



Graph sampling with determinantal processes

Nicolas Tremblay, Pierre-Olivier Amblard, Simon Barthelme

► To cite this version:

Nicolas Tremblay, Pierre-Olivier Amblard, Simon Barthelme. Graph sampling with determinantal processes. EUSIPCO 2017 - 25th European Signal Processing Conference, Aug 2017, Kos Island, Greece. 10.23919/EUSIPCO.2017.8081494 . hal-01483347

HAL Id: hal-01483347

<https://hal.science/hal-01483347>

Submitted on 5 Mar 2017

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.

Graph sampling with determinantal processes

Nicolas Tremblay*, Pierre-Olivier Amblard* and Simon Barthélémy*

*CNRS, GIPSA-lab, Grenoble-Alpes University, Grenoble, France

Abstract—We present a new random sampling strategy for k -bandlimited signals defined on graphs, based on determinantal point processes (DPP). For small graphs, *i.e.*, in cases where the spectrum of the graph is accessible, we exhibit a DPP sampling scheme that enables perfect recovery of bandlimited signals. For large graphs, *i.e.*, in cases where the graph's spectrum is not accessible, we investigate, both theoretically and empirically, a sub-optimal but much faster DPP based on loop-erased random walks on the graph. Preliminary experiments show promising results especially in cases where the number of measurements should stay as small as possible and for graphs that have a strong community structure. Our sampling scheme is efficient and can be applied to graphs with up to 10^6 nodes.

I. INTRODUCTION

Graphs are a central modelling tool for network-structured data. Data on a graph, called graph signals [1], such as individual hobbies in social networks, blood flow of brain regions in neuronal networks, or traffic at a station in transportation networks, may be represented by a scalar per node. Studying such signals with respect to the particular graph topology on which it is defined, is the goal of graph signal processing (GSP). One of the top challenges of GSP is not only to adapt classical processing tools to arbitrary topologies defined by the underlying graph, but to do so with efficient algorithms that can handle the large datasets encountered today. A possible answer to this efficiency challenge is dimension reduction, and in particular sampling. Graph sampling [2]–[6] consists in measuring a graph signal on a reduced set of nodes carefully chosen to optimize a predefined objective, may it be the estimation of some statistics, signal compression via graph filterbanks, signal recovery, etc.

In an effort to generalize Shannon's sampling theorem of bandlimited signals, several authors [3]–[5] have proposed sampling schemes adapted to the recovery of *a priori* smooth graph signals, and in particular k -bandlimited signals, *i.e.*, linear combinations of only k graph Fourier modes. Classically, there are three main types of sampling schemes that one may try to adapt to graphs: periodic, irregular, and random.

Periodic graph sampling is ill-defined unless the graph is not bipartite (or multipartite) and has only been investigated in the design of graph filterbanks where the objective is to sample “one every two nodes” [7], [8].

Irregular graph sampling of k -bandlimited graph signals has been studied by Anis *et al.* [3] and Chen *et al.* [4] who, building upon Pensenson's work [2], look for optimal sampling sets of k -bandlimited graphs, in the sense that they are tight (*i.e.*, they contain only k nodes) and that they maximize noise-robustness of the reconstruction. To exhibit such an optimal set, one needs i) to compute the first k eigenvectors of the

Laplacian which becomes prohibitive at large scale, and ii) to optimize over all subsets of k nodes, which is a large combinatorial problem. Using greedy algorithms and spectral proxies [3] to sidestep some of the computational bottlenecks, these methods enable to efficiently perform optimal sampling on graphs of size up to a few thousands nodes.

Loosening the tightness constraint on sampling sets, random graph sampling [5] shows that if $O(k \log k)$ nodes are drawn with replacement according to a particular probability distribution p^* that depends on the first k eigenvectors of the Laplacian, then recovery is guaranteed with high probability. p^* was shown to be closely related to *leverage scores*, an important concept in randomized numerical linear algebra [9]. Moreover, an efficient estimation of p^* , bypassing expensive spectral computations, enables to sample very large graphs; as was shown in [10], where experiments were performed on graphs with $N \simeq 10^6$ and $k \simeq 200$.

In this paper, we detail another form of random graph sampling, motivated by a simple experimental context. Consider a graph composed of two almost disconnected communities of equal size. Its first Fourier mode is constant (as always), and its second one is positive in one community, and negative in the other, while its absolute value is more or less constant on the whole graph. Therefore, a 2-bandlimited graph signal will approximately be equal to a constant in one community, and to another constant in the other: any sampling set composed of one node drawn from each community will enable to recover the whole signal. Unfortunately, random graph sampling as proposed in [5] samples independently with replacement: it has a fair chance of sampling only in one community, thereby missing half of the information. To avoid such situations, we propose to take advantage of determinantal point processes (DPP), random processes that introduce negative correlations to force diversity within the sampling set.

Contributions. If the first k eigenvectors of the Laplacian are computable, we exhibit a DPP that always samples sets of size k perfectly embedding k -bandlimited graph signals. Moreover, in the case the graph is too large to compute its first k eigenvectors, and building upon recent results linking determinantal processes and random rooted spanning forests [11], we show to what extent a simple algorithm based on loop-erased random walks is an acceptable approximation.

II. BACKGROUND

Consider an undirected weighted graph $\mathcal{G}$ of N nodes represented by its adjacency matrix W , where $W_{ij} = W_{ji} \geq 0$ is the weight of the connection between nodes i and j . Denote by $L = D - W$ its Laplacian matrix, where D is the diagonal

matrix with entries $D_{ii} = \sum_j W_{ij}$. L is real and symmetrical, therefore diagonalizable as $U\Lambda U^\top$ with $U = (\mathbf{u}_1 | \dots | \mathbf{u}_N) \in \mathbb{R}^{N \times N}$ the graph Fourier basis and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_N)$ the diagonal matrix of the associated frequencies, that we sort: $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_N$. For any function $h : \mathbb{R}^+ \rightarrow \mathbb{R}$, we write $h(\Lambda) = \text{diag}(h(\lambda_1), \dots, h(\lambda_N))$. The graph filter associated to the frequency response h reads $U h(\Lambda) U^\top$.

Given the frequency interpretation of the graph Fourier modes, one may define “smooth” signals as linear combinations of the first few low-frequency Fourier modes. Writing $U_k = (\mathbf{u}_1 | \dots | \mathbf{u}_k) \in \mathbb{R}^{N \times k}$, we have the formal definition:

Definition II.1 (k -bandlimited signal on $\mathcal{G}$). A signal $\mathbf{x} \in \mathbb{R}^N$ defined on the nodes of the graph $\mathcal{G}$ is k -bandlimited with $k \in \mathbb{N}^*$ if $\mathbf{x} \in \text{span}(U_k)$, i.e., $\exists \boldsymbol{\alpha} \in \mathbb{R}^k$ such that $\mathbf{x} = U_k \boldsymbol{\alpha}$.

Sampling consists in selecting a subset $\mathcal{A} = (\omega_1, \dots, \omega_m)$ of m nodes of the graph and measuring the signal on it. To each possible sampling set, we associate a measurement matrix $M = (\boldsymbol{\delta}_{\omega_1} | \boldsymbol{\delta}_{\omega_2} | \dots | \boldsymbol{\delta}_{\omega_m})^\top \in \mathbb{R}^{m \times N}$ where $\boldsymbol{\delta}_{\omega_i}(j) = 1$ if $j = \omega_i$, and 0 otherwise. Given a k -bandlimited signal $\mathbf{x}$, its (possibly noisy) measurement on $\mathcal{A}$ reads:

$$\mathbf{y} = M\mathbf{x} + \mathbf{n} \in \mathbb{R}^m, \quad (1)$$

where $\mathbf{n}$ models any kind of measurement noise and/or the fact that $\mathbf{x}$ may not exactly be in $\text{span}(U_k)$. In this latter case, $\mathbf{x} = U_k U_k^\top \mathbf{x} + \mathbf{e}$, where $U_k U_k^\top \mathbf{x}$ is the closest signal to $\mathbf{x}$ that is in $\text{span}(U_k)$; and $M\mathbf{x} = M U_k U_k^\top \mathbf{x} + \mathbf{n}$ with $\mathbf{n} = M\mathbf{e}$.

A. If U_k is known

If U_k is known, the measurement on $\mathcal{A}$ reads $\mathbf{y} = M U_k \boldsymbol{\alpha} + \mathbf{n}$ and one recovers the signal by solving:

$$\mathbf{x}_{\text{rec}} = \underset{\mathbf{z} \in \text{span}(U_k)}{\text{argmin}} \|\mathbf{M}\mathbf{z} - \mathbf{y}\|^2 = U_k (M U_k)^\dagger \mathbf{y}, \quad (2)$$

where $(M U_k)^\dagger \in \mathbb{R}^{k \times m}$ is the Moore-Penrose pseudo-inverse of $M U_k \in \mathbb{R}^{m \times k}$ and $\|\cdot\|$ the Euclidean norm. Of course, perfect reconstruction is impossible if $m < k$, as the system would be underdetermined. Hereafter, we thus suppose $m \geq k$. Denote by $\sigma_1 \leq \dots \leq \sigma_k$ the singular values of $M U_k$.

Proposition II.1 (Classical result). If $\sigma_1 > 0$, then perfect reconstruction (up to the noise level) is possible.

Proof. If $\sigma_1 > 0$, then $U_k^\top M^\top M U_k$ is invertible; and since $\mathbf{x} = U_k \boldsymbol{\alpha}$, we have:

$$\begin{aligned} \mathbf{x}_{\text{rec}} &= U_k (M U_k)^\dagger \mathbf{y} = U_k (U_k^\top M^\top M U_k)^{-1} U_k^\top M^\top (M\mathbf{x} + \mathbf{n}) \\ &= \mathbf{x} + U_k (U_k^\top M^\top M U_k)^{-1} U_k^\top M^\top \mathbf{n}. \end{aligned} \quad \square$$

In the following, in the cases where U_k is known, we fix $m = k$. Chen et al. [4] showed that a sample of size k always exists s.t. $\sigma_1 > 0$. In fact, in general, many possible such subsets exist. Finding the optimal one is a matter of definition. Authors in [3], [4] propose two optimality definitions via noise robustness. A first option is to minimize the worst-case error, which translates to finding the subset that maximizes σ_1^2 :

$$\mathcal{A}^{\text{WCE}} = \underset{\mathcal{A} \text{ s.t. } |\mathcal{A}|=k}{\text{argmax}} \sigma_1^2. \quad (3)$$

Algorithm 1 Greedy algorithm to approximate (3), (4), or (5)

Input: U_k , objective = WCE, MSE or MV.

$\mathcal{Y} \leftarrow \emptyset$

while $|\mathcal{Y}| \neq k$ **do:**

for $n = 1, 2, \dots, N$ **do:**

$\mathcal{Y}_n \leftarrow \mathcal{Y} \cup \{n\}$

 • Compute $\{\sigma_i\}_{i=1, \dots, |\mathcal{Y}_n|}$ the singular values of the restriction of U_k to the lines indexed by $\mathcal{Y}_n$.

 • Store the gain on the objective of adding n

end for

 Add to $\mathcal{Y}$ the node for which the gain is maximal.

end while

Output: $\mathcal{Y}$.

A second option is to find the subset that minimizes the mean square error (assuming $\mathbb{E}(\mathbf{n}\mathbf{n}^\top)$ is the identity) :

$$\mathcal{A}^{\text{MSE}} = \underset{\mathcal{A} \text{ s.t. } |\mathcal{A}|=k}{\text{argmin}} \text{tr} [(U_k^\top M^\top M U_k)^{-1}] = \underset{\mathcal{A} \text{ s.t. } |\mathcal{A}|=k}{\text{argmin}} \sum_{i=1}^k \frac{1}{\sigma_i^2}, \quad (4)$$

where tr is the trace operator.

We argue that another possible definition of optimal subset is such that the determinant of $U_k^\top M^\top M U_k$ is maximal, i.e.:

$$\mathcal{A}^{\text{MV}} = \underset{\mathcal{A} \text{ s.t. } |\mathcal{A}|=k}{\text{argmax}} \prod_{i=1}^k \sigma_i^2, \quad (5)$$

where MV stands for Maximum Volume, as the determinant is proportional to the volume spanned by the k sampled lines of U_k . This definition of optimality does not have a straight-forward noise robustness interpretation, even though it is related to the previous definitions (in particular, increasing σ_1 necessarily increases the determinant). Nevertheless, such volume sampling as it is sometimes called [12] has been proved optimal in the related problem of low-rank approximation [13], and is largely studied in diverse contexts [14].

In all three cases, finding the optimal subset is a very large combinatorial problem (it implies a search over all combinations of k nodes): in practice, greedy algorithms such as Alg. 1 are used to find approximate solutions. In the following, we refer to $\hat{\mathcal{A}}^{\text{WCE}}$, $\hat{\mathcal{A}}^{\text{MSE}}$ and $\hat{\mathcal{A}}^{\text{MV}}$ the (deterministic) solutions obtained by greedy optimization of the above 3 objectives. We refer to $\hat{\mathcal{A}}^{\text{maxvol}}$ the solution to the `maxvol` algorithm proposed in [15], which is another approximation of $\mathcal{A}^{\text{MV}}$.

B. If U_k is unknown

If U_k is unknown, one may use a proxy that penalizes high frequencies in the solution, and thus recover an approximation of the original signal by solving the regularized problem:

$$\mathbf{x}_{\text{rec}} = \underset{\mathbf{z} \in \mathbb{R}^N}{\text{argmin}} \|\mathbf{M}\mathbf{z} - \mathbf{y}\|^2 + \gamma \mathbf{z}^\top \mathbf{L}^r \mathbf{z}, \quad (6)$$

where $\gamma > 0$ is the regularization parameter and r the power of the Laplacian that controls the strength of the high frequency penalization. By differentiating with respect to $\mathbf{z}$, one has:

$$(\mathbf{M}^\top \mathbf{M} + \gamma \mathbf{L}^r) \mathbf{x}_{\text{rec}} = \mathbf{M}^\top \mathbf{y}, \quad (7)$$

which may be solved by direct inversion if N is not too large, or –thanks to the objective’s convexity– by iterative methods such as gradient descent.

C. Reweighting in the case of random sampling

In random sampling, the sample $\mathcal{A} = (\omega_1, \dots, \omega_m)$ is a random variable, and therefore so is the measurement $\mathbf{y} = \mathbf{M}\mathbf{x}$. Consider two signals $\mathbf{x}_1$ and $\mathbf{x}_2 \neq \mathbf{x}_1$, and their measurements $\mathbf{y}_1$ and $\mathbf{y}_2$. If $\mathbf{y}_1 = \mathbf{y}_2$, perfect reconstruction is impossible. To prevent this, one possible solution is to reweight the measurement such that the expected norm of the reweighted measurement equals the signal’s norm. For instance, in [5], m nodes are drawn independently with replacement. At each draw, the probability to sample node i is $p_i^* = \|\mathbf{U}_k^\top \boldsymbol{\delta}_i\|^2 / k$. One has $\mathbb{E}(\|\mathbf{P}^{-1/2} \mathbf{M}\mathbf{x}\|^2) = \|\mathbf{x}\|^2$ if

$$\mathbf{P} = \text{diag}(mp_{\omega_1}^* | \dots | mp_{\omega_m}^*). \quad (8)$$

This guarantees that if $\|\mathbf{x}_1 - \mathbf{x}_2\| > 0$, then, for a large enough m , the reweighted measures will necessarily be distinct in the measurement space, i.e., $\|\mathbf{P}^{-1/2}(\mathbf{y}_1 - \mathbf{y}_2)\| > 0$, ensuring that $\mathbf{x}_1$ and $\mathbf{x}_2$ have a chance of being recovered.

Recovery with known $\mathbf{U}_k$ (Eq. (2)) thus transforms into

$$\mathbf{x}_{\text{rec}} = \underset{\mathbf{z} \in \text{span}(\mathbf{U}_k)}{\text{argmin}} \left\| \mathbf{P}^{-1/2} (\mathbf{M}\mathbf{z} - \mathbf{y}) \right\|^2; \quad (9)$$

and recovery with unknown $\mathbf{U}_k$ (Eq. (6)) transforms into

$$\mathbf{x}_{\text{rec}} = \underset{\mathbf{z} \in \mathbb{R}^N}{\text{argmin}} \left\| \mathbf{P}^{-1/2} (\mathbf{M}\mathbf{z} - \mathbf{y}) \right\|^2 + \gamma \mathbf{z}^\top \mathbf{L}^r \mathbf{z}. \quad (10)$$

III. DETERMINANTAL PROCESSES

Denote by $[N]$ the set of all subsets of $\{1, 2, \dots, N\}$.

Definition III.1 (Determinantal Point Process). *Consider a point process, i.e., a process that randomly draws an element $\mathcal{A} \in [N]$. It is determinantal if, for every $\mathcal{S} \subseteq \mathcal{A}$,*

$$\mathbb{P}(\mathcal{S} \subseteq \mathcal{A}) = \det(\mathbf{K}_{\mathcal{S}}),$$

where $\mathbf{K} \in \mathbb{R}^{N \times N}$, a semi-definite positive matrix s.t. $0 \preceq \mathbf{K} \preceq \mathbf{I}$, is called the marginal kernel; and $\mathbf{K}_{\mathcal{S}}$ is the restriction of $\mathbf{K}$ to the rows and columns indexed by the elements of $\mathcal{S}$.

A. Sampling from a DPP and signal recovery

Given the constraints on $\mathbf{K}$, its eigendecomposition $\mathbf{K} = \sum_{i=1}^N \mu_i \mathbf{v}_i \mathbf{v}_i^\top$ always exists, and Alg. 2 provides a sample from the associated DPP [14]. Note that $|\mathcal{A}|$, the number of elements in $\mathcal{A}$, is distributed as the sum of N Bernoulli trials of probability μ_i (see the “for loop” of Alg. 2). In particular:

$$\mathbb{E}(|\mathcal{A}|) = \text{tr}(\mathbf{K}) = \sum_{i=1}^N \mu_i \text{ and } \text{Var}(|\mathcal{A}|) = \sum_{i=1}^N \mu_i (1 - \mu_i). \quad (11)$$

Consider $\mathcal{A}$, the random variable of a DPP with kernel $\mathbf{K}$; and $\mathbf{M}$ its associated measurement matrix. Denote by $\pi_i = \mathbf{K}_{ii}$ the marginal probability that $\{i\} \subseteq \mathcal{A}$ and write:

$$\mathbf{P} = \text{diag}(\pi_{\omega_1} | \dots | \pi_{\omega_m}) \quad (12)$$

Algorithm 2 Sampling a DPP with marginal kernel $\mathbf{K}$ [14]

Input: Eigendecomposition of $\mathbf{K}$: $\{\mu_i, \mathbf{V} = (\mathbf{v}_1 | \dots | \mathbf{v}_N)\}$.
 $\mathcal{J} \leftarrow \emptyset$
for $n = 1, 2, \dots, N$ **do**:
 $\mathcal{J} \leftarrow \mathcal{J} \cup \{n\}$ with probability μ_n
end for
 $\mathbf{V} \leftarrow \{\mathbf{v}_n\}_{n \in \mathcal{J}}$
 $\mathbf{Y} \leftarrow \emptyset$
while $|\mathbf{U}| > 0$ **do**:
 • Select i from $\mathcal{Y}$ with $\text{Pr}(i) = \frac{1}{|\mathbf{V}|} \sum_{\mathbf{v} \in \mathbf{V}} (\mathbf{v}^\top \mathbf{e}_i)^2$
 • $\mathcal{Y} \leftarrow \mathcal{Y} \cup \{i\}$
 • $\mathbf{V} \leftarrow \mathbf{V}_{\perp}$, an orthonormal basis for the subspace of $\mathbf{V}$
 $\perp$ to $\mathbf{e}_i$.
end while
Output: $\mathcal{Y}$.

Proposition III.1. $\forall \mathbf{x} \in \mathbb{R}^N \quad \mathbb{E}_{\mathcal{A}} \left(\|\mathbf{P}^{-1/2} \mathbf{M}\mathbf{x}\|^2 \right) = \|\mathbf{x}\|^2$.

Proof. Let $\mathbf{x} \in \mathbb{R}^N$. In the following, we need the probability $\mathbb{P}(\mathcal{A})$ that $\mathcal{A}$ is sampled (instead of marginal probabilities). One can show that $\mathbb{P}(\mathcal{A})$ is proportional to the determinant of the restriction to $\mathcal{A}$ of a matrix called L -ensemble [14]. Note that: $\forall i \quad \sum_{\mathcal{A} \ni i} \mathbb{P}(\mathcal{A}) = \mathbb{P}(\{i\} \text{ is sampled}) = \pi_i$. Thus:

$$\begin{aligned} \mathbb{E}_{\mathcal{A}} \left(\|\mathbf{P}^{-1/2} \mathbf{M}\mathbf{x}\|^2 \right) &= \mathbb{E}_{\mathcal{A}} \left(\sum_{i \in \mathcal{A}} \frac{x_i^2}{\pi_i} \right) = \sum_{\mathcal{A}} \mathbb{P}(\mathcal{A}) \sum_{i \in \mathcal{A}} \frac{x_i^2}{\pi_i} \\ &= \sum_{i=1}^N x_i^2 \sum_{\mathcal{A} \ni i} \frac{\mathbb{P}(\mathcal{A})}{\pi_i} = \sum_{i=1}^N x_i^2 = \|\mathbf{x}\|^2 \quad \square \end{aligned}$$

This result is the first step to prove a restricted isometry property (RIP) for $\mathbf{P}^{-1/2} \mathbf{M}$ – currently work in progress. In this paper’s context, it shows how to reweight the measurement in the case of sampling from a DPP: recovery should be performed via optimization of (9) or (10) with $\mathbf{P}$ as in (12).

B. If $\mathbf{U}_k$ is known: the ideal low-pass marginal kernel

Consider the DPP defined by the following marginal kernel:

$$\mathbf{K}_k = \mathbf{U}_k \mathbf{U}_k^\top = \mathbf{U} h_{\lambda_k}(\Lambda) \mathbf{U}^\top, \quad (13)$$

with h_{λ_k} s.t. $h_{\lambda_k}(\lambda) = 1$ if $\lambda \leq \lambda_k$ and 0 otherwise. In GSP words, $\mathbf{K}_k$ is the ideal low-pass with cutting frequency λ_k .

Proposition III.2. *A sample from the DPP with marginal kernel $\mathbf{K}_k$ defined as in (13) is of size k and $\forall \mathbf{x} \in \text{span}(\mathbf{U}_k)$, the measurement $\mathbf{y} = \mathbf{M}\mathbf{x} + \mathbf{n} \in \mathbb{R}^k$ enables perfect reconstruction up to the noise level.*

Proof. $\mathbf{K}_k = \sum_{i=1}^N \mu_i \mathbf{u}_i \mathbf{u}_i^\top$, with $\mu_1 = \dots = \mu_k = 1$ and $\mu_{k+1} = \dots = \mu_N = 0$. According to (11), $\mathbb{E}(|\mathcal{A}|) = k$ and $\text{Var}(|\mathcal{A}|) = 0$, implying that any sample $\mathcal{A}$ has size k . Also, by definition of a DPP, the fact that $\mathcal{A}$ is sampled implies:

$$\det(\mathbf{K}_{\mathcal{A}}) > 0, \text{ i.e., } \det(\mathbf{M} \mathbf{U}_k \mathbf{U}_k^\top \mathbf{M}^\top) = \prod_{i=1}^k \sigma_i^2 > 0, \quad (14)$$

implying $\sigma_1 > 0$. Using Prop. II.1 completes the proof. $\square$

Algorithm 3 Wilson’s algorithm (see Sec. III-C)

Input: Adjacency matrix W of the graph and $q > 0$
 $\mathcal{V} \leftarrow \emptyset, \mathcal{V} \leftarrow \emptyset$
while $\mathcal{V} \neq \{1, \dots, N\}$ **do**:
 • Start a random walk from any node $i \in \{1, \dots, N\} \setminus \mathcal{V}$
 until it reaches either Δ , or a node in $\mathcal{V}$.
 • Erase all the loops of the trajectory and denote by $\mathcal{S}$
 the set of all new visited nodes (after erasure).
 • $\mathcal{V} \leftarrow \mathcal{V} \cup \mathcal{S}$
if the last node of the trajectory is Δ **do**:
 • Denote by l the last visited node before Δ
 • $\mathcal{V} \leftarrow \mathcal{V} \cup \{l\}$
Output: $\mathcal{V}$.

The most probable sample from this DPP is the sample $\mathcal{A}$ of size k that maximizes $\det(\mathbf{M}\mathbf{U}_k\mathbf{U}_k^\top\mathbf{M}^\top)$ i.e., $\mathcal{A}^{\text{MV}}$, the solution to the maximal volume optimisation problem (5).

C. If $\mathbf{U}_k$ is unknown: Wilson’s marginal kernel

First, add a node Δ to the graph, that we will call “absorbing state” and that we connect to all existing N nodes with a weight q . We will propagate random walks on this extended graph. Note that the probability to jump from any node $i \neq \Delta$:

- to any other node $j \neq \Delta$ is equal to $W_{ij}/(D_{ii} + q)$;
- to Δ is $q/(D_{ii} + q)$.

Once the random walk reaches Δ , it cannot escape from it. Wilson and Propp [16], [17] introduced Algorithm 3 in the context of spanning tree sampling. This Algorithm outputs a set $\mathcal{V}$ of nodes by propagating loop-erased random walks.

Proposition III.3 (Corollary of Thm 3.4 in [11]). $\mathcal{V}$ is a sample from a DPP with marginal kernel $K_q = \mathbf{U}g_q(\Lambda)\mathbf{U}^\top$, with $g_q(\lambda) = \frac{q}{q+\lambda}$.

Proof. Given that the weight q we added is finite and constant over all nodes, Thm 3.4 and Lemma 3.3 in [11] state that $\mathcal{V}$ is distributed as a DPP with marginal kernel equal to $(\mathbf{L} + q\mathbf{I})^{-1}q\mathbf{I}$ where $\mathbf{I}$ is the identity matrix (note that L in [11] refers to $-\mathbf{L}$ in our paper). Decomposing $\mathbf{L}$ as $\mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top$ finishes the proof. $\square$

Prop. III.2 shows that sampling with K_k is optimal in the sense it enables perfect reconstruction with only $m = k$ measurements. Unfortunately, as N or k increase, computing K_k becomes prohibitive. Considering g_q as an approximation of h_{λ_k} , one may see sampling with K_q as an approximation of optimal sampling with K_k . Prop. III.3 has a beautiful consequence: one does not need to explicitly compute K_q (which would also be too expensive) to sample from it. In fact, Alg. 3 efficiently samples from K_q . Thus, Alg. 3 approximates sampling from K_k without any spectral computation of $\mathbf{L}$.

IV. EXPERIMENTS

The Stochastic Block Model (SBM). We consider random community-structured graphs drawn from the SBM. We specifically look at graphs with k communities of same size N/k . In the SBM, the probability of connection between any

Algorithm 4 Estimation of $\pi_i = \sum_{j=1}^N g_q(\lambda_j)\mathbf{u}_j(i)^2$

Input: Laplacian $\mathbf{L}$, $g_q : \mathbb{R} \rightarrow \mathbb{R}^+$, parameters d and n
 • Compute λ_N the largest eigenvalue of $\mathbf{L}$
 • Approximate $s = \sqrt{g_q}$ on $[0, \lambda_N]$ with a polynomial of order d i.e., $S = \mathbf{U}s(\Lambda)\mathbf{U}^\top \simeq \mathbf{U}\sum_{l=1}^d \beta_l \Lambda^l \mathbf{U}^\top = \sum_{l=1}^d \beta_l \mathbf{L}^l$
 • Generate $\mathbf{R} \in \mathbb{R}^{N \times n}$ with $\forall(i, j) \quad R_{ij} = \mathcal{N}(0, 1/n)$
 • Compute $\mathbf{SR} \simeq \sum_{l=1}^d \beta_l \mathbf{L}^l \mathbf{R}$ recursively
Output: $\forall i \quad \hat{\pi}_i = \|\delta_i^\top \mathbf{SR}\|^2$

two nodes i and j is q_1 if they are in the same community, and q_2 otherwise. One can show that the average degree reads $c = q_1 (\frac{N}{k} - 1) + q_2 (N - \frac{N}{k})$. Thus, instead of providing the probabilities (q_1, q_2) , one may characterize a SBM by considering $(\epsilon = \frac{q_2}{q_1}, c)$. The larger ϵ , the fuzzier the community structure. In fact, authors in [18] show that above the critical value $\epsilon_c = (c - \sqrt{c})/(c + \sqrt{c}(k - 1))$, community structure becomes undetectable in the large N limit. In the following, $N = 100$, $k = 2$, $c = 16$ and we let ϵ vary.

To generate a k -bandlimited signal, we draw k realisations of a normal distribution with mean 0 and variance 1, to build the vector $\alpha \in \mathbb{R}^k$. Then, renormalize to make sure that $\|\alpha\|^2 = 1$, and write $\mathbf{x} = \mathbf{U}_k \alpha$. In our experiments, the measurement noise is Gaussian with mean 0 and $\sigma_n = 10^{-4}$.

First series: $\mathbf{U}_k$ known. We compare DPP sampling from K_k followed by recovery using (9); to deterministic sampling with $\hat{\mathcal{A}}^{\text{WCE}}, \hat{\mathcal{A}}^{\text{MSE}}, \hat{\mathcal{A}}^{\text{MV}}$ and $\hat{\mathcal{A}}^{\text{maxvol}}$ presented in Sec. II-A followed by recovery using (2). We show in Fig. 1a) the recovery performance with respect to ϵ/ϵ_c : all greedy methods perform similarly, and on average slightly outperform the DPP-based method, especially as the graphs become less structured.

Second series: $\mathbf{U}_k$ unknown. We compare negatively correlated sampling from K_q (using Wilson’s algorithm) followed by recovery using (10) with $\mathbf{P}$ as in (12); to the uncorrelated random sampling [5] where nodes are sampled independently with replacement from p^* with $p_i^* = \|\mathbf{U}_k^\top \delta_i\|^2 / k$, followed by recovery using (10) with $\mathbf{P}$ as in (8). Several comments are in order. i) the number of samples from Wilson’s algorithm is not known in advance, but we know its expected value $\mathbb{E}(|\mathcal{A}|) = \sum_i q/(q + \lambda_i)$. In order to explore the behavior around the critical number of measurements $m = k$, one needs to choose q s.t. $\sum_i q/(q + \lambda_i) \simeq k$, but without computing the λ_i ! To do so, we follow the proposition of [11] (see discussion around Eq. (4.14)) based on a few runs of Alg. 3. Once a sample of size approximately equal to k is exhibited, and for fair comparison, one uses the same number of nodes for the uncorrelated sampling. ii) exact computation of p^* for the uncorrelated sampling requires $\mathbf{U}_k$. We follow the efficient Alg. 1 of [5] to approximate p^* . iii) recovery using $\mathbf{P}$ as in (12) requires to know $\pi_i = \sum_{j=1}^N g_q(\lambda_j)\mathbf{u}_j(i)^2$ for each node i . To estimate π_i without spectral decomposition of $\mathbf{L}$, an elegant solution is via fast graph filtering of random signals, as in Alg. 4. Building upon the Johnson-Lindenstrauss lemma, one can show that $\hat{\pi}_i$ concentrates around π_i for $n = O(\log N)$ (see a similar proof in [19]). We fix $n = 20 \log N$ and $d = 30$.




Fig. 1: a) Comparison of the recovery performance after noisy measurements of k -bandlimited signals between DPP sampling with K_k , and the 4 sampling methods presented in Sec. II-A (grouped here as “greedy”: they perform equivalently). The intervals represent the 10- and 90-percentiles; the lines the averages, computed over 10^4 different signals defined on 100 different SBM realisations for each value of ϵ . b) Reconstruction performance versus the number of measurements m for different values of the regularization parameter γ , for $\epsilon = \epsilon_c/5$. c) Performance versus m for different values of ϵ . In b) and c), the full (resp. dashed) lines represent results obtained after DPP sampling via Wilson’s algorithm (resp. after uncorrelated sampling such as in [5]). For each value of m , we average over 3500 bandlimited signals defined on 100 different SBM realisations.

The reconstruction’s parameters are γ and r . Following [5], r is fixed to 4. Fig. 1b) shows that the reconstruction performance saturates at $\gamma = 10^{-5}$. With these values of r and γ , we compare DPP sampling with K_q vs uncorrelated sampling in Fig. 1c) with respect to the number of measurements, for different values of ϵ . The lower ϵ , *i.e.* the stronger the community structure, the better DPP sampling is compared to uncorrelated sampling. In terms of computation time, a precise comparison of both methods versus the different parameters is out of this paper’s scope. Nevertheless, to give an idea of the method’s scalability, for a SBM with $N = 10^5$ (resp. 10^6), $k = 2$, $c = 16$ and $\epsilon = \epsilon_c/5$, and with $q = 5 \cdot 10^{-4}$, Wilson’s algorithm outputs in average 5 (resp. 36) samples in a mean time of 7 (resp. 90) seconds, using Python on a laptop.

V. CONCLUSION

We first show that sampling from a DPP with marginal kernel $K_k = U_k U_k^T$ *i.e.* the projector onto the first k graph Fourier modes, outputs a subset of size k that enables perfect reconstruction. The robustness to noise of the reconstruction is comparable to state-of-the-art greedy sampling algorithms. Moreover, in the important and fairly common case where the first k eigenvectors of the Laplacian are not computable, we show that Wilson’s algorithm (Alg. 3) may be leveraged as an efficient graph sampling scheme that approximates sampling from K_k . Preliminary experiments on the SBM suggest that in the cases where the number of samples m needs to stay close to k and where the community structure of the graph is strong, sampling with Wilson’s algorithm enables a better reconstruction than the state-of-the-art uncorrelated random sampling. Further investigation, both theoretical and experimental, is necessary to better specify which families of graphs are favorable to DPP sampling, and which are not.

REFERENCES

- [1] D. Shuman, S. Narang, P. Frossard, A. Ortega, and P. Vandergheynst, “The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains,” *Signal Processing Magazine, IEEE*, vol. 30, no. 3, pp. 83–98, May 2013.
- [2] I. Pesenson, “Sampling in Paley-Wiener spaces on combinatorial graphs,” *Transactions of the AMS*, vol. 360, no. 10, 2008.
- [3] A. Anis, A. Gadde, and A. Ortega, “Efficient Sampling Set Selection for Bandlimited Graph Signals Using Graph Spectral Proxies,” *IEEE Transactions on Signal Processing*, vol. 64, no. 14, pp. 3775–3789, 2016.
- [4] S. Chen, R. Varma, A. Sandryhaila, and J. Kovachević, “Discrete signal processing on graphs: Sampling theory,” *IEEE Transactions on Signal Processing*, vol. 63, no. 24, pp. 6510–6523, 2015.
- [5] G. Puy, N. Tremblay, R. Gribonval, and P. Vandergheynst, “Random sampling of bandlimited signals on graphs,” *Applied and Computational Harmonic Analysis*, 2016, in press.
- [6] F. Gama, A. G. Marques, G. Mateos, and A. Ribeiro, “Rethinking sketching as sampling: Linear transforms of graph signals,” in *50th Asilomar Conf. Signals, Systems and Computers*, 2016.
- [7] D. I. Shuman, M. J. Faraji, and P. Vandergheynst, “A Multiscale Pyramid Transform for Graph Signals,” *IEEE Transactions on Signal Processing*, vol. 64, no. 8, pp. 2119–2134, Apr. 2016.
- [8] S. Narang and A. Ortega, “Compact Support Biorthogonal Wavelet Filterbanks for Arbitrary Undirected Graphs,” *Signal Processing, IEEE Transactions on*, vol. 61, no. 19, pp. 4673–4685, Oct. 2013.
- [9] P. Drineas and M. W. Mahoney, “RandNLA: randomized numerical linear algebra,” *Communications of the ACM*, vol. 59, no. 6, 2016.
- [10] N. Tremblay, G. Puy, R. Gribonval, and P. Vandergheynst, “Compressive spectral clustering,” in *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, vol. 48, 2016, pp. 1002–1011.
- [11] L. Avena and A. Gaudillière, “On some random forests with determinantal roots,” *arXiv preprint arXiv:1310.1723*, 2013.
- [12] A. Deshpande and L. Rademacher, “Efficient volume sampling for row/column subset selection,” in *51st Annual IEEE Symposium on the Foundations of Computer Science (FOCS)*, 2010, pp. 329–338.
- [13] S. A. Goreinov and E. E. Tyrtyshnikov, “The maximal-volume concept in approximation by low-rank matrices,” *Contemporary Mathematics*, vol. 280, pp. 47–52, 2001.
- [14] A. Kulesza and B. Taskar, “Determinantal Point Processes for Machine Learning,” *Found. and Trends in Mach. Learn.*, vol. 5, no. 2–3, pp. 123–286, 2012.
- [15] S. A. Goreinov, I. V. Oseledets, D. V. Savostyanov, E. E. Tyrtyshnikov, and N. L. Zamarashkin, “How to find a good submatrix,” *Research Report 08-10, ICM HKBU, Kowloon Tong, Hong Kong*, pp. 08–10, 2008.
- [16] D. B. Wilson, “Generating random spanning trees more quickly than the cover time,” in *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*. ACM, 1996, pp. 296–303.
- [17] J. G. Propp and D. B. Wilson, “How to get a perfectly random sample from a generic markov chain and generate a random spanning tree of a directed graph,” *Journal of Algorithms*, vol. 27, no. 2, 1998.
- [18] A. Decelle, F. Krzakala, C. Moore, and L. Zdeborová, “Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications,” *Phys. Rev. E*, vol. 84, no. 6, 2011.
- [19] N. Tremblay, G. Puy, P. Borgnat, R. Gribonval, and P. Vandergheynst, “Accelerated Spectral Clustering Using Graph Filtering Of Random Signals,” in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 4094–4098.