



HAL
open science

Exchangeable Random Measures for Sparse and Modular Graphs with Overlapping Communities

Adrien Todeschini, Francois Caron

► **To cite this version:**

Adrien Todeschini, Francois Caron. Exchangeable Random Measures for Sparse and Modular Graphs with Overlapping Communities. 2016. hal-01270854

HAL Id: hal-01270854

<https://hal.science/hal-01270854>

Preprint submitted on 8 Feb 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Exchangeable Random Measures for Sparse and Modular Graphs with Overlapping Communities

Adrien Todeschini and François Caron

INRIA & Institut de Mathématiques de Bordeaux, France
e-mail: Adrien.Todeschini@inria.fr

Department of Statistics, University of Oxford, UK
e-mail: caron@stats.ox.ac.uk

Abstract: We propose a novel statistical model for sparse networks with overlapping community structure. The model is based on representing the graph as an exchangeable point process, and naturally generalizes existing probabilistic models with overlapping block-structure to the sparse regime. Our construction builds on vectors of completely random measures, and has interpretable parameters, each node being assigned a vector representing its level of affiliation to some latent communities. We develop methods for simulating this class of random graphs, as well as to perform posterior inference. We show that the proposed approach can recover interpretable structure from two real-world networks and can handle graphs with thousands of nodes and tens of thousands of edges.

Keywords and phrases: Networks, Random Graphs, Multiview Networks, Multigraphs, Completely Random Measures, Lévy measure, Multivariate Subordinator, Sparsity, Non-Negative Factorization, Exchangeability, Point Processes.

1. Introduction

There has been a growing interest in the analysis, understanding and modeling of network data over the recent years. A network is composed of a set of nodes, or vertices, with connections between them. Network data arise in a wide range of fields, and include social networks, collaboration networks, communication networks, biological networks, food webs and are a useful way of representing interactions between sets of objects. Of particular importance is the elaboration of random graph models, which can capture the salient properties of real-world graphs. Following the seminal work of Erdős and Rényi (1959), various network models have been proposed; see the overviews of Newman (2003b, 2009), Kolaczyk (2009), Bollobás (2001), Goldenberg et al. (2010), Fienberg (2012) or Jacobs and Clauset (2014). In particular, a large body of the literature has concentrated on models that can capture some modular or community structure within the network. The first statistical network model in this line of research is the popular stochastic block-model (Holland et al., 1983; Snijders and Nowicki, 1997; Nowicki and Snijders, 2001). The stochastic block-model assumes that each node belongs to one of p latent communities, and the probability of connection between two nodes is given by a $p \times p$ connectivity matrix. This model has been extended in various directions, by introducing degree-correction parameters (Karrer and Newman, 2011), by allowing the number of communities to grow with the size of the network (Kemp et al., 2006), or by considering overlapping communities (Airoldi et al., 2008; Miller et al., 2009; Latouche et al., 2011; Palla et al., 2012; Yang and Leskovec, 2013). Stochastic block-models and their extensions have shown to offer a very flexible modeling framework, with interpretable parameters, and have been successfully used for the analysis of numerous real-world networks. However, as outlined by Orbanz and Roy (2015), when one makes the usual assumption that the ordering of the nodes is irrelevant in the definition of the statistical network model, the Bayesian probabilistic versions of those models lead to dense networks¹: that means that the number of edges grows quadratically with the number of nodes. This property is rather undesirable, as many real-world networks are believed to be sparse.

Recently, Caron and Fox (2014) proposed an alternative framework for statistical network modeling. The framework is based on representing the graph as an exchangeable random measure on the plane. More precisely, the nodes are embedded at some location $\theta_i \in \mathbb{R}_+$ and, for simple graphs, a connection exists

¹We refer to graphs whose number of edges scales quadratically with the number of nodes as dense, and sparse if it scales sub-quadratically.

between two nodes i and j if there is a point at locations (θ_i, θ_j) and (θ_j, θ_i) . An undirected simple graph is therefore represented by a symmetric point process Z on the plane

$$Z = \sum_{i,j} z_{ij} \delta_{(\theta_i, \theta_j)} \quad (1)$$

where $z_{ij} = z_{ji} = 1$ if i and j are connected, 0 otherwise; see Figure 1 for an illustration. Caron and Fox (2014) noted that jointly exchangeable random measures, a notion to be defined in Eq. (21), admit a representation theorem due to Kallenberg (1990), providing a general construction for exchangeable random measures hence random graphs represented by such objects. This connection is further explored by Veitch and Roy (2015) and Borgs et al. (2016), who provide a detailed description and extensive theoretical analysis of the associated class of random graphs, which they name *Kallenberg exchangeable graphs* or *graphon processes*. Within this class of models, Caron and Fox (2014) consider in particular the following simple generative model, where two nodes $i \neq j$ connect with probability

$$\Pr(z_{ij} = 1 | (w_\ell)_{\ell=1,2,\dots}) = 1 - e^{-2w_i w_j} \quad (2)$$

where the $(w_i, \theta_i)_{i=1,2,\dots}$ are the points of a Poisson point process on \mathbb{R}_+^2 . The parameters $w_i > 0$ can be interpreted as sociability parameters. Depending on the properties of the mean measure of the Poisson process, the authors show that it is possible to generate both dense and sparse graphs, with potentially heavy-tailed degree distributions, within this framework. The construction (2) is however rather limited in terms of capturing structure in the network. Herlau et al. (2015) proposed an extension of (2), which can accommodate a community structure. More precisely, introducing latent community membership variables $c_i \in \{1, \dots, p\}$, two nodes $i \neq j$ connect with probability

$$\Pr(z_{ij} = 1 | (w_\ell, c_\ell)_{\ell=1,2,\dots}, (\eta_{k\ell})_{1 \leq k, \ell \leq p}) = 1 - e^{-2\eta_{c_i c_j} w_i w_j} \quad (3)$$

where the $(w_i, c_i, \theta_i)_{i=1,2,\dots}$ are the points of a (marked) Poisson point process on $\mathbb{R}_+ \times \{1, \dots, p\} \times \mathbb{R}_+$ and $\eta_{k\ell}$ are positive random variables parameterizing the strength of interaction between nodes in community k and nodes in community ℓ . The model is similar in spirit to the degree-corrected stochastic block-model (Karrer and Newman, 2011), but within the point process framework (1), and can thus accommodate both sparse and dense networks with community structure. The model of Herlau et al. (2015) however shares the limitations of the (degree-corrected) stochastic block-model, in the sense that it cannot model overlapping community structures, each node being assigned to a single community; see Latouche et al. (2011) and Yang and Leskovec (2013) for more discussion along these lines. Other extensions with block structure or mixed membership block structure are also suggested by Borgs et al. (2016).

In this paper, we consider that each node i is assigned a set of latent non-negative parameters w_{ik} , $k = 1, \dots, p$, and that the probability that two nodes $i \neq j$ connect is given by

$$\Pr(z_{ij} = 1 | (w_{\ell 1}, \dots, w_{\ell p})_{\ell=1,2,\dots}) = 1 - e^{-2 \sum_{k=1}^p w_{ik} w_{jk}}. \quad (4)$$

These non-negative weights can be interpreted as measuring the level of affiliation of node i to the latent communities $k = 1, \dots, p$. For example, in a friendship network, these communities can correspond to colleagues, family, or sport partners, and the weights measure the level of affiliation of an individual to each community. Note that as individuals can have high weights in different communities, the model can capture overlapping communities. The link probability (4) builds on a non-negative factorization; it has been used by other authors for network modeling (Yang and Leskovec, 2013; Zhou, 2015) and is also closely related to the model for multigraphs of Ball et al. (2011). The main contribution of this paper is to use the link probability (4) within the point process framework of Caron and Fox (2014). To this aim, we consider that the node locations and weights $(w_{i1}, \dots, w_{ip}, \theta_i)_{i=1,2,\dots}$ are drawn from a Poisson point process on \mathbb{R}_+^{p+1} with a given mean measure ν . The construction of such multivariate point process relies on vectors of completely random measures (or equivalently multivariate subordinators). In particular, we build on the flexible though tractable construction recently introduced by Griffin and Leisen (2014).

The proposed model generalizes that of Caron and Fox (2014) by allowing the model to capture more structure in the network, while retaining its main features, and is shown to have the following properties:




FIG 1. Representation of an undirected graph via a point process Z . Each node i is embedded in \mathbb{R}_+ at some location θ_i and is associated with a set of positive attributes (w_{i1}, \dots, w_{ip}) . An edge between nodes θ_i and θ_j is represented by a point at locations (θ_i, θ_j) and (θ_j, θ_i) in \mathbb{R}_+^2 .

- **Interpretability:** each node is assigned a set of positive parameters, which can be interpreted as measuring the levels of affiliation of a node to latent communities; once those parameters are learned, they can be used to uncover the latent structure in the network.
- **Sparsity:** we can generate graphs whose number of edges grows subquadratically with the number of nodes.
- **Exchangeability:** in the sense of [Kallenberg \(1990\)](#).

Additionally, we develop a Markov chain Monte Carlo (MCMC) algorithm for posterior inference with this model, and show experiments on two real-world networks with a thousand of nodes and tens of thousands of edges.

The article is organized as follows. The class of random graph models is introduced in Section 2. Properties of the class of graphs and simulation are described in Section 3. We derive a scalable MCMC algorithm for posterior inference in Section 4. In Section 5 we provide illustrations of the proposed method on simulated data and on two networks: a network of citations between political blogs and a network of connections between US airports. We show that the approach is able to discover interpretable structure in the data.

2. Sparse graph models with overlapping communities

In this section, we present the statistical model for simple graphs. The construction builds on vectors of completely random measures (CRM, [Kingman, 1967](#)). We only provide here the necessary material for the definition of the network model; please refer to Appendix A for additional background on vectors of CRMs. The model described in this section can also be extended to bipartite graphs; see Appendix D.

2.1. General construction using vectors of CRMs

We consider that each node i is embedded at some location $\theta_i \in \mathbb{R}_+$, and has some set of positive weights $(w_{i1}, \dots, w_{ip}) \in \mathbb{R}_+^p$. The points $(w_{i1}, \dots, w_{ip}, \theta_i)_{i=1, \dots, \infty}$ are assumed to be drawn from a Poisson process with mean measure

$$\nu(dw_1, \dots, dw_p, d\theta) = \rho(dw_1, \dots, dw_p)\lambda(d\theta) \quad (5)$$

where λ is the Lebesgue measure and ρ is a σ -finite measure on \mathbb{R}_+^p , concentrated on $\mathbb{R}_+^p \setminus \{\mathbf{0}\}$, which satisfies

$$\int_{\mathbb{R}_+^p} \min\left(1, \sum_{k=1}^p w_k\right) \rho(dw_1, \dots, dw_p) < \infty. \quad (6)$$




FIG 2. An example of (a) the restriction on $[0, 1]^2$ of the two atomic measures D_1 and D_2 , (b) the corresponding multiview directed multigraphs (top: view 1; bottom: view 2) and (c) corresponding undirected graph.

Under this condition (Skorohod, 1991; Barndorff-Nielsen et al., 2001), we can describe the set of weights and locations using a vector of completely random measures (W_1, \dots, W_p) on \mathbb{R}_+ :

$$W_k = \sum_{i=1}^{\infty} w_{ik} \delta_{\theta_i}, \text{ for } k = 1, \dots, p. \quad (7)$$

We simply write

$$(W_1, \dots, W_p) \sim \text{CRM}(\rho, \lambda). \quad (8)$$

Mimicking the hierarchical construction of Caron and Fox (2014), we introduce integer-valued random measures D_k on \mathbb{R}_+^2 , $k = 1, \dots, p$,

$$D_k = \sum_{i=1}^{\infty} \sum_{j=1}^{\infty} n_{ijk} \delta_{(\theta_i, \theta_j)} \quad (9)$$

where the n_{ijk} are natural integers. The vector of random measures (D_1, \dots, D_p) can be interpreted as representing a multiview (a.k.a. multiplex or multi-relational) directed multigraph (Verbrugge, 1979; Salter-Townshend and McCormick, 2013), where n_{ijk} represents the number of interactions from node i to node j in the view k ; see Figure 2 for an illustration. Conditionally on the vector of CRMs, the measures D_k are independently drawn from a Poisson process² with mean measure $W_k \times W_k$

$$D_k | (W_1, \dots, W_p) \sim \text{Poisson}(W_k \times W_k) \quad (10)$$

that is, the n_{ijk} are independently Poisson distributed with rate $w_{ik}w_{jk}$.

Finally, the point process Z representing the graph (Eq. (1)) is deterministically obtained from (D_1, \dots, D_p) by setting $z_{ij} = 1$ if there is at least one directed connection between i and j in any view, and 0 otherwise, therefore $z_{ij} = \min(1, \sum_{k=1}^p n_{ijk} + n_{jik})$. To sum up, the graph model is described as follows:

$$\begin{aligned} W_k &= \sum_{i=1}^{\infty} w_{ik} \delta_{\theta_i} & (W_1, \dots, W_p) &\sim \text{CRM}(\rho, \lambda) \\ D_k &= \sum_{i=1}^{\infty} \sum_{j=1}^{\infty} n_{ijk} \delta_{(\theta_i, \theta_j)} & D_k | W_k &\sim \text{Poisson}(W_k \times W_k) \\ Z &= \sum_{i=1}^{\infty} \sum_{j=1}^{\infty} \min(1, \sum_{k=1}^p n_{ijk} + n_{jik}) \delta_{(\theta_i, \theta_j)}. \end{aligned} \quad (11)$$

²Note that we consider a generalized definition of a Poisson process, where the mean measure is allowed to have atoms; see e.g. Daley and Vere-Jones (2008a, Section 2.4).




FIG 3. An example, for $p = 2$, of (a) the product measures $W_k \times W_k$, (b) a draw of the directed multigraph measures $D_k \mid W_k \sim \text{Poisson}(W_k \times W_k)$ and (c) corresponding undirected measure $Z = \sum_{i=1}^{\infty} \sum_{j=1}^{\infty} \min(1, \sum_{k=1}^p n_{ijk} + n_{jik}) \delta_{(\theta_i, \theta_j)}$.

The model construction is illustrated in Figure 3. Integrating out the measures D_k , $k = 1, \dots, p$, the construction can be expressed as, for $i \leq j$

$$z_{ij} | (w_{\ell 1}, \dots, w_{\ell p})_{\ell=1,2,\dots} \sim \begin{cases} \text{Ber}(1 - \exp(-2 \sum_{k=1}^p w_{ik} w_{jk})) & i \neq j \\ \text{Ber}(1 - \exp(-\sum_{k=1}^p w_{ik}^2)) & i = j \end{cases} \quad (12)$$

and $z_{ji} = z_{ij}$; see Figure 1.

Graph Restrictions. Except in trivial cases, we have $W_k(\mathbb{R}_+) = \infty$ a.s. and therefore $Z(\mathbb{R}_+^2) = \infty$ a.s., so the number of points over the plane is infinite a.s. For $\alpha > 0$, we consider restrictions of the measures W_k , $k = 1, \dots, p$, to the interval $[0, \alpha]$ and of the measures D_k and Z to the box $[0, \alpha]^2$, and write respectively $W_{k\alpha}$, $D_{k\alpha}$ and Z_α these restrictions. Note that condition (6) ensures that $W_{k\alpha}([0, \alpha]) < \infty$ a.s. hence $D_{k\alpha}([0, \alpha]^2) < \infty$ and $Z_\alpha([0, \alpha]^2) < \infty$ a.s. As a consequence, for a given $\alpha > 0$, the model yields a finite number of edges a.s., even though there may be an infinite number of points $(w_i, \theta_i) \in \mathbb{R}_+ \times [0, \alpha]$; see Section 3.

Remark 1 The model defined above can also be used for random multigraphs, where $n_{ij} = \sum_{k=1}^p n_{ijk}$ is the number of directed interactions between i and j . Then we have

$$n_{ij} | (w_{\ell 1}, \dots, w_{\ell p})_{\ell=1,2,\dots} \sim \text{Poisson} \left(\sum_{k=1}^p w_{ik} w_{jk} \right)$$

which is a Poisson non-negative factorization (Lee, 1999; Cemgil, 2009; Psorakis et al., 2011; Ball et al., 2011; Gopalan et al., 2015).

Remark 2 The model defined by Eq. (12) allows to model networks which exhibit assortativity (Newman, 2003a), meaning that two nodes with similar characteristics (here similar set of weights) are more likely to connect than nodes with dissimilar characteristics. The link function can be generalized to (see e.g. Zhou, 2015)

$$z_{ij} \sim \text{Ber} \left(1 - \exp \left(- \sum_{k=1}^p \sum_{\ell=1}^p \eta_{k\ell} w_{ik} w_{j\ell} \right) \right)$$




FIG 4. Graph sampled from the model with three latent communities, identified by colors red, green, blue. For each node, the intensity of each color is proportional to the value of the associated weight in that community. Pure red/green/blue color indicates the node is only strongly affiliated to a single community. A mixture of those colors indicates balanced affiliations to different communities. Graph generated with the software Gephi (Bastian et al., 2009).

where $\eta_{k\ell} \geq 0$, in order to be able to capture both assortative and disassortative mixing in the network. In particular, setting larger values off-diagonal than on the diagonal of the matrix $(\eta_{k\ell})_{1 \leq k, \ell \leq p}$ allows to capture disassortative mixing. The properties and algorithms for simulation and posterior inference can trivially be extended to this more general case. In order to keep the notations as simple as possible, we focus here on the simpler link function (12).

2.2. Particular model based on compound CRMs

The key component in our statistical network model is the multivariate Lévy measure ρ in (8). Various approaches have been developed for constructing multivariate Lévy measures (Tankov, 2003; Cont and Tankov, 2003; Kallsen and Tankov, 2006; Barndorff-Nielsen et al., 2001; Skorohod, 1991), or more specifically vectors of completely random measures (Epifani and Lijoi, 2010; Leisen and Lijoi, 2011; Leisen et al., 2013; Griffin et al., 2013; Lijoi et al., 2014). We will in this paper consider the following particular form:

$$\rho(dw_1, \dots, dw_p) = e^{-\sum_{k=1}^p \gamma_k w_k} \int_0^\infty w_0^{-p} F\left(\frac{dw_1}{w_0}, \dots, \frac{dw_p}{w_0}\right) \rho_0(dw_0) \quad (13)$$

where $F(d\beta_1, \dots, d\beta_p)$ is some score probability distribution on \mathbb{R}_+^d , with moment generating function $M(t_1, \dots, t_p)$, ρ_0 is a base Lévy measure on \mathbb{R}_+ and $\gamma_k \geq 0$ are exponentially tilting parameters for $k = 1, \dots, p$. The model defined by (5) and (13) is a special case of the compound completely random measure (CCRM) model proposed by Griffin and Leisen (2014). It admits the following hierarchical construction, which makes interpretability, characterization of the conditionals and analysis of this class of models particularly easy. Let

$$W_0 = \sum_{i=1}^{\infty} w_{i0} \delta_{\theta_i} \sim \text{CRM}(\tilde{\rho}_0, \lambda) \quad (14)$$

where $\tilde{\rho}_0$ is a measure on \mathbb{R}_+ defined by $\tilde{\rho}_0(dw_0) = M(-w_0\gamma_1, \dots, -w_0\gamma_p) \rho_0(dw_0)$, and for $k = 1, \dots, p$ and $i = 1, 2, \dots$

$$w_{ik} = \beta_{ik} w_{i0}$$

where the scores β_{ik} have the following joint distribution

$$(\beta_{i1}, \dots, \beta_{ip}) | w_{i0} \stackrel{\text{ind}}{\sim} H(\cdot | w_{i0}) \quad (15)$$

with H is an exponentially tilted version of F :

$$H(d\beta_1, \dots, d\beta_p | w_0) = \frac{e^{-w_0 \sum_{k=1}^p \gamma_k \beta_k} F(d\beta_1, \dots, d\beta_p)}{\int_{\mathbb{R}_+^p} e^{-w_0 \sum_{k=1}^p \gamma_k \tilde{\beta}_k} F(d\tilde{\beta}_1, \dots, d\tilde{\beta}_p)}. \quad (16)$$

Additionally, the set of points $(w_{i0}, \beta_{i1}, \dots, \beta_{ip})_{i=1,2,\dots}$ is a Poisson point process with mean measure

$$e^{-w_0 \sum_{k=1}^p \gamma_k \beta_k} F(d\beta_1, \dots, d\beta_p) \rho_0(dw_0). \quad (17)$$

Dependence between the different CRMs is both tuned by the shared scaling parameter w_{i0} and potential dependency between the scores $(\beta_{i1}, \dots, \beta_{ip})$. The hierarchical construction has the following interpretation:

- The weight w_{i0} is an individual scaling parameter for node i whose distribution is tuned by the base Lévy measure ρ_0 . It can be considered as a degree correction, as often used in network models (Karrer and Newman, 2011; Zhao et al., 2012; Herlau et al., 2015). As shown in Section 3, ρ_0 tunes the overall sparsity properties of the network.
- The community-related scores β_{ik} tune the level of affiliation of node i to community k ; this is controlled by both the score distribution F and the tilting coefficients γ_k . These parameters tune the overlapping block-structure of the network.

An example of such a graph with three communities is displayed in Figure 4.

Specific choices for F and ρ_0 . We now give here specific choices of score distribution F and base Lévy measure ρ_0 , which lead to scalable inference algorithms. As in Griffin and Leisen (2014), we consider that F is a product of independent gamma distributions

$$F(d\beta_1, \dots, d\beta_p) = \prod_{k=1}^p \beta_k^{a_k-1} e^{-b_k \beta_k} \frac{b_k^{a_k}}{\Gamma(a_k)} d\beta_k \quad (18)$$

where $a_k > 0, b_k > 0, k = 1, \dots, p$, which leads to

$$H(dw_1, \dots, dw_p | w_0) \propto \prod_{k=1}^p w_k^{a_k-1} e^{-\frac{b_k w_k}{w_0} - \gamma_k w_k} dw_k$$

which is also a product of gamma distributions.

ρ_0 is set to be the mean measure of the jump part of a generalized gamma process (Hougaard, 1986; Brix, 1999), which has been extensively used in BNP models due to its generality, the interpretability of its parameters and its attractive conjugacy properties (James, 2002; Lijoi et al., 2007; Saeedi and Bouchard-Côté, 2011; Caron, 2012; Caron et al., 2014). The Lévy measure in this case is

$$\rho_0(dw_0) = \frac{1}{\Gamma(1-\sigma)} w_0^{-1-\sigma} \exp(-w_0 \tau) dw_0 \quad (19)$$

where the parameters (σ, τ) verify

$$\sigma \in (0, 1), \tau \geq 0 \quad \text{or} \quad \sigma \in (-\infty, 0], \tau > 0. \quad (20)$$

The gamma process ($\sigma = 0$), the inverse Gaussian process ($\sigma = \frac{1}{2}$) and the stable process ($\sigma \in (0, 1), \tau = 0$) are special cases. Using (18) and (19), the multivariate Lévy measure has the following analytic form

$$\rho(dw_1, \dots, dw_p) = \frac{2e^{-\sum_{k=1}^p \gamma_k w_k}}{\Gamma(1-\sigma)} \left[\prod_{k=1}^p \frac{w_k^{a_k-1} b_k^{a_k}}{\Gamma(a_k)} \right] \left(\frac{\tau}{\sum_{k=1}^p b_k w_k} \right)^{-\frac{\kappa}{2}} K_\kappa \left(2 \sqrt{\tau \sum_{k=1}^p b_k w_k} \right) dw_1 \dots dw_p$$

where $\kappa = \sigma + \sum_{k=1}^p a_k$ and K is the modified Bessel function of the second kind.

3. Properties and Simulation

3.1. Exchangeability

The point process Z defined by (11) is jointly exchangeable in the sense of Kallenberg (1990, 2005). For any $h > 0$ and any permutation π of \mathbb{N}

$$(Z(A_i \times A_j)) \stackrel{d}{=} (Z(A_{\pi(i)} \times A_{\pi(j)})) \text{ for } (i, j) \in \mathbb{N}^2 \quad (21)$$

where $A_i = [h(i-1), hi]$. This follows directly from the fact that the vector of CRMs (W_1, \dots, W_p) has independent and identically distributed increments, hence

$$(W_1(A_i), \dots, W_p(A_i)) \stackrel{d}{=} (W_1(A_{\pi(i)}), \dots, W_p(A_{\pi(i)})). \quad (22)$$

The model thus falls into the general representation theorem for exchangeable point processes (Kallenberg, 1990).

3.2. Sparsity

In this section, following the asymptotic notations of Janson (2011), we derive the sparsity properties of our graph model, first for the general construction of Section 2.1, then for the specific construction on compound CRMs of Section 2.2. Similarly to the notations in Caron and Fox (2014), let Z_α be the restriction of Z to the box $[0, \alpha]^2$. Let $(N_\alpha)_{\alpha \geq 0}$ and $(N_\alpha^{(e)})_{\alpha \geq 0}$ be counting processes respectively corresponding to the number of nodes and edges in Z_α :

$$\begin{aligned} N_\alpha &= \text{card}(\{\theta_i \in [0, \alpha] | Z(\{\theta_i\} \times [0, \alpha]) > 0\}) \\ N_\alpha^{(e)} &= Z(\{(x, y) \in \mathbb{R}_+^2 | 0 \leq x \leq y \leq \alpha\}). \end{aligned}$$

Note that in the propositions below, we discard the trivial case $\int_{\mathbb{R}_+^p} \rho(dw_1, \dots, dw_p) = 0$ which implies $N_\alpha^{(e)} = N_\alpha = 0$ a.s.

General construction. The next proposition characterizes the sparsity properties of the random graph depending on the properties of the Lévy measure ρ . In particular, if

$$\int_{\mathbb{R}_+^p} \rho(dw_1, \dots, dw_p) = \infty \quad (23)$$

then, for any $\alpha > 0$, there is a.s. an infinite number of $\theta_i \in [0, \alpha]$ for which $\sum_k w_{ik} > 0$ and the vector of CRMs is called infinite-activity. Otherwise, it is finite-activity.

Proposition 3 Assume that, for any $k = 1, \dots, p$,

$$\int_{\mathbb{R}_+^p} w_k \rho(dw_1, \dots, dw_p) < \infty \quad (24)$$

Then

$$N_\alpha^{(e)} = \begin{cases} \Theta(N_\alpha^2) & \text{if } (W_1, \dots, W_p) \text{ is finite-activity} \\ o(N_\alpha^2) & \text{otherwise} \end{cases}$$

a.s. as α tends to ∞ .

The proof is given in Appendix B.

Construction based on CCRMs. For the CCRM Lévy measure (13), the sparsity properties are solely tuned by the base Lévy measure ρ_0 . Ignoring trivial degenerate cases for the score distribution F , it is easily shown that the CCRM model defined by (5) and (13) is infinite-activity iff the Lévy measure ρ_0 verifies

$$\int_0^\infty \rho_0(dw) = \infty. \quad (25)$$

In this case all CRMs W_0, W_1, \dots, W_p are infinite-activity. Otherwise they are all finite-activity and the vector of CRMs is finite-activity. In the particular case of a CCRM with independent gamma distributed scores (18) and generalized gamma process base measure (19), the condition (25) is satisfied whenever $\sigma \geq 0$. The next proposition characterizes the sparsity of the network depending on the properties of the base Lévy measure ρ_0 .

Proposition 4 *Assume that*

$$\int_0^\infty w \rho_0(dw) < \infty \quad (26)$$

and F is not degenerated at 0. Then

$$N_\alpha^{(e)} = \begin{cases} \Theta(N_\alpha^2) & \text{if } \int_0^\infty \rho_0(dw) < \infty \\ o(N_\alpha^2) & \text{otherwise} \end{cases}$$

a.s. as α tends to ∞ . Furthermore, if the tail Lévy intensity $\bar{\rho}_0$ defined by

$$\bar{\rho}_0(x) = \int_x^\infty \rho_0(dw), \quad (27)$$

is a regularly varying function, i.e.

$$\frac{\bar{\rho}_0(x)}{x^{-\sigma} \ell(1/x)} \longrightarrow 1 \text{ as } x \rightarrow 0$$

for some $\sigma \in (0, 1)$ where ℓ is a slowly varying function verifying $\lim_{t \rightarrow \infty} \ell(at)/\ell(t) = 1$ for any $a > 0$ and $\lim_{t \rightarrow \infty} \ell(t) > 0$, then

$$N_\alpha^{(e)} = O(N_\alpha^{2/(1+\sigma)})$$

a.s. as α tends to ∞ . In the particular case of a CCRM with independent gamma distributed scores (18) and generalized gamma process base measure (19), condition (26) is equivalent to having $\tau > 0$. In this case, we therefore have

$$N_\alpha^{(e)} = \begin{cases} \Theta(N_\alpha^2) & \text{if } \sigma < 0 \\ o(N_\alpha^2) & \text{if } \sigma \geq 0 \\ O(N_\alpha^{2/(1+\sigma)}) & \text{if } \sigma \in (0, 1). \end{cases}$$

The proof is given in Appendix B. Figure 5(a) provides an empirical illustration of Proposition 4 for a CCRM with independent gamma scores and generalized gamma based Lévy measure. Figure 5(b) shows empirically that the degree distribution also exhibits a power-law behaviour when $\sigma \in (0, 1)$.

3.3. Simulation

The point process Z is defined on the plane. We describe in this section how to sample realizations of restrictions Z_α of Z to the box $[0, \alpha]^2$.

General construction. The hierarchical construction given by Eq. (11) suggests a direct way to sample from the model:

1. Sample $(w_{i1}, \dots, w_{ip}, \theta_i)_{i=1,2,\dots}$ from a Poisson process with mean measure $\nu(dw_1, \dots, dw_p, d\theta) \mathbb{1}_{\theta \in [0, \alpha]}$.
2. For each pair of points, sample z_{ij} from (12).




FIG 5. Empirical analysis of the properties of CCRM based graphs generated with parameters $p = 2$, $\tau = 1$, $a_k = 0.2$, $b_k = \frac{1}{p}$ and averaging over various α . (a) Number of edges versus the number of nodes and (b) degree distributions on a log-log scale for various σ : one finite-activity CCRM ($\sigma = -0.5$) and three infinite-activity CCRMs ($\sigma = 0.2$, $\sigma = 0.5$ and $\sigma = 0.8$). In (a) we note growth at a rate $\Theta(N_\alpha^2)$ for $\sigma = -0.5$ and $O(N_\alpha^{2/(1+\sigma)})$ for $\sigma \in (0, 1)$.

There are two caveats to this strategy. First, for infinite-activity CRMs, the number of points in $\mathbb{R}_+^p \times [0, \alpha]$ is almost surely infinite; even for finite-activity CRMs, it may be so large that it is not practically feasible. We need therefore to resort to an approximation, by sampling from a Poisson process with an approximate mean measure $\nu^\varepsilon(dw_1, \dots, dw_p, d\theta) \mathbb{1}_{\theta \in [0, \alpha]} = \rho^\varepsilon(dw_1, \dots, dw_p) \lambda(d\theta) \mathbb{1}_{\theta \in [0, \alpha]}$ where

$$\int_{\mathbb{R}_+^p} \rho^\varepsilon(dw_1, \dots, dw_p) < \infty$$

with $\varepsilon > 0$ controlling the level of approximation. The approximation is specific to the choice of the mean measure, and we describe such an approximation for CCRMs below.

The second caveat is that, for applying Eq. (12), we need to consider all pairs $i \leq j$, which can be computationally problematic. We can instead, similarly to Caron and Fox (2014), use the hierarchical Poisson construction as follows:

1. Sample $(w_{i1}, \dots, w_{ip}, \theta_i)_{i=1,2,\dots,K}$ from a Poisson process with mean measure $\nu^\varepsilon(dw_1, \dots, dw_p, d\theta) \mathbb{1}_{\theta \in [0, \alpha]}$. Let $W_{k,\alpha}^\varepsilon = \sum_{i=1}^K w_{ik} \delta_{\theta_i}$ be the associated truncated CRMs and $W_{k,\alpha}^{\varepsilon*} = \sum_{i=1}^K w_{ik}$ their total masses.
2. For $k = 1, \dots, p$, sample $D_{k,\alpha}^* | W_{k,\alpha}^{\varepsilon*} \sim \text{Poisson}((W_{k,\alpha}^{\varepsilon*})^2)$.
3. For $k = 1, \dots, p$, $\ell = 1, \dots, D_{k,\alpha}^*$, $j = 1, 2$, sample $U_{k\ell j} | W_{k,\alpha}^\varepsilon \stackrel{\text{ind}}{\sim} \frac{W_{k,\alpha}^\varepsilon}{W_{k,\alpha}^{\varepsilon*}}$.
4. Set $D_{k,\alpha}^\varepsilon = \sum_{\ell=1}^{D_{k,\alpha}^*} \delta_{U_{k\ell 1}, k\ell 2}$.
5. Obtain Z from (D_1, \dots, D_p) as in (11).

Construction based on CCRMs. The hierarchical construction of compound CRMs suggests an algorithm to simulate a vector of CRMs. We consider the following (truncated) mean measure

$$\rho^\varepsilon(dw_1, \dots, dw_p) = e^{-\sum_{k=1}^p \gamma_k w_k} \int_\varepsilon^\infty w_0^{-p} F\left(\frac{dw_1}{w_0}, \dots, \frac{dw_p}{w_0}\right) \rho_0(dw_0) \quad (28)$$

with $\varepsilon \geq 0$. We can sample from the (truncated) CCRM as follows

1. (a) Sample $(w_{i0}, \theta_i)_{i=1,\dots,K}$ from a Poisson point process with mean measure $\tilde{\rho}_0(dw_0) \lambda(d\theta) \mathbb{1}_{\{w_0 > \varepsilon, \theta \in [0, \alpha]\}}$.
- (b) For $i = 1, \dots, K$ and $k = 1, \dots, p$, set $w_{ik} = \beta_{ik} w_{i0}$ where $(\beta_{i1}, \dots, \beta_{ip}) | w_{i0}$ is drawn from (15).

The truncation level ε is set to 0 for finite-activity CCRMs, and $\varepsilon > 0$ otherwise. We explain in Appendix C how to perform step 1.(a) in the case of a tilted generalized gamma process.

4. Posterior inference

In this section, we describe a MCMC algorithm for posterior inference of the model parameters and hyperparameters in the statistical network model defined in Section 2. We first describe the data augmentation scheme and characterization of conditionals. We then describe the sampler for a general Lévy measure ρ , and finally derive the sampler for compound CRMs.

4.1. Characterization of conditionals and data augmentation

Assume that we have observed a set of connections $(z_{ij})_{1 \leq i, j \leq N_\alpha}$, where N_α is the number of nodes with at least one connection. We aim at inferring the positive parameters $(w_{i1}, \dots, w_{ip})_{i=1, \dots, N_\alpha}$ associated to the nodes with at least one connection. We also want to estimate the positive parameters associated to the other nodes with no connection. The number of such nodes may be large, and even infinite for infinite-activity CRMs; but under our model, these parameters are only identifiable through their sum, denoted (w_{*1}, \dots, w_{*p}) . Note that the node locations θ_i are not likelihood identifiable, and we will not try to infer them. We assume that there is a set of unknown hyperparameters ϕ of the mean intensity ρ , with prior $p(\phi)$. We assume that the Lévy measure ρ is absolutely continuous with respect to the Lebesgue measure on \mathbb{R}^d , and write simply $\rho(dw_1, \dots, dw_p; \phi) = \rho(w_1, \dots, w_p; \phi)dw_1 \dots dw_p$. The parameter α is also assumed to be unknown, with some prior $\alpha \sim \text{Gamma}(a_\alpha, b_\alpha)$ with $a_\alpha > 0, b_\alpha > 0$. We therefore aim at approximating $p((w_{1k}, \dots, w_{N_\alpha k}, w_{*k})_{k=1, \dots, p}, \phi, \alpha | (z_{ij})_{1 \leq i, j \leq N_\alpha})$.

As a first step, we characterize the conditional distribution of the restricted vector of CRMs $(W_{1\alpha}, \dots, W_{p\alpha})$ given the restricted measures $(D_{1\alpha}, \dots, D_{p\alpha})$. Proposition 5 below extends Theorem 12 in Caron and Fox (2014) to the multivariate setting.

Proposition 5 *Let $(\theta_1, \dots, \theta_{N_\alpha})$, $N_\alpha \geq 0$ be the support points of $(D_{1\alpha}, \dots, D_{p\alpha})$, with*

$$D_{k\alpha} = \sum_{1 \leq i, j \leq N_\alpha} n_{ijk} \delta_{(\theta_i, \theta_j)}.$$

The conditional distribution of $(W_{1\alpha}, \dots, W_{p\alpha})$ given $(D_{1\alpha}, \dots, D_{p\alpha})$ is equivalent to the distribution of

$$\left(\widetilde{W}_1 + \sum_{i=1}^{N_\alpha} w_{i1} \delta_{\theta_i}, \dots, \widetilde{W}_p + \sum_{i=1}^{N_\alpha} w_{ip} \delta_{\theta_i} \right) \quad (29)$$

*where $(\widetilde{W}_1, \dots, \widetilde{W}_p)$ is a vector of discrete random measures, which depends on $(D_{1\alpha}, \dots, D_{p\alpha})$ only through the total masses $w_{*k} = \widetilde{W}_k([0, \alpha])$.*

*The set of weights $(w_{ik})_{i=1, \dots, N_\alpha; k=1, \dots, p}$ and $(w_{*k})_{k=1, \dots, p}$ are dependent, with joint conditional distribution*

$$p((w_{1k}, \dots, w_{N_\alpha k}, w_{*k})_{k=1, \dots, p} | (n_{ijk})_{1 \leq i, j \leq N_\alpha; k=1, \dots, p}, \phi, \alpha) \\ \propto \left[\prod_{i=1}^{N_\alpha} \prod_{k=1}^p w_{ik}^{m_{ik}} \right] e^{-\sum_{k=1}^p (w_{*k} + \sum_{i=1}^{N_\alpha} w_{ik})^2} \left[\prod_{i=1}^{N_\alpha} \rho(w_{i1}, \dots, w_{ip}; \phi) \right] \alpha^{N_\alpha} g_{*\alpha}(w_{*1}, \dots, w_{*p}; \phi) \quad (30)$$

where $m_{ik} = \sum_{j=1}^{N_\alpha} n_{ijk} + n_{jik}$ and $g_{\alpha}(w_{*1}, \dots, w_{*p}; \phi)$ is the probability density function of the random vector $(W_1([0, \alpha]), \dots, W_p([0, \alpha]))$.*

The proof can be straightforwardly adapted from that of Caron and Fox (2014), or from Proposition 5.2 of James (2014) and is omitted here. It builds on other posterior characterizations in Bayesian nonparametric models (Prünster, 2002; James, 2002, 2005; James et al., 2009).

Data augmentation. Similarly to Caron and Fox (2014), we introduce latent count variables $\tilde{n}_{ijk} = n_{ijk} + n_{jik}$ with

$$(\tilde{n}_{ij1}, \dots, \tilde{n}_{ijp}) | w, z \sim \begin{cases} \delta_{(0, \dots, 0)} & \text{if } z_{ij} = 0 \\ \text{tPoisson}(2w_{i1}w_{j1}, \dots, 2w_{ip}w_{jp}) & \text{if } z_{ij} = 1, i \neq j \\ \left(\frac{\tilde{n}_{ij1}}{2}, \dots, \frac{\tilde{n}_{ijp}}{2} \right) & \text{if } z_{ij} = 1, i = j \end{cases} \quad (31)$$

where $\text{tPoisson}(\lambda_1, \dots, \lambda_p)$ is the multivariate Poisson distribution truncated at zero, whose pmf is

$$\text{tPoisson}(x_1, \dots, x_p; \lambda_1, \dots, \lambda_p) = \frac{\prod_{k=1}^p \text{Poisson}(x_k; \lambda_k)}{1 - \exp(-\sum_{k=1}^p x_k \lambda_k)} \mathbb{1}_{\{\sum_{k=1}^p x_k > 0\}}.$$

One can sample from this distribution by first sampling $x = \sum_{k=1}^p x_k$ from a zero-truncated Poisson distribution with rate $\sum_{k=1}^p \lambda_k$, and then $(x_1, \dots, x_p) | (\lambda_1, \dots, \lambda_p), x \sim \text{Multinomial}\left(x, \left(\frac{\lambda_1}{\sum_{k=1}^p \lambda_k}, \dots, \frac{\lambda_p}{\sum_{k=1}^p \lambda_k}\right)\right)$.

4.2. Markov chain Monte Carlo algorithm: General construction

Using the data augmentation scheme together with the posterior characterization (30), we can derive the following MCMC sampler, which uses Metropolis-Hastings (MH) and Hamiltonian Monte Carlo (HMC) updates within a Gibbs sampler, and iterates as described in Algorithm 1.

Algorithm 1 Markov chain Monte Carlo sampler for posterior inference.

At each iteration

1. Update (w_{i1}, \dots, w_{ip}) , $i = 1, \dots, N_\alpha$ given the rest using MH or HMC.
 2. Update hyperparameters (ϕ, α) and total masses (w_{*1}, \dots, w_{*p}) given the rest using MH.
 3. Update the latent variables given the rest using (31).
-

In general, if the Lévy intensity ρ can be evaluated pointwise, one can use a MH update for step 1, but it would scale poorly with the number of nodes. Alternatively, if the Lévy intensity ρ is differentiable, one can use a Hamiltonian Monte Carlo update (Duane et al., 1987; Neal, 2011).

The challenging part of the Algorithm 1 is Step 2. From Eq. (30) we have

$$p((w_{*k})_{k=1, \dots, p}, \phi, \alpha | \text{rest}) \propto p(\phi) p(\alpha) e^{-\sum_{k=1}^p (w_{*k} + \sum_{i=1}^{N_\alpha} w_{ik})^2} \left[\prod_{i=1}^{N_\alpha} \rho(w_{i1}, \dots, w_{ip}; \phi) \right] \alpha^{N_\alpha} g_{*\alpha}(w_{*1}, \dots, w_{*p}; \phi).$$

This conditional distribution is not of standard form and involves the multivariate pdf $g_{*\alpha}(w_{*1}, \dots, w_{*p})$ of the random vector $(W_1([0, \alpha]), \dots, W_p([0, \alpha]))$ for which there is typically no analytical expression. All is available is its Laplace transform, which is given by

$$\mathbb{E} \left[e^{-\sum_{k=1}^p t_k W_k([0, \alpha])} \right] = e^{-\alpha \psi(t_1, \dots, t_p; \phi)} \quad (32)$$

where

$$\psi(t_1, \dots, t_p; \phi) = \int_{\mathbb{R}_+^p} \left(1 - e^{-\sum_{k=1}^p t_k w_k} \right) \rho(dw_1, \dots, dw_p; \phi) \quad (33)$$

is the multivariate Laplace exponent, which involves a p -dimensional integral. We propose to use a Metropolis-Hastings step, with proposal

$$q(\tilde{w}_{*1:p}, \tilde{\phi}, \tilde{\alpha} | w_{*1:p}, \phi, \alpha) = q(\tilde{w}_{*1:p} | w_{*1:p}, \tilde{\phi}, \tilde{\alpha}) \times q(\tilde{\phi} | \phi) \times q(\tilde{\alpha} | \alpha, \tilde{\phi}, w_{*1:p}) \quad (34)$$

where

$$q(\tilde{\alpha} | \alpha, \tilde{\phi}, w_{*1:p}) = \text{Gamma}(\tilde{\alpha}; a_\alpha + N_\alpha, b_\alpha + \psi(\lambda_1, \dots, \lambda_p; \tilde{\phi})) \quad (35)$$

and the proposal for $w_{*1:p}$ is an exponentially tilted version of $g_{*\alpha}$

$$q((\tilde{w}_{*k})_{k=1, \dots, p} | (w_{*k})_{k=1, \dots, p}, \tilde{\phi}) = \frac{e^{-\sum_{k=1}^p \lambda_k \tilde{w}_{*k}} g_{*\tilde{\alpha}}(\tilde{w}_1, \dots, \tilde{w}_p; \tilde{\phi})}{e^{-\tilde{\alpha} \psi(\lambda_1, \dots, \lambda_p; \tilde{\phi})}} \quad (36)$$

where $\lambda_k = w_{*k} + 2 \sum_{i=1}^{N_\alpha} w_{ik}$ and $q(\tilde{\phi}|\phi)$ can be freely specified by the user. This leads to the following acceptance ratio

$$r = \frac{p(\tilde{\phi})q(\phi|\tilde{\phi})}{p(\phi)q(\tilde{\phi}|\phi)} \frac{\rho(\tilde{w}_1, \dots, \tilde{w}_p; \tilde{\phi})}{\rho(w_1, \dots, w_p; \phi)} \left(\frac{b_\alpha + \psi(\tilde{\lambda}_1, \dots, \tilde{\lambda}_p; \tilde{\phi})}{b_\alpha + \psi(\lambda_1, \dots, \lambda_p; \phi)} \right)^{a_\alpha + N_\alpha} e^{\sum_{k=1}^p [w_{*k}^2 - \tilde{w}_{*k}^2]}$$

where $\tilde{\lambda}_k = \tilde{w}_{*k} + 2 \sum_{i=1}^{N_\alpha} w_{ik}$. This acceptance ratio involves evaluating the multivariate exponent (33).

In the general case, the MCMC algorithm 1 thus requires to be able to

- (a) evaluate pointwise the Lévy intensity ρ , and potentially differentiate it,
- (b) evaluate pointwise the Laplace exponent (33) and
- (c) sample from the exponentially tilted distribution (36).

Regarding point (c), the random variable with pdf (36) has the same distribution as the random vector $(W'_1([0, \alpha]), \dots, W'_p([0, \alpha]))$ where $(W'_1, \dots, W'_p) \sim \text{CRM}(\rho', \lambda)$ with ρ' is an exponentially tilted version of ρ

$$\rho'(w_1, \dots, w_p) = e^{-\sum_k \lambda_k w_k} \rho(w_1, \dots, w_p). \quad (37)$$

By considering an approximate tilted intensity $\rho^{\varepsilon'}(w_1, \dots, w_p)$, one can approximately sample from (36) by simulating points from a Poisson process with mean measure $\alpha \rho^{\varepsilon'}(w_1, \dots, w_p)$ and summing them up.

4.3. Markov chain Monte Carlo algorithm: Compound CRMs

The hierarchical construction of CCRMs enables to derive a certain number of simplifications in the algorithm described in the previous section. Using the construction $w_{ik} = \beta_{ik} w_{i0}$ where the points $(w_{i0}, \beta_{i1}, \dots, \beta_{ip})_{i=1,2,\dots}$ have Lévy measure (17), we aim at approximating the posterior

$$p((w_{10}, \dots, w_{N_\alpha 0}), (\beta_{1k}, \dots, \beta_{N_\alpha k}, w_{*k})_{k=1,\dots,p}, \phi, \alpha | (z_{ij})_{1 \leq i, j \leq N_\alpha}). \quad (38)$$

Conditional on the latent count variables defined in (31), we have the following conditional characterization, similar to (30)

$$\begin{aligned} & p((w_{10}, \dots, w_{N_\alpha 0}), (\beta_{1k}, \dots, \beta_{N_\alpha k}, w_{*k})_{k=1,\dots,p} | (n_{ijk})_{1 \leq i, j \leq N_\alpha; k=1,\dots,p}, \phi, \alpha) \\ & \propto \left[\prod_{i=1}^{N_\alpha} w_{i0}^{m_i} \prod_{k=1}^p \beta_{ik}^{m_{ik}} \right] e^{-\sum_{k=1}^p (w_{*k} + \sum_{i=1}^{N_\alpha} w_{ik})^2 - \sum_{i=1}^{N_\alpha} w_{i0} (\sum_{k=1}^p \gamma_k \beta_{ik})} \\ & \times \left[\prod_{i=1}^{N_\alpha} f(\beta_{i1}, \dots, \beta_{ip}; \phi) \rho_0(w_{i0}; \phi) \right] \alpha^{N_\alpha} g_{*\alpha}(w_{*1}, \dots, w_{*p}; \phi) \end{aligned} \quad (39)$$

where $m_i = \sum_{k=1}^p m_{ik}$ and f and ρ_0 are densities of F and ρ_0 with respect to the Lebesgue measure.

If f and ρ_0 are differentiable, one can use a HMC update for Step 1 of Algorithm 1. In particular, when they take the form (18) and (19), we obtain the following simple expressions for the gradient:

$$\begin{aligned} \frac{\partial U(q)}{d(\log w_{i0})} &= m_i - \sigma - w_{i0} \left[\tau + 2 \sum_{k=1}^p \beta_{ik} \left(w_{*k} + \sum_{j=1}^{N_\alpha} w_{j0} \beta_{jk} \right) \right], \quad i = 1, \dots, N_\alpha, \\ \frac{\partial U(q)}{d(\log \beta_{ik})} &= m_{ik} + a_k - \beta_{ik} \left[b_k + 2w_{i0} \left(w_{*k} + \sum_{j=1}^{N_\alpha} w_{j0} \beta_{jk} \right) \right], \quad i = 1, \dots, N_\alpha, \quad k = 1, \dots, p, \end{aligned}$$

where $U(q) = \log p(q|\text{rest})$ with $q = (\log w_{i0}, \log \beta_{i1}, \dots, \log \beta_{ip})_{i=1,\dots,N_\alpha}$.

Regarding Step 2 of Algorithm 1, the Laplace exponent for CCRM takes the simple form

$$\psi(t_1, \dots, t_p) = \int_0^\infty [M(-w_0\gamma_1, \dots, -w_0\gamma_p) - M(-w_0(t_1 + \gamma_1), \dots, -w_0(t_p + \gamma_p))] \rho_0(dw_0) \quad (40)$$

which only requires evaluating a one-dimensional integral, whatever the number p of communities, and this can be done numerically. For the specific model defined by (18) and (19), we obtain

$$\psi(t_1, \dots, t_p) = \frac{1}{\Gamma(1-\sigma)} \int_0^\infty \left[1 - \prod_{k=1}^p \left(1 + \frac{w_0 t_k}{b_k + w_0 \gamma_k} \right)^{-a_k} \right] \left[\prod_{k=1}^p \left(1 + \frac{w_0 \gamma_k}{b_k} \right)^{-a_k} \right] w_0^{-1-\sigma} e^{-w_0 \tau} dw_0.$$

Finally, we need to sample total masses (w_{*1}, \dots, w_{*p}) from (36), and this can be done by simulating points $(w_{i0}, \beta_{i1}, \dots, \beta_{ip})_{i=1,2,\dots}$ from a Poisson process with exponentially tilted Lévy intensity

$$\alpha e^{-w_0 \sum_{k=1}^p (\gamma_k + \lambda_k) \beta_k} f(\beta_1, \dots, \beta_p) \rho_0(w_0) \quad (41)$$

and summing up the weights $w_{*k} = \sum_{i=1,2,\dots} w_{i0} \beta_{ik}$ for $k = 1, \dots, p$. For infinite-activity CRMs, this is not feasible, and we suggest to resort to the approximation of Cohen and Rosinski (2007). More precisely, we write

$$(w_{*1}, \dots, w_{*p}) = X_\varepsilon + X^\varepsilon$$

where the random vectors $X_\varepsilon \in \mathbb{R}_+^p$ and $X^\varepsilon \in \mathbb{R}_+^p$ are defined as $X_\varepsilon = \sum_{i|w_{i0} < \varepsilon} w_{i0}(\beta_{i1}, \dots, \beta_{ip})$ and $X^\varepsilon = \sum_{i|w_{i0} > \varepsilon} w_{i0}(\beta_{i1}, \dots, \beta_{ip})$. We can sample a realization of the random vector X^ε exactly by simulating the points of a Poisson process with mean intensity

$$\alpha e^{-w_0 \sum_{k=1}^p (\gamma_k + \lambda_k) \beta_k} f(\beta_1, \dots, \beta_p) \rho_0(w_0) \mathbf{1}_{w_0 > \varepsilon} \quad (42)$$

See Section 3.3 and Appendix C for details. The positive random vector X_ε is approximated by a truncated Gaussian random vector with mean μ_ε and variance Σ_ε such that

$$\begin{aligned} \mu_\varepsilon &= \alpha \int_{\mathbb{R}_+^p} w_{1:p} \rho_\varepsilon(dw_1, \dots, dw_p) \\ \Sigma_\varepsilon &= \alpha \int_{\mathbb{R}_+^p} w_{1:p} w_{1:p}^T \rho_\varepsilon(dw_1, \dots, dw_p) \end{aligned}$$

where

$$\rho_\varepsilon(dw_1, \dots, dw_p) = e^{-\sum_{k=1}^p (\gamma_k + \lambda_k) w_k} \int_0^\varepsilon w_0^{-p} F\left(\frac{dw_1}{w_0}, \dots, \frac{dw_p}{w_0}\right) \rho_0(dw_0).$$

Note that μ_ε and Σ_ε can both be expressed as one-dimensional integrals using the gradient and Hessian of the moment generating function M of F . Theorem 6 in Appendix E, which is an adaptation of the results of Cohen and Rosinski (2007) to CCRM, gives the conditions on the parameters of CCRM under which

$$\Sigma_\varepsilon^{-1/2} (X_\varepsilon - \mu_\varepsilon) \xrightarrow{d} \mathcal{N}(0, I_p) \text{ as } \varepsilon \rightarrow 0$$

and thus the approximation is asymptotically valid. The Gaussian approximation is in particular asymptotically valid for the CCRM defined by (18) and (19) when $\sigma \in (0, 1)$, hence is valid for all infinite-activity cases except $\sigma = 0$.

Note that due to the Gaussian approximation in the proposal distribution for $(w_{*\alpha})$, Algorithm 1 does not actually admit the posterior distribution (38) as invariant distribution, and is an approximation of an exact MCMC algorithm targeting this distribution. We observe in the experimental section that this approximation provides very reasonable results for the examples considered.

5. Experiments

5.1. Simulated data

We first study the convergence of the MCMC algorithm on synthetic data simulated from the CCRM based graph model described in Section 2 where F and ρ_0 take the form (18) and (19). We generate an undirected graph with $p = 2$ communities and parameters $\alpha = 200$, $\sigma = 0.2$, $\tau = 1$, $b_k = b = \frac{1}{p}$, $a_k = a = 0.2$ and $\gamma_k = \gamma = 0$. The sampled graph has 1121 nodes and 12,180 edges. For the inference, we consider that b and γ are known and we assume a vague prior $\text{Gamma}(0.01, 0.01)$ on the unknown parameters α and $\phi = (1 - \sigma, \tau, a)$. We run 3 parallel MCMC chains with different initial values. Each chain starts with 10,000 iterations using our model with only one community where the scores β are fixed to 1, which is equivalent to the model of Caron and Fox (2014). We then run 200,000 iterations using our model with p communities. We use $\varepsilon = 10^{-3}$ as a truncation level for simulating $w_{*1:p}$ and $L = 10$ leapfrog steps for the HMC. The stepsizes of both the HMC and the random walk MH on $(\log(1 - \sigma), \log \tau, \log a)$ are adapted during the first 50,000 iterations so as to target acceptance ratios of 0.65 and 0.23 respectively. The computations take around 5h20 using Matlab on a standard desktop computer. Trace plots of the parameters $\log \alpha$, σ , τ , a and $\bar{w}_* = \frac{1}{p} \sum_{k=1}^p w_{*k}$ and histograms based on the last 50,000 iterations are given in Figure 6. Posterior samples clearly converge around the sampled value. Choosing a threshold value $\epsilon \ll 10^{-3}$ does not lead to any noticeable change in the MCMC histograms, suggesting that the target distribution of our approximate MCMC is very close to the posterior distribution of interest.

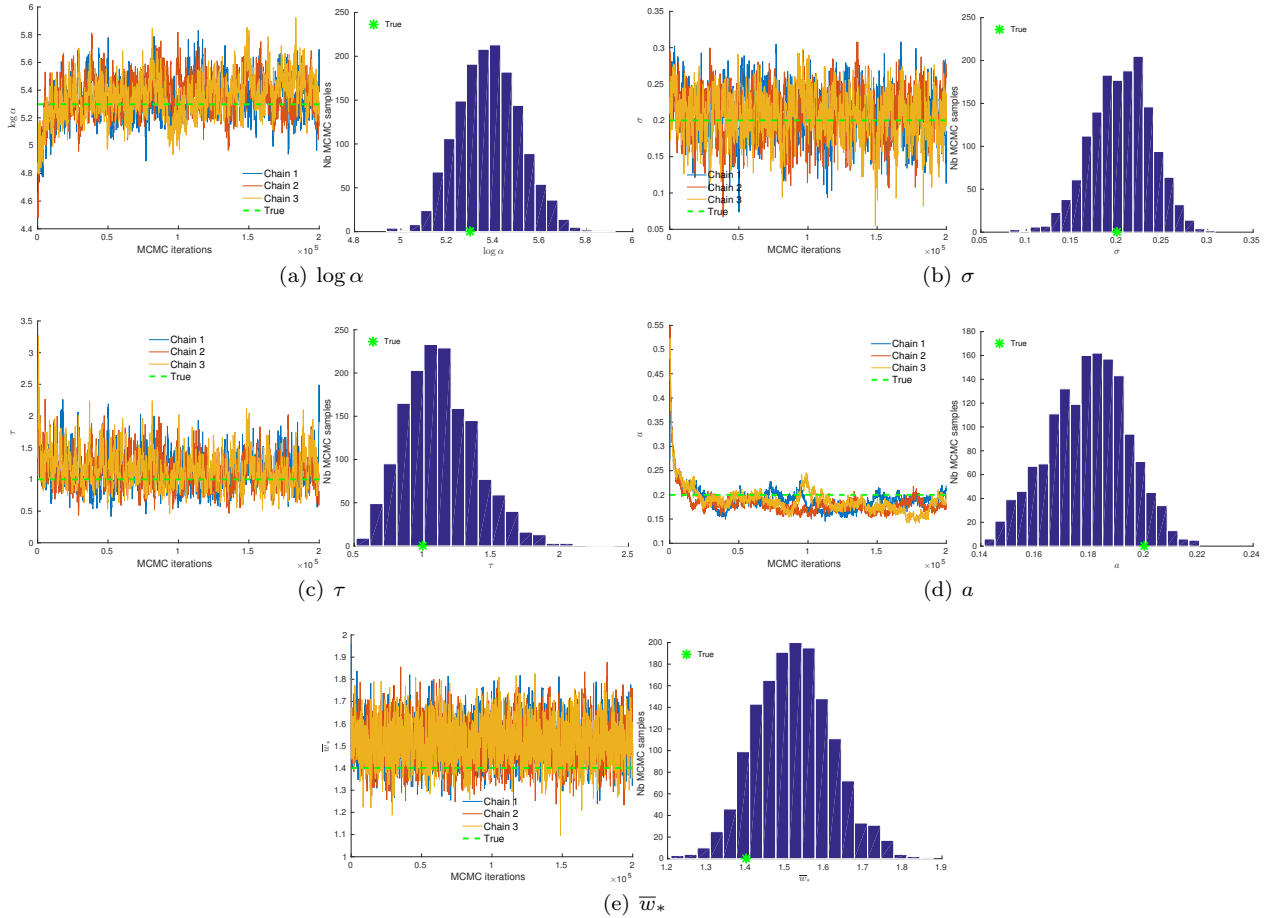


FIG 6. MCMC trace plots (left) and histograms (right) of parameters (a) $\log \alpha$, (b) σ , (c) τ , (d) a and (e) \bar{w}_* for a graph generated with parameters $p = 2$, $\alpha = 200$, $\sigma = 0.2$, $\tau = 1$, $b = \frac{1}{p}$, $a = 0.2$ and $\gamma = 0$.

Our model is able to accurately recover the mean parameters of both low and high degree nodes and

to provide reasonable credible intervals, as shown in Figure 7(a-b). By generating 5000 graphs from the posterior predictive we assess that our model fits the empirical power-law degree distribution of the sparse generated graph as shown in Figure 7(c). We demonstrate the interest of our nonparametric approach by comparing these results to the ones obtained with the parametric version of our model. To achieve this, we fix $w_{*k} = 0$ and force the model to lie in the finite-activity domain by assuming $\sigma \in (-\infty, 0)$ and using the prior distribution $-\sigma \sim \text{Gamma}(0.01, 0.01)$. Note that in this case, the model is equivalent to that of Zhou (2015). As shown in Figure 8(a-b), the parametric model is able to recover the mean parameters of nodes with high degrees, and credible intervals are similar to that obtained with the full model; however, it fails to provide reasonable credible intervals for nodes with low degree. In addition, as shown in Figure 8(c), the posterior predictive degree distribution does not fit the data, illustrating the inability of this parametric model to capture power-law behaviour.

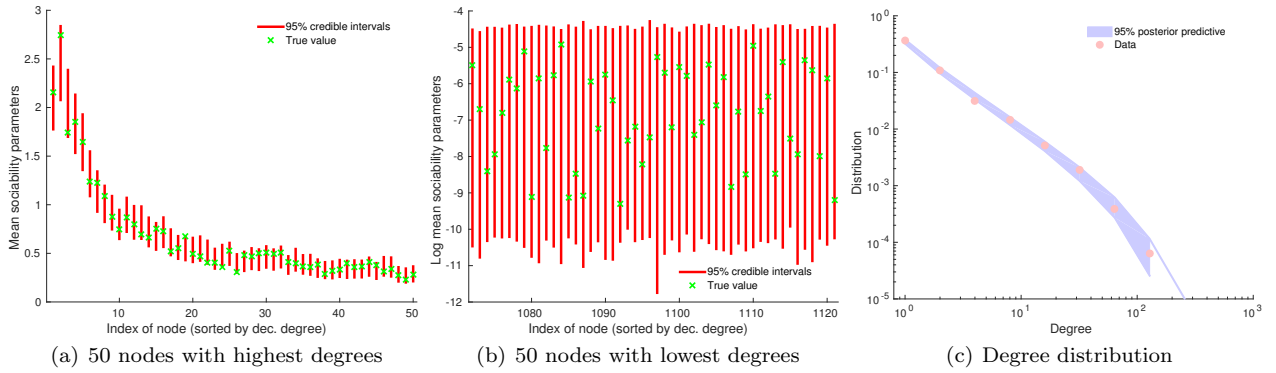


FIG 7. 95% posterior credible intervals and true values of (a) the mean parameters $\bar{w}_i = \frac{1}{p} \sum_{k=1}^p w_{ik}$ of the 50 nodes with highest degrees and (b) the log mean parameters $\log \bar{w}_i$ of the 50 nodes with lowest degrees. (c) Empirical degree distribution and 95% posterior predictive credible interval. Results obtained for a graph generated with parameters $p = 2$, $\alpha = 200$, $\sigma = 0.2$, $\tau = 1$, $b = \frac{1}{p}$, $a = 0.2$ and $\gamma = 0$.




FIG 8. See Figure 7. Results obtained on the same generated graph by inferring a finite-activity model with $w_{*k} = 0$ and $\sigma \leq 0$.

5.2. Real-world graphs

We now apply our methods to learn the latent communities of two real-world undirected simple graphs. The first network to be considered, the **polblogs** network (Adamic and Glance, 2005), is the network of the American political blogosphere in February 2005³. Two blogs are considered as connected if there is at

³<http://www.cise.ufl.edu/research/sparse/matrices/Newman/polblogs>

least one hyperlink from one blog to the other. Additional information on the political leaning of each blog (left/right) is also available. The second network, named **USairport**, is the network of connections between US airports in 2010⁴.

TABLE 1
Size of the networks, number of communities and computational time.

Name	Nb nodes	Nb edges	Nb communities p	Time
polblogs	1224	33,430	2	50m
USairport	1867	44,304	4	5h20m

The sizes of the different networks are given in Table 1. We consider $\gamma_k = 0$ is known and we assume a vague prior Gamma(0.01, 0.01) on the unknown parameters α , $1 - \sigma$, τ , a_k and b_k . We take $p = 2$ communities for **polblogs** and $p = 4$ communities for **USairport**. We run 3 parallel MCMC chains, each with 10,000 + 200,000 iterations, using the same procedure as used for the simulated data; see Section 5.1. Computation times are reported in Table 1. The simulation of $w_{*1:p}$ requires more computational time when $\sigma \geq 0$ (infinite-activity case). This explain the larger computation times for **USairport** compared to **polblogs**.




FIG 9. Empirical degree distribution (red) and posterior predictive (blue) of the (a) *polblogs* and (b) *USairport* networks.

We interpret the communities based on the minimum Bayes risk point estimate where the cost function is a permutation-invariant absolute loss on the weights $w = (w_{ik})_{i=1,\dots,N_\alpha;k=1,\dots,p}$. Let \mathcal{S}_p be the set of permutations of $\{1, \dots, p\}$ and consider the cost function

$$C(w, w^*) = \min_{\pi \in \mathcal{S}_p} \left[\sum_{k=1}^p \sum_{i=1}^{N_\alpha} |w_{i\pi(k)} - w_{ik}^*| + \sum_{k=1}^p |w_{*\pi(k)} - w_{*k}^*| \right]$$

whose evaluation requires solving a combinatorial optimization problem in $O(p^3)$ using the Hungarian method. We therefore want to solve

$$\hat{w} = \arg \min_{w^*} \mathbb{E}[C(w, w^*) | Z]$$

where $\mathbb{E}[C(w, w^*) | Z] \simeq \frac{1}{N} \sum_{t=1}^N C(w^{(t)}, w^*)$ and $(w^{(t)})_{t=1,\dots,N}$ are from the MCMC output. For simplicity, we limit the search of \hat{w} to the set of MCMC samples giving

$$\hat{w} = \arg \min_{w^* \in \{w^{(1)}, \dots, w^{(N)}\}} \frac{1}{N} \sum_{t=1}^N C(w^{(t)}, w^*).$$

Table 2 reports the nodes with highest weights in each community for the **polblogs** network. Figure 10 also shows the weight associated to each of the two community alongside the true left/right class for each

⁴http://www.transtats.bts.gov/DL_SelectFields.asp?Table_ID=292

TABLE 2

Nodes with highest weight in each community for the *polblogs* network. Blog URLs are followed by known political leaning: (L) for left-wing and (R) for right-wing.

Community 1: “Liberal”	Community 2: “Conservative”
dailykos.com (L)	instapundit.com (R)
atrios.blogspot.com (L)	blogsforbush.com (R)
talkingpointsmemo.com (L)	powerlineblog.com (R)
washingtonmonthly.com (L)	drudgereport.com (R)
liberaloasis.com (L)	littlegreenfootballs.com/weblog (R)
talkleft.com (L)	michellemalkin.com (R)
digbysblog.blogspot.com (L)	lashawnbarber.com (R)
newleftblogs.blogspot.com (L)	wizbangblog.com (R)
politicalstrategy.org (L)	hughhewitt.com (R)
juancole.com (L)	truthlaidbear.com (R)

TABLE 3

Nodes with highest weights in each community for the *USAirport* network.

Community 1: “Hub”	Community 2: “East”	Community 3: “West”	Community 4: “Alaska”
Miami, FL	Cleveland, OH	Denver, CO	Anchorage, AK
New York, NY	Detroit, MI	Las Vegas, NV	Fairbanks, AK
Newark, NJ	Nashville, TN	Los Angeles, CA	Bethel, AK
Los Angeles, CA	Chicago, IL	Salt Lake City, UT	St. Mary’s, AK
Atlanta, GA	Knoxville, TN	Seattle, WA	King Salmon, AK
Washington, DC	Atlanta, GA	Burbank, CA	McGrath, AK
Chicago, IL	Louisville, KY	Phoenix, AZ	Unalakleet, AK
Boston, MA	Indianapolis, IN	Oakland, CA	Galena, AK
Houston, TX	Memphis, TN	Portland, OR	Aniak, AK
Orlando, FL	Charlotte, NC	Albuquerque, NM	Kotzebue, AK

blog. The two learned communities, which can be interpreted as “Liberal” and “Conservative”, clearly recover the political leaning of the blogs. Figure 12 shows the adjacency matrices obtained by reordering the nodes by community membership, where each node is assigned to the community whose weight is maximum, clearly showing the block-structure of this network. The obtained clustering yields a 93.95% accuracy when compared to the ground truth classification. Figure 11(a) shows the relative community proportions for a subset of the blogs. *dailykos.com* and *washingtonmonthly.com* are clearly described as liberal while *blogsforbush.com*, *instapundit.com* and *drudgereport.com* are clearly conservative. Other more moderate blogs such as *danieldrezner.com/blog* and *andrewsullivan.com* have more balanced values in both communities. Figure 9(a) shows that the posterior predictive degree distribution provides a good fit to the data.

For *USAirport*, the four learned communities can also be easily interpreted, as seen in Table 3. The first community, labeled “Hub”, represents highly connected airports with no preferred location, while the three others, labeled “East”, “West” and “Alaska”, are communities based on the location of the airport. In Figure 11(b), we can see that some airports have a strong level of affiliation in a single community: New York and Miami for “Hub”, Lansing for “East”, Seattle for “West” and Bethel and Anchorage for “Alaska”. Other airports have significant weights in different communities: Raleigh/Durham and Los Angeles are hubs with strong regional connections, Nashville and Minneapolis share a significant number of connections with both East and West of the USA. Anchorage has a significant “Hub” weight, while most airports in Alaska are disconnected from the rest of the world as can be seen in Figure 12(b). “Alaska” appears as a separate block while substantial overlaps are observed between the “Hub”, “East” and “West” communities. Figure 9(b) shows that the posterior predictive degree distribution also provides a good fit to the data.

Acknowledgments

The authors thank George Deligiannidis for pointing out the article of [Asmussen and Rosiński \(2001\)](#). FC acknowledges the support of the European Commission under the Marie Curie Intra-European Fellowship Programme. Part of this work has been supported by the BNPSI ANR project no ANR-13-BS-03-0006-01.




FIG 10. Level of affiliation of each blog of the polblogs network to the communities identified as “Liberal” (top) and “Conservative” (bottom). The names of the blogs are grouped according to the left-right wing ground truth. Left-wing blogs are represented in blue on the left, right-wing blogs in red on the right.




FIG 11. Relative values of the weights in each community for a subset of the nodes of the (a) polblogs and (b) USairport networks.

Appendix A: Background on completely random measures

A.1. Completely random measures

Completely random measures (CRM) were introduced by Kingman (1967, 1993) and are now standard tools for constructing flexible Bayesian nonparametric (BNP) models; see for example the surveys of Lijoi and Prünster (2010) or Daley and Vere-Jones (2008b, Section 10.1).

A CRM W on \mathbb{R}_+ is a random measure such that, for any collection of disjoint measurable subsets A_1, \dots, A_n of \mathbb{R}_+ , $W(A_1), \dots, W(A_n)$ are independent. A CRM can be decomposed into a sum of three independent parts: a non-random measure, a countable collection of atoms with random masses at fixed locations, and a countable collection of atoms with random masses and random locations. Here, we will only consider CRMs with random masses and random locations, which take the form

$$W = \sum_{i=1}^{\infty} w_i \delta_{\theta_i} \tag{43}$$

where the $w_i \in \mathbb{R}_+$ are the masses and $\theta_i \in \mathbb{R}_+$ are the locations. The law of W can actually be characterized by a Poisson point process $N = \{(w_i, \theta_i)_{i=1,2,\dots}\}$ on \mathbb{R}_+^2 with mean measure $\nu(dw, d\theta)$. We focus here on

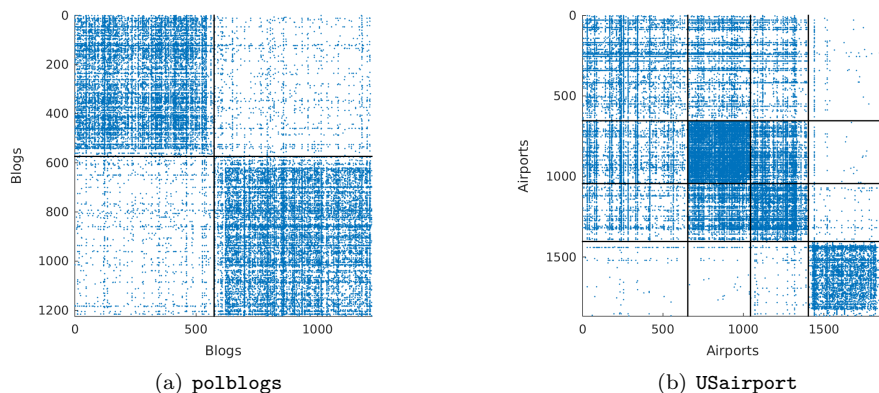


FIG 12. Adjacency matrices of the (a) *polblogs* and (b) *USairport* networks, reordered by associating each node to the community where it has the highest weight.

the case where the CRM is homogeneous with independent increments. This implies that the location θ_i are independent of the weights w_i and the mean measure decomposes as $\nu(dw, d\theta) = \rho(dw)\lambda(d\theta)$ where λ is the Lebesgue measure and ρ is a measure on \mathbb{R}_+ such that

$$\int_0^\infty (1 - e^{-w})\rho(dw) < \infty. \quad (44)$$

We write $W \sim \text{CRM}(\rho, \lambda)$. Note that $W([0, T]) < \infty$ a.s. for any real T while $W(\mathbb{R}_+) = \infty$ a.s. if ρ is not degenerate at 0. If

$$\int_0^\infty \rho(dw) = \infty \quad (45)$$

then there will be a.s. an infinite number of jumps in any interval $[0, T]$ and we refer to the CRM as infinite-activity. Otherwise, it is called finite activity. Let $\bar{\rho}$ be the tail Lévy intensity defined as

$$\bar{\rho}(x) = \int_x^\infty \rho(dw) \quad (46)$$

for $x > 0$. This function corresponds to the expected number of points (w_i, θ_i) such that $w_i > x$ and $\theta_i \in [0, 1]$, and its asymptotic properties play an important role in the characterization of the graph properties.

A.2. Vectors of CRMs

Multivariate extensions of CRMs have been proposed recently by various authors (Epifani and Lijoi, 2010; Leisen and Lijoi, 2011; Leisen et al., 2013; Griffin et al., 2013; Lijoi et al., 2014). These models are closely related to Lévy copulas (Tankov, 2003; Cont and Tankov, 2003; Kallsen and Tankov, 2006) and multivariate subordinators on cones (Barndorff-Nielsen et al., 2001; Skorohod, 1991). A vector of CRMs (W_1, \dots, W_p) on \mathbb{R}_+ is a collection of random measures W_k , $k = 1, \dots, p$, such that, for any collection of disjoint measurable subsets A_1, \dots, A_n of \mathbb{R}_+ , the vectors $(W_1(A_1), \dots, W_p(A_1)), (W_1(A_2), \dots, W_p(A_2)), \dots, (W_1(A_n), \dots, W_p(A_n))$ are mutually independent. We only consider here vectors of CRMs with both random weights and locations. In this case, the measures W_k , $k = 1, \dots, p$, are a.s. discrete and take the form

$$W_k = \sum_{i=1}^{\infty} w_{ik} \delta_{\theta_i}. \quad (47)$$

The law of the vector of CRMs can be characterized by a Poisson point process on \mathbb{R}_+^{p+1} with mean measure $\nu(dw_1, \dots, dw_p, d\theta)$. We focus again on homogeneous vectors of CRMs with independent increments where the mean measure can be written as

$$\nu(dw_1, \dots, dw_p, d\theta) = \rho(dw_1, \dots, dw_p)\lambda(d\theta). \quad (48)$$

where ρ is a measure on \mathbb{R}_+^p , concentrated on $\mathbb{R}_+^p \setminus \{\mathbf{0}\}$, which satisfies

$$\int_{\mathbb{R}_+^p} \min \left(1, \sum_{k=1}^p w_k \right) \rho(dw_1, \dots, dw_p) < \infty. \quad (49)$$

We use the same notation as for (scalar) CRMs and write simply $(W_1, \dots, W_p) \sim \text{CRM}(\rho, \lambda)$. A key quantity is the multivariate Laplace exponent defined by

$$\psi(t_1, \dots, t_p) := -\log \mathbb{E} \left[e^{-\sum_{k=1}^p t_k W_k([0,1])} \right] \quad (50)$$

$$= \int_{\mathbb{R}_+^p} \left(1 - e^{-\sum_{k=1}^p t_k w_k} \right) \rho(dw_1, \dots, dw_p) \quad (51)$$

Note that this quantity involves a p -dimensional integral which may not be analytically computable, and may be expensive to evaluate numerically. As for CRMs, if

$$\int_{\mathbb{R}_+^p} \rho(dw_1, \dots, dw_p) = \infty \quad (52)$$

then there will be an infinite number of $\theta_i \in [0, T]$ for which $\sum_k w_{ik} > 0$ and the vector of CRMs is called infinite-activity. Otherwise, it is called finite-activity. Note that some (but not all) CRMs may still be marginally finite-activity.

Appendix B: Proof of Propositions 3 and 4

The proof follows the lines of the sparsity proof in Caron and Fox (2014), and we only provide a sketch of it. First, as Z is a jointly exchangeable point process verifying (21) and under the moment condition (24), it follows from the law of large numbers that

$$N_\alpha^{(e)} = \Theta(\alpha^2) \text{ a.s. as } \alpha \rightarrow \infty.$$

Finite-activity case. If the vector of CRMs is finite-activity, the jump locations arise from an homogeneous Poisson process with finite rate, and $N_\alpha = \Theta(\alpha)$ a.s. It follows that

$$N_\alpha^{(e)} = \Theta(N_\alpha^2) \text{ a.s. as } \alpha \rightarrow \infty.$$

Infinite-activity case. Consider now the infinite-activity case. Following Caron and Fox (2014), one can lower bound the node counting process N_α by a counting process \tilde{N}_α which is conditionally Poisson with mean measure $\lambda(S_\alpha^1)\psi(W_1(S_\alpha^2), \dots, W_p(S_\alpha^2))$ where (S_α^1, S_α^2) is a partition of $[0, \alpha]$ such that $\lambda(S_\alpha^1) = \lambda(S_\alpha^2) = \frac{\alpha}{2}$. As $\psi(W_1(S_\alpha^2), \dots, W_p(S_\alpha^2)) \rightarrow \infty$ a.s. in the infinite-activity case, it follows that $N_\alpha = \Omega(\alpha)$ a.s., and therefore

$$N_\alpha^{(e)} = o(N_\alpha^2) \text{ a.s. as } \alpha \rightarrow \infty.$$

Finally, for compound CRMs with regularly varying ρ_0 with exponent σ , Proposition 7 in Appendix F implies that $\psi(W_1(S_\alpha^2), \dots, W_p(S_\alpha^2)) = \Theta(\alpha^\sigma)$ a.s. hence $N_\alpha = \omega(\alpha^{1+\sigma})$ a.s. and

$$N_\alpha^{(e)} = O(N_\alpha^{2/(1+\sigma)}) \text{ a.s. as } \alpha \rightarrow \infty.$$

Appendix C: Simulation from a tilted truncated generalized gamma process

We want to sample points from a Poisson process with truncated mean measure

$$\rho^\varepsilon(dw) = h(w)w^{-1-\sigma}e^{-\tau w}\mathbf{1}_{w>\varepsilon}dw \quad (53)$$

where h is a monotone decreasing and bounded function, and (τ, σ) verify either $\tau \geq 0$ and $\sigma \in (0, 1)$, or $\tau > 0$ and $\sigma \in (-1, 0]$. We will resort to adaptive thinning (Lewis and Shedler, 1979; Ogata, 1981; Favaro and Teh, 2013).

For $\tau > 0$, consider the family of adaptive bounds

$$g_t(s) = h(t)t^{-1-\sigma} \exp(-\tau s)$$

with $g_t(s) > \rho(s)$ for $s > t$. We have,

$$\begin{aligned} G_t(s) &= \int_t^s g_t(s') ds' \\ &= \frac{h(t)}{\tau} t^{-1-\sigma} (\exp(-\tau t) - \exp(-\tau s)) \end{aligned}$$

and

$$G_t^{-1}(r) = -\frac{1}{\tau} \log \left(\exp(-\tau t) - \frac{r\tau}{t^{-1-\sigma}h(t)} \right).$$

For $\tau = 0$, we consider bounds

$$g_t(s) = h(t)s^{-1-\sigma}$$

and we obtain

$$\begin{aligned} G_t(s) &= \frac{h(t)}{\sigma} (t^{-\sigma} - s^{-\sigma}) \\ G_t^{-1}(r) &= \left[t^{-\sigma} - \frac{r\sigma}{h(t)} \right]^{-1/\sigma}. \end{aligned}$$

The adaptive thinning sampling scheme is as follows:

1. Set $N = \emptyset, t = \varepsilon$
2. iterate until termination
 - (a) Draw $r \sim \text{Exp}(1)$
 - (b) If $r > G_t(\infty)$, terminate; else set $t' = G_t^{-1}(r)$;
 - (c) with probability $\rho^\varepsilon(t')/g_t(t')$ accept sample t' and set $N = N \cup \{t'\}$
 - (d) set $t = t'$ and continue
3. Return N a draw from the Poisson random measure with intensity ρ^ε on $[\varepsilon, +\infty)$

The efficiency of this approach depends on the acceptance probability, which is given, for $\tau > 0$, by

$$\frac{\rho^\varepsilon(s)}{g_t(s)} = \frac{h(s)s^{-1-\sigma}}{h(t)t^{-1-\sigma}} < 1$$

for $s > t$.

Appendix D: Bipartite networks

It is possible to use a construction similar to that of Section 2 to model bipartite graphs, and extend the model of Caron (2012). A bipartite graph is a graph with two types of nodes, where only connections between nodes of different types are allowed. Nodes of the first type are embedded at locations $\theta_i \in \mathbb{R}_+$, and nodes of the second type at location $\theta'_j \in \mathbb{R}_+$. The bipartite graph will be represented by a (non-symmetric) point process

$$Z = \sum_{i,j} z_{ij} \delta_{(\theta_i, \theta'_j)} \tag{54}$$

where $z_{ij} = 1$ if there is an edge between node i of type 1 and node j of type 2.

Statistical Model. We consider the model

$$\begin{aligned} W_1, \dots, W_p &\sim \text{CRM}(\rho, \lambda) \\ W'_1, \dots, W'_p &\sim \text{CRM}(\rho', \lambda) \end{aligned}$$

and for $k = 1, \dots, p$,

$$\begin{aligned} D_k | W_k, W'_k &\sim \text{Poisson}(W_k \times W'_k) \\ D_k &= \sum_{i,j} n_{ijk} \delta_{(\theta_i, \theta'_j)} \end{aligned}$$

and $z_{ij} = \min(1, \sum_{k=1}^p n_{ijk})$.

Posterior inference. We derive here the inference algorithm when (W_1, \dots, W_p) and (W'_1, \dots, W'_p) are compound CRMs with F and ρ_0 taking the form (18) and (19).

Assume that we observe a set of connections $z = (z_{ij})_{i=1, \dots, N_\alpha; j=1, \dots, N'_\alpha}$. We introduce latent variables n_{ijk} , for $1 \leq i \leq N_\alpha$, $1 \leq j \leq N'_\alpha$, $k = 1, \dots, p$,

$$(n_{ij1}, \dots, n_{ijp}) | w, w', z \sim \begin{cases} \delta_{(0, \dots, 0)} & \text{if } z_{ij} = 0 \\ \text{tPoisson}(w_{i1} w'_{j1}, \dots, w_{ip} w'_{jp}) & \text{if } z_{ij} = 1 \end{cases}.$$

We want to approximate

$$p((w_{10}, \dots, w_{N_\alpha 0}), (\beta_{1k}, \dots, \beta_{N_\alpha k}, w_{*k})_{k=1, \dots, p}, (w'_{10}, \dots, w'_{N'_\alpha 0}), (\beta'_{1k}, \dots, \beta'_{N'_\alpha k}, w'_{*k})_{k=1, \dots, p}, \phi, \alpha, \phi', \alpha' | z)$$

Denote $m_{ik} = \sum_{j=1}^{N'_\alpha} n_{ijk}$ and $m_i = \sum_{k=1}^p m_{ik}$. The MCMC algorithm iterates as follows:

1. Update $(\alpha, \phi) | \text{rest}$ using a Metropolis-Hastings step.
2. Update

$$w_{i0} | \text{rest} \sim \text{Gamma} \left(m_i - \sigma, \tau + \sum_{k=1}^p \beta_{ik} \left[\gamma_k + \left(\sum_{j=1}^{N'_\alpha} w'_{jk} \right) + w'_{*k} \right] \right).$$

3. Update

$$\beta_{ik} | \text{rest} \sim \text{Gamma} \left(a_k + m_{ik}, b_k + w_{i0} \left[\gamma_k + \left(\sum_{j=1}^{N'_\alpha} w'_{jk} \right) + w'_{*k} \right] \right).$$

4. Update $(w_{*1}, \dots, w_{*p}) | \text{rest}$.
5. Update the latent variables $n_{ijk} | \text{rest}$.
6. Repeat steps 1-4 to update (α', ϕ') , $(w'_{10}, \dots, w'_{N'_\alpha 0})$, $(\beta'_{1k}, \dots, \beta'_{N'_\alpha k})_{k=1, \dots, p}$ and $(w'_{*1}, \dots, w'_{*p})$.

Appendix E: Gaussian approximation of the sum of small jumps

Theorem 6 Consider the multivariate random variable $X_\varepsilon \in \mathbb{R}_+^p$ with moment generating function

$$\mathbb{E}[e^{-t^T X_\varepsilon}] = \exp \left[-\alpha \int_{\mathbb{R}_+^p} \left(1 - e^{-\sum_{k=1}^p t_k w_k} \right) \rho_\varepsilon(dw_1, \dots, dw_p) \right]$$

where $\alpha > 0$ and

$$\rho_\varepsilon(dw_1, \dots, dw_p) = e^{-\sum_{k=1}^p \gamma_k w_k} \int_0^\varepsilon w_0^{-p} F \left(\frac{dw_1}{w_0}, \dots, \frac{dw_p}{w_0} \right) \rho_0(dw_0)$$

with $\varepsilon > 0$, ρ_0 is a Lévy measure on \mathbb{R}_+ and F is a probability distribution on \mathbb{R}_+^p with density f verifying

$$\int_0^\infty f(zu_1, \dots, zu_p) dz > 0 \text{ } U\text{-almost everywhere}$$

$$\int_{\mathbb{R}_+^p} \|\beta_{1:p}\|^2 f(\beta_1, \dots, \beta_p) d\beta_{1:p} < \infty$$

where U is the uniform distribution on the unit sphere S^{p-1} . Then if ρ_0 is a regularly varying Lévy measure with exponent $\sigma \in (0, 1)$, i.e.

$$\int_x^\infty \rho_0(dw_0) \stackrel{x \downarrow 0}{\sim} x^{-\sigma} \ell(1/x)$$

where $\ell : (0, \infty) \rightarrow (0, \infty)$ is a slowly varying function then

$$\Sigma_\varepsilon^{-1/2}(X_\varepsilon - \mu_\varepsilon) \xrightarrow{d} \mathcal{N}(0, I_p)$$

as $\varepsilon \rightarrow 0$, where

$$\mu_\varepsilon = \alpha \int_{\mathbb{R}_+^p} w \rho_\varepsilon(dw_1, \dots, dw_p)$$

$$\Sigma_\varepsilon = \alpha \int_{\mathbb{R}_+^p} ww^T \rho_\varepsilon(dw_1, \dots, dw_p)$$

with

$$\mu_\varepsilon \sim \alpha \mathbb{E}[\beta] \frac{\sigma}{1-\sigma} \varepsilon^{1-\sigma} \ell(1/\varepsilon) \text{ as } \varepsilon \rightarrow 0$$

$$\Sigma_\varepsilon \sim \alpha \mathbb{E}[\beta\beta^T] \frac{\sigma}{2-\sigma} \varepsilon^{2-\sigma} \ell(1/\varepsilon) \text{ as } \varepsilon \rightarrow 0$$

where β is distributed from F .

Proof. We write the model in spherical form. Let $r = \sqrt{\sum w_k^2}$ and $u_k = \frac{w_k}{r}$ for $k = 1, \dots, p-1$. The determinant of the Jacobian is $\frac{r^{p-1}}{\sqrt{1-\sum_{k=1}^{p-1} u_k^2}}$ and so

$$\tilde{\rho}_\varepsilon(r, u_1, \dots, u_{p-1}) = \frac{r^{p-1}}{u_p} e^{-r \sum_{k=1}^p \gamma_k u_k} \int_0^\varepsilon w_0^{-p} f\left(\frac{ru_1}{w_0}, \dots, \frac{ru_p}{w_0}\right) \rho_0(dw_0) dr du_{1:p-1}$$

$$:= \mu_\varepsilon(dr|u_{1:p-1}) U(du_{1:p-1})$$

where $u_p = \sqrt{1 - \sum_{k=1}^{p-1} u_k^2}$, $\mu_\varepsilon(dr|u) = r^{p-1} e^{-r \sum_{k=1}^p \gamma_k u_k} \int_0^\varepsilon w_0^{-p} f\left(\frac{ru_1}{w_0}, \dots, \frac{ru_p}{w_0}\right) \rho_0(dw_0) dr$ and $U(du) = \frac{1}{u_p} du_{1:p}$ is the uniform distribution on the unit sphere S^{p-1} .

In order to apply Theorem 2.4 in [Cohen and Rosinski \(2007\)](#) (see also [Asmussen and Rosiński, 2001](#)), we need to show that there exists a function $b_\varepsilon : (0, 1] \rightarrow (0, +\infty)$ such that

$$\lim_{\varepsilon \rightarrow 0} \frac{\sigma_\varepsilon(u)}{b_\varepsilon} > 0, \text{ } U\text{-almost everywhere} \quad (55)$$

where

$$\sigma_\varepsilon^2(u) = \int_0^\infty r^2 \mu_\varepsilon(dr|u)$$

and for every $\kappa > \varepsilon$

$$\lim_{\varepsilon \rightarrow 0} \frac{1}{b_\varepsilon^2} \int_{\|w_{1:p}\| > \kappa b_\varepsilon} \|w_{1:p}\|^2 \rho_\varepsilon(dw_1, \dots, dw_p) = 0 \quad (56)$$

Assume that $\int_0^\infty f(zu_1, \dots, zu_p) dz > 0$ U -almost everywhere. With the change of variable $z = \frac{r}{w_0}$, and the dominated convergence theorem we obtain

$$\begin{aligned} \sigma_\varepsilon^2(u) &= \int_0^\infty z^{p+1} f(zu_1, \dots, zu_p) \left[\int_0^\varepsilon e^{-zw_0 \sum_{k=1}^p \gamma_k u_k} w_0^2 \rho_0(dw_0) \right] dz \\ &\sim \left(\int_0^\infty z^{p+1} f(zu_1, \dots, zu_p) dz \right) \left(\int_0^\varepsilon w_0^2 \rho_0(dw_0) \right) \text{ as } \varepsilon \rightarrow 0 \\ &\sim \left(\int_0^\infty z^{p+1} f(zu_1, \dots, zu_p) dz \right) \frac{\sigma}{2-\sigma} \varepsilon^{2-\sigma} \ell(1/\varepsilon) \text{ as } \varepsilon \rightarrow 0 \end{aligned}$$

Let $b_\varepsilon = \varepsilon^{1-\sigma/2} \sqrt{\ell(1/\varepsilon)}$, we have

$$\lim_{\varepsilon \rightarrow 0} \frac{\sigma_\varepsilon^2(u)}{b_\varepsilon^2} = \left(\int_0^\infty z^{p+1} f(zu_1, \dots, zu_p) dz \right) \frac{\sigma}{2-\sigma} > 0, \text{ } U\text{-almost everywhere} \quad (57)$$

Now consider, for any $\kappa > 0$,

$$\begin{aligned} I_\varepsilon &= \int_{\|w_{1:p}\| > \kappa b_\varepsilon} \|w_{1:p}\|^2 \nu_\varepsilon(dw_1, \dots, dw_p) \\ &= \int_0^\varepsilon \int_{\|\beta_{1:p}\| > \frac{\kappa b_\varepsilon}{w_0}} w_0^2 \|\beta_{1:p}\|^2 e^{-w_0 \sum_{k=1}^p \gamma_k \beta_k} f(\beta_1, \dots, \beta_p) \rho_0(dw_0) d\beta_{1:p} \end{aligned}$$

For $w_0 \in (0, \varepsilon)$, we have $\frac{\kappa b_\varepsilon}{w_0} \geq \frac{\kappa b_\varepsilon}{\varepsilon} = \varepsilon^{-\sigma/2} \ell(1/\varepsilon) > \kappa_2 \varepsilon^{-\sigma/4}$ for ε small enough as $t^\delta \ell(t) \rightarrow 0$ for any $\delta > 0$ as $t \rightarrow \infty$. So for ε small enough

$$\begin{aligned} I_\varepsilon &> \int_0^\varepsilon \int_{\|\beta_{1:p}\| > \kappa_2 \varepsilon^{-\sigma/4}} w_0^2 \|\beta_{1:p}\|^2 e^{-w_0 \sum_{k=1}^p \gamma_k \beta_k} f(\beta_1, \dots, \beta_p) \rho_0(dw_0) d\beta_{1:p} \\ &> \left[\int_{\|\beta_{1:p}\| > \kappa_2 \varepsilon^{-\sigma/4}} \|\beta_{1:p}\|^2 f(\beta_1, \dots, \beta_p) d\beta_{1:p} \right] \left[\int_0^\varepsilon w_0^2 \rho_0(dw_0) \right] \end{aligned}$$

As $\left[\int_0^\varepsilon w_0^2 \rho_0(dw_0) \right] \sim \frac{\sigma}{2-\sigma} b_\varepsilon^2$ when $\varepsilon \rightarrow 0$, we conclude that

$$\lim_{\varepsilon \rightarrow 0} I_\varepsilon = \lim_{\varepsilon \rightarrow 0} \frac{\sigma}{2-\sigma} \int_{\|\beta_{1:p}\| > \kappa_2 \varepsilon^{-\sigma/4}} \|\beta_{1:p}\|^2 f(\beta_1, \dots, \beta_p) d\beta_{1:p} = 0 \quad (58)$$

Equations (57) and (58) with Theorem 2.4 of [Cohen and Rosinski \(2007\)](#) yield

$$\Sigma_\varepsilon^{-1/2} (X_\varepsilon - \mu_\varepsilon) \xrightarrow{d} \mathcal{N}(0, I_p)$$

as $\varepsilon \rightarrow 0$, where

$$\begin{aligned} \mu_\varepsilon &= \alpha \int_{\mathbb{R}_+^p} w_{1:p} \rho_\varepsilon(dw_1, \dots, dw_p) \\ &= \alpha \int_{\mathbb{R}_+^p} \int_0^\varepsilon w_0 \beta_{1:p} e^{-w_0 \sum_{k=1}^p \gamma_k \beta_k} \rho_0(dw_0) f(\beta_1, \dots, \beta_p) d\beta_{1:p} \\ &\sim \alpha \mathbb{E}[\beta_{1:p}] \frac{\sigma}{1-\sigma} \varepsilon^{1-\sigma} \ell(1/\varepsilon) \text{ as } \varepsilon \rightarrow 0 \end{aligned}$$

and

$$\begin{aligned} \Sigma_\varepsilon &= \alpha \int_{\mathbb{R}_+^p} w_{1:p} w_{1:p}^T \rho_\varepsilon(dw_1, \dots, dw_p) \\ &\sim \alpha \mathbb{E}[\beta_{1:p} \beta_{1:p}^T] \frac{\sigma}{2-\sigma} \varepsilon^{2-\sigma} \ell(1/\varepsilon) \text{ as } \varepsilon \rightarrow 0 \end{aligned}$$

using the dominated convergence theorem and lemmas 8 and 9. ■

Appendix F: Technical lemmas

Proposition 7 Let ν be a Lévy measure defined by Eq. (5) and (13) and ψ be its multivariate Laplace exponent. Assume that $\bar{\rho}_0$ is a regularly varying function with exponent $\sigma \in (0, 1)$:

$$\bar{\rho}_0 \stackrel{x \downarrow 0}{\sim} x^{-\sigma} \ell(1/x) \quad (59)$$

Then ψ is (multivariate) regularly varying (Resnick, 2013), with exponent σ . More precisely, for any $(x_1, \dots, x_p) \in (0, \infty)^p$, we have

$$\begin{aligned} \psi(tx_1, \dots, tx_p) &= \int_{\mathbb{R}_+^p} \left(1 - e^{-t \sum_{k=1}^p x_k w_k}\right) \nu(dw_1, \dots, dw_p) \\ &\stackrel{t \uparrow \infty}{\sim} t^\sigma \Gamma(1 - \sigma) \ell(t) \mathbb{E} \left[\left(\sum_{k=1}^p x_k \beta_k \right)^\sigma \right]. \end{aligned}$$

Proof.

$$\begin{aligned} \psi(tx_1, \dots, tx_p) &= \int_{\mathbb{R}_+^p} \left(1 - e^{-t \sum_{k=1}^p x_k w_k}\right) \nu(dw_1, \dots, dw_p) \\ &= \int_{\mathbb{R}_+^p} \left(1 - e^{-t \sum_{k=1}^p x_k w_k}\right) \nu(dw_1, \dots, dw_p) \\ &= \int_{\mathbb{R}_+^p} f(\beta_1, \dots, \beta_p) \left[\int_0^\infty \left(1 - e^{-w_0 t \sum_{k=1}^p x_k \beta_k}\right) e^{-w_0 \sum_{k=1}^p \gamma_k \beta_k} \rho_0(dw_0) \right] d\beta_{1:p} \end{aligned}$$

which gives, using Lemmas 8, 9, and the dominated convergence theorem

$$\psi(tx_1, \dots, tx_p) \stackrel{t \uparrow \infty}{\sim} t^\sigma \Gamma(1 - \sigma) \ell(t) \int_{(0, \infty)^p} \left(\sum_{k=1}^p x_k \beta_k \right)^\sigma f(\beta_1, \dots, \beta_p) d\beta_{1:p}.$$

■

Lemma 8 If

$$\int_x^\infty \rho(dw) \stackrel{x \downarrow 0}{\sim} x^{-\sigma} \ell(1/x)$$

Then

$$\int_x^\infty e^{-cw} \rho(dw) \stackrel{x \downarrow 0}{\sim} x^{-\sigma} \ell(1/x)$$

Proof.

$$\begin{aligned} \int_x^\infty e^{-cw} \rho(dw) &= \int_x^\infty \rho(dw) - \int_x^\infty (1 - e^{-cw}) \rho(dw) \\ &\stackrel{x \downarrow 0}{\sim} x^{-\sigma} \ell(1/x) \end{aligned}$$

as $\int_0^\infty (1 - e^{-cw}) \rho(dw) < \infty$ for any $c > 0$. ■

Lemma 9 (Gnedin et al., 2007; Bingham et al., 1989). Let ρ be a Lévy measure with regularly varying tail Lévy intensity

$$\int_x^\infty \rho(dw) \stackrel{x \downarrow 0}{\sim} x^{-\sigma} \ell(1/x) \quad (60)$$

where $\sigma \in (0, 1)$ and ℓ is a slowly varying function (at infinity). Then (60) is equivalent to

$$\int_0^x w^k \rho(dw) \stackrel{x \downarrow 0}{\sim} \frac{\sigma}{k - \sigma} x^{k - \sigma} \ell(1/x) \quad (61)$$

$$\int_0^\infty (1 - e^{-tw}) \rho(dw) \stackrel{t \uparrow \infty}{\sim} \Gamma(1 - \sigma) t^\sigma \ell(t) \quad (62)$$

for any $k \geq 1$.

References

- L. A. Adamic and N. Glance. The political blogosphere and the 2004 US election: divided they blog. In *Proceedings of the 3rd international workshop on Link discovery*, pages 36–43. ACM, 2005.
- E. M. Airoldi, D. Blei, S. E. Fienberg, and E. Xing. Mixed membership stochastic blockmodels. *The Journal of Machine Learning Research*, 9:1981–2014, 2008.
- S. Asmussen and J. Rosiński. Approximations of small jumps of Lévy processes with a view towards simulation. *Journal of Applied Probability*, pages 482–493, 2001.
- B. Ball, B. Karrer, and M. E. J. Newman. Efficient and principled method for detecting communities in networks. *Physical Review E*, 84(3):036103, 2011.
- O. E. Barndorff-Nielsen, J. Pedersen, and K.-I. Sato. Multivariate subordination, self-decomposability and stability. *Advances in Applied Probability*, 33:160–187, 2001.
- M. Bastian, S. Heymann, and M. Jacomy. Gephi: An open source software for exploring and manipulating networks. In *International AAAI Conference on Weblogs and Social Media*, 2009.
- N. H. Bingham, C. M. Goldie, and J. L. Teugels. *Regular variation*, volume 27. Cambridge university press, 1989.
- B. Bollobás. *Random graphs*, volume 73. Cambridge University Press, 2001.
- C. Borgs, J. T. Chayes, H. Cohn, and N. Holden. Sparse exchangeable graphs and their limits via graphon processes. *ArXiv preprint arXiv:1601.07134*, 2016.
- A. Brix. Generalized gamma measures and shot-noise Cox processes. *Advances in Applied Probability*, 31(4): 929–953, 1999.
- F. Caron. Bayesian nonparametric models for bipartite graphs. In F. Pereira, C.J.C. Burges, L. Bottou, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 2051–2059. Curran Associates, Inc., 2012.
- F. Caron and E. B. Fox. Sparse graphs using exchangeable random measures. *arXiv preprint arXiv:1401.1137*, 2014.
- F. Caron, Y.W. Teh, and T.B. Murphy. Bayesian nonparametric Plackett-Luce models for the analysis of preferences for college degree programmes. *The Annals of Applied Statistics*, 8(2):1145–1181, 2014.
- A. T. Cemgil. Bayesian inference for nonnegative matrix factorisation models. *Computational Intelligence and Neuroscience*, 2009, 2009.
- S. Cohen and J. Rosinski. Gaussian approximation of multivariate Lévy processes with applications to simulation of tempered stable processes. *Bernoulli*, 13(1):195–210, 2007.
- R. Cont and P. Tankov. *Financial modelling with jump processes*, volume 2. CRC press, 2003.
- D.J. Daley and D. Vere-Jones. *An Introduction to the Theory of Point Processes. Volume I: Elementary Theory and Methods*. Springer Verlag, 2nd edition edition, 2008a.
- D.J. Daley and D. Vere-Jones. *An Introduction to the Theory of Point Processes. Volume II: General Theory and Structure*. Springer Verlag, 2nd edition edition, 2008b.
- S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth. Hybrid Monte Carlo. *Physics Letters B*, 195(2): 216–222, 1987.
- I. Epifani and A. Lijoi. Nonparametric priors for vectors of survival functions. *Statistica Sinica*, pages 1455–1484, 2010.
- P. Erdős and A. Rényi. On random graphs. *Publicationes Mathematicae*, 6:290–297, 1959.
- S. Favaro and Y.W. Teh. MCMC for normalized random measure mixture models. *Statistical Science*, 28(3): 335–359, 2013.
- S. E. Fienberg. A brief history of statistical models for network analysis and open challenges. *Journal of Computational and Graphical Statistics*, 21(4):825–839, 2012.
- A. Gnedin, B. Hansen, and J. Pitman. Notes on the occupancy problem with infinitely many boxes: general asymptotics and power laws. *Probab. Surv*, 4(146-171):88, 2007.
- A. Goldenberg, A.X. Zheng, S.E. Fienberg, and E.M. Airoldi. A survey of statistical network models. *Foundations and Trends in Machine Learning*, 2(2):129–233, 2010.
- P. Gopalan, J. M. Hofman, and D. M. Blei. Scalable recommendation with hierarchical Poisson factorization. In *Conference on Uncertainty in Artificial Intelligence*, 2015.
- J. E. Griffin and F. Leisen. Compound random measures and their use in Bayesian nonparametrics. *arXiv preprint arXiv:1410.0611*, 2014.

- J. E. Griffin, M. Kolossiatis, and M. F. J. Steel. Comparing distributions by using dependent normalized random-measure mixtures. *Journal of the Royal Statistical Society: Series B*, 75(3):499–529, 2013.
- T. Herlau, M. N. Schmidt, and M. Mørup. Completely random measures for modelling block-structured sparse networks. Technical report, arXiv preprint arXiv:1507.02925, 2015.
- P. W. Holland, K. B. Laskey, and S. Leinhardt. Stochastic blockmodels: First steps. *Social networks*, 5(2):109–137, 1983.
- P. Hougaard. Survival models for heterogeneous populations derived from stable distributions. *Biometrika*, 73(2):387–396, 1986.
- A. Z. Jacobs and A. Clauset. A unified view of generative models for networks: models, methods, opportunities and challenges. Technical report, arXiv:1411.4070, 2014.
- L. James. Bayesian Poisson process partition calculus with an application to Bayesian Lévy moving averages. *The Annals of Statistics*, pages 1771–1799, 2005.
- L. James. Poisson latent feature calculus for generalized Indian buffet processes. Technical report, arXiv:1411.2936, 2014.
- L. F. James. Poisson process partition calculus with applications to exchangeable models and bayesian nonparametrics. *arXiv preprint math/0205093*, 2002.
- L.F. James, A. Lijoi, and I. Prünster. Posterior analysis for normalized random measures with independent increments. *Scandinavian Journal of Statistics*, 36(1):76–97, 2009.
- S. Janson. Probability asymptotics: notes on notation. Technical report, arXiv:1108.3924, 2011.
- O. Kallenberg. Exchangeable random measures in the plane. *Journal of Theoretical Probability*, 3(1):81–136, 1990.
- O. Kallenberg. *Probabilistic symmetries and invariance principles*. Springer, 2005.
- J. Kallsen and P. Tankov. Characterization of dependence of multidimensional Lévy processes using Lévy copulas. *Journal of Multivariate Analysis*, 97(7):1551–1572, 2006.
- B. Karrer and M. E. J. Newman. Stochastic blockmodels and community structure in networks. *Physical Review E*, 83(1):016107, 2011.
- C. Kemp, J. B. Tenenbaum, T. L. Griffiths, T. Yamada, and N. Ueda. Learning systems of concepts with an infinite relational model. In *AAAI*, volume 21, page 381, 2006.
- J.F.C. Kingman. Completely random measures. *Pacific Journal of Mathematics*, 21(1):59–78, 1967.
- J.F.C. Kingman. *Poisson processes*. Oxford University Press, USA, 1993.
- E. D. Kolaczyk. *Statistical Analysis of Network Data: Methods and Models*. Springer, 2009.
- P. Latouche, E. Birmelé, and C. Ambroise. Overlapping stochastic block models with application to the French political blogosphere. *The Annals of Applied Statistics*, pages 309–336, 2011.
- H. S. Lee, D. D. and Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.
- F. Leisen and A. Lijoi. Vectors of two-parameter Poisson–Dirichlet processes. *Journal of Multivariate Analysis*, 102(3):482–495, 2011.
- F. Leisen, A. Lijoi, and D. Spanó. A vector of Dirichlet processes. *Electronic Journal of Statistics*, 7:62–90, 2013.
- P. A. Lewis and G. S. Shedler. Simulation of nonhomogeneous Poisson processes by thinning. *Naval Research Logistics Quarterly*, 26(3):403–413, 1979.
- A. Lijoi and I. Prünster. Models beyond the Dirichlet process. In N. L. Hjort, C. Holmes, P. Müller, and S.G. Walker, editors, *Bayesian nonparametrics*, volume 28, page 80. Camb. Ser. Stat. Probab. Math, 2010.
- A. Lijoi, R. H. Mena, and I. Prünster. Controlling the reinforcement in Bayesian non-parametric mixture models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(4):715–740, 2007.
- A. Lijoi, B. Nipoti, and I. Prünster. Bayesian inference with dependent normalized completely random measures. *Bernoulli*, 20(3):1260–1291, 2014.
- K. Miller, T. Griffiths, and M. Jordan. Nonparametric latent feature models for link prediction. In *NIPS*, 2009.
- R. M. Neal. MCMC using Hamiltonian dynamics. In S. Brooks, A. Gelman, G. Jones, and X.-L. Meng, editors, *Handbook of Markov Chain Monte Carlo*, volume 2. Chapman & Hall / CRC Press, 2011.
- M. Newman. *Networks: an introduction*. OUP Oxford, 2009.
- Mark EJ Newman. Mixing patterns in networks. *Physical Review E*, 67(2):026126, 2003a.
- M.E.J. Newman. The structure and function of complex networks. *SIAM review*, pages 167–256, 2003b.

- K. Nowicki and T. Snijders. Estimation and prediction for stochastic blockstructures. *Journal of the American Statistical Association*, 96(455):1077–1087, 2001.
- Y. Ogata. On Lewis’ simulation method for point processes. *IEEE Transactions on Information Theory*, 27(1):23–31, 1981.
- P. Orbanz and D. M. Roy. Bayesian models of graphs, arrays and other exchangeable random structures. *IEEE Trans. Pattern Anal. Mach. Intelligence (PAMI)*, 37(2):437–461, 2015.
- K. Palla, D. A. Knowles, and Z. Ghahramani. An infinite latent attribute model for network data. In *ICML*, 2012.
- I. Prünster. *Random probability measures derived from increasing additive processes and their application to Bayesian statistics*. PhD thesis, University of Pavia, 2002.
- I. Psorakis, S. Roberts, M. Ebden, and B. Sheldon. Overlapping community detection using Bayesian non-negative matrix factorization. *Physical Review E*, 83(6):066114, 2011.
- S. Resnick. *Extreme values, regular variation and point processes*. Springer, 2013.
- A. Saeedi and A. Bouchard-Côté. Priors over recurrent continuous time processes. In *Advances in Neural Information Processing Systems*, pages 2052–2060, 2011.
- M. Salter-Townshend and T. H. McCormick. Latent space models for multiview network data. Technical report, Technical Report, 2013.
- A. V. Skorohod. *Random processes with independent increments*, volume 47. Springer, 1991.
- T. A. B. Snijders and K. Nowicki. Estimation and prediction for stochastic blockmodels for graphs with latent block structure. *Journal of classification*, 14(1):75–100, 1997.
- P. Tankov. Dependence structure of spectrally positive multidimensional Lévy processes. *Unpublished manuscript*, 2003.
- V. Veitch and D. M. Roy. The class of random graphs arising from exchangeable random measures. *arXiv preprint arXiv:1512.03099*, 2015.
- L. M. Verbrugge. Multiplexity in adult friendships. *Social Forces*, 57(4):1286–1309, 1979.
- J. Yang and J. Leskovec. Overlapping community detection at scale: a nonnegative matrix factorization approach. In *Proceedings of the sixth ACM international conference on Web search and data mining*, pages 587–596. ACM, 2013.
- Y. Zhao, E. Levina, and J. Zhu. Consistency of community detection in networks under degree-corrected stochastic block models. *The Annals of Statistics*, 40(4):2266–2292, 2012.
- M. Zhou. Infinite edge partition models for overlapping community detection and link prediction. In *Artificial Intelligence and Statistics (AISTATS2015), JMLR W&CP*, volume 38, 2015.