of all the other particles.

The MARCUS-LUSHNIKOV model is a non-spatial (i.e., a mean-field) coagulation model [MARCUS 1968], [GILLESPIE 1972], [LUSHNIKOV 1978]:
Continuous-time Markov process of vectors of particle masses at time $t \in [0, \infty)$:

$$M_1^{(N)}(t) \geq M_2^{(N)}(t) \geq M_3^{(N)}(t) \geq \cdots \geq M_{n(t)}^{(N)}(t) \geq 1, \qquad \sum_{i=1}^{n(t)} M_i^{(N)}(t) = N.$$

We start with $M_i^{(N)}(0) = 1$ for any i.

Coagulation mechanism:

Particles with masses m and $\widetilde{m}$ coagulate after an exponential random time with parameter $K_N(m, \widetilde{m})$ (the coagulation kernel) independently of all the other particles.

> Our main question:
>
> Is there a gelation phase transition at some fixed time $t_c \in (0, \infty)$ in the limit $N \to \infty$?

That is, is there a (deterministic) time after which a gel emerges, i.e., a particle with size $M_1^{(N)}(t) \asymp N$?

Here, we make the special choice of the multiplicative kernel:

$$K_N(m, \widetilde{m}) = \frac{m\,\widetilde{m}}{N}.$$

Advantage:

The model is now a function of a time-dependent version of the well-known ERDŐS-RÉNYI random graph model. Indeed, the vector $(M_i^{(N)}(t))_{i=1}^{n(t)}$ is in distribution equal to the collection of sizes of all the connected components of the graph $\mathcal{G}(N, 1 - \mathrm{e}^{-t/N})$.

Explanation:

Equip each unordered pair $\{i, j\}$ of different numbers in $\{1, \dots, N\}$ independently with an exponentially distributed random time $\mathrm{e}_{i,j}$ with expected value N. After the elapsure of $\mathrm{e}_{i,j}$, there is a bond created between i and j. Then, at time t, for each pair, the probability to have a bond between them is equal to $1 - \mathrm{e}^{-t/N}$.

Imagine the component containing i (with size m) and the one containing j (with size $\widetilde{m}$) have been turned into a new component (with size $m + \widetilde{m}$). Then the new component inherits all the active $m\widetilde{m}$ exponential random times of the two earlier components.

Our purpose

From now, we stick to the sparse Erdős–Rényi graph on $[N] = \{1, \dots, N\}$.

Goal 1: Explicit joint large-deviation principle for the statistics of all the component sizes k, distinguished into microscopic ($k \asymp 1$), mesoscopic ($1 \ll k \ll N$) and the macroscopic ($k \asymp N$) sizes. Explicit identification of the gelation phase transition as a consequence.

Goal 2: The same for a "spatial" version, the *inhomogeneous Erdős-Rényi graph*.

Earlier works on LDPs for sparse random graphs:

- [O'CONNELL 1998]: LDP for size of larges component and number of isolated points
- [ENGEL, MONASSON, HARTMANN 2004]: LDP for free energy of a tilted version with weights on the number of components, connections with Potts model.
- [BORDENAVE, CAPUTO 2015]: LDP for the microscopic connected subgraphs
- [PUHALSKII 2005]: LDP for the number of components, number of macroscopic components, number of excess edges in them. (Proof ansatz and rate function very different from ours).

Fix $t > 0$ and consider the standard Erdős–Rényi graph $\mathcal{G}(N, \frac{t}{N})$ with components of sizes $S_1^{(N)} \geq S_2^{(N)} \geq \cdots \geq S_n^{(N)} \geq 1$.

Microscopic and macroscopic empirical measures of the particle sizes:

$$\mathrm{Mi}^{(N)} = \frac{1}{N} \sum_{i=1}^{n} \delta_{S_i^{(N)}} \qquad \text{and} \qquad \mathrm{Ma}^{(N)} = \sum_{i=1}^{n} \delta_{\frac{1}{N} S_i^{(N)}}.$$

Fix $t > 0$ and consider the standard Erdős–Rényi graph $\mathcal{G}(N, \frac{t}{N})$ with components of sizes $S_1^{(N)} \geq S_2^{(N)} \geq \cdots \geq S_n^{(N)} \geq 1$.

Microscopic and macroscopic empirical measures of the particle sizes:

$$\mathrm{Mi}^{(N)} = \frac{1}{N}\sum_{i=1}^{n}\delta_{S_i^{(N)}} \qquad \text{and} \qquad \mathrm{Ma}^{(N)} = \sum_{i=1}^{n}\delta_{\frac{1}{N}S_i^{(N)}}.$$

Then $\mathrm{Mi}^{(N)}$ is a random member of the set $\mathcal{N} = \mathcal{N}(1)$, where

$$\mathcal{N}(c) = \Big\{\lambda \in [0,\infty)^{\mathbb{N}} \colon \sum_{k\in\mathbb{N}} k\lambda_k = c\Big\} \qquad \text{(coordinatewise top.).}$$

$\mathrm{Ma}^{(N)}$ is a random element of $\mathcal{M}_{\mathbb{N}_0} = \mathcal{M}_{\mathbb{N}_0}(1)$, where

$$\mathcal{M}_{\mathbb{N}_0}(c) = \Big\{\alpha \in \mathcal{M}_{\mathbb{N}_0}((0,1]) \colon \int_{(0,1]} x\,\alpha(\mathrm{d}x) = c\Big\} \qquad \text{(vague top.).}$$

and $\mathcal{M}_{\mathbb{N}_0}((0,1])$ is the set of all measures on $(0,1]$ with values in $\mathbb{N}_0$.

Note that the total masses

$$c_\lambda = \sum_{k\in\mathbb{N}} k\lambda_k \qquad \text{and} \qquad c_\alpha = \int_{(0,1]} x\,\alpha(\mathrm{d}x)$$

are discontinuous functions of λ resp. α.

LDP for the micro- and macroscopic parts

As $N \to \infty$, the pair $(\mathrm{Mi}^{(N)}, \mathrm{Ma}^{(N)})$ satisfies an LDP with rate function

$$I(\lambda, \alpha; t) = \begin{cases} I_{\mathrm{Mi}}(\lambda; t) + I_{\mathrm{Ma}}(\alpha; t) + (1 - c_\lambda - c_\alpha)\Big(\frac{t}{2} - \log t\Big), & \text{if } c_\lambda + c_\alpha \leq 1, \\ \infty & \text{otherwise,} \end{cases}$$

where

$$\begin{aligned} I_{\mathrm{Mi}}(\lambda; t) &= \sum_{k=1}^{\infty} \lambda_k \log \frac{k! t \lambda_k}{\mathrm{e}\, k^{k-2}} + c_\lambda \Big(1 + \frac{t}{2} - \log t\Big), \\ I_{\mathrm{Ma}}(\alpha; t) &= \int_0^1 \Big[x \log \frac{x}{1 - \mathrm{e}^{-tx}} + \frac{t}{2} x(1-x)\Big]\, \alpha(\mathrm{d}x). \end{aligned}$$

Microscopic total mass phase transition

1. For $t \in (0,1)$, the minimum of micro-part of the rate function is attained precisely at

$$\lambda_k^*(c;t) = \frac{k^{k-2} c^k t^{k-1} \mathrm{e}^{-ctk}}{k!}, \qquad k \in \mathbb{N},$$

and the minimum of the micro total mass is attained precisely at $c = 1$. Therefore the infimum of the joint rate function $I(\cdot,\cdot;t)$ is attained at $(\lambda,\alpha) = (\lambda_k^*(1;t), \mathbf{0})$.

Microscopic total mass phase transition

1. For $t \in (0,1)$, the minimum of micro-part of the rate function is attained precisely at

$$\lambda_k^*(c;t) = \frac{k^{k-2} c^k t^{k-1} \mathrm{e}^{-ctk}}{k!}, \qquad k \in \mathbb{N},$$

and the minimum of the micro total mass is attained precisely at $c = 1$. Therefore the infimum of the joint rate function $I(\cdot,\cdot;t)$ is attained at $(\lambda, \alpha) = (\lambda_k^*(1;t), \mathbf{0})$.

2. For $t \in (1,\infty)$, the minimum of the micro total mass rate function is attained precisely at $c = \beta_t$ for some $\beta_t \in (0, \frac{1}{t})$, given as the smallest positive solution to $\log \beta_t = t\beta_t - t$. The infimum is attained precisely at $(\lambda, \alpha) = (\lambda^*(\beta_t;t), (1-\beta_t, 0, \dots))$.

Microscopic total mass phase transition

1. For $t \in (0,1)$, the minimum of micro-part of the rate function is attained precisely at

$$\lambda_k^*(c;t) = \frac{k^{k-2} c^k t^{k-1} \mathrm{e}^{-ctk}}{k!}, \qquad k \in \mathbb{N},$$

   and the minimum of the micro total mass is attained precisely at $c = 1$. Therefore the infimum of the joint rate function $I(\cdot,\cdot;t)$ is attained at $(\lambda, \alpha) = (\lambda_k^*(1;t), \mathbf{0})$.

2. For $t \in (1,\infty)$, the minimum of the micro total mass rate function is attained precisely at $c = \beta_t$ for some $\beta_t \in (0, \frac{1}{t})$, given as the smallest positive solution to $\log \beta_t = t\beta_t - t$. The infimum is attained precisely at $(\lambda, \alpha) = (\lambda^*(\beta_t;t), (1-\beta_t, 0, \dots))$.

Hence, $t_c = 1$ is the gelation transition time. On a linear level, we can say:

- Before time 1, all particles are finitely large, and the statistics of their sizes follow the Borel distribution.
- After time 1, there is precisely one macroscopic particle of size $\sim (1-\beta_t)N$, and a Borel-distributed statistics of remaining particle sizes.

LDP for mesoscopic total mass

Fix $t \in [0, \infty)$ and $\varepsilon > 0$ and $R \in \mathbb{N}$. Then the mesoscopic (ε, R)-total mass,

$$\overline{\mathrm{Me}}_{R,\varepsilon}^{(N)} = \frac{1}{N} \sum_{i:\ R < M_i^{(N)}(t) < \varepsilon N} S_i^{(N)}.$$

satisfies an LDP with some rate function $\mathcal{J}_{\mathrm{Me}}^{(\varepsilon,R)}$ whose limit for $\varepsilon \downarrow 0$ and $R \to \infty$ is equal to

$$\mathcal{J}_{\mathrm{Me}}(c) = (1-c)\Big(\log(1-c)t - \frac{(1-c)t}{2}\Big) + \frac{t}{2} - \log t.$$

- $\mathcal{J}_{\mathrm{Me}}$ is strictly increasing and has a unique zero at $c = 0$.
- We also proved that $\overline{\mathrm{Me}}_{R_N,\varepsilon_N}^{(N)}$ satisfies an LDP with rate function $\mathcal{J}_{\mathrm{Me}}$ if $1 \ll R_N \leq N\varepsilon_N \ll N$.

Let $\mathbb{P}_{k,p}$ be the probability measure for $\mathcal{G} \sim \mathcal{G}(k,p)$. Put

$$\mu_k(p) = \mathbb{P}_{k,p}\big(\mathcal{G} \text{ is connected}\big),$$

then we have

Distribution of statistics

For any N and any $\ell = (\ell_k)_k \in \mathbb{N}_0^{\mathbb{N}}$ satisfying $\sum_k k\ell_k = N$, write

$$A_N(\ell) = \bigcap_{k\in\mathbb{N}} \{\#\{i \colon S_i^{(N)} = k\} = \ell_k\},$$

then

$$\mathbb{P}_{N,p}(A_N(\ell)) = N! \prod_k \frac{\mu_k(p)^{\ell_k}(1-p)^{\frac{1}{2}k(N-k)\ell_k}}{k!^{\ell_k}\,\ell_k!}.$$

Proof: elementary combinatorics.

Micro and macro asymptotics [STEPANOV 1970]

- $$(1-p)^{\frac{1}{2}(k-1)(k-2)} \le \frac{\mu_k(p)}{k^{k-2}p^{k-1}} \le 1, \qquad k \in \mathbb{N}.$$

  In particular, if $k = o(\sqrt{N})$, then

  $$\mu_k(\tfrac{t}{N}) = k^{k-2}(\tfrac{t}{N})^{k-1}, \qquad N \to \infty.$$

- $$\mu_{\lfloor \alpha N}(\tfrac{t}{N}) \sim \Big(1 - \frac{\alpha t}{e^{\alpha t}-1}\Big)\big(1-e^{-\alpha t}\big)^{\alpha N}, \qquad \alpha \in (0,1).$$

Consequences for the coagulation model:

- In [SMOLUCHOWSKI 1916] a system of ODEs is introduced for the evolution of the (microscopic) particle sizes:

$$\frac{\mathrm{d}}{\mathrm{d}t}\lambda_k(t) = \frac{1}{2}\sum_{m,\widetilde{m}\in\mathbb{N}\colon\, m+\widetilde{m}=k} \lambda_m(t)\lambda_{\widetilde{m}}(t)K(m,\widetilde{m}) - \lambda_k(t)\sum_{m\in\mathbb{N}}\lambda_m(t)K(k,m),$$

  where $K = \lim_{N\to\infty} NK_N$, and
  $\lambda_m(t) = \lim_{N\to\infty}\frac{1}{N}\#\{\text{particles at time } t \text{ of size } k\}$.
- One can check that the minimizers λ_k^* of our variational formula satisfy them.
- Convergence of stochastic coagulation processes towards these ODEs was expected for long time, but the first rigorous proof was given only in [LANG, NGUYEN 1980].
- In [LUSHNIKOV 1978] the formation of a gel is realized and explained.
- Pathwise large deviations appear cumbersome, but doable.
- Such LDPs have been derived by [MIELKE et. al. (2017)] for general chemical reactions, following a Freidlin-Wentsel approach, but the rate function is rather inexplicit and not easy to evaluate at a fixed time.

Consider the non-interacting Bose gas in the thermodynamic limit at temperature $1/\beta \in (0, \infty)$ with particle density $\rho \in (0, \infty)$. Then the partition function is given by

$$Z_N(\beta, \rho) = \sum_{(\ell_k)_{k\in\mathbb{N}} \in \mathbb{N}_0^{\mathbb{N}} :\ \sum_k k\ell_k = N} \prod_k \frac{N^{\ell_k}}{\ell_k!\, k^{\ell_k}} [\rho(4\pi\beta k)^{\frac{d}{2}}]^{-\ell_k}.$$

Consider the non-interacting Bose gas in the thermodynamic limit at temperature $1/\beta \in (0,\infty)$ with particle density $\rho \in (0,\infty)$. Then the partition function is given by

$$Z_N(\beta,\rho) = \sum_{(\ell_k)_{k\in\mathbb{N}}\in\mathbb{N}_0^{\mathbb{N}}\colon \sum_k k\ell_k = N} \prod_k \frac{N^{\ell_k}}{\ell_k!\, k^{\ell_k}} [\rho(4\pi\beta k)^{\frac{d}{2}}]^{-\ell_k}.$$

The free energy per particle is then

$$f(\beta,\rho) = \lim_{N\to\infty} \frac{1}{N} \log Z_N(\beta,\rho) = -\inf_{\lambda\in\mathcal{N}(\rho)} I(\lambda), \quad \text{where} \quad I(\lambda) = \sum_{k\in\mathbb{N}} \lambda_k \log \frac{\lambda_k k}{(4\pi\beta k)^{\frac{d}{2}}\mathrm{e}}.$$

Consider the non-interacting Bose gas in the thermodynamic limit at temperature $1/\beta \in (0,\infty)$ with particle density $\rho \in (0,\infty)$. Then the partition function is given by

$$Z_N(\beta,\rho) = \sum_{(\ell_k)_{k\in\mathbb{N}}\in\mathbb{N}_0^{\mathbb{N}}\colon \sum_k k\ell_k = N} \prod_k \frac{N^{\ell_k}}{\ell_k!\, k^{\ell_k}} [\rho(4\pi\beta k)^{\frac{d}{2}}]^{-\ell_k}.$$

The free energy per particle is then

$$f(\beta,\rho) = \lim_{N\to\infty} \frac{1}{N} \log Z_N(\beta,\rho) = -\inf_{\lambda\in\mathcal{N}(\rho)} I(\lambda), \quad \text{where} \quad I(\lambda) = \sum_{k\in\mathbb{N}} \lambda_k \log \frac{\lambda_k k}{(4\pi\beta k)^{\frac{d}{2}} \mathrm{e}}.$$

Comparison: In Lushnikov's model, we face roughly

$$t^N \mathrm{e}^{\frac{t}{2}N} \sum_{(\ell_k)_{k\in\mathbb{N}}\colon \sum_k k\ell_k = N} \prod_k \frac{k^{(k-2)\ell_k} t^{-\ell_k}}{\ell_k!\, k!^{\ell_k}}.$$

The two respective minimizers are

$$k\lambda_k^{(\mathrm{Lush})}(c;t) = \frac{1}{t} \frac{(ct\mathrm{e}^{-ct})^k}{k^{1-k}\, k!} \quad \text{and} \quad k\lambda_k^{(\mathrm{BEC})}(\alpha;t) = \frac{1}{\rho(4\pi\beta)^{\frac{d}{2}}} \frac{\mathrm{e}^{-\alpha k}}{k^{\frac{d}{2}}},$$

where c and α control the value of $\sum_k k\lambda_k$ (note that $k^{1-k}\, k! \asymp k^{3/2}$).

Comparison to Bose-Einstein condensation (II)

One difference: In the non-interacting Bose gas, the macroscopic part gives no energetic contribution, while in the Lushnikov model it does.

One difference: In the non-interacting Bose gas, the macroscopic part gives no energetic contribution, while in the Lushnikov model it does.

In the Bose gas, increasing ρ drives more and more particles into the finite cycles. There is a natural threshold, the critical inverse temperature β_c, characterised by

$$(4\pi\beta)^{\frac{d}{2}}\rho = \sum_{k\in\mathbb{N}} k^{-\frac{d}{2}}.$$

Only when all finite cycles are filled entirely, the first “infinite” cycle arises.

The BEC is a saturation transition.

In contrast, in the Lushnikov model, increasing t makes each particle larger, until some decide to make the jump to infinity. However, the other micro particles keep growing (recall that $\beta_t < \frac{1}{t}$).

The gelation phase transition is an explosion transition.

- type space: compact metric space $\mathcal{S}$
- vertex distribution: probability measure μ on $\mathcal{S}$
- connectivity probability function: positive symmetric irreducible kernel κ from $\mathcal{S}$ to $\mathcal{S}$.
- vertex set: $[N] = \{1, \dots, N\}$. Vertex i has the type $x_i \in \mathcal{S}$. Type vector $x = (x_1, \dots, x_N) \in \mathcal{S}^N$.

$\mathcal{G}_N = \mathcal{G}([N], x, \frac{1}{N}\kappa)$ is the graph on $[N] = \{1, \dots, N\}$, having a bond $\{i, j\}$ with probability $\frac{1}{N}\kappa(x_i, x_j) \wedge 1$, independently over all pairs (i, j) with $i \neq j$.

- type space: compact metric space $\mathcal{S}$
- vertex distribution: probability measure μ on $\mathcal{S}$
- connectivity probability function: positive symmetric irreducible kernel κ from $\mathcal{S}$ to $\mathcal{S}$.
- vertex set: $[N] = \{1, \dots, N\}$. Vertex i has the type $x_i \in \mathcal{S}$. Type vector $x = (x_1, \dots, x_N) \in \mathcal{S}^N$.

$\mathcal{G}_N = \mathcal{G}([N], x, \frac{1}{N}\kappa)$ is the graph on $[N] = \{1, \dots, N\}$, having a bond $\{i, j\}$ with probability $\frac{1}{N}\kappa(x_i, x_j) \wedge 1$, independently over all pairs (i, j) with $i \neq j$.

- There is a spatial coagulation model that can be mapped onto this graph model.
- The 120-pages article [Bollobas, Janson, Riordan 2007] derived a sufficient and necessary criterion for the phase transition of the existence of a giant component in $\mathcal{G}_N$. The main tool is a multitype branching process.
- We will prove the extension to an LDP for the statistics of the microscopic and the macroscopic components of $\mathcal{G}_N$, and obtain this criterion independently in a different way. The main tool is the explicit identification of the joint distribution of the statistics of all the connected components according to their multi-types.

For a measure ν on $\mathcal{S}$, introduce the operator

$$T_{\kappa,\nu}\colon L^2(\nu) \to L^2(\nu), \qquad T_{\kappa,\nu} f(x) = \int_{\mathcal{S}} \kappa(x,y) f(y)\, \nu(\mathrm{d}y),$$

and its norm

$$\Sigma(\kappa,\nu) = \|T_{\kappa,\nu}\|_{L^2(\nu)} = \sup_{f\in L^2(\nu)\colon \|f\|_{L^2(\nu)}=1} \|T_{\kappa,\nu} f\|_{L^2(\nu)}.$$

Existence of a giant component

- If $\Sigma(\kappa,\mu) \leq 1$, then the largest component of $\mathcal{G}_N$ has size $o(N)$ as $N\to\infty$ with high probability (in fact, $O(\log N)$).
- If $\Sigma(\kappa,\mu) > 1$, then it has size $\asymp N$. More precisely, if $\rho\colon \mathcal{S}\to[0,\infty)$ denotes the maximal solution of

$$\rho = 1 - \mathrm{e}^{-T_{\kappa,\mu}\rho},$$

  then the size of the largest component of $\mathcal{G}_N$ is $\sim N \int_{\mathcal{S}} \rho(x)\,\mu(\mathrm{d}x)$.

The sizes of the microscopic clusters are characterized in terms of the distribution of the sizes of the offspring of the multitype branching process, in which each particle of type $x\in\mathcal{S}$ has offspring with distribution that is a Poisson process with intensity $\kappa(x,y)\,\mu(\mathrm{d}y)$.

Assume that $\mathcal{S}$ is a finite set. Assume that $\mu_N = \frac{1}{N}\sum_{i=1}^N \delta_{x_i} \Longrightarrow \mu$ as $N \to \infty$.

We denote by $(\mathcal{C}_i)_i$ the collection of the connected components of $\mathcal{G}_N$.

Let $\eta_r(A)$ be the number of type-r sites in $A \subset [N]$, and $\eta(A) = (\eta_r(A))_{r\in\mathcal{S}}$.

For $k \in \mathbb{N}_0^{\mathcal{S}}$ let $\mathrm{Mi}_N(k) := \frac{1}{N}\sum_i \delta_{\eta(\mathcal{C}_i)}(k)$, then $\sum_k \mathrm{Mi}_N(k) k_r = \mu_N(r)$.

For $\alpha \in (0,1]^{\mathcal{S}}$ let $\mathrm{Ma}_N(k) := \sum_i \delta_{\frac{1}{N}\eta(\mathcal{C}_i)}(k)$, then $\int_{(0,1]^{\mathcal{S}}} \mathrm{Ma}_N(\mathrm{d}y)\, y_r = \mu_N(r)$.

Joint distribution of the cluster types

For any $l = (l_k)_{k\in\mathbb{N}_0^{\mathcal{S}}} \in \mathbb{N}_0^{(\mathbb{N}_0^{\mathcal{S}})}$ satisfying $\sum_k l_k k_r = N\mu_N(r)$ for any $r \in \mathcal{S}$,

$$\mathbb{P}(N\mathrm{Mi}_N(k) = l_k\ \forall k) = \Big(\prod_{r\in\mathcal{S}} (N\mu_N(r))!\Big)$$

$$\times \prod_{k\in\mathbb{N}^{\mathcal{S}}} \frac{p_N(k)^{l_k}}{l_k!\,(k_r!)^{l_k}} \left(\prod_{r,s\in\mathcal{S}} \Big(1 - \frac{\kappa(r,s)}{N}\Big)^{k_s(N\mu_N(r)-k_r)/2}\right)^{l_k}$$

with $p_N(k)$ the connection probability for the graph $\mathcal{G}(|k|, x, \frac{1}{N}\kappa)$ for any k-compatible x.

Define

$$\tau(k) := \sum_{T\in\mathcal{T}(k,x)} \prod_{\{i,j\}\in E(T)} \kappa\,(x_i, x_j)\,, \qquad k \in \mathbb{N}_0^{\mathcal{S}},$$

where $x \in \mathcal{S}^{|k|}$ is k-compatible, and $\mathcal{T}(k,x)$ is the set of spanning trees on $[|k|]$.

Notable extension of [STEPANOV 1970]:

Asymptotics of $p_N(k)$ as $N \to \infty$

$$p_N(k) \sim N^{1-|k|}\tau(k), \qquad k \in \mathbb{N}_0^{\mathcal{S}}.$$

and

$$\frac{1}{N}\log p_N(\lfloor Ny \rfloor) \to \sum_{r\in\mathcal{S}} y_r \log\left(1 - \mathrm{e}^{-\sum_{s\in\mathcal{S}}\kappa(r,s)y_s}\right), \qquad y \in (0,1]^{\mathcal{S}}.$$

The second assertion is of independent interest and is also proved for $\mathcal{S}$ a compact metric space. The technical problem is that giant clusters can be connected with just one bond, whose probabily is not seen on the exponential scale.

Denote $c_r(\lambda) = \sum_{k \in \mathbb{N}_0^{\mathcal{S}}} \lambda_k k_r$ and $c_r(\alpha) = \int_{(0,1]^{\mathcal{S}}} \alpha(\mathrm{d}y)\, y_r$.

The LDP

As $N \to \infty$, the pair $(\mathrm{Mi}_N, \mathrm{Ma}_N)$ satisfies a large-deviations principle with rate function

$$I(\lambda, \alpha) = I_{\mathrm{Mi}}(\lambda) + I_{\mathrm{Ma}}(\alpha) + I_{\mathrm{Me}}(\mu - c(\lambda) - c(\alpha)),$$

where

$$I_{\mathrm{Mi}}(\lambda) = \sum_{k \in \mathbb{N}^{\mathcal{S}}} \lambda_k \log \frac{\lambda_k}{\tau(k) \prod_{s \in \mathcal{S}} \frac{\mu_s^{k_s}}{k_s!}} + \sum_{k \in \mathbb{N}^{\mathcal{S}}} \lambda_k (|k| - 1) + \frac{1}{2} \langle c(\lambda), \kappa\mu \rangle,$$

$$I_{\mathrm{Ma}}(\alpha) = \int_{[0,1]^{\mathcal{S}}} \alpha(\mathrm{d}y) \left(\left\langle y, \log \frac{y}{1 - \mathrm{e}^{-\kappa * y}} \right\rangle + \frac{1}{2} \langle y, \kappa * (\mu - y) \rangle \right),$$

$$I_{\mathrm{Me}}(\nu) = \left\langle \nu, \log \frac{\nu}{(\kappa * \nu)\, \mu} \right\rangle + \frac{1}{2} \langle \nu, \kappa * \mu \rangle.$$

- entropies $\Longleftrightarrow$ combinatorics
- terms with $\frac{1}{2}$ $\Longleftrightarrow$ non-connection probabilities
- term τ times Poisson $\Longleftrightarrow$ reference process, conditioned on being connected

- We indeed prove this also for $\mathcal{S}$ a compact metric space. The lift from discrete $\mathcal{S}$ to continuous $\mathcal{S}$ is a cumbersome and technical work in the spirit of the DAWSON–GÄRTNER theorem.
- We do not know about earlier work in that direction.
- One application is to i.i.d. random $x_1, \dots, x_N \Longrightarrow$ quenched LDP. Annealed versions follow easily.
- Standard consequences are contracted separate LDPs for Mi_N and Ma_N. ($\Longrightarrow$ interesting conditional phase transition, see later)
- We abstained from formulating an LDP for the mesoscopic part.

Discussion of minimizers

Given $c = (c_r)_{r\in\mathbb{N}_0^{\mathcal{S}}}$, the Euler-Lagrange equations for minimizers λ of I_{Mi} subject to $c(\lambda) = c$, i.e., $\sum_k \lambda_k k_r = c_r$ for $r \in \mathcal{S}$, are

$$\lambda_k = \tau(k) \prod_{r\in\mathbb{N}_0^{\mathcal{S}}} \frac{t_r^{k_r}}{k_r!}, \qquad k \in \mathbb{N}_0^{\mathcal{S}}.$$

Given $c = (c_r)_{r \in \mathbb{N}_0^{\mathcal{S}}}$, the Euler-Lagrange equations for minimizers λ of I_{Mi} subject to $c(\lambda) = c$, i.e., $\sum_k \lambda_k k_r = c_r$ for $r \in \mathcal{S}$, are

$$\lambda_k = \tau(k) \prod_{r \in \mathbb{N}_0^{\mathcal{S}}} \frac{t_r^{k_r}}{k_r!}, \qquad k \in \mathbb{N}_0^{\mathcal{S}}.$$

The only candidate is $t_r(c) = c_r \, \mathrm{e}^{-\kappa * c(r)}$. Call the solution $\lambda^*(c)$ if it exists.

Given $c = (c_r)_{r\in\mathbb{N}_0^{\mathcal{S}}}$, the Euler-Lagrange equations for minimizers λ of I_{Mi} subject to $c(\lambda) = c$, i.e., $\sum_k \lambda_k k_r = c_r$ for $r \in \mathcal{S}$, are

$$\lambda_k = \tau(k) \prod_{r\in\mathbb{N}_0^{\mathcal{S}}} \frac{t_r^{k_r}}{k_r!}, \qquad k \in \mathbb{N}_0^{\mathcal{S}}.$$

The only candidate is $t_r(c) = c_r\,\mathrm{e}^{-\kappa * c(r)}$. Call the solution $\lambda^*(c)$ if it exists.

Existence of $\lambda^*(c)$

- $t(c)$ is a solution $\Longleftrightarrow \tilde{c} \mapsto t(\tilde{c})$ is invertible and the inverse map is analytic in $t(c)$,
- this invertibility is true if and only if $\Sigma(\kappa, c) < 1$.

Discussion of minimizers

Given $c = (c_r)_{r \in \mathbb{N}_0^{\mathcal{S}}}$, the Euler-Lagrange equations for minimizers λ of I_{Mi} subject to $c(\lambda) = c$, i.e., $\sum_k \lambda_k k_r = c_r$ for $r \in \mathcal{S}$, are

$$\lambda_k = \tau(k) \prod_{r \in \mathbb{N}_0^{\mathcal{S}}} \frac{t_r^{k_r}}{k_r!}, \qquad k \in \mathbb{N}_0^{\mathcal{S}}.$$

The only candidate is $t_r(c) = c_r \, \mathrm{e}^{-\kappa * c(r)}$. Call the solution $\lambda^*(c)$ if it exists.

Existence of $\lambda^*(c)$

- $t(c)$ is a solution $\Longleftrightarrow$ $\tilde{c} \mapsto t(\tilde{c})$ is invertible and the inverse map is analytic in $t(c)$,
- this invertibility is true if and only if $\Sigma(\kappa, c) < 1$.

Minimizers of LDP rate function

- If $\Sigma(\kappa, \mu) \le 1$, then $\lambda^*(\mu)$ exists, and $(\lambda^*(\mu), 0)$ is the minimizer of I. No giant component arises.
- If $\Sigma(\kappa, \mu) > 1$, then the optimal microcluster distribution c^* is characterised by $c_r = \mu_r \, \mathrm{e}^{\kappa * (\mu - c)(r)}$, and the minimizer of I is equal to $(\lambda(c^*), \delta_{\mu - c^*})$. The latter corresponds to a giant cluster with $\sim N(\mu_r - c_r^*)$ vertices of multitype $r \in \mathbb{N}_0^{\mathcal{S}}$ for any r.

Recall the multitype branching process, in which each particle of type $x \in \mathcal{S}$ has offspring with distribution that is a Poisson process with intensity $\kappa(x,y)\,\mu(\mathrm{d}y)$.

Denote by $\Xi(\mathrm{d}r)$ the entire progeneity (total offspring) of type $r \in \mathcal{S}$ the process. Let $\mathtt{P}_r$ denote the probability measure when the process starts with just one particle of type r.

Then

$$\mu(\mathrm{d}r)\mathtt{P}_r(\Xi \in \mathrm{d}k) = \lambda_\mu(\mathrm{d}k)k(\mathrm{d}r), \qquad k \in \mathcal{M}_{\mathbb{N}_0}(\mathcal{S}), r \in \mathcal{S}.$$

In words: the empirical statistics of the microscopic components in $\mathcal{G}_N$ in the subcritical case approximate the distribution of the total offspring of the characteristic branching process.

Contraction principle $\Longrightarrow \mathrm{Ma}_N$ satisfies an LDP with rate function

$$\mathcal{I}_{\mathrm{Ma}}(\alpha) = \inf_{\lambda} I(\lambda, \alpha) = I_{\mathrm{Ma}}(\alpha) + J(\mu - c_\alpha),$$

where

$$J(c) = \begin{cases} I_{\mathrm{Mi}}(\lambda_c) & \text{if } \Sigma(\kappa, c) \leq 1, \\ I_{\mathrm{Mi}}(\lambda_{b*}) + I_{\mathrm{Me}}(c - b^*) & \text{if } \Sigma(\kappa, c) > 1, \end{cases}$$

and $b^* = b^*(c) \in \mathcal{M}(\mathcal{S})$ is the minimal solution $\neq c$ of the characteristic equation

$$\kappa(c - b^*)(r)\, b^*(\mathrm{d}r) = (c - b^*)(\mathrm{d}r), \qquad b^* \leq c,$$

and b^* is saturated in the sense that $\Sigma(\kappa, b^*) = 1$.

Contraction principle $\Longrightarrow \mathrm{Ma}_N$ satisfies an LDP with rate function

$$\mathcal{I}_{\mathrm{Ma}}(\alpha) = \inf_{\lambda} I(\lambda, \alpha) = I_{\mathrm{Ma}}(\alpha) + J(\mu - c_\alpha),$$

where

$$J(c) = \begin{cases} I_{\mathrm{Mi}}(\lambda_c) & \text{if } \Sigma(\kappa, c) \leq 1, \\ I_{\mathrm{Mi}}(\lambda_{b*}) + I_{\mathrm{Me}}(c - b^*) & \text{if } \Sigma(\kappa, c) > 1, \end{cases}$$

and $b^* = b^*(c) \in \mathcal{M}(\mathcal{S})$ is the minimal solution $\neq c$ of the characteristic equation

$$\kappa(c - b^*)(r)\, b^*(\mathrm{d}r) = (c - b^*)(\mathrm{d}r), \qquad b^* \leq c,$$

and b^* is saturated in the sense that $\Sigma(\kappa, b^*) = 1$.

Hence, conditional on $\{\mathrm{Ma}_N \approx \alpha\}$, we have, as $N \to \infty$,

$$\mathrm{Mi}_N \overset{N\to\infty}{\Longrightarrow} \begin{cases} \lambda_{\mu - c_\alpha} & \text{if } \Sigma(\kappa, \mu - c_\alpha) < 1 \quad \Longrightarrow \text{no mesoscopic part,} \\ \lambda_{b^*} & \text{if } \Sigma(\kappa, \mu - c_\alpha) \geq 1 \quad \Longrightarrow \text{mesoscopic part.} \end{cases}$$

$\Longrightarrow$ saturation phase transition: If the macroscopic part α is fixed, and more and more bonds are trown in, then first the microscopic part increases until λ_{b^*} is attained, then it is frozen, and only the mesoscopic part increases. ($\Longrightarrow$ frozen percolation)