



HAL
open science

Inference problems in large random graphs

Ludovic Stephan

► **To cite this version:**

Ludovic Stephan. Inference problems in large random graphs. General Mathematics [math.GM]. Sorbonne Université, 2021. English. NNT : 2021SORUS497 . tel-03720004

HAL Id: tel-03720004

<https://theses.hal.science/tel-03720004>

Submitted on 11 Jul 2022

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.



Sorbonne Université
Ecole Doctorale Sciences Mathématiques de Paris Centre

*Equipe DYOGENE
Inria Paris, DI ENS*

Inférence dans des grands graphes aléatoires

Ludovic Stephan

Thèse de doctorat en Mathématiques Appliquées
Sous la direction de **Laurent Massoulié**

Présentée et soutenue publiquement le 22 octobre 2021, devant un jury composé de

M.	Couillet	Romain	Professeur	Rapporteur
Mme	Péché	Sandrine	Professeure	Examinatrice
M.	Broutin	Nicolas	Professeur	Examinateur
Mme	Zdeborová	Lenka	Professeure associée	Examinatrice
M.	Massoulié	Laurent	Directeur de Recherche	Directeur de thèse

Remerciements

I don't know half of you half as well as I should like; and I like less than half of you half as well as you deserve.

Bilbo Baggins, *The Lord of the Rings*

Je remercie tout d'abord mon directeur de thèse, Laurent Massoulié, que ce soit pour avoir accepté de m'accueillir ou pour son accompagnement tout au long de ces trois années. Que ce soit scientifiquement ou humainement, je n'aurais pu demander un meilleur directeur, et je lui suis très reconnaissant de m'avoir si bien guidé dans mon apprentissage de la recherche.

M'étant retrouvé dans une situation administrative complexe, il m'a fallu bien de l'assistance pour réussir à m'en dépêtrer et travailler dans les meilleures conditions. De ce point de vue, Hélène Bessin-Rousseau (pour l'Inria), Sophie Jaudon (pour l'ENS) et Corentin Lacombe (pour l'ED) ont été d'une grande patience et d'une aide précieuse, et ce travail aurait été bien moins conséquent sans eux.

Un laboratoire ne vaut que par les personnes qui le composent, et leur bonne humeur ; de ce côté là, l'équipe Dyogene est très bien fournie, et je remercie Simon, Luca, Bastien, Hélène (encore !), Marc, Mathieu, Maxime, Matthieu et Ilia, pour toutes les discussions mathématiques ou non, et pour m'avoir donné l'envie d'aller au travail même quand ma motivation me faisait défaut.

Enfin, mes derniers remerciements vont à mes ami.e.s et ma famille, qui m'ont supporté pendant ces dernières années avec énormément de bienveillance et d'amour. J'ai l'immense chance de ne pas avoir la place de les lister ici, mais j'ai bon espoir qu'ils et elles se reconnaîtront, et savent toute l'affection que j'ai pour eux. Merci enfin à Sarah, pour avoir partagé ma vie et mes galères pendant tout ce temps, ainsi qu'à Ron et Magellan pour leur amour et leur fluff inconditionnels.

Résumé en français

On s'intéresse dans ce manuscrit à divers problèmes d'inférence dans des grands graphes aléatoires. Sauf mention expresse du contraire, les graphes étudiés sont « creux », c'est-à-dire que le degré moyen d'un sommet du graphe ne dépend pas de la taille du graphe. Ce régime de faible densité est connu pour sa difficulté d'étude, en raison notamment d'une proportion non négligeable de sommets isolés et d'une forte variabilité dans les degrés des sommets.

On s'intéresse dans un premier temps à une variation du problème connu de la clique plantée. Considérons le modèle suivant : étant donné un arbre Γ de taille k et de forme supposée connue, on insère une copie de Γ dans un graphe d'Erdős-Rényi $\mathcal{G}(n, \lambda/n)$. Les questions classiques de l'inférence dans les graphes se posent alors : premièrement, est-il possible de différencier ce modèle d'un graphe d'Erdős-Rényi de mêmes paramètres, et deuxièmement, est-il possible de retrouver la copie cachée de Γ , ou tout du moins une fraction non négligeable d'icelle ? On se pose aussi la question de l'existence d'algorithmes polynômiaux qui réalisent ces deux tâches, respectivement appelées détection et reconstruction.

Étant donné la difficulté d'obtenir des résultats sur des arbres Γ généraux, ce manuscrit se restreint à des formes d'arbres spécifiques : un chemin de longueur k , ou un arbre d -régulier de hauteur h (et donc de taille d^h), déjà plus difficile à étudier. Les résultats obtenus diffèrent fortement du paysage habituel des problèmes d'inférence similaires, comme la clique plantée.

La première particularité est l'absence de phase « difficile » : dès que la détection ou la reconstruction sont possibles, il existe un algorithme polynômial pour ces tâches. Cela diffère de beaucoup de problèmes d'inférence étudiés précédemment, qui comportent souvent une phase où la détection et la reconstruction sont possibles, mais où l'on conjecture qu'aucun algorithme polynômial n'en est capable.

La seconde différence notable est le « découplage » de la détection et de la reconstruction : pour un large choix de paramètres, la détection est trivialement possible (via un comptage d'arêtes), mais il est impossible de reconstruire ne serait-ce qu'une fraction non-négligeable des sommets de Γ lorsque n tend vers $+\infty$. Là encore, c'est un départ notable du paysage attendu, puisque dans l'immense majorité des cas la détection implique au moins qu'une fraction de la structure plantée est retrouvable.

Une explication possible pour ces différences est la notion de « rang » de la structure plantée, que l'on définit comme le rang de sa matrice d'adjacence : la majorité des problèmes d'inférence étudiés historiquement concernent des structures de faible rang, ce qui n'est pas le cas des arbres considérés. Cette supposition est supportée par d'autres travaux récents sur des problèmes d'inférences de fort rang, qui présentent eux aussi des comportements inhabituels.

On s'intéresse dans un deuxième chapitre au problème de la détection de communautés : étant donné un graphe dont on suppose qu'il est partitionné en plusieurs communautés, et que cette partition a une influence sur la présence ou l'absence de liens, le but est de reconstruire ces communautés le plus fidèlement possible. Pour l'étude théorique de ce problème, le modèle le plus classique est le modèle stochastique par blocs : il s'agit d'un graphe aléatoire G , dont les sommets sont partitionnés en r communautés de taille comparable, et telle que la probabilité de présence d'une arête ne dépend que des communautés auxquelles appartiennent ses extrémités.

Ici encore, on s'intéresse à la version peu dense de ce modèle, où les probabilités de connexion sont d'ordre $1/n$. Dans ce régime, il est impossible de reconstruire intégralement les communautés, et on s'intéresse donc uniquement à la reconstruction partielle. Plusieurs classes d'algorithmes ont été proposés pour répondre à ce problème : parmi eux, les algorithmes spectraux, qui consistent à considérer les vecteurs propres d'une matrice associée à G pour en tirer de l'information sur les communautés. La première contribution de ce chapitre est la conception, ainsi que la démonstration du bon fonctionnement, d'une nouvelle méthode spectrale. Celle-ci est basée sur la matrice des distances $D^{(\ell)}$, dont le coefficient (x, y) vaut 1 si et seulement si x et y sont à distance exactement ℓ .

La deuxième question à laquelle on s'intéresse dans ce chapitre est celle de la robustesse des méthodes : étant donné que le modèle stochastique par blocs ne correspond pas forcément aux graphes réels, on étudie l'influence de petites perturbations (ajouts ou retractions d'arêtes) sur la performance des algorithmes. Les méthodes spectrales sont connues pour être extrêmement sensibles à ce type de modification, et échouent très rapidement. Par contraste, on montre que la nouvelle méthode basée sur la matrice $D^{(\ell)}$ est beaucoup plus robuste aux perturbations: elle continue de fonctionner malgré l'ajout de jusqu'à n^ε arêtes, où $\varepsilon > 0$. Ce résultat n'atteint pas la robustesse des algorithmes dits « semi-définis », mais notre algorithme a d'autres avantages : il est plus rapide, et peut détecter des communautés dans une plus large gamme de paramètres.

Les deux derniers chapitres de ce manuscrit portent sur une généralisation du modèle stochastique par blocs. En effet, celui-ci présente plusieurs limites dans son adéquation avec les réseaux réels, et de nombreuses variantes ont été développées : rajouter des informations supplémentaires aux arêtes, corriger la distribution des degrés ou encore travailler sur des graphes dirigés. Nous introduisons donc dans ce manuscrit un modèle très général de graphe aléatoire pondéré qui inclut toutes les variations susmentionnées ; les seules conditions imposées sont l'indépendance des arêtes, la faible densité du graphe, ainsi que le faible rang de la matrice d'adjacence, qui sont toutes vérifiées dans les modèles qui nous intéressent.

Dans le cas non dirigé et peu dense (3ème chapitre), il est connu que le spectre de la matrice d'adjacence ne donne pas d'information sur la structure sous-jacente du graphe : les vecteurs propres sont localisés autour des sommets de forts degrés. C'est entre autres pour cette raison que nous avons utilisé des matrices auxiliaires telles que $D^{(\ell)}$ précédemment. On étudie dans ce chapitre la matrice « non-backtracking » B , déjà utilisée dans le cadre de la détection de communautés. En particulier, on montre que son spectre suit un comportement déjà observé dans le chapitre précédent : ses plus grandes valeurs propres suivent celles de l'espérance de la matrice d'adjacence A , tandis que le reste du spectre est confiné dans un cercle de rayon connu. Ce résultat est valide dans une vaste gamme de degrés, y compris dans des régimes plus denses.

Lorsque les degrés du graphe G sont suffisamment grands, ce résultat sur la matrice B a une conséquence particulièrement intéressante. En effet, via la formule d'Ihara-Bass reliant le spectre de B et celui de A , on retrouve des résultats très précis (avec un terme d'erreur tendant vers 0) sur les valeurs propres de A , ainsi que sur la corrélation entre les vecteurs propres de A et ceux de son espérance. Ces résultats reflètent une transition de phase connue sous le nom de Baik-Ben Arous-Péché, l'étendant ainsi des matrices de Wigner aux graphes aléatoires.

Dans le cas dirigé (4ème chapitre), les efforts des algorithmes de classification se portent habituellement sur des symétrisations de la matrice d'adjacence, par exemple sur le calcul de ses valeurs singulières (qui correspondent aux valeurs propres de A^*A). Par contraste, ce manuscrit montre que la même méthode que celle utilisée pour la matrice B s'applique directement à la matrice d'adjacence asymétrique A ; des simulations numériques semblent montrer que l'utilisation des vecteurs propres à gauche et à droite de A dans un algorithme de classification résulte en une performance similaire aux algorithmes de l'état de l'art.

Enfin, dans les deux cas, des expériences numériques semblent montrer que la distribution empirique des vecteurs propres de A (resp. B) tend vers un mélange de gaussiennes de paramètres calculables théoriquement. Cela expliquerait l'efficacité des algorithmes développés dans le dernier chapitre, qui sont entre autres basés sur l'estimation de ces mélanges.

Contents

1	Introduction	1
1.1	Inference of planted subgraphs	1
1.1.1	From NP-hard to planted problems	1
1.1.2	Hard phase in the planted clique problem	2
1.1.3	The planted tree problem	3
1.1.4	Counting subgraphs in random graphs	3
1.1.5	Phase transition landscape of the planted tree problem	4
1.2	Spectral algorithms in community detection	6
1.2.1	Simulating real-world networks	6
1.2.2	Sparse SBM : definitions	7
1.2.3	Tree approximation and fundamental limits	8
1.2.4	Classifying reconstruction algorithms	9
1.2.5	Spectral algorithms	10
1.2.6	Robustness of community detection	11
1.3	Spectra of weighted inhomogeneous random graphs	13
1.3.1	Beyond the classical SBM	13
1.3.2	A generic random graph model	14
1.3.3	Aside : Phase transitions in random matrix theory	15
1.3.4	Non-backtracking matrix of undirected graphs	15
1.3.5	Adjacency matrix of directed graphs	17
1.3.6	Perspectives : eigenvector limit distribution	18
2	Planting trees in graphs	19
2.1	Introduction	19
2.2	Related work	20
2.3	Model and main results	20
2.4	Preliminary results	22
2.5	Proof strategy for planted paths	23
2.5.1	Impossibility of detection	23
2.5.2	Easiness of detection and reconstruction, sparse case	24
2.5.3	Impossibility of reconstruction, dense case	24
2.6	Proof strategy for planted D -ary trees	25
2.6.1	Study of the Galton-Watson process	26
2.6.2	Coupling and application to planted trees	27
2.6.3	Likelihood ratio and detection for $h \leq \underline{h}$	27
2.6.4	Reconstruction for large h	28
2.7	Proof of preliminary results	28
2.7.1	Proof of Lemma 1	28
2.7.2	Proof of Theorem 1	28

2.7.3	Proof of Theorem 4	29
2.7.4	Proof of Theorem 5	30
2.8	Detailed proofs for planted paths	32
2.8.1	Proof of Lemma 2	32
2.8.2	Proof of Lemma 4	32
2.8.3	Proof of Lemma 3	33
2.8.4	Proof of Lemma 5	34
2.8.5	Proof of Theorem 7	35
2.8.6	Proof of Lemma 6, Equation (2.8)	36
2.8.7	Proof of of Lemma 6, Equation (2.9)	38
2.9	Proofs for planted D -ary trees	40
2.9.1	Proof of Lemma 7	40
2.9.2	Proof of Theorem 8	40
2.9.3	Proof of Lemma 8	41
2.9.4	Proof of Theorem 9	42
2.9.5	Proof of Lemma 9	42
2.9.6	Proof of Theorem 11	44
3	Robustness of community detection in sparse networks	46
3.1	Introduction	46
3.1.1	Background	46
3.1.2	Summary of main results	47
3.1.3	Detailed setting and results	49
3.1.4	Notations and outline of the paper	52
3.2	Spectral structure of $B^{(\ell)}$	52
3.2.1	A theorem on eigenspace perturbation	52
3.2.2	Strategy of proof	52
3.2.3	A new reconstruction algorithm	53
3.3	Study of the matrix $D^{(\ell)}$	54
3.3.1	From $B^{(\ell)}$ to $D^{(\ell)}$	54
3.3.2	Stability to graph perturbation	55
3.4	Proof of Propositions 2 and 3	56
3.4.1	Outline of the proof and similarities with [29]	56
3.4.2	Local structure of G	56
3.4.3	From local neighbourhoods to the matrix $B^{(\ell)}$	57
3.4.4	Ramanujan property of $B^{(\ell)}$	58
3.4.5	Proof of Proposition 3	60
3.5	Proofs for Theorem 16	62
3.5.1	Proof of Lemma 14	62
3.5.2	Proof of limit (3.15)	64
3.5.3	Explicit bounds on K	64
3.6	Proof of Lemma 15	67
3.7	Proof of Lemma 16	68
3.8	Proof of Proposition 5	68
3.9	Proof of Theorem 19	69

4	Non-backtracking spectrum of inhomogeneous random graphs	71
4.1	Introduction	71
4.2	Detailed setting and results	72
4.2.1	Notations	72
4.2.2	Defining the convergence parameters	73
4.2.3	Main theorem	74
4.3	Applications	75
4.3.1	Phase transition in random graphs	75
4.3.2	Community detection in random networks	78
4.3.3	Extension to gaussian weights	81
4.4	A Bauer-Fike type bound for almost orthogonal diagonalization	81
4.4.1	A custom perturbation lemma for almost diagonalizable matrices	82
4.4.2	Matrix power perturbation and phase perturbation control	86
4.5	Proof of Theorem 22	88
4.5.1	Structure of the matrices U and V	89
4.5.2	Proof of the perturbation bounds	90
4.6	Preliminary computations	92
4.7	Local study of G	96
4.7.1	Setting and definitions	96
4.7.2	Growth properties of trees and graphs	96
4.7.3	Local tree-like structure	97
4.7.4	Coupling between rooted graphs and trees	98
4.8	Near eigenvectors of G	101
4.8.1	Functionals on (T, o)	101
4.8.2	Spatial averaging of graph functionals	102
4.8.3	Structure of near eigenvectors	104
4.9	Proof of Theorem 30	105
4.9.1	A telescopic trick: proof of (4.29)	105
4.9.2	Bounding $\ B^\ell P_{H^\perp}\ $	106
4.10	Trace method: proof of Proposition 21	109
4.11	Applications of Theorem 22	114
4.11.1	Proof of Proposition 11	114
4.11.2	Proof of Theorem 23	114
4.11.3	Proof of Theorem 24	117
4.11.4	Proof of Theorem 25 and Proposition 12	117
4.11.5	Proof of Theorem 26	118
4.12	Computing functionals on trees	118
4.12.1	Study of compound Poisson processes	119
4.12.2	Decomposing the tree functionals	120
4.12.3	Edge functionals	121
4.13	Near eigenvectors: computations	122
4.14	Proofs for Theorem 30	125
4.14.1	Proof of (4.26)-(4.28)	125
4.14.2	Bounding $\ B^\ell\ $: proof of (4.31)	126
4.14.3	Proof of Lemma 36	126
4.15	Norm bounds: additional proofs	127
4.15.1	Bound (4.69) on $\ MB^k\ $	127
4.15.2	Bound (4.70) on $\ \Delta^{(t-1)}\tilde{M}B^{k-t-1}\ $	128
4.15.3	Bound (4.71) on $R_t^{(\ell)}$	128

5	Spectral clustering in directed networks	130
5.1	Introduction	130
5.2	The Master Theorem	132
5.3	Spectral embeddings of directed SBM	134
5.3.1	Model definition	134
5.3.2	SBM with a pathwise structure	135
5.4	Geometric clustering and community detection	137
5.4.1	Algorithm and measure of performance	137
5.4.2	Choice of the clustering algorithm	138
5.4.3	Numerical validation of our results	139
5.5	Conclusion and future prospects	139
5.6	Gaussian Mixture clustering and Gaussian fluctuations	142
5.7	A bird's eye view on the proof of Theorem 31	145
5.7.1	Warmup: notations	145
5.7.2	The Pseudo-Master Theorem: A is nearly diagonalized by pseudo-eigenvectors	146
5.7.3	Master Theorem = Pseudo-Master Theorem + perturbation theory	147
5.8	Proof of the Pseudo-Master Theorem	148
5.8.1	The graph has few short cycles, and small neighbourhood growth	148
5.8.2	The graph is locally approximated by trees	149
5.8.3	Concentration of linear functionals	150
5.8.4	Pseudo-eigenvectors on the random tree	151
5.8.5	The martingale equation	151
5.8.6	Proving (5.16)-(5.17)-(5.18)	152
5.8.7	Proving (5.19): martingale correlations	152
5.8.8	The trace method	154
5.9	Master Theorem for the stochastic block model	157
5.9.1	Spectral decomposition of P	157
5.9.2	Master Theorem for SBM	158
5.10	Pathwise SBM	159
5.10.1	Motivation	160
5.10.2	Model density	160
5.10.3	Eigendecomposition of tridiagonal Toeplitz matrices.	160
5.10.4	Digression: the full threshold	161
5.10.5	Computations for the two-block case	161
5.10.6	Computations when there are two blocks	163
5.11	Convergence of eigenvectors	163
5.11.1	Convergence of $\tilde{U}_i(x)$	163
5.11.2	Convergence of the pseudo-eigenvectors	165
5.11.3	Convergence of the eigenvector	165
	References	167

Chapter 1

Introduction

1.1 Inference of planted subgraphs

In this section, we focus on the study of a specific planted subgraph problem, where a tree Γ with known shape is planted inside a sparse Erdős-Rényi graph with mean degree λ . All results and proofs are gathered in Chapter 2 of this thesis.

1.1.1 From NP-hard to planted problems

We begin this exposition with a motivation for the study of inference problems as a random instance of NP-hard problems. Let $G = (V, E)$ be a graph on n vertices; some questions that we could consider on G are as follows:

- the *max-clique* problem : find the size of the maximum fully connected subset (or clique) of V . We can also consider the associated decision problem : given k , find a clique of size k in G (or determine that there are no such cliques).
- the *min-bisection* problem : find a partition of V in two subsets of size $n/2$ with the minimum number of edges between those subsets. A more general problem asks for a partition of V in $k > 2$ parts of equal size with as few edges between the partitions as possible.
- the *Hamiltonian path* problem : find a path in G that goes through each vertex of V exactly once. Again, this can be generalized to finding a spanning tree of G with maximum degree lower than a constant k , or finding the path of maximum length in G .

All of those problems, and their generalizations, are known to be NP-hard [75, 61], which means that *in the worst case* they are very hard to solve. However, we can reasonably expect that the real-world instances of those problems do not match those worst-case scenarii : if we are looking for the maximum clique in a real network, it is likely because we assume that there is one that stands out. Conversely, we may want to know if the found maximum clique is statistically significant, or if it arises from random noise.

For those reasons, we consider here the planted version of those problems, where a specific structure (a somewhat large clique, a hamiltonian path, or a partition structure) is “hidden” inside a random graph. An often convenient assumption is to consider an Erdős-Rényi model $\mathcal{G}(n, p)$, where $p = p(n)$ is adapted to the problem at hand. As an example, for the planted clique problem the most widely used setting is $p = 1/2$ independently from n , while in the planted hamiltonian cycle problem the setting used in [32] is $p = d/n$ for constant d .

The two main tasks we want to perform on our inference models are usually the following:

- *Detection*: given a graph G , is there a (preferably polynomial) algorithm that can detect with high probability whether G was drawn from $\mathcal{G}(n, p)$ or whether G contains a planted structure ?
- *Reconstruction*: given a graph G drawn from the planted model, is there a (polynomial) algorithm that can recover the vertices in the planted structure with high probability, or at least a positive fraction of them ?

An important thing to note, however, is that a problem does not need to be NP-hard to yield insightful planted variants ; indeed, the planted matching problem [115] has a polynomially feasible optimization counterpart (the celebrated blossom algorithm [52]), but a rich phase transition landscape.

1.1.2 Hard phase in the planted clique problem

We now temporarily focus on the planted clique problem, where a clique of size k is randomly hidden inside a $\mathcal{G}(n, 1/2)$ graph; we denote the resulting random model by $\mathcal{PC}(n, k)$. In [93], the author shows that the maximum clique in a $\mathcal{G}(n, p)$ graph has size

$$s(p) = 2 \frac{\log(n)}{\log(1/p)} + O(\log \log(n)).$$

This implies that whenever $k \geq (1 + \varepsilon)s(1/2) = 2(1 + \varepsilon) \log_2(n)$, w.h.p a graph drawn according to $\mathcal{G}(n, 1/2)$ has no clique of size k , whereas one drawn according to $\mathcal{PC}(n, k)$ has (by definition) at least one. This illustrates a general method for planted subgraph problems: whenever the underlying random graph contains no copy of the planted subgraph w.h.p, detection is easily shown to be possible !

In the same regime, for the specific planted clique model, it can be shown that w.h.p. the only k -clique in $G \sim \mathcal{PC}(n, k)$ is the planted one; as a result we have that reconstruction is possible as soon as $k \geq (2 + \varepsilon) \log_2(n)$.

However, the two algorithms outlined above are non-polynomial : finding the planted k -clique or proving its non-existence requires inspecting $n^{\Omega(\log(n))}$ vertices. In fact, no polynomial algorithms for either detection or reconstruction are known unless $k = \Omega(n)$; the state of the art is a quasi-linear algorithm when $k > \sqrt{n}/e$ [49], as well as a "doubling trick" by Alon et al. [10] that allows any algorithm to be expanded to $k > \varepsilon \sqrt{n}$ for any $\varepsilon > 0$ at the cost of increasing the complexity: the resulting algorithm is still polynomial, but with a degree proportional to $\log_2(\varepsilon)$.

This is the core of what is known as the *information-computation gap*: when $\log(n) \ll k \ll \sqrt{n}$, there is theoretically enough information to detect and recover the planted clique, but with no polynomial algorithm known. It is conjectured, and supported by some theoretical evidence (see for example [72, 15]) that it is impossible to recover the planted clique in polynomial time when $k = o(\sqrt{n})$. We can therefore pinpoint three regimes of parameters in this problem :

- the *impossible* phase ($k < (1 - \varepsilon) \log(n)$): it is provably impossible to perform either detection or reconstruction with nontrivial success probability,
- the *hard* phase ($\log(n) \ll k \ll \sqrt{n}$): an exhaustive (and non-polynomial) search is guaranteed to find the planted clique with high probability, but no polynomial algorithm is known to work,
- and the *easy* phase, where both detection and reconstruction are polynomially feasible.

This type of phase transition landscape has been observed in several inference problems, including sparse PCA [24] or community detection [98]. We refer to [31] for a systematic study of this phenomenon. Notable exceptions include the planted Hamiltonian path [12] or the planted matching [115] problems ; a possible explanation for this diverging behavior is given at the end of this section.

1.1.3 The planted tree problem

The focus of this manuscript is on inference problems in sparse random graphs, i.e. graphs where the expected degree is independent from n . This translates, for homogeneous random graphs, in an Erdős-Rényi model with parameters n and λ/n for fixed λ . The main reason for this focus is the structure of real-world networks: as can be seen in Table 1.1, those networks are extremely sparse, and are thus better modeled by random graphs with lower density.

To match the sparsity of the underlying graph, the planted structures need to be sparser as well; for this reason, we focus on the case of planted trees. This choice can also be motivated by a problem in network security: assuming that a population of k malicious agents intend to coordinate to launch some kind of attack, the most inconspicuous way for them to do so is to form a minimal number of connections; consequently, their connection network will form a tree.

Without further restrictions, our model is not specific enough (see [105] for the number of graph on k vertices). In this manuscript, except for a digression on planted stars (which are related to the degree distribution of a random graph), we focus on a few families of trees:

- a path of length k ,
- a complete D -ary tree (with $D \geq 2$ fixed) of height h .

Those choices fit the framework of NP-hard relaxations discussed at the beginning of this section: both the decisions problems of finding a path or a D -ary tree of a given size in a graph are known to be NP-complete [40].

In order to quantify the performance of our reconstruction algorithms, we use the following metric: given the set of planted vertices $S \subseteq V$, we define the *overlap* of an estimator \hat{S} as

$$\text{ov}(\hat{S}) = \frac{1}{|S|} \mathbf{E} \left[|S \cap \hat{S}| \right].$$

This diverges from the usual convention of searching for an estimator \hat{S} that achieves reconstruction with high probability, and not in expectation. However, these definitions match in the regimes that are of interest in this problem: in particular, if $\text{ov}(\hat{S}) = o(1)$, then it is straightforward to show that

$$\limsup |S \cap \hat{S}| = o(1) \quad \text{w.h.p.},$$

and both definitions for the impossibility of reconstruction coincide. The only regime where a significant difference can occur is that of *partial* reconstruction, where $0 < \text{ov}(\hat{S}) < 1$, but our results in this regime can be extended to the high probability case without much effort.

1.1.4 Counting subgraphs in random graphs

In the following, we denote the planted structure by $\Gamma = \Gamma(k)$. The most intuitive method to study our planted tree model is to simply count the copies of Γ in an Erdős-Rényi graph. As such, we define the random variable

$$X_\Gamma(G) = \text{number of copies of } \Gamma \text{ in } G.$$

Network	Nodes	Edges	Average degree
DBLP collaboration	317,080	1,049,866	6.62
arXiv CondMat collaboration	23,133	93,497	8.08
Enron emails	36,692	183,831	10.02
Web graph (Google)	875,513	5,105,039	11.66

Table 1.1: Number of nodes, edges, and average degree of several real-world networks. Data is taken from the SNAP dataset collection [85]

As we mentioned above, whenever $X_\Gamma(G) = 0$, then a planted copy of Γ in G is easy to detect, although not necessarily easy to reconstruct. This yields the following heuristic:

Whenever $\mathbf{E}(X_\Gamma) = o(1)$ in an Erdős-Rényi graph, then planted copy of Γ is detectable with high probability.

The implied detection procedure need not be polynomial, as testing the existence of a copy of Γ in an arbitrary graph can be complex. It also does not imply anything about reconstruction: even though there is no copy of Γ in G prior to adding the planted one, this addition may create many “decoy” copies that prevent reconstruction.

A more surprising fact is that the distribution of X_Γ also allows to prove results about the impossibility of detection! More precisely, the following holds:

Proposition 1. *If in an Erdős-Rényi graph $\mathcal{G}(n, \lambda)$ one has*

$$\text{Var}(X_\Gamma) = o(\mathbf{E}(X_\Gamma)),$$

then detection is impossible: no algorithm can distinguish from the Erdős-Rényi and planted subgraph models with nontrivial probability.

Further, if

$$\mathbf{E}(X_\Gamma) = \omega(1),$$

then reconstruction is also impossible: no algorithm can recover a positive fraction of the planted subgraph’s vertices with nonvanishing probability.

In other words, the study of detection and (sometimes) reconstruction often reduces to computing the first and second moment of X_Γ for some specific Γ . For fixed-size subgraphs, as $n \rightarrow \infty$, this was extensively studied by Bollobás [26]; those methods can be adapted to our setting when the size of Γ grows slowly enough.

1.1.5 Phase transition landscape of the planted tree problem

The main results achieved in Chapter 2 for the planted path and D -ary tree are summarized in Figures 1.1 and 1.2. We obtain a near-complete phase transition landscape for the planted path problem, with the exception of the case where $k = \Omega(n)$; for the planted D -ary tree, the results are less comprehensive, although we provide several conjectures to fill the gaps in the landscape. Already though, this planted tree problem exhibits some key differences with the usual landscape of inference problems:

Absence of hard phase A common staple of inference problems, e.g. planted clique or community detection (see [7] for the latter) is the presence of a so-called *hard phase*, where detection and reconstruction are possible but no polynomial algorithm is known. A tentative

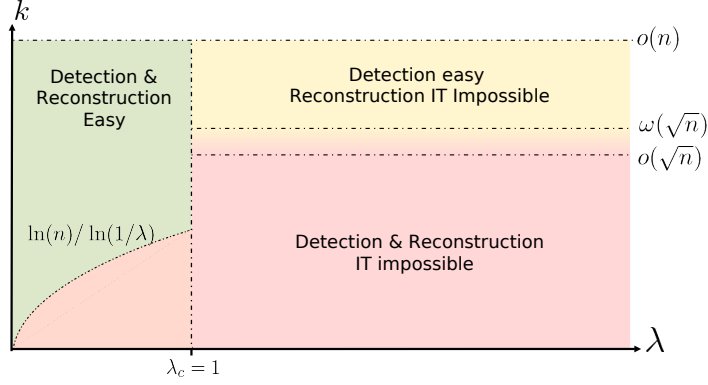


Figure 1.1: Phase transition diagram for the planted path problem, in terms of k and λ .

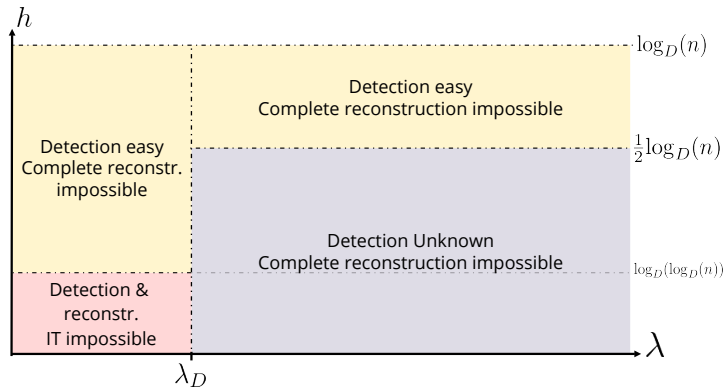


Figure 1.2: Phase transition diagram for the planted D -ary tree problem ($D \geq 2$), in terms of the height h and λ . The threshold λ_D corresponds to the emergence of the D -core in $\mathcal{G}(n, \lambda/n)$ [110, 99].

systematic study of such phenomena can be found in [31], drawing parallels with known NP-hard problems. By contrast, no such phase exists in the studied planted tree models : as soon as detection or reconstruction are possible, they can be performed by a polynomial algorithm, often even pseudo-linear in the number of edges.

Detection without reconstruction Another feature usually observed is the strong dependency between detection and reconstruction. Although ad-hoc counter-examples exist [14], reconstruction is almost always impossible whenever detection is; and conversely, whenever detection is possible, it is often feasible to recover at least a portion of the planted structure. In the planted path model, we instead observed a large region, from $k = \omega(\sqrt{n})$ to $o(n)$, where detection was feasible (and polynomially so, per the previous point), but it is impossible to recover more than a positive fraction of the planted path. The impossibility of complete reconstruction for the planted D -ary tree is a more common behavior, encountered for example in sparse community detection [1].

A possible explanation for those behaviors may be related to the concept of *low-rank* planted structure : viewing the graph through its adjacency matrix, in most mentioned inference problems the planted structure is (up to diagonal terms) a low-rank matrix. As an example, the adjacency matrix of a clique of size k has rank 1, up to its diagonal terms. To the contrary,

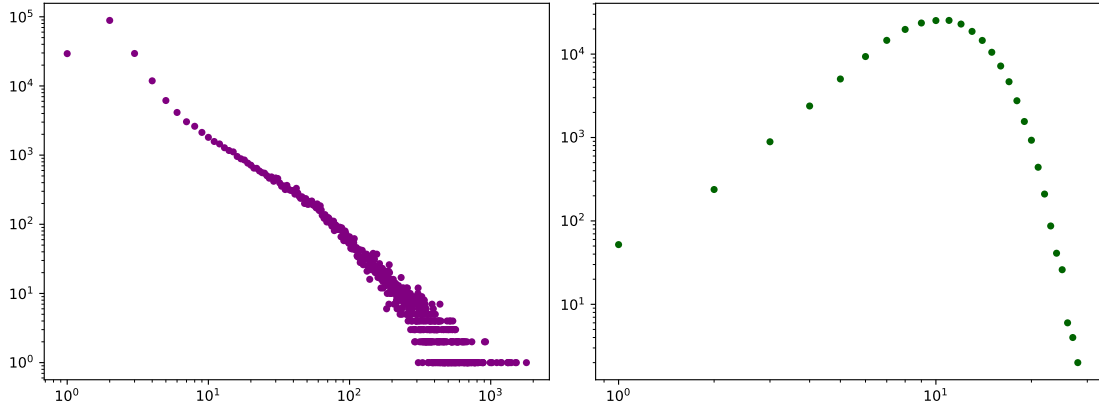


Figure 1.3: Degree distribution of a real-world (left) and random (right) graph, in log-log space. To the left is the graph of Google+ friendship circles, whose degree distribution roughly exhibits a power law, as in [38]. To the contrary, the distribution of an Erdős-Rényi graph with the same degree (right) is close to a Poisson distribution, which can be seen to be wildly different.

a tree of size k can have rank up to k , which is in particular the case for the planted path. Other examples of high-rank inference problems include the planted Hamiltonian path in [12], for which a polynomial algorithm also matches the information-theoretic lower bound up to negligible factors.

1.2 Spectral algorithms in community detection

In this second section, we move on to community detection and the study of the stochastic block model. More specifically, we focus on *sparse* stochastic block models, where as in the last section the connection probabilities are in $O(1/n)$. All original proofs can be found in chapter 3 of this thesis.

1.2.1 Simulating real-world networks

Although Erdős-Rényi models are a popular way to simulate real-world networks, they lack several desirable features. One notable example is the degree distribution, which is known to be wildly different between the Erdős-Rényi model and many known networks (see Figure 1.3). Another point of interest, which will be our main focus, is the presence of inhomogeneities: we expect the presence of underlying communities – geographic circles in social networks, similarity of function in protein interaction networks – to influence the formation of edges. As such, the homogeneity of the Erdős-Rényi model, where each edge has the same probability of appearing, is an obstacle to the accuracy of modeling.

Whenever a network exhibits this sort of community structure, a natural problem to consider is the possibility, and accuracy, of recovering those latent communities: the design and analysis of algorithms suited for this task is called *community detection*. The main obstacle in this endeavor is the following: except in some rare cases (see e.g. Figure 1.4), real-world networks with ground truth communities are rare, and thus the performance analysis of such algorithms is hard to define.

To address this shortcoming, the stochastic block model (or SBM in short) was proposed in [69]; it first consists in separating the vertex set of a random graph in *blocks* (modeling the underlying communities). Given this partition, each edge is then drawn independently from all others, with a probability of presence depending only on the communities of its endpoints.

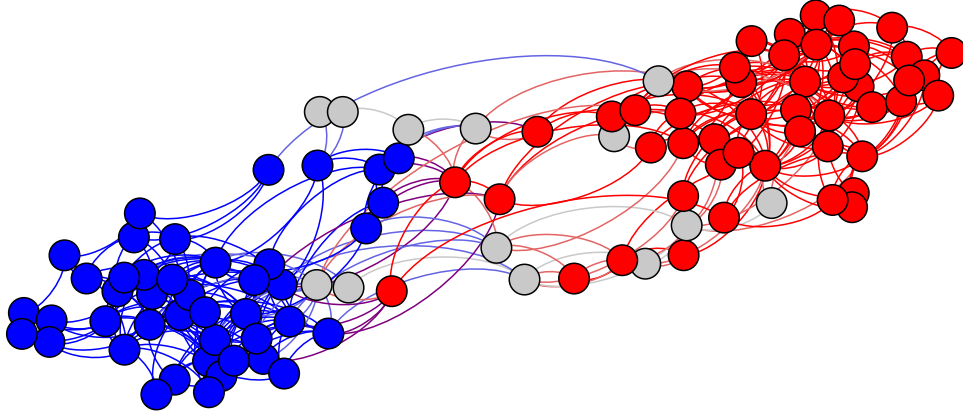


Figure 1.4: A network with ground truth communities : political books written on the 2004 U.S presidential election, classified by partisan leaning : democrat (blue), republican (red) and neutral (grey). Edges between two books indicate frequent co-purchases.

Usually – in the so-called *assortative* models – this gives rise to denser groups, with fewer edges connecting different groups. Several related models have been devised to address other issues, such as setting an arbitrary degree distribution [76], overlapping communities [9], or a geometric model to bring it closer to real-world networks [59].

In light of the aforementioned Table 1.1, sparse stochastic block models will be of particular interest: in those models, the average degrees of the vertices are constant, or grow only very slightly faster.

1.2.2 Sparse SBM : definitions

During the rest of this section, unless specifically mentioned, all stochastic block models studied will be sparse. We first elaborate on the SBM definition, as follows. Given a community assignment $\sigma : [n] \rightarrow [r]$, and a symmetric connectivity matrix P of size $r \times r$ with positive entries, we build a random graph $G = (V, E)$ such that $V = [n]$, and each edge (i, j) is present in E independently from all others, with probability

$$\mathbf{P}((i, j) \in E) = \frac{P_{\sigma(i)\sigma(j)}}{n}. \quad (1.1)$$

A special attention is given to the two-block symmetrical case, with two communities of equal size $n/2$ and the connectivity matrix having the form

$$P = \begin{pmatrix} a & b \\ b & a \end{pmatrix}.$$

In the general case, we shall assume that the communities are linear in size, i.e.

$$\frac{|\{x \in [n] \mid \sigma(x) = i\}|}{n} = \pi_i + o(1), \quad (1.2)$$

and we define the *mean progeny matrix* M as

$$M_{ij} = P_{ij}\pi_j. \quad (1.3)$$

Informally, the matrix M represents the expected number of neighbours of a vertex with type i that have type j . Finally, we will assume that the expected degree of each vertex is the same,

which translates to

$$\sum_{j \in [r]} M_{ij} = d \quad \text{for all } i \in [r]. \quad (1.4)$$

This condition is actually not a restriction: indeed, if (1.4) does not hold, a simple clustering on the vertex degrees gives an easy reconstruction procedure. It implies, in particular, that the Perron eigenvalue of M is the average degree d , with associated eigenvector $\mathbf{1}$.

In order to measure algorithm performance, we shall need a measure of how close an estimated community assignment is to the ground truth. The most commonly used metric is the *overlap*; given a community assignment $\sigma : [n] \rightarrow [r]$ and an estimator $\hat{\sigma}$, we can define an unadjusted overlap as

$$\text{uov}(\sigma, \hat{\sigma}) = \max_{\tau \in \mathfrak{S}_r} \frac{1}{n} \sum_{i \in [n]} \mathbf{1}_{\hat{\sigma}(i) = \tau \circ \sigma(i)},$$

where the maximum is taken over all permutations of the communities. This overlap is then rescaled to set the “naive” estimator (consisting of assigning every vertex to the largest community) to have an overlap of zero, while still ranging from 0 to 1, which yields the following definition:

$$\text{ov}(\sigma, \hat{\sigma}) = \frac{\text{uov}(\sigma, \hat{\sigma}) - \pi_{\max}}{1 - \pi_{\max}},$$

where π_{\max} is the maximum entry of π . This rescaled version is the one we’ll use in the following.

Our interest lies in the limit $n \rightarrow +\infty$; we classify the reconstruction performance of an estimator $\hat{\sigma}$ in three categories:

- **Exact recovery:** $\mathbf{P}(\text{ov}(\hat{\sigma}, \sigma) = 1) = 1 - o(1)$,
- **Almost exact recovery:** $\mathbf{P}(\text{ov}(\hat{\sigma}, \sigma) = 1 - o(1)) = 1 - o(1)$,
- **Partial recovery:** $\mathbf{P}(\text{ov}(\hat{\sigma}, \sigma) > \alpha) = 1 - o(1)$ for some $\alpha > 0$.

Due to the presence of a positive proportion of isolated vertices in sparse SBMs, exact and almost exact recovery are unachievable; as such, our efforts only focus on attaining partial recovery.

1.2.3 Tree approximation and fundamental limits

It is a well-known fact (see [68]) that when $n \rightarrow \infty$ a sparse Erdős-Rényi random graph $\mathcal{G}(n, d/n)$ converges, in the sense of Benjamini-Schramm [23], to a Galton-Watson tree with offspring distribution $\text{Poi}(d)$. This result can be extended to the stochastic block model, defining a *multitype Galton-Watson* tree as follows:

- the root vertex has a random type drawn according to the probability vector π ,
- each vertex of type i has a number $\text{Poi}(M_{ij})$ of children of type j , for all $j \in [r]$, where M is the matrix in (1.1).

Then, the SBM converges (in the same sense) to the random tree defined above [68]. This seems to indicate a strong relationship between the properties of the two models; and indeed many papers [103, 91, 29, 102] make use of couplings to translate results from one model to the other.

Of particular interest in the study of such tree processes is the *census reconstruction problem*: given only the number of vertices of each type at depth R , is it possible to recover the type of the root vertex with better accuracy than random chance, when $R \rightarrow \infty$?

Although not described this way, a paper from Kesten and Stigum [79] provides an affirmative answer to this question, provided

$$\lambda_2(M)^2 > d \tag{1.5}$$

where $\lambda_2(M)$ is the second largest (in magnitude) eigenvalue of M . For this reason, equation (1.5) is often dubbed the *Kesten-Stigum (KS) threshold*. In the case of the two-blocks symmetric SBM, this threshold becomes

$$(a - b)^2 > 2(a + b),$$

and is known to be sharp: if the other (strict) inequality is true, it is impossible to estimate the type of the root vertex from the long distance census (see [54]). This complete result for $r = 2$ is leveraged in [100] (for the negative part) and [91, 102] (for the positive part) to show that those results translate to the possibility or impossibility of detection and reconstruction in the symmetric SBM.

For larger values of r , the picture is more complicated: following a conjecture from [47], the methods in [91] have been extended to more than two communities in [29, 6], when above the KS threshold. However, a major difference is the possibility of reconstruction below this threshold, shown in [7], although with non-polynomial algorithms. This is thus another instance of the information-computation gap mentioned in the previous section. More information, as well as a survey on denser SBMs, can be found in [1].

1.2.4 Classifying reconstruction algorithms

Algorithms for community detection in the SBM can roughly be classified into three categories: semi-definite programming relaxations, message passing algorithms, and spectral methods. We briefly touch on the first two, and will dedicate the rest of this section to the third.

Semi-definite programming In the two-block symmetric case, labelling the groups with $\{-1, 1\}$, finding the maximum likelihood estimator for σ is equivalent to the following:

$$\text{maximize } \sigma^* A \sigma \quad \text{subject to } \sigma \in \{-1, 1\}^n, \sum_i \sigma_i = 0, \tag{1.6}$$

where A is the adjacency matrix of G . This is a reformulation of the min bisection problem discussed in the previous section, and is therefore NP-hard to solve. This optimization problem is therefore relaxed into another, more easily tractable one, the main choice being [64, 97]:

$$\text{maximize } \left\langle A - \frac{d}{n} \mathbf{1}\mathbf{1}^*, X \right\rangle \quad \text{subject to } X \succeq 0, \quad X_{ii} = 1 \quad \forall i \in [n]$$

where $X \succeq 0$ means that X is symmetric positive-semidefinite, d is an estimate of the mean degree of G , and the scalar product is $\langle A, B \rangle = \text{Tr}(A^*B)$. One then hopes that the found maximum is close to $X = \sigma\sigma^*$, so that the community labels can be recovered from the argmax X . This often becomes false when the parameters a, b become close to the KS threshold, and as a consequence no SDP algorithm is known to work up to the precise threshold, only approaching it as $a + b$ diverges [97]. The paper [88] presents a generalization of the SDP method to multiple communities, but also misses the KS threshold by a (possibly large) constant.

Belief propagation and message-passing Assuming that the community labels are drawn according to a probability vector π (reminiscent of (1.2)), the probability of a community assignment σ is

$$\mathbf{P}(\sigma \mid G; P; \pi) \propto \prod_{i < j} P_{\sigma(i), \sigma(j)}^{A_{ij}} \left(1 - \frac{P_{\sigma(i), \sigma(j)}}{n} \right)^{1 - A_{ij}} \prod_{i \in [n]} \pi_{\sigma(i)}. \tag{1.7}$$

Disregarding the correlations corresponding to non-edges, Equation (1.7) corresponds to a Markovian random field; as such, a natural method of inference is the use of belief propagation (BP) algorithms, whose building block is the sum-product algorithm. Despite promising numerical evidence [47], no proof of correctness for this algorithm is known to work in the SBM down to the KS threshold. However, linearized versions of this algorithm, either an adaptation of the cavity method [95] or the ABP algorithm in [6], have been proven to work down to the KS threshold. Those proofs usually involve linking the properties of those linearized algorithms to the spectrum of specific matrices that depend on the graph G , hence an important connection with the spectral algorithms discussed later.

Even without proof of its correctness, the classic BP algorithm has another use: it allows to amplify the signal found by other algorithms, often up to the optimal overlap; this improvement is especially discussed in [101], for the two-block case, and [6] in the general setting.

1.2.5 Spectral algorithms

The third class of algorithms, and the one studied in this thesis, are spectral algorithms. They hinge on a simple observation: the non-zero eigenvalues of the expected adjacency matrix $\mathbf{E}[A]$ are the same as those of M , with associated eigenvectors constant on the clusters. Writing

$$A = \mathbf{E}[A] + (A - \mathbf{E}[A]),$$

we could expect that using arguments from perturbation theory we are able to show that those eigenvalues are reflected within the spectrum of A as well, so that clustering on e.g. the second eigenvector of A yields non-trivial reconstruction. This ansatz has been verified incrementally for a wide regime of degrees (see [1] and references therein), and [55, 19] shows that the spectrum of A retains information from the one of $\mathbf{E}[A]$ up to the regime

$$d \gg \sqrt{\frac{\log(n)}{\log \log(n)}}.$$

Unfortunately, this property fails to hold when d is below this threshold. The reason for this failure is that for very sparse random graphs, the support of the eigenvectors of A associated with the highest eigenvalues concentrate around the vertices with high degrees, instead of capturing the community structure (see Figure 1.5 for an example of this phenomenon). Special care therefore has to be taken to limit the influence of those vertices.

The first efforts in this direction were done in [39], by trimming those vertices from the graph; however, this algorithm does not achieve positive overlap down to the KS threshold. The most successful approaches involve replacing the adjacency matrix A by another matrix that depends on G , but is less influenced by vertex degrees. Examples include:

- the *self-avoiding matrix* $B^{(\ell)}$, such that $B_{ij}^{(\ell)}$ counts the number of self-avoiding paths between i and j . This matrix was the first to reach the KS threshold in [91], and was extended to the case of multiple communities in [117] (see Chapter 3),
- the *non-backtracking matrix* B , indexed by the oriented edges of G , such that

$$B_{(x,y),(z,t)} = \mathbf{1}_{y=z} \mathbf{1}_{x \neq t};$$

powers of B thus count the number of paths between oriented edges that do not backtrack in successive steps. This matrix was suggested in [47], stemming from a linearization of the BP algorithm. Its spectral properties were then studied in [29], which showed the possibility of reconstruction above the KS threshold upon a small spectral condition on M .

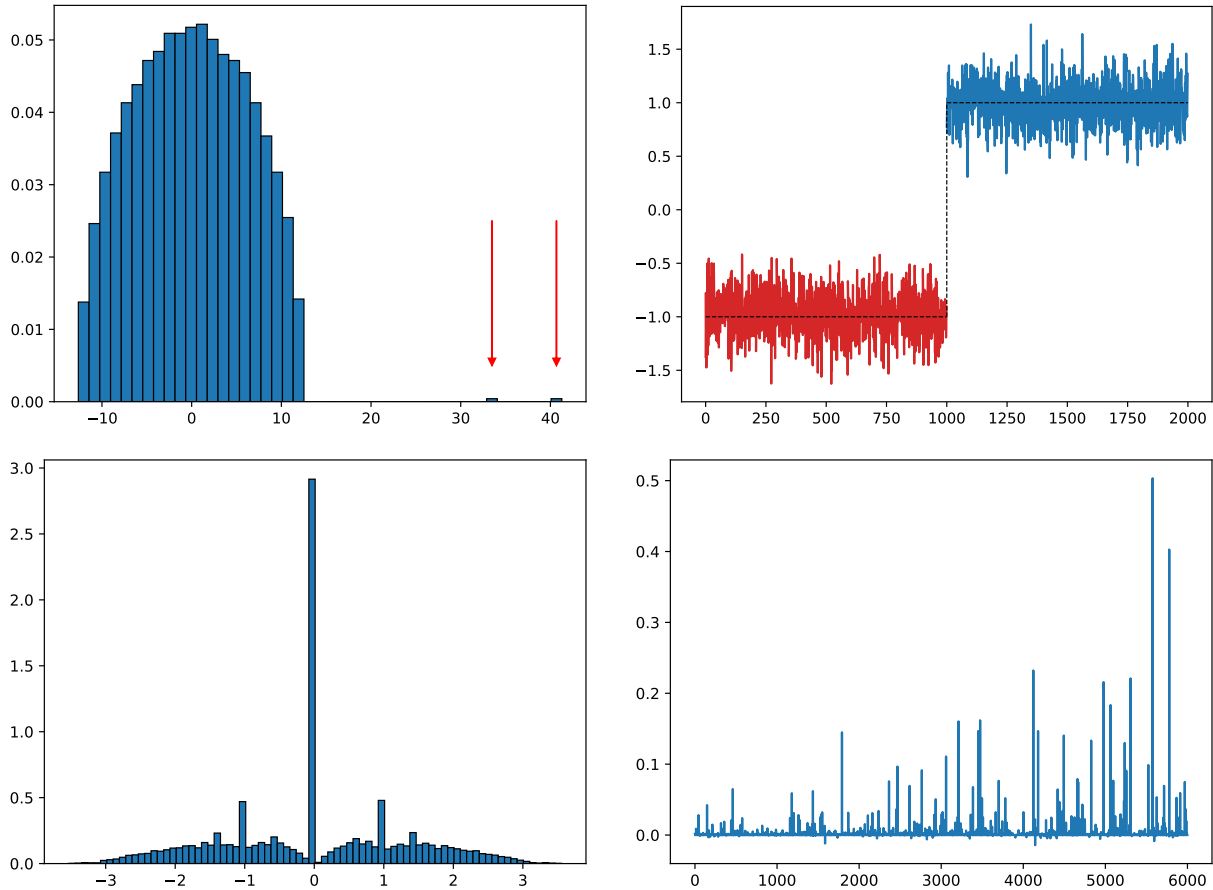


Figure 1.5: Eigenvalue histogram (left) and plot of the second eigenvector (right) of the adjacency matrix of a somewhat dense two-block SBM ($n = 2000, d = 40$, top), and a sparser one ($n = 6000, d = 1.7$, bottom). In the dense case, eigenvalues of $\mathbf{E}[A]$ are reflected in the spectrum of A (red arrows), and the associated eigenvector separates the communities. To the contrary, the sparse case has no outlier, and the top eigenvectors are not correlated with the community structure.

The non-backtracking matrix is of particular interest, for many reasons: it doesn't require tuning a path length parameter ℓ (even though the proof in [29] actually considers the matrix B^ℓ), and the non-backtracking operator has importance in other areas of graph theory, such as Ramanujan graphs [87] or the graph zeta function [66].

We also introduce in [117] another matrix suited for spectral community detection: the *distance matrix* $D^{(\ell)}$, such that

$$D_{ij}^{(\ell)} = \mathbf{1}_{d(i,j)=\ell}$$

where $d(i, j)$ is the usual graph distance between i and j . While very close (and thus having similar properties) as the matrix $B^{(\ell)}$ defined above in the typical SBM-generated graph, it is more robust to edge addition, as we will see below.

1.2.6 Robustness of community detection

The problem of robustness draws its root in the inaccuracies of the SBM to model real-world networks. More precisely, graphs generated according to a sparse SBM are known to be locally tree-like, with very few cycles and no cliques of size greater than four. In contrast, real-world

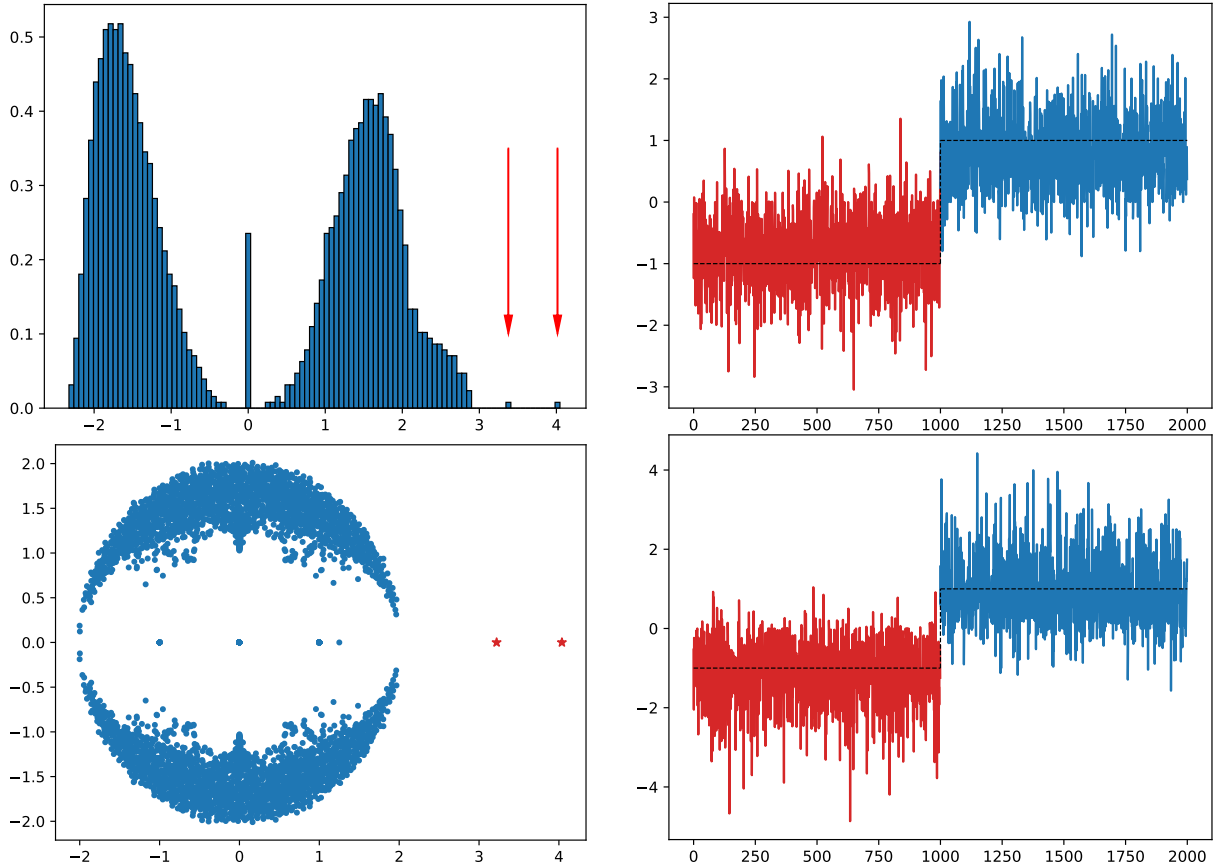


Figure 1.6: Eigenvalues and eigenvector plot of the self-avoiding (top) and non-backtracking (bottom) matrices of a sparse two-block SBM ($n = 2000, d = 4$). Since the non-backtracking matrix is not symmetric, its eigenvalues are complex; however, the outlier eigenvalues (as well as the associated eigenvectors) are real. Notice that the communities are less separated in the eigenvector plots than in Figure 1.5, a result of the model sparsity.

networks often contain cliques or denser subgraphs of moderate size, such as friend groups in social network graphs. As a result, several adaptations have been proposed to rectify the SBM, the two major ones being:

- the *geometric block models*, where each vertex u is assigned a position X_u in \mathbb{R}^k (or some subset thereof), and then two vertices u and v are connected iff their distance $\|X_u - X_v\|$ is less than a predetermined threshold. Typically, either the position X_u (in [3]) or the thresholds (in [59]) will depend on the community memberships.
- a perturbation of the classical SBM, where m edges are added either randomly or adversarially to a graph G drawn according to a sparse SBM. Generally, the adversarial version is preferred, or a “mixed” model where small cliques are added randomly (see [71] for such an example).

Unfortunately, classical spectral and message-passing algorithms do not fare well in those settings: they often rely on fragile properties of the SBM (such as the local tree approximation), which are quickly lost. On the other hand, SDP algorithms, while slower and not working up to the KS threshold, are very robust: they guarantee reconstruction even in the presence of a perturbation of up to $o(n)$ edges. By comparison, the spectrum of $B^{(\ell)}$ or B can be greatly perturbed by the addition of a clique of size $o(\log(n))$.

We show in [117] that the aforementioned distance matrix $D^{(\ell)}$ fares much better than other spectral methods: it can sustain a perturbation affecting at most $O(n^\varepsilon)$ vertices, where ε is a small (but nonzero) constant. We also show that this bound is tight in the fully adversarial context: adding a clique of size n^ε in a well chosen manner will make the spectral algorithm fail.

Unfortunately, this upper bound does not match the stellar performance of the SDP algorithms, and working out the precise constant ε shows that it is fairly small. However, this is a first step in showing that the perceived fragility of spectral methods is not a foregone conclusion, and that an algorithm that combines the best of both worlds is possible.

1.3 Spectra of weighted inhomogeneous random graphs

We now focus on more general community detection problems, including community detection in directed graphs. In particular, we study the spectrum of the non-backtracking matrix and directed adjacency matrix of random graphs under a very general model. All proofs and details are in Chapters 4 and 5 of this manuscript.

1.3.1 Beyond the classical SBM

We first expand on several variants of the stochastic block model, that will be the motivation for the general graph model introduced later.

Labeled stochastic block model Real-world networks often carry more information than simply the presence or absence of edges. For this reason, a common adaptation of the SBM includes edge information: given an array $\mathbf{P}_{i,j}$ of probability distributions on a common label space \mathcal{L} , the classical SBM is augmented by a label $\ell_{x,y}$ on each edge (x,y) drawn according to the distribution $\mathbf{P}_{\sigma(x),\sigma(y)}$. Such a model was defined in [67], where the authors conjectured an analogue of the KS threshold in the two-block case, and furthered in [84]. In particular, one of the proposed reconstruction methods involves choosing a function $w : \mathcal{L} \rightarrow \mathbb{R}$ and using it to transform G into a weighted graph. Note that this model encompasses the popular censored block model of [2, 113].

Degree-Corrected SBM To remediate the mismatch between the degree distributions of real-world networks and the SBM (see Figure 1.3) the degree-corrected SBM was proposed in [46]: it consists in an additional *affinity vector* $\theta \in \mathbb{R}^n$, such that the probability of presence of an edge (i,j) in G is

$$\mathbf{P}((i,j) \in E) = \frac{\theta_i P_{ij} \theta_j}{n}.$$

Depending on the models, the vector θ can be either deterministic [63, 82] or i.i.d [62]. In the case where P represents a SBM with constant expected degree, and θ is random, it is easy to see that the degree distribution of G is a rescaling of the one of the θ_i , which allows for matching arbitrary graphs more closely. This model can of course be augmented with labels, in the sense described above.

Directed SBM Although less studied than its undirected version, community detection in directed networks is also a topic of research (see [89] for a survey). One obstacle to reconstruction is that the adjacency matrix of a directed graph is *a priori* not diagonalizable; as such, commonly used reconstruction techniques use symmetrizations of the adjacency or Laplacian matrices, e.g. the singular value decomposition [114, 37].

The SBM can be readily extended to the directed setting [123]; in this case, the matrix P need not be symmetric, and we can assign different *types* σ and σ' to the start and end vertices of each edge, so that

$$\mathbf{P}((i, j) \in E) = \frac{P_{\sigma(i)\sigma'(j)}}{n}.$$

Assuming that, as in the undirected case,

$$\frac{|\{x \in [n] \mid \sigma(x) = j, \sigma'(x) = i\}|}{n} = \Pi_{ij} + o(1)$$

the mean progeny matrix M defined in Equation (1.3) has the form

$$M = P\Pi,$$

and we assume that this matrix is diagonalizable. As with the undirected version, the directed SBM can be extended with label or degree corrections.

1.3.2 A generic random graph model

Although specific adaptations exist when the affinity vector θ is i.i.d (see [62]), in the general DCSBM case the expected adjacency matrix is not constant by blocks. We therefore need a much more general model, outlined below.

Given an $n \times n$ matrix P with entries in $[0, 1]$ and a random matrix W , we define a random graph G as follows : each edge (i, j) is present in G with probability P_{ij} , and holds a weight W_{ij} . Depending on the model (directed or undirected), the matrices P and W may be required to be symmetric. It is also very likely that the results of this manuscript can be extended to an interpolating model, where the presence of an edge between (i, j) and (j, i) is correlated, but where the resulting graph is not necessarily undirected.

The expectation of the adjacency matrix $Q := \mathbf{E}[A]$ has a simple expression:

$$Q = P \circ \mathbf{E}[W],$$

where \circ is the elementwise (or Hadamard) product. The matrix Q therefore serves as the *signal* matrix, that we want to recover from the observation of G .

Given the application scope of our model, we make several assumptions on Q and W :

- the graph G is sparse: $P_{xy} \leq d/n$ for all x, y ,
- Q has low rank r ,
- the eigenvectors of Q are delocalized, in the sense that $\|\varphi\|_\infty \leq b\sqrt{n}$ for any normed eigenvector φ ,
- the entries of W are uniformly bounded by a constant L .

Note that those restrictions are actually fairly lax: our analysis allows for $r, b = n^{o(1)}$, and $d, L = \exp\left(o\left(\sqrt{\log(n)}\right)\right)$.

Note that this random model encompasses more than simply the SBM variants discussed above: as an example, it can be used as a model for very sparse (noisy) matrix completion – see [28] for an excellent account on this topic.

1.3.3 Aside : Phase transitions in random matrix theory

We temporarily place ourselves in the undirected case, where A and Q are symmetric. Writing the adjacency matrix of G as

$$A = Q + (A - Q),$$

where the second term has an expected value of zero, we can rephrase the problem of characterizing the spectrum of G into a problem of random matrix theory. The usual questions that we would like to answer are the following:

- (i) what is the relationship between the spectrum of A and the one of $\mathbf{E}[A]$?
- (ii) what is the limiting spectral distribution of A ?

Those questions have a rich history in random matrix theory, especially when (as in the examples discussed above) the matrix Q has low rank. A celebrated result in this topic is the Baik-Ben Arous-Péché (BBP) phase transition [13], on the largest eigenvalue of empirical covariance matrices of Gaussian vectors. It was the first example of a widespread phenomenon on random matrix models, that can be summarized as follows:

- if the largest eigenvalue μ_1 of $\mathbf{E}[M]$ is below some threshold ϑ , then the largest eigenvalue of M is equal to $\rho + o(1)$, where both ϑ and ρ depend on the model,
- if, on the other hand, the largest eigenvalue μ_1 of $\mathbf{E}[M]$ satisfies $\mu_1 > \vartheta$, then the largest value of M is of the form $f(\mu_1) + o(1)$, with f a deterministic function.

Such phase transitions are called *BBP phase transitions*, and have been proven to occur in various models, e.g. Wigner [107] matrices; more recently, [21] has expanded this result to a larger class of random matrices, as well as to eigenvalues beyond the first.

In this manuscript, we show that part of this phase transition occurs in our weighted random graph model, assuming the weighted degrees

$$d_x = \sum_{y \sim x} W_{xy}^2 \tag{1.8}$$

concentrate around a common value d_0 independent from x . Under this condition, if μ is an eigenvalue of Q such that $\mu > \sqrt{d_0}$, then there is an eigenvalue λ of the adjacency matrix A such that

$$\lambda = \mu + \frac{d_0}{\mu} + o(1).$$

This matches the theoretical results for Wigner matrices in [107, 21], which is in accordance with other results on the limit distribution of A [51]. We also obtain asymptotic results on the scalar products between the eigenvectors of Q and those of A , again matching [21]. However, the concentration hypothesis (1.8) is quite restrictive: in the unweighted case, it requires the common average degree d to be at least $\omega(\log(n))$.

1.3.4 Non-backtracking matrix of undirected graphs

The fact that degree concentration is needed for the spectral properties of A observed in the previous paragraph is not a proof artefact: as with the more specific community detection, in the undirected case, the spectrum of A is very sensitive to sparsity. In the case of unweighted graphs with $d \ll \log(n)$, the behaviour of the largest eigenvalues has been studied in [19]; the authors showed that the largest eigenvalues of A only correspond to vertices of highest degrees, with

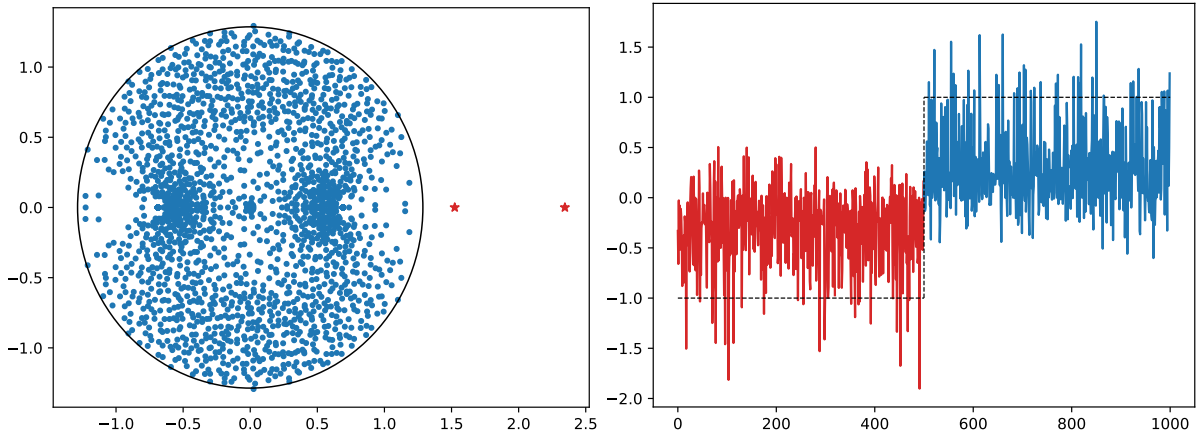


Figure 1.7: (Left) Spectrum of the non-backtracking matrix B of a weighted random graph, and (Right) Embedding of the second eigenvector of B in n -dimensional space. The underlying graph is a symmetric labeled two-block SBM ($n = 1000, a = 8.4, b = 3.6$); we generate spins according to the censored block model in [113] ($\varepsilon = 0.3$), and assign weights according to equation (4.13) in Chapter 4. The parameters have been chosen so that reconstruction is impossible using only either the unweighted edges or the spins, but possible with a combination of both. Note that the shape of the spectrum inside the bulk is very distinct from the unweighted case (Figure 1.6). The embedding used to generate the right figure is the one in Equation (1.9).

associated eigenvectors concentrated around those vertices. The more precise phase transition around $d \sim \log(n)$ has also been studied in [11], this time on a uniform weighted Erdős-Rényi graph. The spectrum of A also has a shape that is far from that of a Wigner matrix: in particular, it has an atom (of known mass) at zero (see [41] for details).

The solution developed in this manuscript follows the lines of [29], and uses the weighted non-backtracking matrix B defined as

$$B_{(x,y),(z,t)} = W_{zt} \mathbf{1}_{y=z} \mathbf{1}_{x \neq t}.$$

It enjoys similar desirable properties as its non-weighted counterpart, namely:

- it squashes the contribution of high-degree vertices,
- it does not require tuning a length parameter ℓ ,
- it allows to recover information about the adjacency matrix A through the Ihara-Bass formula [17, 124].

This relative stability allows us to obtain results similar to those of [29]: all eigenvalues of Q above a certain threshold ϑ are reflected in the spectrum of B , with the other eigenvalues of B being enclosed in a circle of radius $\vartheta(1 + o(1))$. This holds for a large range of parameters: the rank of Q can be up to $n^{o(1)}$, and the max edge probability $\|P\|_\infty$ up to $\exp(o(\sqrt{\log(n)}))$.

To perform reconstruction on the various inference models mentioned in this section, we need a vector of length n instead of $|\vec{E}|$. This is done through local averages: for an eigenvector χ of B , we define the associated embedded vector φ as

$$\xi(x) = \sum_{e: e_2=x} \chi(e). \quad (1.9)$$

The main drawback of clustering using the non-backtracking matrix is its size: instead of clustering on a $n \times n$ matrix, we have to compute the eigendecomposition of a $dn \times dn$ matrix,

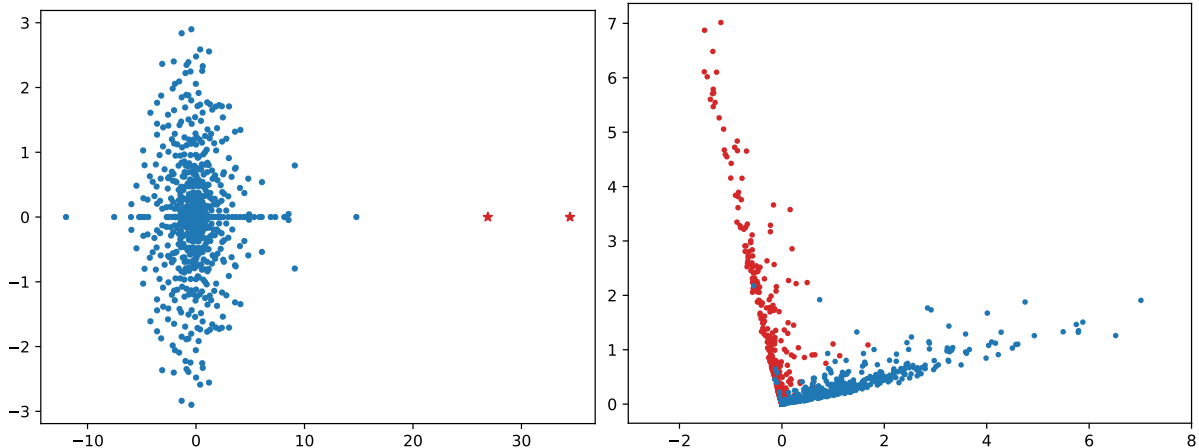


Figure 1.8: Left: Eigenvalues of the adjacency matrix A of a real-world network: the political blogs dataset [8]. While the bulk is less circular than in stochastic networks, the outlier eigenvalues are clearly visible – and real. Right: Scatter plot of the first two eigenvectors of A , with the true labels (blue for Democrat and red for Republican) revealed. The clusters are clearly visible.

where d is the mean degree of the graph. This can incur a large computational cost whenever the mean degree of the graph diverges (slowly). An alternative to the non-backtracking matrix that does not suffer from this issue is the Bethe Hessian matrix, defined as

$$H(r) = (r^2 - 1)I - rA + D$$

where D is the diagonal matrix of vertex degrees. Several heuristics for the choice of r have been proposed [112, 45], which empirically provide good results. The Bethe Hessian is closely related to the non-backtracking matrix; in particular, the eigenvector of B associated to an eigenvalue λ corresponds to a singular vector of $H(\lambda)$ via the transformation (1.9). Studying the non-backtracking matrix thus also yields insights about the behavior of the Bethe Hessian, as evidenced in [43].

1.3.5 Adjacency matrix of directed graphs

As previously mentioned, most community detection methods in the directed case make use of symmetrized versions of common matrices: a symmetric laplacian [37], singular values of A [128], or even hermitian matrices [44]. In contrast, following a line of work on the use of asymmetric matrices in inference problems started in [36], we show that using the unprocessed adjacency matrix already yields a plethora of useful information, at least in the directed SBM and its variants.

In particular, we show the exact same results as in the undirected case, with the non-backtracking matrix B being replaced by the adjacency matrix of G : the eigenvalues of Q above a certain threshold ϑ are reflected in the spectrum of A , with the other eigenvalues confined in a bulk of radius ρ . Such a similarity between the adjacency and non-backtracking is no coincidence; defining the $V \times E$ *start* and *terminal* matrices S and T as

$$S_{x,e} = \mathbf{1}_{e_1=x} \quad \text{and} \quad T_{x,e} = \mathbf{1}_{e_2=x},$$

the matrices A and B satisfy the relationships

$$A = ST^* \quad \text{and} \quad B = T^*S - J,$$

where J is the edge reversal operator:

$$J(e_1, e_2) = (e_2, e_1).$$

In the sparsity regimes we consider, the matrix J is extremely sparse, so that we can write $B \approx T^*S$ and the spectra of A and B are nearly identical.

The main strong point of our spectral analysis is that it encompasses various models inside a single algorithm: simply compute the spectrum of the non-backtracking matrix B (resp. the adjacency matrix A for directed graphs), and cluster on the eigenvectors that are outside the bulk. This generality allows us to apply this algorithm to most real-world graphs and expect decent results (see Figure 1.8). The optimal clustering algorithm is discussed in Chapter 5.

1.3.6 Perspectives : eigenvector limit distribution

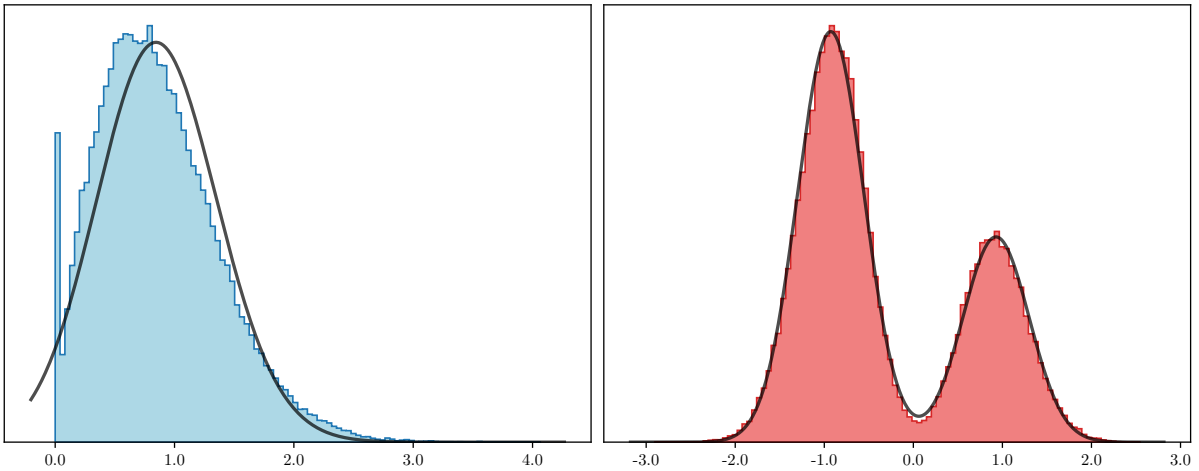


Figure 1.9: Eigenvector plots of a sparse (Left, $d \approx 4.7$) and denser (Right, $d = 20$) directed stochastic block model. Overlaid are the corresponding Gaussian mixture distributions with matching model parameters.

The results of the empirical determination of the best clustering procedure yields a very interesting insight: gaussian mixture models seem to perform better than the much more widely used k -means algorithm.

We present in this manuscript a possible explanation for this behavior: when the mean degree of G goes to infinity, the eigenvectors of B (resp. A) seem to behave exactly like gaussian mixtures.

As can be seen in Figure 1.9, the empirical distribution matches quite closely the theoretical gaussian mixture, even for fairly low degrees ($d = 20$). We prove in Chapter 5 that the empirical distribution of the i -th eigenvector of A converges as $n \rightarrow \infty$ to a mixture of random variables X_{ij} with known mean and variance, that are obtained as limits of martingale processes on a multitype Galton-Watson tree. In particular, for the directed Erdős-Rényi model, the limit variable $X = X_{11}$ satisfies the recursion

$$X = \frac{1}{d} \sum_{i=1}^N Z_i,$$

where N is a $\text{Poi}(d)$ random variable and the Z_i are i.i.d with the same distribution as X . This is reminiscent of a random version of the central limit theorem, and provides an intuitive justification for the observed convergence.

Chapter 2

Planting trees in graphs

This chapter is based on the paper [92], a joint work with L. Massoulié and D. Towsley.

2.1 Introduction

This paper is concerned with the detection of additional structures planted in a graph initially without structure (such as an Erdős-Rényi graph) and, in case such a structure is detected, with the reconstruction of the corresponding structure. We focus on planted structures that consist in a superimposed graph, and more specifically on superimposed trees.

A first motivation for this focus stems from the following application scenario. Assume that the original graph without planted structure represents normal communications among agents, while the superimposed graph represents communications among a subset of *attackers* who, when active, connect directly among themselves to coordinate their activity. Detection then amounts to estimating whether an attack occurs, while reconstruction amounts to identifying the attackers in case of an attack.

A second motivation is theoretical: previous work reviewed in Section 2.2 has shown that detection and reconstruction of planted structures in graphs displays rich and intriguing behaviour, with phases where the task is either impossible, computationally hard, or easy. It is important to understand what causes such phases, and whether phases for detection always coincide with the corresponding phase for reconstruction. Our present study sheds light on these questions, by showing that in the cases of planted tree structures we consider, no hard phase occurs, while feasibility phases of detection and reconstruction differ widely. In contrast, the latter property does not hold for previously studied low rank planted structures.

More specifically, our contributions are as follows. In the particular case of planted line graphs, we determine the complete phase diagram for detection and reconstruction: In a low density region where the average degree λ of the original graph is below some critical value λ_c , both detection and reconstruction go from impossible to easy as the line length K crosses some critical value $K^* = f(\lambda) \ln(n)$, where n is the number of nodes in the graph. In a high density region where $\lambda > \lambda_c$, detection goes from impossible to easy as K goes from $o(\sqrt{n})$ to $\omega(\sqrt{n})$. In contrast, reconstruction remains impossible so long as $K = o(n)$.

We then consider the case of D -ary trees for fixed $D > 1$, of height h . For these our results provide a similar picture with significant differences. Specifically, there exists a limit height $h_* = \ln \ln(D) + O(1)$ such that detection is impossible if $h < h_* - \ln(h_*)$, and easy for $h > h_* + \Omega(1)$. In that latter case, non-trivial reconstruction is feasible, but it must fail on a non-vanishing fraction of the K attack nodes. In a high-density region $\lambda > \lambda_D$, we have again that detection is easy for $K = \omega(\sqrt{n})$, and that reconstruction must fail at least on a fraction of nodes.

The paper is organized as follows. We review related work in Section 2.2. We describe our model and main results in Section 2.3. The proofs for planted lines and planted D -ary trees are in Sections 2.5 and 2.6 respectively, with detailed proofs of auxiliary results in the Appendix.

2.2 Related work

Planted clique detection and reconstruction has been the object of many works, see e.g. [48], [49], [15] for recent results and surveys. A central result in that context is that detection appears hard (i.e. no algorithm is known to succeed at detection in polynomial time) for cliques of size $o(\sqrt{n})$ planted in $G(n, 1/2)$. IT thresholds for planted dense subgraph detection are developed in [121].

Computational hardness of planted clique is used in reduction arguments to show that other planted structure detection problems are hard, eg sparse PCA [24], and dense subgraph detection [65]. The latter also displays IT-impossible phases, hard phases and easy phases. A systematic development of such reductions between problems with planted structure is initiated in [31].

Community detection and reconstruction has also been thoroughly studied, the seminal article [47] introducing several conjectures on feasibility of detection and reconstruction for the stochastic block model. Almost all conjectures in [47] have been verified in subsequent works, in particular [100, 91, 102, 7]. It is particularly relevant for the present article that these works identify regimes where both detection and reconstruction are jointly either impossible, possible but computationally hard, or computationally easy.

Presence of specific subgraphs in random graphs has been thoroughly studied, see e.g. [70]. We leverage the corresponding techniques in our study of low density regions, for which detection feasibility corresponds to absence of copies of the planted graph structure in the original random graph.

Most planted structures considered so far were typically of “low rank” (e.g. planted dense graph’s expected adjacency matrix is, up to diagonal terms, a rank one perturbation); in contrast, adjacency matrices of trees and lines are not close to a low rank matrix. One notable exception is the planted Hamiltonian cycle reconstruction addressed in [12]. Interestingly, for that problem too there is no hard phase for reconstruction; but in contrast to our scenario, no large parameter region is identified where detection is easy while reconstruction is impossible.

2.3 Model and main results

A total population of n agents interconnects according to one of the following two modalities. Under the null hypothesis H_0 the interconnection does not display any specific structure. We assume that the corresponding graph G is an Erdős-Rényi $\mathcal{G}(n, p)$ graph, with edge probability $p \in [0, 1]$ taken equal to λ/n for some fixed $\lambda > 0$. We thus focus on sparse random graphs with average degree $O(1)$. Under the alternative hypothesis H_1 , the graph G is the union of a base graph G_0 distributed according to $\mathcal{G}(n, p)$, with another graph G' connecting a distinguished subset \mathcal{K} of nodes. Specifically, for a fixed graph Γ on node set $[K]$ with edge set \mathcal{E} , and an injective map $\sigma : [K] \rightarrow [n]$ chosen uniformly at random and independently of G_0 , G' consists of the nodes $\mathcal{K} = \{\sigma(i), i \in [K]\}$ and edges $\{(\sigma(i), \sigma(j)), (i, j) \in \mathcal{E}\}$.

We shall mostly focus on tree graphs Γ , and more specifically on D -ary trees, i.e. trees with a distinguished root, or depth-0 node, and for each $\ell \in [h-1]$, D^ℓ depth- ℓ nodes being connected to one parent at depth $\ell-1$ and D children at depth $\ell+1$. The two extreme cases are a line graph for $D=1$ and a star for $D=K-1$.

We are interested in answering, on the basis of an observed graph G , the following questions:

Q1 (Detection): For a given planted graph shape Γ (e.g. line, star, D -ary tree, ...), under what parameter regimes specified by λ and K is there a test that distinguishes H_0 from H_1 with error probabilities of both kinds going to zero as $n \rightarrow \infty$? This is an information-theoretic property characterized by the likelihood ratio $\frac{\mathbb{P}_1(G)}{\mathbb{P}_0(G)}$, where \mathbb{P}_i denotes the distribution of G under H_i , $i = \{0, 1\}$. Indeed by the Neyman-Pearson lemma, among tests with given probability of correctly deciding H_1 , there is one which minimizes probability of erroneously rejecting H_0 which decides H_1 if and only if the likelihood ratio $L(G) := \frac{\mathbb{P}_1(G)}{\mathbb{P}_0(G)}$ is larger than some threshold τ . We can ask the same question as Q1 when we restrict ourselves to tests that can be implemented in polynomial time. This then corresponds to a computational property, and it is of interest to identify whether the two thresholds (informational and computational) coincide or not.

Q2 (Reconstruction): Can one reconstruct the planted structure G' , or at least a subset of its constituent nodes? Several metrics of reconstruction accuracy are possible. We shall focus on the following *overlap* metric, which we now define for estimation procedures that produce a set $\hat{\mathcal{K}}$ of K nodes in $[n]$, aimed to estimate at best the actual set \mathcal{K} of K nodes involved in the attack.

Definition 1. *The overlap of a set $\hat{\mathcal{K}}$ estimating the actual ground truth \mathcal{K} is by definition the expected size of their intersection, i.e.*

$$\text{ov}(\hat{\mathcal{K}}) := \sum_{i \in [n]} \mathbb{P}(i \in \hat{\mathcal{K}} \cap \mathcal{K}).$$

We say that a particular reconstruction $\hat{\mathcal{K}}$ of size K fails if $\text{ov}(\hat{\mathcal{K}}) = o(K)$, succeeds if $\text{ov}(\hat{\mathcal{K}}) = K(1 - o(1))$, and partially succeeds if $\text{ov}(\hat{\mathcal{K}}) = cK(1 - o(1))$ for some $c \in (0, 1)$.

Reconstruction (respectively, partial reconstruction) is then deemed *feasible* if there exists an estimator $\hat{\mathcal{K}}$ that is successful (respectively, partially successful). These properties are of an information-theoretic nature. Indeed the best possible overlap is achieved by the so-called Maximum a Posteriori (MAP) estimation procedure, and these properties are therefore determined by the overlap of the MAP estimator. One can, as for detection, consider a computational version of reconstruction: reconstruction (respectively, partial reconstruction) is *easy* when it can be achieved by an estimator $\hat{\mathcal{K}}$ that is efficiently computable.

Before stating our results for planted lines and D -ary trees, we first consider planted star graphs, for which a simpler picture holds:

Theorem 1. *For any fixed $\lambda > 0$, a planted star of size $K = \ln(n)/\ln(\ln(n))[1 - \omega(1/\ln(\ln(n)))]$ is not detectable, while both detection and reconstruction of a planted star of size $K = \ln(n)/\ln(\ln(n))[1 + \omega(1/\ln(\ln(n)))]$ are easy.*

The result for line graphs, summarized in Table 2.1, is

Theorem 2 (Line graphs). *In the low-density region $\lambda < \lambda_c = 1$, detection and reconstruction are impossible if $K = \ln(n)/\ln(1/\lambda) - \omega(\ln(\ln(n)))$, while both detection and reconstruction are easy if $K = \ln(n)/\ln(1/\lambda) + \omega(1)$ and $K = o(n/\ln(n))$.*

In the high-density region $\lambda > \lambda_c = 1$, detection and reconstruction are impossible if $K = o(\sqrt{n})$, detection is easy if $K = \omega(\sqrt{n})$, while reconstruction is impossible for $K = o(n)$.

For D -ary trees, the results are similar. However the critical parameter λ_D defined in (2.11) is the threshold for emergence of the D -core (see [99]), and only partial reconstruction is possible in the subcritical regime $\lambda < \lambda_D$. We consider D -ary trees Γ of depth h with corresponding size $K = \frac{D^{h+1}-1}{D-1}$; the main results (in terms of h) are summarized in Table 2.2.

$K < \ln(n)/\ln(1/\lambda)$	$\ln(n)/\ln(1/\lambda) < K \ll n/\ln(n)$
Detection & reconstruction IT impossible	Detection & reconstruction easy

(a) Subcritical regime : $\lambda < 1$

$K \ll \sqrt{n}$	$\sqrt{n} \ll K \ll n$
Detection & reconstruction IT impossible	Detection easy reconstruction IT impossible

(b) Supercritical regime : $\lambda > 1$

Table 2.1: Summary of results for planted line graph

$h < \log_D \log_D(n)$	$\log_D \log_D(n) < h < \log_D(n)$
Detection & reconstruction IT impossible	Detection easy complete reconstruction impossible

(a) Subcritical regime : $\lambda < \lambda_D$

$h < \log_D(n)/2$	$\log_D(n)/2 < h < \log_D(n)$
Detection unknown complete reconstruction impossible	Detection easy complete reconstruction impossible

(b) Supercritical regime : $\lambda > \lambda_D$ Table 2.2: Summary of results for planted D -ary tree

Theorem 3 (D-ary trees). *In the low-density region $\lambda < \lambda_D$, there exist two parameters \underline{h} and \bar{h} such that the following holds.*

$\bar{h} = \ln \ln(n)/\ln(D) + \Theta(1)$, and $\underline{h} = \bar{h} - 1$ for almost all λ .

When $h \leq \underline{h} - O(\ln(\underline{h}))$, both detection and reconstruction are impossible with high probability.

Detection is easy whenever $h \geq \bar{h} + O(1)$.

For any $\lambda > 0$, hence in both low-density and high-density regions, detection is easy whenever $K = \omega(\sqrt{n})$ while complete reconstruction is impossible for $K = o(n)$.

2.4 Preliminary results

We now state three results that hold for arbitrary planted structures, and that will be used extensively. The first is a characterization of the likelihood ratio $\frac{\mathbb{P}_1}{\mathbb{P}_0}$:

Lemma 1. *The likelihood ratio $L(G) = \frac{\mathbb{P}_1(G)}{\mathbb{P}_0(G)}$ is given by $L(G) = \frac{X_\Gamma}{\mathbf{E}_0(X_\Gamma)}$, where X_Γ denotes the number of copies of Γ in G .*

The second gives a generic detection process that succeeds for K large enough, and all planted graph structures Γ that are connected.

Theorem 4. *Assume that $\lambda > 0$, $K = \omega(\sqrt{n})$, and the hidden graph is any connected subgraph on K nodes, not necessarily a line. Then the total variation distance $|\mathbb{P}_1 - \mathbb{P}_0|_{\text{var}}$ between \mathbb{P}_0 and \mathbb{P}_1 goes to 1 as $n \rightarrow \infty$.*

Let A_i , $i \in \{1, 2, 3\}$ denote the number of size i -connected components in G , $\hat{\lambda} = (nA_3)/(A_1A_2)$, and $\hat{k} = n - e^{\hat{\lambda}}A_1$. The test that decides H_1 if $\hat{k} \geq t_n := \sqrt{K\sqrt{n}}$, and H_0 otherwise is polynomial-time computable and distinguishes with high probability graphs sampled from \mathbb{P}_1 or \mathbb{P}_0 .

Remark. When λ is known, a simpler test based on the number of edges in the graph also succeeds. The test in Theorem 4 still applies even when λ is unknown. The proof further implies that under \mathbb{P}_1 , G can be distinguished from $\mathcal{G}(n, \lambda'/n)$ for any λ' not necessarily equal to λ .

Finally, it is important to note that, as evidenced in [14], impossibility of detection does not imply immediately that of reconstruction. Fortunately, in our setting, the following result will imply the latter as soon as the former is proved :

Theorem 5. Assume that $K = o(\sqrt{n})$ that $\mathbf{E}_0(X_\Gamma) = \omega(1)$ and that $\mathbf{E}_0(L^2) = 1 + o(1)$. Then, for every estimator $\hat{\mathcal{K}}$ of the planted set \mathcal{K} , we have

$$\text{ov}(\hat{\mathcal{K}}) = o(K),$$

that is, reconstruction fails as well.

2.5 Proof strategy for planted paths

We say that the ordered set $\{i_1, \dots, i_K\}$ of K distinct nodes in $[n]$ is a K -path in G if and only if the edges $(i_\ell, i_{\ell+1})$ are present in G for all $\ell = 1, \dots, K-1$. The previous Lemma 1 yields, in the case where Γ is the line graph, the following result, whose proof is in the appendix:

Lemma 2. For planted K -path, the likelihood ratio reads

$$L(G) := \frac{\mathbb{P}_1(G)}{\mathbb{P}_0(G)} = \frac{1}{n(n-1)\cdots(n-K+1)} |\{K\text{-paths in } G\}| \left(\frac{\lambda}{n}\right)^{-K+1}. \quad (2.1)$$

Moreover one has

$$\mathbf{E}_0(L^2) = \mathbf{E}_0(x^S), \quad (2.2)$$

where $x = n/\lambda$, and S is a random variable counting the number of edges common to the K -path $(1 - 2 - \dots - K)$ and a random K -path π chosen uniformly at random among the $n(n-1)\cdots(n-K+1)$ possible ones on node set $[n]$.

2.5.1 Impossibility of detection

We have the following

Theorem 6. Assume that $\lambda > 1$ and $K = o(\sqrt{n})$, or alternatively that $\lambda < 1$ and $K = \ln(n)/\ln(1/\lambda) - \omega(\ln(\ln(n)))$. Then the total variation distance $|\mathbb{P}_1 - \mathbb{P}_0|_{\text{var}}$ between \mathbb{P}_0 and \mathbb{P}_1 goes to zero as $n \rightarrow \infty$. Thus for any arbitrary test $T(G) \in \{0, 1\}$, $\mathbb{P}_1(T(G) = 1) - \mathbb{P}_0(T(G) = 1) \rightarrow 0$ as $n \rightarrow \infty$.

By a standard argument, the variation distance $|\mathbb{P}_1 - \mathbb{P}_0|_{\text{var}}$ is upper-bounded by $\sqrt{\mathbf{E}_0(L^2) - 1}$, and thus the Theorem is a direct consequence of the following

Lemma 3. Assume that $\lambda > 1$ and $K = o(\sqrt{n})$, or alternatively that $\lambda < 1$ and $K = \ln(n)/\ln(1/\lambda) - \omega(\ln(\ln(n)))$. Then $\lim_{n \rightarrow \infty} \mathbf{E}_0(L^2) = 1$.

The proof of Lemma 3 (details in the Appendix) is based on an analysis of expression (2.2). Set $Z_t = 1$ if edge (I_t, I_{t+1}) is part of path $(1 \cdots K)$, $Z_t = 0$ if it is not part of that path, but $I_{t+1} \in [K]$, and finally $Z_t = -1$ if $I_{t+1} \notin [K]$, so that

$$\mathbf{E}_0(L^2) = \mathbf{E}_0(x^{\sum_{t=1}^{K-1} Z_t^+}) \quad (2.3)$$

In order to upper-bound this expression, a key step is the following Lemma, which exhibits a tractable upper bound involving a Markov chain:

Lemma 4. *Let $n' := n - K$. The Markov chain $\{Z_t'\}_{t \geq 1}$ taking values in $\{-1, 0, 1\}$ with transition probability matrix*

$$P := \begin{pmatrix} 1 - K/n' & K/n' & 0 \\ 1 - K/n' & (K-2)/n' & 2/n' \\ 1 - K/n' & (K-1)/n' & 1/n' \end{pmatrix} \quad (2.4)$$

can be constructed jointly with process $\{Z_t\}_{t \geq 1}$ so that, for all $m \geq 1$, one has

$$\mathbf{E}_0(x^{\sum_{t=1}^m Z_t^+}) \leq \mathbf{E}_0(x^{\sum_{t=1}^m Z_t'^+}). \quad (2.5)$$

Its proof is in the appendix, together with the analysis of the right-hand side of (2.5). The latter relies on spectral analysis of a matrix derived from P in (2.4), which leverages perturbation arguments as $K/n \rightarrow 0$. It concludes the proof of Lemma 3 by showing that $\mathbf{E}_0(L^2) = 1 + o(1)$ under the Lemma's assumptions.

2.5.2 Easiness of detection and reconstruction, sparse case

Assume $\lambda < 1$ and $K = \ln(n)/\ln(1/\lambda) + \omega(1)$. Detection is then easy: under \mathbb{P}_0 , the expected number of K -paths in the graph is $o(1)$. A test which decides \mathbb{P}_1 if there is a K -path and \mathbb{P}_0 otherwise thus discriminates the two hypotheses with high probability. Presence of a K -path can moreover be determined in polynomial time by running depth-first searches from each node in G .

For reconstruction, we need the following

Lemma 5. *For $\lambda < 1$, $K = \ln(n)/\ln(1/\lambda) + \omega(1)$ and $K = o(n/\ln(n))$, let C be the connected component of the graph containing the longest path. Apply \sqrt{K} times a peeling operation to C , which consists in removing all degree one nodes, to obtain set C' . Under \mathbb{P}_1 , set C' and its intersection with the planted path both have with high probability size $K \pm o(K)$.*

The Lemma readily implies a polynomial-time algorithm for reconstruction that achieves overlap $K - o(K)$: set C' can be obtained in polynomial time. By adding / removing $o(K)$ nodes to it one obtains a set of size K with overlap $K - o(K)$.

2.5.3 Impossibility of reconstruction, dense case

We assume $\lambda > 1$ and $K = \omega(\sqrt{n})$. We have seen that with high probability, observation of G allows to determine whether or not an attack has taken place. We now assume that an attack has indeed happened. We have the following result, showing the impossibility of efficient planted structure reconstruction:

Theorem 7. *Given $\lambda > 1$, $K = \omega(\sqrt{n})$, $K = o(n)$, and a realization G of the graph under \mathbb{P}_1 , any estimator $\hat{\mathcal{K}}$ of the ground truth achieves negligible overlap, i.e. $\text{ov}(\hat{\mathcal{K}}) = o(K)$.*

Its proof structure is as follows. Fix an arbitrary integer $\tau \geq 1$. We shall establish that necessarily

$$\text{ov}(\mathcal{K}) \leq K/(\tau + 1) + o(K). \quad (2.6)$$

Fix

$$L = C \ln(n) \text{ for some suitable constant } C, D \gg L \text{ and } D^2 \ll \frac{n}{\ln(n)}. \quad (2.7)$$

Condition on the event the attack path is precisely $k_1, \dots, k_K =: k_1^K$. Chop the attack path into $K/(L + D)$ contiguous segments, each of length $M := L + D$.

Consider the ℓ -th segment $\{k_{(\ell-1)M+1}, \dots, k_{\ell M}\}$. We shall construct, for some $I(\ell) \in [(\ell - 1)M + 1, (\ell - 1)M + L]$, τ random paths of edges in the graph G of the form

$$k_{I(\ell)}, I_2(t, \ell), I_3(t, \ell), \dots, I_D(t, \ell), k_{I(\ell)+D}$$

for $t \in [\tau]$ such that the nodes $I_2(t, \ell), \dots, I_D(t, \ell)$ are all distinct, none of them belongs to the attack path, and such that the paths $(k_1, \dots, k_K) =: k_1^K$ and $k_1^{I(\ell)}, I_2^D(t, \ell), k_{I(\ell)+D}^K$ are statistically indistinguishable. More precisely, we have the following:

Lemma 6. *There is a construction, for any $\ell \in [K/M]$, of τ random paths*

$$k_{I(\ell)}, I_2(t, \ell), I_3(t, \ell), \dots, I_D(t, \ell), k_{I(\ell)+D}, t \in [\tau],$$

such that for any $i \in [(\ell - 1)M + 1, (\ell - 1)M + L]$, any τ disjoint ordered sets of $D - 1$ distinct nodes $i_2^D(t)$, $t \in [\tau]$ in $[n] \setminus k_1^K$, we have

$$d_{\text{var}}(\mathbb{P}_1(G \in \cdot | \mathcal{K} = k_1^K, I(\ell) = i, (I_2^D(t, \ell))_t = (i_2^D(t))_t), \mathbb{P}_0(G \in \cdot | k_1^K \in G, (k_i, i_2^D(t), k_{i+D})_t \in G)) = \epsilon = o(1). \quad (2.8)$$

This construction moreover verifies the following property. There is an event \mathcal{E} such that $\mathbb{P}_1(\mathcal{E}) = 1 - o(1)$, and such that, denoting $|(\cup_{t \in [\tau]} I_2^D(t, \ell)) \cap (\cup_{t \in [\tau]} I_2^D(t, \ell'))|$ the number of common points between the node sets $\cup_{t \in [\tau]} I_2^D(t, \ell)$ and $\cup_{t \in [\tau]} I_2^D(t, \ell')$, one has:

$$\forall \ell \neq \ell' \in [K/M], \quad \mathbf{E}_1 (|(\cup_{t \in [\tau]} I_2^D(t, \ell)) \cap (\cup_{t \in [\tau]} I_2^D(t, \ell'))| \mathbf{1}_{\mathcal{E}}) = O\left(\frac{D^2}{n}\right). \quad (2.9)$$

The Lemma's proof idea is as follows. The τ non-overlapping alternative path segments, that we refer to as a τ -path, are obtained by selecting uniformly at random one such τ -path among all present in the graph. Then (2.8) is established by showing that the number of τ -paths concentrates. In turn, this concentration is established by bounding the variance of the number of τ -paths. This is done using the Markov chain bounding technique used in Lemma 4. The second part of the Lemma, (2.9), requires further concentration results on the numbers of τ -paths, that follow from applying Janson's inequality ([30], p. 205, Theorem 6.31).

The proof idea of Theorem 7 (detailed in the appendix) is then as follows. The τ -paths of Lemma 6 provide τ alternative K -paths to the actual planted path. These are "lures" for the optimal MAP reconstruction algorithm, that must return on average as many points of each of these lure paths as of the planted path. Since all these $\tau + 1$ paths have intersection of negligible size, the overlap achieved by MAP must necessarily be at most $K/(\tau + 1)$.

2.6 Proof strategy for planted D -ary trees

We assume here that Γ is a complete D -ary tree of size K and depth h , with $D > 1$ a fixed constant.

Under \mathbb{P}_0 , the neighbourhood of a given vertex in G is close to a Galton-Watson process with offspring law $\text{Poi}(\lambda)$. The probability of the existence of an infinite D -ary subtree in this process is the largest non-negative root $p_*(D, \lambda)$ of the equation

$$p = \psi_D(\lambda p), \quad (2.10)$$

where

$$\psi_D(\mu) := \mathbb{P}(\text{Poi}(\mu) \geq D), \quad \mu \geq 0.$$

The behavior of the random graph differs based on whether the above probability is zero or not. We define the *critical* threshold λ_D as

$$\lambda_D = \sup \{ \lambda > 0 \mid p_*(D, \lambda) = 0 \} \quad (2.11)$$

In the following, we focus on *subcritical* λ , that is whenever $\lambda < \lambda_D$.

2.6.1 Study of the Galton-Watson process

Let (T, o) be a rooted Galton-Watson tree with offspring law $\text{Poi}(\lambda)$, with $\lambda < \lambda_D$. The following Theorem characterizes the distribution of the maximum height of a D -ary tree rooted in o .

Theorem 8. *Let (T, o) be a Galton-Watson tree as above, and $n > 0$. Let p_h be the probability that a D -ary tree of height h rooted in o is contained in T . Then, for almost all λ , there exists h_* such that*

$$p_{h_*+1} = o\left(\frac{1}{n}\right) \quad (2.12)$$

$$p_{h_*} = \Omega(n^{-c}) \text{ for some } c < 1 \quad (2.13)$$

Moreover, as $n \rightarrow \infty$ one has $h_* = \frac{\ln \ln(n)}{\ln(D)} + O(1)$.

Thus h_* depends on λ only through terms of lower (constant) order. The Theorem's proof, detailed in the appendix, relies on the following

Lemma 7. *The sequence p_h satisfies the recurrence relation*

$$\begin{aligned} p_1 &= 1 \\ p_{h+1} &= \psi_D(\lambda p_h) \text{ for all } h \geq 1. \end{aligned}$$

Necessarily $0 \leq p_{h+1} \leq p_h$ for all h (since a tree of height $h+1$ contains a tree of height h), and therefore by continuity of ψ_D , p_h converges as $h \rightarrow \infty$ to the largest fixed point of (2.10). By definition of λ_D , the only solution of this equation is $p_\infty = 0$, and thus

$$\lim_{h \rightarrow \infty} p_h = 0 \quad (2.14)$$

Now, $\psi_D(x) \sim \frac{x^D}{D!}$ as $x \rightarrow 0$, which implies that for h large enough, $p_{h+1} \simeq C p_h^D$, and thus $p_h \simeq C \varepsilon^{D^h}$ for some small $\varepsilon > 0$. A more rigorous version of this argument, as well as its use in the proof of Theorem 8, is presented in the Appendix.

2.6.2 Coupling and application to planted trees

Following the insights from the previous section, we define the two thresholds \bar{h} and \underline{h} by :

$$\bar{h} = \inf \left\{ h > 0 \mid p_h < \frac{1}{n} \right\}, \quad \underline{h} = \sup \left\{ h > 0 \mid p_h > \frac{\ln(n)}{n} \right\}.$$

Theorem 8 implies that $\bar{h} \sim \frac{\ln \ln(n)}{\ln(D)}$, and that for almost all λ , $\bar{h} = \underline{h} + 1$, and otherwise $\bar{h} = \underline{h} + 2$. Also, $p_{\bar{h}} = o(\frac{1}{n})$ and $p_{\underline{h}} = \Omega(n^{-c})$ for some $c > 1$. The following Theorem connects the study from section 2.6.1 to our planted tree problem:

Theorem 9. *Let G be a graph drawn according to \mathbb{P}_0 , and $h > 0$. Then with high probability:*

1. *For $h \leq \underline{h}$, there are $\omega(1)$ D -ary trees of height h in G .*
2. *For $h \geq \bar{h} + C$, where C is a large enough constant, there are no D -trees of height h in G .*

The second part of this theorem yields an easy detection algorithm whenever $h \geq \bar{h} + \Omega(1)$.

Corollary 1. *Assume that Γ is a complete D -ary tree of height h , with $h \geq \bar{h} + \Omega(1)$. Then w.h.p under \mathbb{P}_0 , $X_\Gamma = 0$, and therefore the test $T(G) = 1$ iff $X_\Gamma > 0$ discriminates between H_0 and H_1 correctly with high probability.*

The two statements of Theorem 9 are a consequence of the following coupling lemma, whose proof, as well as the full proof of the theorem, is deferred to the appendix :

Lemma 8. *For a graph G and a vertex v in G , denote by $(G, v)_\ell$ the ℓ -neighbourhood of v in G . Similarly, let $(T, o)_\ell$ be the ℓ -neighbourhood of o in the Galton-Watson process described above.*

Then, under \mathbb{P}_0 , assuming that $\ell = o(\log(n))$, the total distance variation between the law of $(G, v)_\ell$ and that of $(T, o)_\ell$ goes to 0 as a negative power of n when $n \rightarrow \infty$.

Furthermore, for $\lambda' > \lambda$, and (T', o') a GW process with parameter λ' , then, provided the ℓ -neighbourhood of v is cycle-free, there exists a coupling between $(G, v)_\ell$ and $(T', o')_\ell$ such that $(G, v)_\ell \subseteq (T', o')_\ell$ with probability 1.

There is therefore a sharp cutoff in the probability of presence of tree of height h in G , and we have already seen in Corollary 1 that it can be leveraged to obtain a detection algorithm when $h \leq \underline{h}$. It remains however to study two aspects of the problem: reconstruction for $h \geq \bar{h}$, as well as the possibility (or lack thereof) of detection when $h \leq \underline{h}$.

2.6.3 Likelihood ratio and detection for $h \leq \underline{h}$

We conjecture, as is the case when $D = 1$, that when $h = \underline{h} - \omega(1)$, then the total variation distance $|\mathbb{P}_1 - \mathbb{P}_0|_{\text{var}}$ goes to 0 when $n \rightarrow \infty$. However, the Markov chain bounds used for lines cannot be easily adapted to the current setting, and we only prove this result for $h \leq \underline{h} - \Omega(\ln \ln \ln(n))$:

Theorem 10. *Assume that Γ is a D -ary tree of height h , with $D > 1$ and*

$$h \leq \underline{h} - \frac{\ln(\underline{h})}{\ln(D)} + \frac{\ln(1 - \frac{1}{D})}{\ln(D)}.$$

Then, the total variation distance $|\mathbb{P}_1 - \mathbb{P}_0|_{\text{var}}$ goes to zero as $n \rightarrow \infty$. Thus, for any test $T(G) \in \{0, 1\}$, $\mathbb{P}_1(T(G) = 1) - \mathbb{P}_0(T(G) = 1) \rightarrow 0$ as $n \rightarrow \infty$.

As before this is deduced from the following Lemma, shown in the Appendix:

Lemma 9. *Under the same assumptions as Theorem 10, $\mathbf{E}_0(L^2) \rightarrow 1$ as $n \rightarrow \infty$.*

We believe the following stronger version of the Theorem to hold:

Conjecture 1. *The result of Theorem 10 holds true for all $h \leq \underline{h}$.*

If true, this conjecture would complete the bottom left part of the phase diagram for D -ary tree, with a sharp threshold between undetectability and detection/reconstruction.

2.6.4 Reconstruction for large h

When $\lambda < \lambda_D$ and $h \geq \bar{h}$, we have shown that under \mathbb{P}_0 there is w.h.p no copy of Γ in G . One could therefore expect to be able to reconstruct Γ with overlap $1 - o(1)$; however, this is not the case :

Theorem 11. *Given $\lambda > 0$, $h \geq \bar{h}$ such that $K = o(n)$, and a realization G of the graph under \mathbb{P}_1 , the overlap achieved by any estimator \hat{K} of the attack is bounded above, i.e $\text{ov}(\hat{K}) \leq (1 - \delta)K$ for some $\delta > 0$.*

The proof is based on the fact that when $D > 1$, the leaves make up a positive proportion of Γ , and they are hard to reconstruct with high precision. On the other hand, since there is no copy of Γ in G w.h.p, one can still reasonably expect to achieve a partial reconstruction. This is the heuristic behind our second conjecture :

Conjecture 2. *For all $h \geq \bar{h}$, there exists a $\delta > 0$ and an estimator (possibly random) \hat{K} such that w.h.p $\text{ov}(\hat{K}) \geq \delta K$.*

2.7 Proof of preliminary results

2.7.1 Proof of Lemma 1

Let $\Gamma_1, \dots, \Gamma_m$ be the copies of Γ in K_n the complete graph on $[n]$, where, denoting $\text{Aut}(\Gamma)$ the automorphism group of Γ , $m = \binom{n}{K} \frac{K!}{|\text{Aut}(\Gamma)|}$. Then, by Bayes' formula, letting $e(G)$ denote the number of edges in graph G , one has for any graph g :

$$\begin{aligned} \mathbb{P}_1(G = g) &= \frac{1}{m} \sum_{i=1}^m \mathbb{P}_0(G = g \mid \Gamma_i \in G) = \frac{1}{m} \sum_{i=1}^m \mathbf{1}_{\Gamma_i \in g} \left(\frac{\lambda}{n}\right)^{e(g) - e(\Gamma_i)} \left(1 - \frac{\lambda}{n}\right)^{\binom{n}{2} - e(g)} \\ &= \frac{1}{m} \left(\frac{\lambda}{n}\right)^{-e(\Gamma)} \sum_{i=1}^m \mathbf{1}_{\Gamma_i \in g} \mathbb{P}_0(G = g) \\ &= \frac{X_\Gamma}{\mathbf{E}_0[X_\Gamma]} \mathbb{P}_0(G), \end{aligned}$$

which completes the proof of Lemma 1.

2.7.2 Proof of Theorem 1

We first prove that planted stars of size $K = \ln(n)/\ln(\ln(n))[1 - \omega(1/\ln(\ln(n)))]$ are undetectable. The number X of K -stars verifies

$$\mathbf{E}_0(X) = n \binom{n-1}{K} \left(\frac{\lambda}{n}\right)^K \sim n \frac{\lambda^K}{K!}.$$

We will have undetectability if $\mathbf{E}_0(L^2) \sim 1$, or equivalently by symmetry arguments, if

$$\mathbf{E}_0(X|\Gamma_1 \in G) \sim \mathbf{E}_0(X),$$

where Γ_1 is an arbitrary K -star, e.g. that made of edges $(i, K+1)$, $i \in [K]$. We decompose $\mathbf{E}_0(X|\Gamma_1 \in G)$ into three terms M_1 , M_2 and M_3 , the expected numbers of K -stars centered respectively: at node $K+1$, at some node $i \in [K]$, and finally at some node $i \in [n] \setminus [K+1]$. Since M_3 is upper-bounded by $\mathbf{E}_0(X)$, it suffices to show that M_1 and M_2 are $o(\mathbf{E}_0(X))$. One has:

$$\begin{aligned} M_2 &= K \left(\binom{n-2}{K-1} \left(\frac{\lambda}{n}\right)^{K-1} + \binom{n-2}{K} \left(\frac{\lambda}{n}\right)^{K-1} \right) \\ &\leq \frac{2K^2}{n} \mathbf{E}_0(X) \\ &\ll \mathbf{E}_0(X). \end{aligned}$$

Also,

$$\begin{aligned} M_1 &= \sum_{\ell=0}^K \binom{K}{\ell} \binom{n-K-1}{K-\ell} \left(\frac{\lambda}{n}\right)^{K-\ell} \\ &\leq \sum_{\ell=0}^K \binom{K}{\ell} \frac{\lambda^\ell}{\ell!} \\ &\leq (1+\lambda)^K. \end{aligned}$$

The desired result $M_1 \ll \mathbf{E}_0(X)$ will follow if

$$\ln(n) + K \ln(\lambda) - \ln(K!) - K \ln(1+\lambda) \rightarrow +\infty.$$

The terms in K are of order at most $\ln(n)/\ln(\ln(n))$. By Stirling's formula, this will therefore hold provided $\ln(n) - K \ln(K) = \omega(\ln(n)/\ln(\ln(n)))$. By assumption,

$$K \ln(K) \leq \frac{\ln(n)}{\ln(\ln(n))} (1 - \omega(1/\ln(\ln(n)))) \ln(\ln(n)) = \ln(n) - \omega(\ln(n)/\ln(\ln(n))),$$

hence the undetectability result.

Similarly for detectability, the assumption that $K = \ln(n)/\ln(\ln(n))[1 + \omega(1/\ln(\ln(n)))]$ entails that

$$\ln(\mathbf{E}_0(X)) = K \ln(\lambda) + \ln(n) - \ln(K!) = -\omega(1).$$

Thus a test which decides H_1 if and only if there is a node in G with degree at least K succeeds with high probability. Moreover, with high probability, only the centre of the planted star has degree at least K . The reconstruction method which consists in choosing, besides the highest degree node, K of its neighbours chosen uniformly at random, achieves an overlap of $K - o(K)$: indeed, conditional on the planted star's centre having initially Y neighbors in the original graph, the expected number of nodes in the reconstructed set will be

$$1 + \frac{K^2}{Y+K} \geq 1 + K(1 - Y/K) = 1 + K - Y.$$

Its expectation is lower-bounded by $K+1-\lambda$, and is thus $K - o(K)$.

2.7.3 Proof of Theorem 4

Let k denote the size of the hidden connected component, with $k = 0$ under \mathbb{P}_0 and $k = K$ under \mathbb{P}_1 . Let A_1 count the number of isolated nodes in G , A_2 the number of connected pairs (i, j) that form an isolated component, and A_3 the number of triplets (i, j, k) that form a connected component.

These quantities satisfy with high probability

$$A_1 = e^{-\lambda}(n-k) + O(\sqrt{n}), \quad A_2 = \frac{(n-k)^2 \lambda}{2} e^{-2\lambda} + O(\sqrt{n}), \quad A_3 = \frac{(n-k)^3 \lambda^2}{2} e^{-3\lambda} + O(\sqrt{n}). \quad (2.15)$$

Indeed, only the $n-k$ nodes that are not part of the hidden connected graph can contribute to counts of connected components of size 1, 2 or 3. (2.15) then follows from evaluation of the expectation and variance of these quantities.

Set $\hat{\lambda} = (nA_3)/(A_1A_2)$. By (2.15), $\hat{\lambda} = \lambda + O(n^{-1/2})$. Now form $\hat{k} = n - e^{\hat{\lambda}}A_1$. Again by (2.15), $\hat{k} = n - (1 - O(n^{-1/2}))(n-k) + O(\sqrt{n}) = k + O(\sqrt{n})$. Our test then decides H_1 if $\hat{k} \geq t_n$ and H_0 otherwise where t_n is such that $\sqrt{n} \ll t_n \ll K$, which is indeed satisfied for $t_n = \sqrt{K\sqrt{n}}$. This ensures that the test discriminates correctly between the two hypotheses with high probability. Necessarily then, the variation distance $|\mathbb{P}_0 - \mathbb{P}_1|_{var}$ goes to 1 as $n \rightarrow \infty$.

2.7.4 Proof of Theorem 5

We first begin by a simple lemma, using the concentration of X_Γ :

Lemma 10. *Let \mathcal{I}_Γ be the proportion of pairs copies of Γ in G whose intersection is nonempty :*

$$\mathcal{I}_\Gamma = \frac{1}{X_\Gamma^2} \sum_{\Gamma', \Gamma'' \in G} \mathbf{1}_{\Gamma' \cap \Gamma'' \neq \emptyset},$$

where Γ' and Γ'' range over all copies of Γ in G .

Then $\mathbf{E}_0(\mathcal{I}_\Gamma) = o(1)$.

Proof. (of Lemma 10). As in the proof of Lemma 1, let $\Gamma_1, \dots, \Gamma_m$ be the copies of Γ in K_n , and let $X_i = \mathbf{1}_{\Gamma_i \in G}$. Write

$$\mathbf{E}_0(X_\Gamma^2) = \sum_{i,j} \mathbf{E}_0(X_i X_j) = \mathbf{E}' + \mathbf{E}'', \quad (2.16)$$

where \mathbf{E}' is the sum over Γ_i, Γ_j having disjoint vertex sets.

We can easily compute \mathbf{E}' :

$$\mathbf{E}' = \binom{n}{K} \binom{n-K}{K} \left(\frac{K!}{|\text{Aut}(\Gamma)|} \right)^2 p^{2K-2} \sim \frac{n^{2K} p^{2K-2}}{|\text{Aut}(\Gamma)|^2} \sim \mathbf{E}_0(X_\Gamma)^2$$

Since $\mathbf{E}_0(L^2) = 1 + o(1)$, it follows that

$$\frac{\mathbf{E}''}{\mathbf{E}_0(X_\Gamma^2)} = o(1). \quad (2.17)$$

Now, it is straightforward to see that

$$\sum_{\Gamma', \Gamma'' \in G} \mathbf{1}_{\Gamma' \cap \Gamma'' \neq \emptyset} = \sum_{\Gamma_i \cap \Gamma_j \neq \emptyset} X_i X_j.$$

Recall that $L = X_\Gamma / \mathbf{E}_0(X_\Gamma)$; we can decompose \mathcal{I}_Γ as follows :

$$\begin{aligned} \mathcal{I}_\Gamma &= \mathcal{I}_\Gamma \mathbf{1}_{L^2 > 1/2} + \mathcal{I}_\Gamma \mathbf{1}_{L^2 < 1/2} \\ &= \frac{\sum_{\Gamma_i \cap \Gamma_j \neq \emptyset} X_i X_j}{\mathbf{E}_0(X_\Gamma^2)} \cdot \frac{1}{L^2} \cdot \mathbf{1}_{L^2 > 1/2} + \mathcal{I}_\Gamma \mathbf{1}_{L < 1/\sqrt{2}} \end{aligned}$$

We can now bound each term separately. The first one is straightforward since $1/L^2 < 2$ whenever the indicator variable is nonzero ; for the second one, notice that $I_\Gamma \leq 1$ and thus

$$\begin{aligned} \mathbf{E}_0(\mathcal{I}_\Gamma) &\leq \frac{\mathbf{E}''}{\mathbf{E}_0(X_\Gamma^2)} \cdot 2 + \mathbb{P}_0\left(L < \frac{1}{\sqrt{2}}\right) \\ &= \frac{\mathbf{E}''}{\mathbf{E}_0(X_\Gamma^2)} \cdot 2 + o(1), \end{aligned}$$

having used the Bienaymé-Chebychev inequality to bound the second term.

Using (2.17) then completes the proof. \square

We can now move on to the proof of Theorem 5 ; we first transform the expression of $\text{ov}(\hat{\mathcal{K}})$ to better suit our needs :

$$\begin{aligned} \text{ov}(\hat{\mathcal{K}}) &= \sum_G \sum_{\mathcal{K}} \mathbb{P}_1(G, \mathcal{K}) \left| \hat{\mathcal{K}} \cap \mathcal{K} \right| \\ &= \sum_G \mathbb{P}_1(G) \sum_{\mathcal{K}} \mathbb{P}_1(\mathcal{K} | G) \left| \hat{\mathcal{K}} \cap \mathcal{K} \right| \end{aligned}$$

where \mathcal{K} ranges over all K -subsets of $[n]$ and G over all graphs on n vertices.

The second sum can be transformed as in the proof of Lemma 1 into :

$$\begin{aligned} \text{ov}(\hat{\mathcal{K}}) &= \sum_G \mathbb{P}_1(G) \sum_{\Gamma' \in G} \frac{|\hat{\mathcal{K}} \cap \Gamma'|}{X_\Gamma} \\ &= \sum_G \mathbb{P}_0(G) \sum_{\Gamma' \in G} \frac{|\hat{\mathcal{K}} \cap \Gamma'|}{X_\Gamma} + o(K), \end{aligned}$$

since the conditions in Theorem 5 imply that $|\mathbb{P}_1 - \mathbb{P}_0|_{\text{var}} = o(1)$ (see the remark after Theorem 6). The sum now ranges over all copies of Γ in G .

This can now be expressed as an expectation :

$$\begin{aligned} \text{ov}(\hat{\mathcal{K}}) &= \mathbf{E}_0 \left[\sum_{\Gamma' \in G} \frac{|\hat{\mathcal{K}} \cap \Gamma'|}{X_\Gamma} \right] + o(K) \\ &= \sum_{i \in [n]} \mathbf{E}_0 \left[\mathbf{1}_{i \in \hat{\mathcal{K}}} \sum_{\Gamma' \in G} \frac{\mathbf{1}_{i \in \Gamma'}}{X_\Gamma} \right] + o(K). \end{aligned}$$

We can now finally use Lemma 10 : indeed,

$$\begin{aligned} \left(\sum_{\Gamma' \in G} \frac{\mathbf{1}_{i \in \Gamma'}}{X_\Gamma} \right)^2 &= \frac{1}{X_\Gamma^2} \sum_{\Gamma', \Gamma'' \in G} \mathbf{1}_{i \in \Gamma'} \mathbf{1}_{i \in \Gamma''} \\ &\leq \frac{1}{X_\Gamma^2} \sum_{\Gamma', \Gamma'' \in G} \mathbf{1}_{\Gamma' \cap \Gamma'' \neq \emptyset} \\ &= \mathcal{I}_\Gamma. \end{aligned}$$

Therefore,

$$\begin{aligned} \text{ov}(\hat{\mathcal{K}}) &\leq \sum_{i \in [n]} \mathbf{E}_0 \left[\mathbf{1}_{i \in \hat{\mathcal{K}}} \sqrt{\mathcal{I}_\Gamma} \right] + o(K) \\ &= K \mathbf{E}_0 \left[\sqrt{\mathcal{I}_\Gamma} \right] + o(K) \\ &= o(K), \end{aligned}$$

using Jensen's inequality as well as Lemma 10. This completes the proof of Theorem 2.17.

2.8 Detailed proofs for planted paths

2.8.1 Proof of Lemma 2

Expression (2.1) follows directly from Lemma 1. In the display below, by $\sum_{(i_1 \dots i_K)}$ we mean summation over all the $n(n-1) \dots (n-K+1)$ oriented paths (i_1, \dots, i_K) of length K over nodes in $[n]$. Write:

$$\begin{aligned} \mathbf{E}_0(L^2) &= \sum_{(i_1 \dots i_K)} \sum_{(j_1 \dots j_K)} \left(\frac{(n/\lambda)^{K-1}}{n \dots (n-K+1)} \right)^2 \mathbb{P}_0(\text{paths } (i_1 \dots i_K) \text{ and } (j_1 \dots j_K) \text{ present in } G) \\ &= \sum_{(i_1 \dots i_K)} \left(\frac{(n/\lambda)^{2(K-1)}}{n \dots (n-K+1)} \right) \mathbb{P}_0(\text{paths } (i_1 \dots i_K) \text{ and } (1 \dots K) \text{ present in } G) \\ &= \left(\frac{n}{\lambda} \right)^{K-1} \mathbb{P}_0(\text{ path } \pi = (I_1 \dots I_K) \text{ present in } G \mid \text{ path } (1 \dots K) \text{ present in } G), \end{aligned}$$

where $\pi = (I_1 \dots I_K)$ is a candidate path chosen uniformly at random from the $n(n-1) \dots (n-K)$ possible length- K paths. In the above we used symmetry to consider a single path $(1 \dots K)$ instead of all paths $(j_1 \dots j_K)$.

Note that conditionally on the event that path $(1 \dots K)$ be present in G and on the path π , the probability that path π is also present in G is given by $(\lambda/n)^{K-1-S}$, where S is the number of edges in common between the two paths π and $(1 \dots K)$. This yields expression (2.2).

2.8.2 Proof of Lemma 4

Let $\mathcal{F}_t = \sigma(I_1, \dots, I_t)$. Recall that $n' = n - K$. It is easily verified that we have the following inequalities for all $t = 2, \dots, K-1$:

$$\mathbb{P}(Z_t = 1 \mid \mathcal{F}_t) \leq \begin{cases} \frac{1}{n'} & \text{if } Z_{t-1} = 1, \\ \frac{2}{n'} & \text{if } Z_{t-1} = 0, \\ 0 & \text{if } Z_{t-1} = -1. \end{cases}$$

Similarly we have

$$\mathbb{P}(Z_t \geq 0 \mid \mathcal{F}_t) \leq \frac{K}{n'}.$$

Moreover it is easily seen that $\mathbb{P}(Z_1 = 1) \leq (K/n')(2/n')$, and $\mathbb{P}(Z_1 \geq 0) \leq K/n'$.

As in Lemma 4, we introduce the Markov chain $\{Z'_t\}_{t \geq 1}$ on state space $\{-1, 0, 1\}$ specified by the initial distribution $\mathbb{P}(Z'_1 = 1) = (K/n')(2/n')$, $\mathbb{P}(Z'_1 \geq 0) = K/n'$ and by the transition probability matrix P in (2.4), that we recall for convenience:

$$P = \begin{pmatrix} 1 - K/n' & K/n' & 0 \\ 1 - K/n' & (K-2)/n' & 2/n' \\ 1 - K/n' & (K-1)/n' & 1/n' \end{pmatrix}$$

The previous inequalities ensure that we can construct by induction over t a coupled version of the two processes $\{Z_t\}$ and $\{Z'_t\}$ such that $Z_1 \leq Z'_1$, and for $t \geq 1$, if $Z'_t = -1$ then $Z_t = -1$, and furthermore we have the following implications:

$$\begin{aligned} Z_t = -1 &\quad \Rightarrow Z_{t+1} \leq Z'_{t+1}, \\ Z_t = Z'_t &\quad \Rightarrow Z_{t+1} \leq Z'_{t+1}, \\ (Z_t, Z'_t) = (1, 0) &\quad \Rightarrow Z_{t+1} \leq Z'_{t+1}. \end{aligned}$$

Thus the only situation when we can have $Z_{t+1} > Z'_{t+1}$ is when $(Z_t, Z'_t) = (0, 1)$. That is to say, for each time $t + 1$ when process Z hits 1 while chain Z' does not, then at time t chain Z' hits 1 while process Z does not.

Because of this, the number of times t at which process Z hits 1 is upper-bounded by the number of times t at which chain Z' does. Thus (2.5) holds, concluding the proof of Lemma 4.

2.8.3 Proof of Lemma 3

By (2.5) and (2.3), $\mathbf{E}_0(L^2)$ is upper bounded by

$$\mathbf{E}_0(L^2) \leq \mathbf{E}_0 x^{\sum_{s=1}^{K-1} Z'_s}. \quad (2.18)$$

To evaluate this term, introduce the row vector $F(t) := \{f_z(t)\}_{z \in \{-1, 0, 1\}}$ where

$$f_z(t) := \mathbf{E}_0 \left[x^{\sum_{s=1}^t Z'_s} \mathbf{1}_{Z'_t=z} \right].$$

We then have

$$F(1) = (\mathbb{P}(Z'_1 = -1), \mathbb{P}(Z'_1 = 0), x\mathbb{P}(Z'_1 = 1)) = (1 - K/n', K/n'(1 - 2/n'), x(K/n')(2/n')), \quad (2.19)$$

together with the recurrence relation

$$F(t+1) = F(t)M, \quad (2.20)$$

where

$$M = \begin{pmatrix} 1 - K/n' & K/n' & 0 \\ 1 - K/n' & K/n' - 2/n' & x2/n' \\ 1 - K/n' & K/n' - 1/n' & x/n' \end{pmatrix}$$

Recall now that $x = n/\lambda$ and $n' = n - K$, so that x/n' is asymptotic to $1/\lambda$. Thus the above matrix M reads

$$M = M_0 + (K/n)M_1,$$

where

$$M_0 = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 0 & 2/\lambda \\ 1 & 0 & 1/\lambda \end{pmatrix}, \quad (2.21)$$

and the entries of matrix M_1 are $O(1)$. Note that M_0 admits eigenvalues $0, 1/\lambda, 1$ with respective left eigenvectors

$$\begin{aligned} u_0 &:= (1, 1, -2), \\ u_{1/\lambda} &:= (-\lambda/(\lambda - 1), 0, 1), \\ u_1 &:= (1, 0, 0). \end{aligned}$$

We shall denote (μ_r, v_r) the (eigenvalue, eigenvector) pair of M obtained by perturbation of the eigenpair (r, u_r) of M_0 , with $r \in \{0, 1/\lambda, 1\}$. By the Bauer-Fike theorem (see [25], Theorem VI.25.1), $|\mu_r - r| = O(K/n)$ for all r .

Moreover Eq. (1.16), p. 67 in [77] implies that a normed left (resp., right) eigenvector of M associated to an eigenvalue μ_r of M differs in norm from a normed left (resp., right) eigenvector of M_0 associated to eigenvalue r by $O(K/n)$. We can thus chose $v_r = u_r + O(K/n)$.

Let the decomposition of vector $F(1)$ in the basis provided by the eigenvectors $\{v_r\}$ be given by:

$$F(1) = \sum_{r \in \{0, 1/\lambda, 1\}} \alpha_r v_r.$$

Denote by e the all-ones 3×1 column vector. The upper bound (2.18) on $\mathbf{E}_0(L^2)$ then gives

$$\begin{aligned} \mathbf{E}_0(L^2) &\leq F(K-1)e \\ &= F(1)M^{K-2}e \\ &= \sum_{r \in \{0, 1/\lambda, 1\}} \alpha_r v_r \mu_r^{K-2} e. \end{aligned} \tag{2.22}$$

By our choice of eigenvectors v_r such that $|v_r - u_r| = O(K/n)$, and the fact that

$$F(1) = (1 + O(K/n))u_1 + O(K/n)u_{1/\lambda} + O(K/n)u_0,$$

corresponding weights α_r verify $\alpha_1 = 1 + O(K/n)$, $\alpha_{1/\lambda} = O(K/n)$, $\alpha_0 = O(K/n)$.

In the case where $\lambda > 1$ and $K = o(\sqrt{n})$, (2.22) yields

$$\mathbf{E}_0(L^2) \leq o(1) + (1 + o(1))\mu_1^{K-2} = (1 + o(1))(1 + O(K/n))^{K-2} \leq (1 + o(1))e^{O(K^2/n)}.$$

The assumption that $K = o(\sqrt{n})$ then allows to conclude.

For $\lambda < 1$ and $K = \ln(n)/\ln(1/\lambda) - \omega(\ln(\ln(n)))$, (2.22) yields

$$\mathbf{E}_0(L^2) \leq (1 + o(1)) (1 + O(K/n))^{K-2} + O(K/n) (1/\lambda + O(K/n))^{K-2}.$$

The first term is $1 + o(1)$ since $K^2/n = o(1)$. The second term's logarithm is equivalent to

$$\ln(K) - \ln(n) + (K-2)\ln(1/\lambda) \leq \ln(\ln(n)) - \ln(\ln(1/\lambda)) - \omega(\ln(\ln(n))),$$

and goes to $-\infty$ by assumption.

2.8.4 Proof of Lemma 5

We place ourselves under \mathbb{P}_1 and condition on the fact that the K -path planted in the original Erdős-Rényi graph G_0 is k_1^K . Denote for each $i \in [K]$ by C_i the connected component of node k_i in G_0 . Denote by \mathcal{E}_i the event that $C_i \cap \{\cup_{j \neq i} C_j\} \neq \emptyset$ and by \mathcal{E}'_i the event that C_i contains a cycle.

A standard construction of connected components based on a random walk exploration implies the existence of a constant $c > 0$ such that for all $\ell \geq 0$,

$$\begin{aligned} \mathbb{P}(\mathcal{E}'_i, |C_i| = \ell) &\leq \frac{\lambda \ell^2}{n} \mathbb{P}(|C_i| = \ell) \leq \frac{\lambda \ell^2}{n} e^{-c\ell}, \\ \mathbb{P}(\mathcal{E}_i, |C_i| = \ell) &\leq \frac{\ell K}{n} e^{-c\ell}, \\ \mathbb{P}(|C_i| \geq \ell) &\leq e^{-c\ell}. \end{aligned} \tag{2.23}$$

The first evaluation implies that with high probability, no C_i contains a cycle (i.e. no \mathcal{E}'_i occurs) when $K = o(n)$. The second evaluation implies that the expected number of $i \in [K]$ such that \mathcal{E}_i occurs and $|C_i| \geq \ell$ is upper bounded, for some constant $c' > 0$, by

$$\sum_{i \in [K]} \mathbb{P}(\mathcal{E}_i, |C_i| \geq \ell) \leq \frac{K^2}{n} e^{-c'\ell}.$$

If $K^2 = o(n)$, then this implies that with high probability, no \mathcal{E}_i occurs. Thus with high probability, there is no cycle in the connected component C . Moreover, the third evaluation in (2.23) ensures that

$$\sum_{i \in K} \mathbb{P}(C_i \geq \sqrt{K}) \leq K e^{-c\sqrt{K}} = o(1).$$

Thus the peeling process applied \sqrt{K} times to C returns exactly the planted K -path, except for \sqrt{K} nodes at each of its ends.

If on the other hand, $K^2 > o(n)$, we choose $\ell^* = \theta \ln(n)$ and deduce from (2.23) that with probability $1 - O(n^{-2})$, say, there is no $i \in [K]$ such that both \mathcal{E}_i and $|C_i| \geq \theta \ln(n)$ hold. The peeling process applied \sqrt{K} times to C then returns the planted path, shortened by no more than \sqrt{K} nodes at each end, plus parts of the neighborhoods C_i for which \mathcal{E}_i occurs. The expected number of nodes returned that do not belong to the planted path is therefore no more than

$$K \mathbb{P}(\mathcal{E}_i) \ell^* = O\left(\frac{K^2}{n}\right) \theta \ln(n).$$

This is $o(K)$ under the assumption that $K = o(n/\ln(n))$. The conclusion of the Lemma follows.

2.8.5 Proof of Theorem 7

We show that Lemma 6 implies (2.6). First, the optimal overlap is achieved by the *Maximum A Posteriori* (MAP) inference procedure, i.e. by putting in $\hat{\mathcal{K}}$ the K nodes with the highest probability, conditional on the observed graph G , of being in \mathcal{K} . The probability that node j belongs to \mathcal{K} conditional on G is proportional to the number of K -paths in G to which j belongs. We denote by \mathcal{K}^* the corresponding set.

Second, when under the alternative distribution $\mathbb{P}_2 := \mathbb{P}_0(G \in \cdot | k_1^K \in G, (k_i, i_2^D, k_{i+D}) \in G)$ in (2.8), the joint distribution of the numbers of K -paths going through the nodes k_1^K or through the nodes in k_1^i, i_2^D, k_{i+D}^K are statistically indistinguishable. Thus, letting N_ℓ (respectively N'_ℓ) denote the number of points of $k_{(\ell-1)M+1}^{\ell M}$ (respectively, $k_{(\ell-1)M+1}^i, i_2^D, k_{i+D}^{\ell M}$) that the MAP estimate selects, one has:

$$\mathbf{E}_2(N_\ell) = \mathbf{E}_2(N'_\ell).$$

Let also $N'_{t,\ell}$ denote the number of points that the MAP estimate selects in $k_{(\ell-1)M+1}^{I(t,\ell)}, I_2^D(t,\ell), k_{I(t,\ell)+D}^{\ell M}$. Since each of these variables is bounded by $M = L+D$, the variation distance bound (2.8) implies

$$\mathbf{E}_1(N_\ell) \leq \mathbf{E}_1(N'_{t,\ell}) + \epsilon M.$$

Summing these inequalities over $\ell \in [K/M]$ and $t \in [\tau]$ yields

$$\tau \sum_{\ell=1}^{K/M} \mathbf{E}_1(N_\ell) = \tau \text{ov}(\mathcal{K}^*) \leq \sum_{t=1}^{\tau} \sum_{\ell=1}^{K/M} \mathbf{E}_1(N'_{t,\ell}) + \epsilon \tau K. \quad (2.24)$$

However, it holds that:

$$\sum_{i=1}^K \mathbf{1}_{k_i \in \mathcal{K}^*} + \sum_{j \in \cup_{t,\ell} I_2^D(t,\ell)} \mathbf{1}_{j \in \mathcal{K}^*} \leq K.$$

This entails (using e.g. Bonferroni's inequality):

$$\sum_{i=1}^K \mathbf{1}_{k_i \in \mathcal{K}^*} + \sum_{\ell=1}^{K/M} \sum_{r=2}^D \sum_{t \in [\tau]} \mathbf{1}_{I_r(t,\ell) \in \mathcal{K}^*} - \sum_{\ell \neq \ell', \ell, \ell' \in [K/M]} |(\cup_{t \in [\tau]} I_2^D(t,\ell)) \cap (\cup_{t \in [\tau]} I_2^D(t,\ell'))| \leq K.$$

Taking expectations and using the last statement (2.9) of the Lemma yields, separating evaluations on event \mathcal{E} and on its complementary set $\bar{\mathcal{E}}$:

$$\text{ov}(\mathcal{K}^*) + \left\{ \sum_{t \in [\tau]} \sum_{\ell=1}^{K/M} \mathbf{E}_1(N'_{t,\ell}) \right\} - \tau L(K/M) - (K/M)^2 O(D^2/n) - \tau K \mathbb{P}_1(\bar{\mathcal{E}}) \leq K.$$

Summed with the previous equation (2.24), this gives:

$$(\tau + 1) \text{ov}(\mathcal{K}^*) \leq K + K\tau (\epsilon + (L/M) + (K/n)(D/M)^2 + \mathbb{P}_1(\bar{\mathcal{E}})).$$

The announced result follows from $\epsilon \ll 1$, $L \ll D$, $K = o(n)$ and $\mathbb{P}_1(\bar{\mathcal{E}}) = o(1)$.

2.8.6 Proof of Lemma 6, Equation (2.8)

We let π_i denote the set of τ candidate paths $(k_i, i_2^D(t, \ell), k_{i+D})_{t \in [\tau]}$ of the graph, where for fixed ℓ , the $\{i_2^D(t, \ell)\}_{t \in [\tau]}$ are distinct and in $[n] \setminus k_1^K$. For $i \in [(\ell - 1)M + 1, (\ell - 1)M + L]$ these can all be used to construct the set of τ alternative paths in the ℓ -th segment of k_1^K . We denote by

$$\pi(\ell) = \cup_{i \in [(\ell-1)M+1, (\ell-1)M+L]} \pi_i$$

the corresponding collection. Our construction simply amounts to choosing a set of τ paths (that we shall call for short a τ -path) uniformly at random from $\pi(\ell)$ in order to construct the alternative τ -path for the ℓ -th segment, and this independently for each segment.

Denote $Z_i = |\pi_i|$. Then

$$\mathbf{E}_1(Z_i) = (n - K)(n - K - 1) \cdots (n - K - \tau(D - 1) + 1) \left(\frac{\lambda}{n}\right)^{\tau D} \sim \frac{1}{n^\tau} \lambda^{\tau D},$$

since we assumed in (2.7) that $D \sim C \ln(n)$. Also, by symmetry,

$$\begin{aligned} \mathbf{E}_1 Z_i^2 &= \sum_{i_2^D(t), j_2^D(t)} \mathbb{P}_1(\forall t \in [\tau], (k_i, i_2^D(t), k_{i+D}) \in G, (k_i, j_2^D(t), k_{i+D}) \in G) \\ &= \mathbf{E}_1(Z_i) \sum_{j_2^D(t)} \mathbb{P}_1(\forall t \in [\tau], (k_i, j_2^D(t), k_{i+D}) \in G | \forall t \in [\tau], (k_i, i_2^D(t), k_{i+D}) \in G), \end{aligned}$$

where in the last expression we fixed an arbitrary choice $(i_2^D(t))_{t \in [\tau]}$. It follows that:

$$\mathbf{E}_1 Z_i^2 = (\mathbf{E}_1(Z_i))^2 \mathbf{E}_1((n/\lambda)^S),$$

where S is the number of common edges between the fixed τ -path $(k_i, i_2^D(t), k_{i+D})_{t \in [\tau]}$ and the τ -path $(k_i, j_2^D(t), k_{i+D})_{t \in [\tau]}$ where $(j_2^D(t))_{t \in [\tau]}$ is chosen uniformly at random among $(\tau(D - 1))$ sequences in $[n] \setminus k_1^K$.

To control this second moment, we will condition on the number of common edges between each path $J_2^D(t)$ in the randomly selected τ -path at its beginning and end with the beginning and end of some of the fixed paths $i_2^D(t')$, that we shall denote by X_t and Y_t . These satisfy the constraints $X_t, Y_t \geq 0$, $X_t + Y_t \leq D$. For $X_t + Y_t < D$, this forces the choice of $X_t + Y_t$ nodes among the $D - 1$ to be chosen for path $J_2^D(t)$; for $X_t + Y_t = D$, this forces all the $D - 1$ choices. Moreover, conditionally on $(X_t, Y_t)_{t \in [\tau]}$, the expectation of the variable $(n/\lambda)^S$ verifies

$$\mathbf{E}_1((n/\lambda)^S | (X_t, Y_t)_{t \in [\tau]}) \leq (n/\lambda)^{\sum_{t \in [\tau]} X_t + Y_t} (1 + O(D/n))^{\tau D},$$

by the Markov chain bounds in Lemma 4. By assumption, $D \ll \sqrt{n}$ so that $(1 + O(D/n))^D = 1 + o(1)$. Thus, accounts for the $(\tau!)^2$ choices of path correspondences between the beginnings and ends of the planted and random paths:

$$\begin{aligned} \mathbf{E}_1 Z_i^2 &\leq (\mathbf{E}_1(Z_i))^2 (\tau!)^2 \left[(n/\lambda)^D n^{D-1} + \sum_{x,y \geq 0, x+y < D} (n/\lambda)^{x+y} n^{-(x+y)} (1 + o(1)) \right]^\tau \\ &\leq (\mathbf{E}_1(Z_i))^2 (1 + o(1)) (\tau!)^2 [n\lambda^{-D} + (\sum_{x \geq 0} \lambda^{-x})^2]^\tau \\ &\leq (\mathbf{E}_1(Z_i))^2 (1 + o(1)) (\tau!)^2 \left(\frac{\lambda}{\lambda-1} \right)^{2\tau}, \end{aligned}$$

where we used that $n\lambda^{-D} = o(1)$.

We now evaluate $\mathbf{E}_1(Z_i Z_j)$ for $i \neq j$. The Markov chain bounding technique of Lemma 4 directly applies to give:

$$\mathbf{E}_1(Z_i Z_j) \leq (\mathbf{E}(Z_i))^2 (1 + o(1)).$$

Finally we obtain:

$$\begin{aligned} \text{Var}(|\pi(\ell)|) &= L \text{Var}(Z_i) + L(L-1) \text{Cov}(Z_i, Z_j) \\ &\leq \mathbf{E}_1(Z_i)^2 \left[L(1 + o(1)) (\tau!)^2 \left(\frac{\lambda}{\lambda-1} \right)^{2\tau} + L^2 o(1) \right] \\ &\leq \mathbf{E}_1(|\pi(\ell)|)^2 \left[\frac{O(1)}{L} + o(1) \right]. \end{aligned}$$

Since by assumption $L \gg 1$, Tchebitchev's inequality implies that the random variable $|\pi(\ell)|$ concentrates: for some suitable $\epsilon = o(1)$, one has

$$\mathbb{P}_1 \left(\left| \frac{|\pi(\ell)|}{\mathbf{E}_1|\pi(\ell)|} - 1 \right| \geq \epsilon \right) \leq \epsilon.$$

Denote by \mathcal{A} the event $\mathcal{A} := \{ \left| \frac{|\pi(\ell)|}{\mathbf{E}_1|\pi(\ell)|} - 1 \right| \leq \epsilon \}$. It thus has probability at least $1 - \epsilon$. Consider a bounded function f of the graph G . This concentration result allows us to establish the variation distance bound (2.8) as follows. For some arbitrary candidate τ -path $(i, i_2^D(t))_{t \in [\tau]}$, omitting for brevity the argument t below, write:

$$\mathbf{E}_1(f(G) | \mathcal{A}, \mathcal{K} = k_1^K, \mathbf{I}(\ell) = i, \mathbf{I}_2^D(\ell) = i_2^D) = \frac{\mathbf{E}_1[f(G) \mathbf{1}_{\mathcal{A}} \mathbf{1}_{(k_i, i_2^D, k_{i+D}) \in G} \frac{1}{|\pi(\ell)|}]}{\mathbf{E}_1(\mathbf{1}_{\mathcal{A}} \mathbf{1}_{(k_i, i_2^D, k_{i+D}) \in G} \frac{1}{|\pi(\ell)|})}.$$

On \mathcal{A} one has

$$\frac{1}{\mathbf{E}_1|\pi(\ell)|} \frac{1}{1 + \epsilon} \leq \frac{1}{|\pi(\ell)|} \leq \frac{1}{\mathbf{E}_1|\pi(\ell)|} \frac{1}{1 - \epsilon}.$$

This yields:

$$\frac{1 - \epsilon}{1 + \epsilon} \frac{\mathbf{E}_1[f(G) \mathbf{1}_{\mathcal{A}} \mathbf{1}_{(k_i, i_2^D, k_{i+D}) \in G}]}{\mathbb{P}_1((k_i, i_2^D, k_{i+D}) \in G)} \leq \mathbf{E}_1(f(G) | \mathcal{A}, \mathcal{K} = k_1^K, \mathbf{I}(\ell) = i, \mathbf{I}_2^D(\ell) = i_2^D) \leq \frac{1 + \epsilon}{1 - \epsilon} \frac{\mathbf{E}_1[f(G) \mathbf{1}_{(k_i, i_2^D, k_{i+D}) \in G}]}{\mathbb{P}_1(\mathcal{A} \cap (k_i, i_2^D, k_{i+D}) \in G)}$$

By symmetry over all τ -paths in $\pi(\ell)$, denoting by Z the total number of possible such τ -paths in it ($Z \sim Ln^{\tau(D-1)}$), one has

$$\mathbb{P}_1(\mathcal{A} \cap (k_i, i_2^D, k_{i+D}) \in G) = \frac{1}{Z} \mathbf{E}_1(|\pi(\ell)| \mathbf{1}_{\mathcal{A}}).$$

However by definition of \mathcal{A} this is no smaller than

$$\frac{1}{Z} (1 - \epsilon) \mathbf{E}_1|\pi(\ell)| \mathbb{P}_1(\mathcal{A}) \geq (1 - \epsilon)^2 \mathbb{P}_1((k_i, i_2^D, k_{i+D}) \in G).$$

Finally we obtain:

$$\begin{aligned} \frac{1-\epsilon}{1+\epsilon} \left[\mathbf{E}_1[f(G) | (k_i, i_2^D, k_{i+D}) \in G] - \|f\|_\infty \epsilon \right] &\leq \mathbf{E}_1(f(G) | \mathcal{A}, \mathcal{K} = k_1^K, \mathbf{I}(\ell) = i, \mathbf{I}_2^D(\ell) = i_2^D) \leq \dots \\ \dots &\leq \frac{1+\epsilon}{(1-\epsilon)^3} \mathbf{E}_1[f(G) | (k_i, i_2^D, k_{i+D}) \in G]. \end{aligned}$$

The result of Equation (2.8) follows.

2.8.7 Proof of Lemma 6, Equation (2.9)

We define the event \mathcal{E} as, for some suitable constant $\alpha = \Omega(1)$:

$$\mathcal{E} := \cap_{\ell \in [K/M]} \mathcal{E}_\ell, \text{ where } \mathcal{E}_\ell := \{|\pi(\ell)| \geq \alpha \mathbf{E}_1 |\pi(\ell)|\}. \quad (2.25)$$

In the below display we let $I_2^D(\ell) = \cup_{t \in [\tau]} I_2^D(t, \ell)$, and $I_2^D(\ell) \cap I_2^D(\ell')$ the intersection of the two corresponding sets of nodes. We then have for arbitrary $\ell \neq \ell' \in [K/M]$:

$$\mathbf{E}_1(|I_2^D(\ell) \cap I_2^D(\ell')| \mathbf{1}_\mathcal{E}) = \sum_i \sum_j \mathbf{E}_1 \left(\frac{1}{|\pi(\ell)| \cdot |\pi(\ell')|} \sum_{i_2^D \in \pi_i} \sum_{j_2^D \in \pi_j} |i_2^D \cap j_2^D| \mathbf{1}_\mathcal{E} \right)$$

where the first summations are over $i \in [M(\ell - 1) + 1, M(\ell - 1) + L]$ and $j \in [M(\ell' - 1) + 1, M(\ell' - 1) + L]$. The expectation in the right-hand side does not depend on i and j , by symmetry. Moreover, on \mathcal{E} we can upper bound the fraction in the expectation by $1/(\alpha \mathbf{E}_1 |\pi(\ell)|)^2$. Thus fixing some arbitrary $i \neq j$:

$$\begin{aligned} \mathbf{E}_1(|I_2^D(\ell) \cap I_2^D(\ell')| \mathbf{1}_\mathcal{E}) &\leq \frac{L^2}{(\alpha \mathbf{E}_1 |\pi(\ell)|)^2} \mathbf{E}_1 \left(\sum_{i_2^D \in \pi_i} \sum_{j_2^D \in \pi_j} |i_2^D \cap j_2^D| \right) \\ &\leq \frac{L^2}{(\alpha \mathbf{E}_1 |\pi(\ell)|)^2} \sum_{i_2^D, j_2^D} \mathbf{E}_1 \left(\mathbf{1}_{(k_i, i_2^D, k_{i+D}) \in G} \mathbf{1}_{(k_j, j_2^D, k_{j+D}) \in G} |i_2^D \cap j_2^D| \right), \end{aligned}$$

where summation is over all pairs of lists i_2^D and j_2^D of $\tau(D - 2)$ distinct elements in $[n] \setminus k_1^K$. Denote by J_2^D one such list selected uniformly at random, and by i_2^D a fixed, arbitrary choice of one such list. One then has, recalling the expression of $\mathbf{E}_1 |\pi(\ell)| = L(\lambda/n)^{\tau D} (n - K) \cdots (n - K - \tau(D - 1) + 1)$:

$$\mathbf{E}_1(|I_2^D(\ell) \cap I_2^D(\ell')| \mathbf{1}_\mathcal{E}) \leq \frac{1}{\alpha^2} \mathbf{E}_1 \left(\left(\frac{n}{\lambda} \right)^S |i_2^D \cap J_2^D| \right), \quad (2.26)$$

where S denotes the number of edges in common between the two τ -paths i_2^D and J_2^D .

As in Lemma 4, we now define the Markov chain $\{Z'_t\}_{t \geq 0}$ on the three states $\{-1, 0, 1\}$, with transition probabilities given by the matrix

$$P := \begin{pmatrix} 1 - D/n' & D/n' & 0 \\ 1 - D/n' & (D - 2)/n' & 2/n' \\ 1 - D/n' & (D - 1)/n' & 1/n' \end{pmatrix},$$

where $n' = n - K - D$, and with initial condition $Z'_0 = -1$. These states are interpreted as follows: $Z'_t = -1$ if $J_{t+1} \notin i_2^D$, $Z'_t = 0$ if $J_t \notin i_2^D$ and $J_{t+1} \in i_2^D$, and $Z'_t = 1$ if $J_t, J_{t+1} \in i_2^D$. The same coupling argument as for Lemma 4 implies, letting $x = n/\lambda$, the following, where the subscript in the second expectation specifies the initial state of the Markov chain $\{Z'_t\}$:

$$\mathbf{E}_1 \left(\left(\frac{n}{\lambda} \right)^S |i_2^D \cap J_2^D| \right) \leq \mathbf{E}_{-1} \left(x^{\sum_{i=1}^{\tau(D-1)} Z'_i} \sum_{j=1}^{\tau(D-1)} \mathbf{1}_{Z'_j \geq 0} \right).$$

We introduce the notation $F_z(t) = (F_{z,-1}(t), F_{z,0}(t), F_{z,1}(t))$, where

$$F_{z,y}(t) := \mathbf{E}_z \left(x^{\sum_{s=1}^t Z'_s} \mathbf{1}_{Z'_t=y} \right).$$

It readily follows that

$$F_z(t) = (\mathbf{1}_{z=-1}, \mathbf{1}_{z=0}, \mathbf{1}_{z=1}) M^t,$$

where

$$M := \begin{pmatrix} 1 - D/n' & D/n' & 0 \\ 1 - D/n' & (D-2)/n' & x * (2/n') \\ 1 - D/n' & (D-1)/n' & x/n' \end{pmatrix}.$$

This matrix M reads, as previously, $M_0 + O(D/n)$ where M_0 is given by (2.21).

Write then, using Markov's property:

$$\begin{aligned} \mathbf{E}_{-1} \left(x^{\sum_{i=1}^{\tau(D-1)} Z_i^+} \sum_{j=1}^{\tau(D-1)} \mathbf{1}_{Z_j \geq 0} \right) &= \sum_{j=1}^{\tau(D-1)} \sum_{z \in \{0,1\}} \mathbf{E}_{-1} \left(x^{\sum_{i=1}^j Z_i^+} \mathbf{1}_{Z_j = z} \right) \mathbf{E}_z \left(x^{\sum_{i=1}^{\tau(D-1)-j} Z_i^+} \right) \\ &= \sum_{j=1}^{\tau(D-1)} \sum_{z \in \{0,1\}} F_{-1,z}(j) \sum_{y=-1,0,1} F_{z,y}(\tau(D-1)-j). \end{aligned}$$

Previously given perturbation results give the existence of coefficients $[\beta_{z,r}]_{z \in \{-1,0,1\}, r \in \{0,1/\lambda,1\}}$ all in $O(1)$ such that

$$F_z(0) = \sum_{r \in \{0,1/\lambda,1\}} \beta_{z,r} v_r.$$

It follows that

$$F_z(\tau(D-1)-j) = \sum_{r \in \{0,1/\lambda,1\}} \beta_{z,r} \mu_r^{\tau(D-1)-j} v_r = O(1),$$

since $|\mu_r| \leq 1 + O(D/n)$ and $D^2 \ll n$. It follows that

$$\begin{aligned} \mathbf{E}_{-1} \left(x^{\sum_{i=1}^{\tau(D-1)} Z_i^+} \sum_{j=1}^{\tau(D-1)} \mathbf{1}_{Z_j \geq 0} \right) &= \sum_{j=1}^{\tau(D-1)} F_{-1}(j) \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix} \times O(1) \\ &= \sum_{j=1}^{\tau(D-1)} \sum_{r \in \{0,1/\lambda,1\}} \beta_{-1,r} \mu_r^j v_r \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix} \times O(1). \end{aligned}$$

Since $F_{-1}(0) = u_1$, it holds that $\beta_{-1,1} = 1 + O(D/n)$, and $\beta_{-1,r} = O(D/n)$ for $r = 0, 1/\lambda$. The terms with $r = 0, 1/\lambda$ in the previous expression thus contribute at most $O(D^2/n)$. The terms with $r = 1$ give

$$\sum_{j=1}^{\tau(D-1)} \beta_{-1,r} \mu_1^j v_1 \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix} \times O(1) = O(D^2/n),$$

by using the fact that $v_1 = (1, 0, 0) + O(D/n)$.

It remains to prove that the event \mathcal{E} defined in (2.25) is such that $\mathbb{P}_1(\mathcal{E}) = 1 - o(1)$. It will suffice to prove that for all $\ell \in [K/M]$, $\mathbb{P}_1(\mathcal{E}_\ell) \geq 1 - o(M/K)$. To show this we shall leverage Janson's inequality, as described in [30], p.205, Theorem 6.31. Applied to the random variable $|\pi(\ell)|$, it guarantees that for all $0 \leq t \leq \mathbf{E}|\pi(\ell)|$ one has

$$\mathbb{P}_1(|\pi(\ell)| \leq \mathbf{E}|\pi(\ell)| - t) \leq e^{-t^2/(2\delta)}, \quad (2.27)$$

where δ is the expected number of ordered pairs of τ -paths (P, Q) in $\pi(\ell)$ that share at least an edge. Paralleling our previous bound on the variance of $|\pi(\ell)|$, we distinguish the pairs of τ -paths (P, Q) according to whether they share the same starting point $i \in [(\ell-1)M+1, (\ell-1)M+L]$ or not to write $\delta = \delta_1 + \delta_2$, and obtain:

$$\begin{aligned} \delta_1 &\leq L \frac{\lambda^{2D\tau}}{n^{2\tau}} (1 + o(1)) (\tau!)^2 \left(\frac{\lambda}{\lambda-1} \right)^{2\tau^2}, \\ \delta_2 &\leq L^2 \frac{\lambda^{2D\tau}}{n^{2\tau}} O\left(\frac{D^2}{n}\right). \end{aligned}$$

We moreover have that $\mathbf{E}|\pi(\ell)| \sim L \frac{\lambda^{D\tau}}{n^\tau}$, so that

$$\frac{(\mathbf{E}|\pi(\ell)|)^2}{\delta} \geq \frac{\Omega(1)}{\frac{1}{L} + \frac{D^2}{n}}.$$

By our choices (2.7) for L and D , this lower bound is also $\Omega(1)L = C\Omega(1)\ln(n)$. Taking $t = (1 - \alpha)\mathbf{E}|\pi(\ell)|$ for some $\alpha \in (0, 1)$ in (2.27), we obtain

$$\mathbb{P}_1(|\pi(\ell)| \leq \alpha\mathbf{E}|\pi(\ell)|) \leq \exp(-\alpha^2 C\Omega(1)\ln(n)/2).$$

It readily follows that, for sufficiently large C , this probability can be made $o(n^{-3})$ (say), which suffices to conclude the proof of the Lemma.

2.9 Proofs for planted D -ary trees

2.9.1 Proof of Lemma 7

Proof. The property $p_1 = 1$ is trivial. For $h \geq 1$, let $Z \text{ Poi}(\lambda)$ be the number of children of the root o . Each of the Z children has independently a probability p_h of being the root of a D -ary tree of height h . Therefore, if we define Z_h to be the number of such children, we have

$$\mathcal{L}(Z_h | Z) \sim \text{Bin}(Z, p_h).$$

By the splitting property of Poisson random variables, Z_h follows the distribution $\text{Poi}(\lambda p_h)$. But T contains a D -ary tree of height h rooted in o if and only if $Z_h \geq D$, and the lemma follows. \square

2.9.2 Proof of Theorem 8

Proof. Let $h_0 > 0$ to be fixed later on ; there exists $\kappa > 0$ such that

$$\frac{(\lambda x)^D}{D!} \leq \psi_D(\lambda x) \leq e^{\kappa(D-1)p_{h_0}} \frac{(\lambda x)^D}{D!} \quad (2.28)$$

for all $x \leq \epsilon$. Therefore, for $h \geq h_0$, one has

$$D \ln(p_h) + (D - 1)c_{\lambda, D} \leq \ln(p_{h+1}) \leq D \ln(p_h) + (D - 1)(c_{\lambda, D} + \kappa p_{h_0}). \quad (2.29)$$

Iterating inequality (2.29), we get that for all $h \geq 0$:

$$D^h (\ln(p_{h_0}) + c_{\lambda, D}) - c_{\lambda, D} \leq \ln(p_{h+h_0}) \leq D^h (\ln(p_{h_0}) + c_{\lambda, D} + \kappa p_{h_0}) - c_{\lambda, D} - \kappa p_{h_0} \quad (2.30)$$

Choose h_0 such that $\alpha := -(\ln(p_{h_0}) + c_{\lambda, D} + \kappa p_{h_0}) > 0$, and let

$$h_* = \left\lceil \frac{\ln\left(\frac{\ln(n)}{\alpha}\right)}{\ln(D)} \right\rceil + h_0$$

Then $h_* + 1 = \frac{\ln\left(\frac{\ln(n)}{\alpha}\right)}{\ln(D)} + h_0 + \delta$ for some $\delta > 0$. Thus, using (2.30), we find

$$\ln(p_{h_*+1}) \leq -D^\delta \ln(n) - c_{\lambda, D} - \kappa p_{h_0}$$

which yields that $p_{h_*+1} = o\left(\frac{1}{n}\right)$ as required.

On the other hand, for almost all λ there is a choice of h_0 such that

$$h_* < \frac{\ln\left(\frac{\ln(n)}{\alpha}\right)}{\ln(D)} + h_0 - \ln\left(\frac{\alpha}{\ln(p_{h_0}) + c_{\lambda,D}}\right)$$

by continuity of the right-hand side. Then, for some $\delta' > 0$, we have

$$\ln(p_{h_*}) \geq -D^{-\delta'} \ln(n) - c_{\lambda,D}$$

which implies the second result of Theorem 8. □

2.9.3 Proof of Lemma 8

This lemma is a classical result in sparse random graph theory (see e.g. [29]) ; we reproduce it here for the sake of self-containedness. First, a result on the size of neighbourhoods in G :

Lemma 11 (Lemma 29 in [29]). *For a vertex v in G , let $S_t(v)$ denote the size of the t -neighbourhood of v . Then there exists a constant C such that with high probability, for every vertex $v \in G$ and $t \geq 0$:*

$$S_t(v) \leq C \ln(n) \alpha^t$$

We'll also use a bound on the number of vertices whose neighbourhood contains a cycle ; its proof, as well as the preceding lemma, can be found in [29].

Lemma 12 (Lemma 30 in [29]). *Assume that $\ell = o(\ln(n))$. Then w.h.p there are at most $\ln(n)\lambda^{2\ell}$ vertices whose ℓ -neighbourhood contains a cycle. Moreover, with high probability the graph G is ℓ tangle-free, i.e. no vertex has more than one cycle in its ℓ -neighbourhood.*

We can now prove the first part of our lemma : consider the classical breadth-first exploration process which starts with $A_0 = \{v\}$ and at step $t \geq 0$, considers (if possible) a vertex $v_t \in A_t$ at minimal distance from v and reveals its neighbors N_{t+1} in $[n] \setminus \bigcup_{s \leq t} A_s$. It then updates A_{t+1} as $A_t \cup N_{t+1}$ and repeats the process. We denote by \mathcal{F}_t the filtration generated by A_0, \dots, A_t .

Proof. (First part of Lemma 8). Let τ be the stopping time at which $(G, v)_\ell$ has been revealed. By the two previous lemmas, with probability at least $1 - c\lambda^{2\ell}/n$, the neighbourhood $(G, v)_\ell$ is a tree. Therefore, we can mirror the discovery process in (T, o) , where at each step we discover the children of v_t . To establish the desired coupling result, we then only need to focus on the number of children of each node.

Given \mathcal{F}_t , the number of discovered neighbors y_{t+1} of the node v_t has distribution $\text{Bin}(n_t, \lambda/n)$, where

$$n_t = n - \sum_{s=0}^t y_s$$

Therefore, given \mathcal{F}_t , the total variation distance between the number of children of v_t in $(G, v)_\ell$ and in $(T, o)_\ell$ is

$$\left| \text{Bin}\left(n_t, \frac{\lambda}{n}\right) - \text{Poi}(\lambda) \right|_{\text{var}}$$

The Stein-Chen method (see for example [16]) yields that

$$\left| \text{Bin}\left(n_t, \frac{\lambda}{n}\right) - \text{Poi}\left(\lambda \frac{n_t}{n}\right) \right|_{\text{var}} \leq \frac{\lambda}{n},$$

and a classical bound for Poisson law (see again [16]) that

$$\left| \text{Poi}\left(\lambda \frac{n_t}{n}\right) - \text{Poi}(\lambda) \right|_{\text{var}} \leq \lambda \left(1 - \frac{n_t}{n}\right)$$

From Lemma 11, we find that $n_t \geq n - C \ln(n) \lambda^\ell$ with probability greater than $1 - 1/n$, and thus

$$|P_{t+1} - Q_{t+1}|_{\text{var}} \leq \frac{\lambda}{n} + \lambda \frac{C \ln(n) \lambda^\ell}{n},$$

where P_{t+1} is the distribution of y_{t+1} given \mathcal{F}_t and Q_{t+1} is a $\text{Poi}(\lambda)$ random variable independent of \mathcal{F}_t . This finishes the proof of the first part of the lemma. \square

For the second part, note that there exists a coupling (X, X') such that $X \sim \text{Poi}(\lambda)$, $X' \sim \text{Poi}(\lambda')$ and $X' > X$ a.s. (take for example $X' = X + Z$ where $Z \sim \text{Poi}(\lambda' - \lambda)$).

The proof is then straightforward : for every vertex v , we produce a coupling between the exploration process of $(G, v)_\ell$ and $(T', o')_\ell$ such that at each step t , the number of neighbors y_t of v_t in G is less than in T' .

2.9.4 Proof of Theorem 9

Proof. We first apply the first part of Lemma 8 to $\ell = \underline{h} = O(\ln \ln(n))$. Then, for at least $n - O(\ln(n)^\alpha)$ vertices v (for some $\alpha > 0$), there is a coupling between $(T, o)_{\underline{h}}$ and $(G, v)_{\underline{h}}$. Since in $(T, o)_{\underline{h}}$, there is a copy of Γ in $(T, o)_{\underline{h}}$ with probability $\Omega(n^{-c})$. It follows that w.h.p there is $\omega(1)$ copies of Γ in G .

Now, assume that $h = \bar{h} + C$, where C is large enough such that for some $\lambda' > \lambda$, there are no trees of height h in (T', o') with probability $1 - o(1/n)$.

For every $v \in G$ such that the h -neighbourhood of v is a tree, we can produce a coupling of $(G, v)_h$ and $(T', o')_h$ such that $(G, v)_h \subseteq (T', o')_h$ with probability 1. Thus, with high probability, no vertex whose h -neighbourhood is a tree contains a copy of Γ in said neighbourhood.

Assume now that there is one cycle in the h -neighbourhood of v . With high probability, there is only one cycle going through v in the neighbourhood. Thus, there are only two vertices in the neighbors of v whose offspring contains a cycle. With probability $1 - O(n^{-c})$, no other neighbour of v is the root of a D -ary tree of height $h - 1$. If $D > 2$, then there is no copy of Γ rooted in v ; if $D = 2$, then both neighbors of v in the cycle must be roots of disjoint binary trees of size $h - 1$, in which case the cycle edge does not help.

To summarize, the probability of presence of a copy of Γ rooted at v is upper bounded by $o(1/n)$ if the h -neighbourhood of v is cycle-free, and by $O(n^{-c})$ if it is not. Since there are $O(\ln(n)^\alpha)$ such vertices, w.h.p there is no copy of Γ in G . \square

2.9.5 Proof of Lemma 9

Proof. In view of Lemma 1, we aim to bound the ratio

$$\mathbf{E}_0(L^2) = \frac{\mathbf{E}_0(X_\Gamma^2)}{\mathbf{E}_0(X_\Gamma)^2}$$

. As before, let $\Gamma_1, \dots, \Gamma_m$ be the copies of Γ in the complete graph K_n , and let $X_i = \mathbf{1}_{\Gamma_i \in G}$.

We follow the proof sketch from [27] : write

$$\mathbf{E}_0(X_\Gamma^2) = \sum_{i,j} \mathbf{E}_0(X_i X_j) = \mathbf{E}' + \mathbf{E}'', \quad (2.31)$$

where \mathbf{E}' is the sum over Γ_i, Γ_j having disjoint vertex sets.

We can easily compute \mathbf{E}' :

$$\mathbf{E}' = \binom{n}{K} \binom{n-K}{K} \left(\frac{K!}{|\text{Aut}(\Gamma)|} \right)^2 p^{2K-2} \sim \frac{n^{2K} p^{2K-2}}{|\text{Aut}(\Gamma)|^2} \sim \mathbf{E}_0(X_\Gamma)^2$$

We therefore need to show that $\mathbf{E}'' = o(\mathbf{E}_0(X_\Gamma)^2)$; to this end, note that if Γ_i and Γ_j are such that $v(\Gamma_i \cup \Gamma_j) = s$, then $e(\Gamma_i \cap \Gamma_j) \leq 2K - s - 1$ (since $\Gamma_i \cap \Gamma_j$ is a forest of size $2K - s$) and therefore $e(\Gamma_i \cup \Gamma_j) \geq s - 1$.

Grouping the terms of \mathbf{E}'' by the size of $\Gamma_i \cup \Gamma_j$, we get

$$\begin{aligned} \mathbf{E}'' &\leq \sum_{s=K}^{2K-1} \binom{n}{s} \binom{s}{s-K, s-K, 2K-s} \left(\frac{K!}{|\text{Aut}(\Gamma)|} \right)^2 \left(\frac{\lambda}{n} \right)^{s-1} \\ &= \frac{n}{\lambda |\text{Aut}(\Gamma)|^2} \sum_{s=K}^{2K-1} \frac{n^s \lambda^s}{n^s} \frac{K!^2}{(s-K)!^2 (2K-s)!} \\ &= \frac{n}{\lambda |\text{Aut}(\Gamma)|^2} \sum_{s=K}^{2K-1} \lambda^s \frac{K!^2}{(s-K)!^2 (2K-s)!} \left(1 + O\left(\frac{K^2}{n}\right) \right) \\ &\leq \frac{n \lambda^{K-1} (1 + o(1))}{|\text{Aut}(\Gamma)|^2} \sum_{u=0}^{K-1} \lambda^u \frac{K!^2}{u!^2 (K-u)!}, \end{aligned}$$

where we made the change of variables $u = s - K$. Now, write

$$\frac{K!^2}{u!^2 (K-u)!} = \binom{K}{u} \frac{K!}{u!} \leq \binom{K}{u} K^{K-u},$$

and we get

$$\begin{aligned} \mathbf{E}'' &\leq \frac{n \lambda^{K-1} K^K}{|\text{Aut}(\Gamma)|^2} (1 + o(1)) \sum_{u=0}^{K-1} \binom{K}{u} \left(\frac{\lambda}{K} \right)^u \\ &\leq \frac{n \lambda^{K-1} K^K}{|\text{Aut}(\Gamma)|^2} (1 + o(1)) \left(1 + \frac{\lambda}{K} \right)^K \\ &\leq \frac{n \lambda^{K-1} K^K e^\lambda}{|\text{Aut}(\Gamma)|^2} (1 + o(1)) \\ &= O\left(\mathbf{E}_0(X_\Gamma)^2 \times \frac{K^K}{n \lambda^K} \right) \end{aligned}$$

When $K \leq \frac{\ln(n)}{\ln \ln(n)}$, we find that $\mathbf{E}'' = o(\mathbf{E}_0[X_\Gamma]^2)$, as requested. But $K = \frac{D^{h+1}-1}{D-1} \leq \frac{\ln(n)}{\ln \ln(n)}$ whenever

$$h \leq \underline{h} - \frac{\ln(\underline{h})}{\ln(D)} + \frac{\ln\left(1 - \frac{1}{D}\right)}{\ln(D)},$$

which is the condition mentioned in Theorem 10. □

2.9.6 Proof of Theorem 11

Proof. For $0 \leq p \leq h$, let L_p be the set of vertices at depth p of Γ , and T_p the set of vertices at depth $\leq p$.

The strategy of proof is as follows : we aim to prove that there exists a universal constant δ such that given G and

$$\mathcal{T} := \sigma(T_{h-1}) \subset G,$$

the location of the first $h - 1$ rows of Γ , we have with high probability on G

$$\mathbb{P}_1 \left((\text{ov}(\hat{\mathcal{K}}) \leq (1 - \delta)K \mid G, \mathcal{T}) = 1 - o(1) \right) \quad (2.32)$$

In what follows, we will consider \mathcal{T} to be fixed, and G drawn under \mathbb{P}_1 .

Let $\varepsilon > 0$ to be adapted later, and consider two cases :

- $|\hat{\mathcal{K}} \cap \mathcal{T}| \leq (1 - \varepsilon)|\mathcal{T}|$: in this case, we easily get

$$\begin{aligned} \text{ov}(\hat{\mathcal{K}}) &\leq D^h + (1 - \varepsilon) \frac{D^h - 1}{D - 1} \\ &= K - \varepsilon \frac{D^h - 1}{D - 1} \\ &= K - \varepsilon \frac{K - 1}{D} \\ &= \left(1 - \frac{\varepsilon}{D}\right)K + o(K), \end{aligned}$$

from which equation (2.32) follows since ε is independent from G and \mathcal{T} .

- if $|\hat{\mathcal{K}} \cap \mathcal{T}| > (1 - \varepsilon)|\mathcal{T}|$, we need the following lemma :

Lemma 13. *Let $\sigma(L_{h-1}) = \{i_1, \dots, i_{D^{h-1}}\}$, and define $n_k = |\mathcal{N}(i_k)|$ and $m_k = |\hat{\mathcal{K}} \cap \mathcal{N}(i_k)|$. Then*

$$\mathbf{E}_1 \left(|\hat{\mathcal{K}} \cap \sigma(L_h)| \mid G, \mathcal{T} \right) = D \sum_k \frac{m_k}{n_k}$$

Proof. (of lemma 13). Given \mathcal{T} , all vertices that are neighbours of a vertex in $\sigma(L_{h-1})$ are equally likely to belong to Γ , since all D -ary trees in G have the same probability of generating G .

Therefore, given G , $\sigma(L_{h-1}) = \{i_1, \dots, i_{D^{h-1}}\}$, the random variable $N_k = |\hat{\mathcal{K}} \cap \mathcal{N}(i_k)|$ follows a hypergeometric law of parameters (n_k, D, m_k) . It follows that

$$\mathbf{E}_1(N_k) = D \frac{m_k}{n_k}$$

Now, with high probability the neighbourhoods $\mathcal{N}(i_k)$ are disjoint and the variables N_k are thus independent. Since $|\hat{\mathcal{K}} \cap \sigma(L_h)| = \sum_k N_k$ whenever the $\mathcal{N}(i_k)$ are disjoint, the lemma follows. \square

We can now prove our main theorem : notice that $|\mathcal{N}(i_k)| \sim D + \text{Poi}(\lambda)$ since $K = o(n)$, so w.h.p a proportion α (for a universal constant α) of the i_k are such that $|\mathcal{N}(i_k)| \geq D + 1$. Moreover,

$$S := \sum_k m_k = K - |\hat{\mathcal{K}} \cap \mathcal{T}| < D^h + \varepsilon|\mathcal{T}| = \left(1 + \frac{\varepsilon}{D-1}\right)D^h + o(D^h)$$

Thus, $S \leq (1 + \varepsilon')D^h$ for some $\varepsilon' > 0$.

Let I_1 be the set of indices such that $n_k = D$; we have

$$\begin{aligned} \sum_k \frac{m_k}{n_k} &= \sum_{k \in I_1} \frac{m_k}{n_k} + \sum_{k \notin I_1} \frac{m_k}{n_k} \\ &\leq \sum_{k \in I_1} \frac{m_k}{D} + \sum_{k \notin I_1} \frac{m_k}{D+1} \end{aligned}$$

Let $S_1 = \sum_{k \in I_1} m_k$; we know that

$$S_1 \leq D|I_1| \leq D(1 - \alpha)D^{h-1},$$

since $m_k \leq n_k = D$ on I_1 , which yields

$$\begin{aligned} \sum_k \frac{m_k}{n_k} &\leq \frac{S_1}{D} + \frac{S - S_1}{D+1} \\ &= \frac{S}{D+1} + \frac{S_1}{D(D+1)} \\ &\leq D^{h-1} \left((1 + \varepsilon) \frac{D}{D+1} + (1 - \alpha) \frac{1}{D+1} \right) \\ &\leq D^{h-1} \left(1 - \frac{\alpha - D\varepsilon}{D+1} \right) \end{aligned}$$

Choosing ε such that $\alpha - D\varepsilon > 0$, we eventually find

$$\mathbf{E}_1 \left(|\hat{\mathcal{K}} \cap L_h| \mid G, \mathcal{T} \right) \leq (1 - \gamma)D^h \quad (2.33)$$

for some $\gamma > 0$.

Finally, we can bound $\hat{\mathcal{K}} \cap \mathcal{K}$:

$$\begin{aligned} \mathbf{E}_1 \left(|\hat{\mathcal{K}} \cap \mathcal{K}| \mid G, \mathcal{T} \right) &\leq |\mathcal{T}| + \mathbf{E}_1 \left(|\hat{\mathcal{K}} \cap L_h| \mid G, \mathcal{T} \right) \\ &\leq (1 - \gamma)D^h + |\mathcal{T}| \\ &\leq K - \gamma D^h + o(D^h) \\ &\leq \left(1 - \gamma \frac{D-1}{D} \right) K + o(K), \end{aligned}$$

which completes the proof of Theorem 11. □

Chapter 3

Robustness of community detection in sparse networks

This chapter is based on the paper [117], a joint work with L. Massoulié.

3.1 Introduction

3.1.1 Background

Community detection is the task of finding large groups of similar items inside a large relationship graph, where it is expected that related items are (in the assortative case) more likely to be linked together. The Stochastic Block Model (abbreviated in SBM) has been designed by [69] to analyze the performance of algorithms for this task; it consists in a random graph G whose edge probabilities depend only on the community membership of their endpoints. Since then, a large number of articles have been devoted to the study of this model; a survey of these results can be found in [1], or in [57] for a more general view on community detection.

The sparse case, when edge probabilities are in $O(1/n)$, is known to be much harder to study than denser models; the existence of a positive portion of isolated vertices makes complete reconstruction impossible, and studies usually focus on partial recovery of the community structure. Insights on this topic often stem from statistical physics; in the two-community case, [47] conjectured the existence of a threshold for reconstruction, as well as its exact value; this conjecture was then proved in [100] for the first part, [91] and [102] for the converse part. Similarly, in the general case, a method was first presented in [80] and then proven to work in [29] – bar a technical condition – and [6].

The main issue in the sparse setting is that the usual method relying on the eigenvectors of the adjacency matrix of G fails due to the lack of separation of the eigenvalues. Consequently, a wide array of alternative spectral methods have been designed, relying on the spectrum of a matrix associated to G . More precisely, the eigenvectors associated to the highest eigenvalues of those matrices will often carry some information about the community structure of G , enough for partial reconstruction. Examples include the path expansion matrix used in [91], or the non-backtracking matrix in [80].

Additionally, other types of methods can be used in this setting: for example, the semi-definite programming (or SDP) algorithm relaxes the problem into a convex optimization one, which can then be approximately solved (see for example [97]).

An important feature of real-life networks that is missing from the SBM is the existence of small-scale regions of higher density, that arise from phenomena unrelated to the community structure. For this reason, a common variant of the SBM is the addition of small cliques to

the generated random graph. Commonly-used spectral methods, for example those relying on the non-backtracking matrix in [29], are known to fail in this setting, due to the apparition of localized eigenvectors, with no ties to the community structure, and corresponding to large eigenvalues – see [127] for a comparison of those methods, as well as a proposed heuristic to deal with those localized vectors by lowering their associated eigenvalues. SDP methods are the most studied for this problem, due to their natural stability; in particular, [88] show a reconstruction algorithm that is robust to the adversarial addition of $o(n)$ edges, in the case of an arbitrary number of communities; this was also shown independently by [96]. However, all the SDP methods mentioned here fail to reach the KS threshold by at least a large constant, with only [97] approaching it as the average degree increases.

After completing this work we became aware of the article of [3]. It establishes results akin to ours on robustness (although with a different definition thereof) of spectral methods for detection in SBM. We use however a slightly different matrix, and our results apply to an arbitrary number of blocks, whereas they only consider SBM with two blocks.

3.1.2 Summary of main results

This article focuses on the Stochastic Block Model, as defined in [69]; we recall here a succinct definition:

Definition 2. *Let $r \in \mathbb{N}$ be fixed, W be a $r \times r$ symmetric matrix with nonnegative entries, and π a probability vector on $[r]$. A random graph $G = (V, E)$ with $|V| = n$ is said to be distributed according to the Stochastic Block Model (or SBM) with r blocks and parameters (W, π) if:*

- (i) *each vertex $v \in V$ is assigned a type $\sigma(v)$ sampled independently from π ,*
- (ii) *any two vertices u, v in V are joined with an edge randomly and independently from every other edge, with probability*

$$\min\left(\frac{W_{\sigma(u), \sigma(v)}}{n}, 1\right).$$

Given a random graph G sampled according to the above model (with the types of each vertex hidden), the aim is to estimate the type assignment σ from the observation of G only. However, since there is a positive proportion of isolated vertices, perfect reconstruction is theoretically impossible; we will thus only focus on retrieving only a positive proportion of the types. Our metric for reconstruction will be the following:

Definition 3. *Let σ be the true type estimation, and $\hat{\sigma}$ a type estimate of σ ; the empirical overlap between σ and $\hat{\sigma}$ is defined as:*

$$\text{ov}(\sigma, \hat{\sigma}) = \max_{\tau \in \mathfrak{S}_r} \frac{1}{n} \sum_{v=1}^n \mathbf{1}_{\hat{\sigma}(v)=\tau(\sigma(v))} - \max_{k \in [r]} \pi_k, \quad (3.1)$$

where \mathfrak{S}_r is the set of permutations of $[r]$. For a given algorithm leading to estimates $\hat{\sigma}$ for all n , we say that this algorithm achieves partial reconstruction if

$$\liminf_{n \rightarrow \infty} \text{ov}(\sigma, \hat{\sigma}) > 0 \quad \text{w.h.p.} \quad (3.2)$$

Spectral methods in denser settings (with average degrees of about $\log(n)$) usually consist in clustering the eigenvectors of the adjacency matrix of G ; however those methods are known to fail in sparser graphs (see [1]). As a result, different (and more complex) matrices are needed:

Definition 4. Let G be any graph, and ℓ be a positive integer. We define two matrices associated with G :

- (i) the path expansion matrix $B^{(\ell)}$ (studied in [91]), whose (i, j) coefficient counts the number of self-avoiding paths (that is, paths that do not go through the same vertex twice) of length ℓ between i and j ,
- (ii) the distance matrix $D^{(\ell)}$, defined by $D_{ij}^{(\ell)} = 1$ if $d(i, j) = \ell$ and 0 otherwise, where d denotes the usual graph distance.

We are now ready to state our first result:

Theorem 12. Assume that $\pi \equiv 1/r$, W is a stochastic positive regular matrix, and that the two highest eigenvalues μ_1, μ_2 of W satisfy the condition:

$$\mu_2^2 > \mu_1.$$

Then there exists an algorithm, based only on an eigenvector of $B^{(\ell)}$ associated with its second highest eigenvalue, that achieves partial reconstruction whenever $\ell \sim \delta \log(n)$ for small enough δ .

The same algorithm also achieves partial reconstruction when applied to $D^{(\ell)}$ instead of $B^{(\ell)}$, with the same conditions on ℓ .

Regardless of the change of matrices, this is already an improvement on [29]; indeed, we managed to remove a technical asymmetry condition on W (namely, the existence of a simple eigenvector associated with a high eigenvalue).

We now move on to study the stability of our algorithm; as opposed to most papers that classify the difficulty of an adversary according to the number of altered edges, ours considers the number of affected vertices.

Definition 5. Let $\gamma := \gamma(n)$ be a positive integer, and G any graph on n vertices. The adversary of strength γ is allowed to arbitrarily add and remove edges at will, as long as the number of vertices affected (i.e. vertices that are endpoints of altered edges) is at most γ .

Our main result on stability is then the following:

Theorem 13. Under the same assumptions as Theorem 12, let G a graph generated via SBM, and \tilde{G}_γ a graph obtained when perturbed by an adversary of strength γ .

Then, assuming

$$\gamma = o\left(\frac{(\mu_2^2/\mu_1)^{\ell/2}}{\log(n)}\right),$$

the algorithm of Theorem 12 still achieves partial reconstruction on \tilde{G}_γ .

The above result on γ is the best possible, up to a factor of $\log(n)$.

Compared to the spectral algorithm in [29], this is a substantial improvement: their algorithm is known (see e.g. [127]) to be highly unstable w.r.t edge addition. In contrast, the above result reaches a perturbation of size a small power of n (since ℓ is of order $\log(n)$). This is sharp, and thus still far from the $o(n)$ bound achieved by various SDP methods (notably [88, 97]); this discrepancy is likely a result of delicate graph properties involved in spectral algorithms that make them more sensitive to perturbations.

However, our result still has several advantages compared to the other cited methods, namely:

- (i) the threshold for partial reconstruction in our method is exactly the KS threshold, whereas SDP-based methods require a slightly stronger condition, especially when the mean degree of G is low.
- (ii) as will be proved later, the running time of our algorithm is at most $O(n^{13/12})$, which is much faster than the usual methods for SDP algorithms.
- (iii) Finally, all the SDP methods mentioned throughout this paper only consider the symmetric case even in the case of multiple communities.

3.1.3 Detailed setting and results

Stochastic block model

Consider the SBM as defined above; following [29], we introduce $\Pi = \text{diag}(\pi_1, \dots, \pi_r)$ and define the *mean progeny matrix* $M = \Pi W$; the eigenvalues of M are the same as those of the symmetric matrix $S = \Pi^{1/2} W \Pi^{1/2}$ and in particular are real. We denote them by

$$\mu_1 \geq |\mu_2| \geq \dots \geq |\mu_r|.$$

We shall make the following regularity assumptions: first,

$$\mu_1 > 1 \quad \text{and } M \text{ is positive regular,} \quad (3.3)$$

i.e. the coefficients of M^t are all positive for some t . Secondly, each type of vertex has the same asymptotic average degree, that is

$$\sum_{i=1}^r M_{ij} = \sum_{i=1}^r \pi_i W_{ij} = \alpha \quad \text{for all } j \in [r]. \quad (3.4)$$

In this case, the matrix $M^* = M/\alpha$ is a stochastic matrix and therefore

$$\mu_1 = \alpha > 1. \quad (3.5)$$

Since $M = \Pi^{-1/2} S \Pi^{1/2}$, M is diagonalizable; let (ϕ_1, \dots, ϕ_r) be a basis of normed left eigenvectors for M , that is

$$\phi_i^\top M = \mu_i \phi_i^\top \quad \text{for all } i \in [r]. \quad (3.6)$$

Condition (3.4) implies that $\phi_1 = \mathbf{1}/\sqrt{r}$, where $\mathbf{1}$ is the all-ones vector.

It has been shown in [29] and [6] that polynomial-time algorithms achieve partial reconstruction when the following condition, called the Kesten-Stigum threshold, is verified:

$$\tau := \mu_2^2 / \mu_1 > 1. \quad (3.7)$$

The quantity τ is commonly referred to as the *signal-to-noise ratio*.

Alternatively, we define r_0 such that

$$\mu_{r_0+1}^2 \leq \mu_1 < \mu_{r_0}^2. \quad (3.8)$$

Therefore, the condition (3.7) is equivalent to $r_0 > 1$.

In the two-community case, the above condition is equivalent to the possibility of reconstruction (see [91], [100]). However, in the general setting ($r > 4$), non-polynomial algorithms can achieve partial reconstruction even below this threshold. This was originally conjectured in [47], and more recently proven in [6].

Path expansion matrix

In [91], an algorithm for partial reconstruction in the two-community case makes use of the path expansion matrix $B^{(\ell)}$. Our first aim is to extend the result from this paper to the general case; we first define for all $k \in [r]$ the vectors χ_k and φ_k by

$$\chi_k(v) = \phi_k(\sigma(v)) \quad \text{and} \quad \varphi_k = \frac{B^{(\ell)}\chi_k}{\|B^{(\ell)}\chi_k\|}. \quad (3.9)$$

Let $\lambda_1(B^{(\ell)}) \geq |\lambda_2(B^{(\ell)})| \geq |\lambda_n(B^{(\ell)})|$ be the eigenvalues of $B^{(\ell)}$ ordered by absolute value; our first theorem is an extension of Theorem 2.1 in [91]:

Theorem 14. *Consider a graph G generated as above, and let $\ell \sim \kappa \log_\alpha(n)$, with $\kappa < 1/12$. Then, with probability going to 1 as n goes to $+\infty$:*

- (i) $\lambda_k(B^{(\ell)}) = \Theta(\mu_k^\ell)$ for $k \in [r_0]$,
- (ii) For $k > r_0$, $\lambda_k(B^{(\ell)}) = O(\log(n)^c \alpha^{\ell/2})$ for some constant $c > 0$.

Furthermore, consider μ such that $\mu^2 > \alpha$ and μ is an eigenvalue of multiplicity d of M . Let $\phi^{(1)}, \dots, \phi^{(d)}$ be an orthonormal basis of eigenvectors of M associated to μ , and $\varphi^{(1)}, \dots, \varphi^{(d)}$ the vectors defined as in (3.9).

There exist orthogonal vectors $\xi^{(1)}, \dots, \xi^{(d)}$ in \mathbb{R}^n such that the following holds:

- (i) for all i , $\xi^{(i)}$ is an eigenvector of $B^{(\ell)}$, with associated eigenvalue $\Theta(\mu^\ell)$
- (ii) there exists an orthogonal matrix $Q \in O(d)$ such that

$$\|\varphi Q - \xi\|_2 = O\left(\alpha^{\ell/2} \mu^{-\ell}\right),$$

where φ (resp. ξ) is the $n \times d$ matrix whose columns are the $\varphi^{(i)}$ (resp. the $\xi^{(i)}$).

The above theorem does not yield immediately an algorithm for community reconstruction; however, adapting the one found in [29], we designed the following:

Algorithm 15. *Let ξ be an eigenvector of $B^{(\ell)}$ associated to the eigenvalue $\lambda_2(B^{(\ell)})$, normalized such that $\|\xi\|^2 = n$, and K an arbitrary large constant. First, partition V in two sets (I^+, I^-) via the following procedure: put v in I^+ with probability*

$$\mathbf{P}(v \in I^+ | \xi) = \frac{1}{2} + \frac{1}{2K} \xi(v) \mathbf{1}_{|\xi(v)| \leq K}$$

Then, assign the label 1 to every vertex in I^+ and label 2 to every vertex in I^- .

We then have the following theorem:

Theorem 16. *Assume that $\pi \equiv 1/r$, and that $r_0 > 1$, i.e. that we are above the Kesten-Stigum threshold. Then Algorithm 15 yields an asymptotically positive overlap when $n \rightarrow \infty$ for some choice of K .*

Note that we don't need the asymmetry condition from [29] anymore; our algorithm can deal with multiple eigenvalues as well. Additionally, an explicit value for K is derived in the appendix, which makes our algorithm easy to implement and eliminates the need for "magic" constants, such as the ones in [127] or [6].

A crucial feature of this algorithm is that it depends only on the second eigenvalue of $B^{(\ell)}$; for any perturbation that leaves the r_0 highest eigenvalues – or even the second highest – unchanged, the result from Theorem 16 will hold.

The distance matrix

We introduce now the *distance matrix* $D^{(\ell)}$, defined by $D_{ij}^{(\ell)} = 1$ if and only if $d(i, j) = \ell$, where d is the distance in G . This matrix, while sparser than $B^{(\ell)}$, retains much of the desired spectral properties. In particular, we have the following theorem:

Theorem 17. *Assume that condition (3.7) holds, and set ℓ such that $\ell \sim \kappa \log_\alpha(n)$, where κ is a constant such that $\kappa < 1/12$. Then the results of Theorem 14 hold with the matrix $B^{(\ell)}$ replaced by $D^{(\ell)}$.*

As a result, Algorithm 15 will still succeed when applied to the matrix $D^{(\ell)}$.

Graph perturbation

As mentioned in the introduction, community detection algorithms have to be resilient to the presence of small cliques (or denser subgraphs) to be useful in practice, since this kind of pattern is often present in real-life networks. We chose to focus here on adversarial perturbations, as defined in the summary, whereas other papers (mainly [3]) focus instead on other random graph models, more prone to small loops and cliques.

As shown in [127], the usual spectral methods do not fare well against adversarial (or even random) perturbation, especially when the added subgraph contains several cliques. This is especially the case for the non-backtracking matrix in [29], but also the path expansion matrix in [91].

However, the distance matrix is more stable to clique addition, since it does not count the number of paths between two vertices – which is affected significantly by small perturbations. We can therefore allow a perturbation of size up to a small power of n , as stated in the following theorem:

Theorem 18. *Let G be an SBM as above, with $\pi_i \equiv 1/r$. Assume that $r_0 > 1$, and recall that $\tau = \mu_2^2/\mu_1 > 1$ is the signal-to-noise ratio.*

Then, if $\gamma = o(\tau^\ell/\log(n))$, then Algorithm 15 based on the distance matrix recovers the original communities with asymptotically positive overlap, even after a perturbation affecting at most γ vertices.

The controls in the above theorem can be shown to be sharp, up to a factor of $\log(n)$:

Theorem 19. *With the same assumptions as above, let $D^{(\ell)}$ be the distance matrix of G and $\tilde{D}^{(\ell)}$ the one of the graph after the adversarial perturbation.*

If $\gamma = \Omega(\tau^\ell)$, then there exists a perturbation of size at most γ such that $\tilde{D}^{(\ell)}$ has an eigenvalue of size $\Omega(\mu_2^\ell)$, with associated eigenvector asymptotically perpendicular to the first r_0 ones of $D^{(\ell)}$.

Therefore, we cannot guarantee the stability of the eigenvectors of $D^{(\ell)}$ when the perturbation affects too many vertices. This means that the best bound we can get on the size of allowed perturbations of the matrix $D^{(\ell)}$ is τ^ℓ , which we can rewrite as

$$\tau^\ell = n^{\kappa \log_\alpha(\tau)}.$$

The spectral method on the distance matrix is thus robust to perturbations of size at most n^ε , with $\varepsilon = \kappa \log_\alpha(\tau)$ going to zero as we approach the KS threshold.

3.1.4 Notations and outline of the paper

Throughout this paper, we will make use of the following notation: for two functions f, g , we say that $f = \tilde{O}(g)$ if there exists a constant c such that $f = O(\log(n)^c \cdot g)$. We similarly define the notations $\tilde{\Theta}$ and $\tilde{\Omega}$.

The next Section is devoted to the study of the spectral structure of $B^{(\ell)}$; we also state there an important theorem on spectral perturbation that will be useful for the study of matrix $D^{(\ell)}$ as well. In Section 3, we study the distance matrix $D^{(\ell)}$ and introduce a method to deal with perturbations of this matrix. We then leverage this method to obtain bounds on the size of allowed perturbations.

3.2 Spectral structure of $B^{(\ell)}$

3.2.1 A theorem on eigenspace perturbation

In the following, we'll need a way to link the operator norm of a matrix perturbation to the consequent perturbation of its eigenvectors. This is provided by the following variant of the Davis-Kahan $\sin \theta$ theorem ([126], Theorem 2):

Theorem 20. *Let $\Sigma, \hat{\Sigma}$ be symmetric $n \times n$ matrices, with eigenvalues $\lambda_1 \geq \dots \geq \lambda_n$ and $\hat{\lambda}_1 \geq \dots \geq \hat{\lambda}_n$ respectively. Fix $1 \leq r \leq s \leq n$ and assume that $\min(\lambda_{r-1} - \lambda_r, \lambda_s - \lambda_{s+1}) > 0$, where we define $\lambda_0 = +\infty$ and $\lambda_{n+1} = -\infty$.*

Let $d = s - r + 1$, and let $V = (v_r, \dots, v_s)$ and $\hat{V} = (\hat{v}_r, \dots, \hat{v}_s)$ have orthonormal columns satisfying $\Sigma v_j = \lambda_j v_j$ and $\hat{\Sigma} \hat{v}_j = \hat{\lambda}_j \hat{v}_j$ for $j \in \{r, \dots, s\}$.

Then there exists an orthogonal matrix $Q \in O(d)$ such that

$$\|VQ - \hat{V}\|_F \leq \frac{2\sqrt{2d}\|\hat{\Sigma} - \Sigma\|_{\text{op}}}{\min(\lambda_{r-1} - \lambda_r, \lambda_s - \lambda_{s+1})}. \quad (3.10)$$

3.2.2 Strategy of proof

We present here the main ideas of the proof, and defer its full version to the appendix. The first step is an adaptation of Proposition 19 from [29]:

Proposition 2. *Let $\ell \sim \kappa \log_\alpha(n)$ with $\kappa < 1/12$. Define, for $k \in [r]$,*

$$\theta_k = \|B^{(\ell)}\varphi_k\| \text{ and } \zeta_k = \frac{B^{(\ell)}\varphi_k}{\theta_k}, \quad (3.11)$$

with φ_k as in (3.9).

Then, with high probability, we have the following estimations for every $\gamma < 1/2$:

- (i) $\theta_k = \Theta(\mu_k^\ell)$ for $k \in [r_0]$,
- (ii) $|\langle \varphi_j, \varphi_k \rangle| = \tilde{O}(\alpha^{3\ell/2} n^{-\gamma/2})$ for $j \neq k \in [r_0]$,
- (iii) $|\langle \zeta_j, \varphi_k \rangle| = \tilde{O}(\alpha^{2\ell} n^{-\gamma/2})$ for $j \neq k \in [r_0]$.

Now, let (z_1, \dots, z_{r_0}) be the Gram-Schmidt orthonormalization of $(\varphi_1, \dots, \varphi_{r_0})$, and define

$$D = \sum_{k=1}^{r_0} \theta_k z_k z_k^\top.$$

The non-zero eigenvalues of D are thus the θ_k , with corresponding eigenvectors z_k . Then, using the asymptotic orthogonality properties of Proposition 2, we prove the following:

Proposition 3. For all $k \in [r_0]$, z_k is asymptotically parallel to φ_k .

Furthermore,

$$\|B^{(\ell)} - D\|_{\text{op}} = \tilde{O}(\alpha^{\ell/2}). \quad (3.12)$$

Theorem 14 then results from a simple application of the Weyl inequality ([125]) and Theorem 20:

Proof. (of Theorem 14): let $\theta_k = 0$ for $k > r_0$; the eigenvalues of D are then exactly the θ_i for $i \leq n$.

By Weyl's inequality, we have for all $i \in [n]$

$$|\lambda_i(B^{(\ell)}) - \theta_i| = \tilde{O}(\alpha^{\ell/2}).$$

Since $\theta_k = \Theta(\mu_k^\ell)$ for $k \in [r_0]$, this implies the statements (i) and (ii) of the Theorem.

We now define $z^{(1)}, \dots, z^{(d)}$ as the z_i associated to $\varphi^{(1)}, \dots, \varphi^{(d)}$, and \mathbf{z} as in Theorem 14. Applying inequality (3.10) to $B^{(\ell)}$ and D yields the existence of an orthogonal matrix $Q \in O(d)$ such that

$$\|\mathbf{z}Q - \boldsymbol{\xi}\| = O\left(\alpha^{\ell/2}\mu^{-\ell}\right),$$

and the proof of Proposition 3 shows that $\|z^{(i)} - \varphi^{(i)}\| = O\left(\alpha^{\ell/2}\mu^{-\ell}\right)$ for all i . Using the triangular inequality (and the fact that Q preserves the norm) completes the proof of Theorem 14. \square

3.2.3 A new reconstruction algorithm

We now sketch the proof for Theorem 16; it hinges on one key lemma, whose proof (adapted from [29]) is in the appendix:

Lemma 14. Let ξ be as in Algorithm 15. For all $i \in [r]$, there exists a random variable X_i such that for every $K > 0$ that is a continuity point of X_i , in probability,

$$\frac{1}{n} \sum_{v \in V} \mathbf{1}_{\sigma(v)=i} \xi(v) \mathbf{1}_{|\xi(v)| \leq K} \rightarrow \pi(i) \mathbf{E}[X_i \mathbf{1}_{|X_i| \leq K}],$$

where the convergence is independent from ξ .

Furthermore, we have

$$\sum_{i \in [r]} \mathbf{E}[X_i] = 0 \quad \text{and} \quad \sum_{i \in [r]} \mathbf{E}[X_i]^2 > c \quad (3.13)$$

for some absolute constant $c > 0$, and for all $\varepsilon > 0$ there exists a choice of M (independent from the chosen eigenvector ξ) such that

$$|\mathbf{E}[X_i \mathbf{1}_{|X_i| \leq K}] - \mathbf{E}[X_i]| < \varepsilon \quad (3.14)$$

In particular, Lemma 14 implies that there exist i and j such that $|\mathbf{E}[X_i] - \mathbf{E}[X_j]| > \sqrt{c}$.

We then use a concentration bound to show that for all i , in probability,

$$\frac{1}{n} \sum_{v \in V} \mathbf{1}_{\sigma(v)=i} \mathbf{1}_{v \in \mathbb{I}^+} \rightarrow \pi(i) \left(\frac{\mathbf{E}[X_i \mathbf{1}_{|X_i| \leq K}] + 1/2}{2K} + 1/2 \right) := \pi(i) \tilde{p}_i \quad (3.15)$$

where the convergence is independent from ξ .

Assume now that $\pi \equiv 1/r$; from Lemma 14, for a large enough M there exists a $\delta > 0$ such that $\tilde{p}_i > \tilde{p}_j + \delta$. Assign label 1 to \mathbb{I}^+ and 2 to \mathbb{I}^- , and let τ be a permutation such that $\tau(i) = 1$ and $\tau(j) = 2$. The overlap achieved by τ is thus

$$\frac{1}{n} \sum_{v \in V} \mathbf{1}_{\sigma(v)=i} \mathbf{1}_{v \in \mathbb{I}^+} + \frac{1}{n} \sum_{v \in V} \mathbf{1}_{\sigma(v)=j} \mathbf{1}_{v \in \mathbb{I}^-} - \frac{1}{r} = \frac{1}{r} (\tilde{p}_i + 1 - \tilde{p}_j) - \frac{1}{r} > \frac{\delta}{r}, \quad (3.16)$$

which completes the proof of Theorem 16.

3.3 Study of the matrix $D^{(\ell)}$

3.3.1 From $B^{(\ell)}$ to $D^{(\ell)}$

The first aim of this section is to prove Theorem 17, i.e. that we can replace matrix $B^{(\ell)}$ by $D^{(\ell)}$ in the algorithm from Theorem 16. Directly proving this theorem is hard, because of the lack of a decomposition such as the one in Lemma 17 for $D^{(\ell)}$. However, in view of the proof of Theorem 14 above, it is sufficient to prove the following proposition:

Proposition 4. *Let G be a SBM as above, and $\ell \sim \kappa \log_\alpha(n)$ with $\kappa < 1/12$. Let $B^{(\ell)}$ be the path expansion matrix of G , and $D^{(\ell)}$ its distance matrix. Then, with high probability:*

$$\rho(B^{(\ell)} - D^{(\ell)}) = \tilde{O}(\alpha^{\ell/2}), \quad (3.17)$$

where ρ is the spectral radius of a matrix.

For ease of notation, let $\delta^{(\ell)} = B^{(\ell)} - D^{(\ell)}$; we first notice that $\delta^{(\ell)}$ is a 0 – 1 matrix:

Lemma 15. *Let $\ell \sim \kappa \log_\alpha(n)$ with $\kappa < 1/12$. For all vertices $i, j \in \{1, \dots, n\}$,*

$$0 \leq \delta_{ij}^{(\ell)} \leq 1. \quad (3.18)$$

Furthermore, if $\delta_{ij}^{(\ell)} = 1$, then there exists a cycle \mathcal{C} such that:

$$d(i, \mathcal{C}) + d(j, \mathcal{C}) \leq \ell. \quad (3.19)$$

Define now a matrix $P^{(\ell)}$ by $P_{ij}^{(\ell)} = 1$ if there is a cycle \mathcal{C} such that $d(i, \mathcal{C}) + d(j, \mathcal{C}) \leq \ell$. By the previous lemma, we have $\delta_{ij}^{(\ell)} \leq P_{ij}^{(\ell)}$ for all (i, j) , and the Perron-Frobenius theorem implies:

$$\rho(\delta^{(\ell)}) \leq \rho(P^{(\ell)}). \quad (3.20)$$

It remains then to bound the spectral radius of $P^{(\ell)}$; the key lemma is the following:

Lemma 16. *For a given cycle \mathcal{C} , let $P_{\mathcal{C}}^{(\ell)}$ be the matrix defined by $P_{\mathcal{C},ij}^{(\ell)} = 1$ if $d(i, \mathcal{C}) + d(j, \mathcal{C}) \leq \ell$, and $V_{\mathcal{C}}$ the set of vertices such that $d(i, \mathcal{C}) \leq \ell$. Then:*

(i) $P_{\mathcal{C}}^{(\ell)}$ is zero outside of $V_{\mathcal{C}} \times V_{\mathcal{C}}$,

(ii) $\rho(P^{(\ell)}) = \max_{\mathcal{C}} \rho(P_{\mathcal{C}}^{(\ell)})$.

By part (ii) of the above lemma, it is sufficient to bound $\rho(P_{\mathcal{C}}^{(\ell)})$ for a given cycle \mathcal{C} in \mathcal{G} ; using part (i) of the above lemma, we can restrict our study to the subspace spanned by the vertices in $V_{\mathcal{C}}$.

Let v be a normed vector of size $|V_{\mathcal{C}}|$ corresponding to the highest eigenvalue of $P_{\mathcal{C}}^{(\ell)}$; as the coefficient (i, j) of $P_{\mathcal{C}}^{(\ell)}$ only depends on the distance of i and j to \mathcal{C} , we likewise group the coefficients of v by their distance t to \mathcal{C} , and write

$$v = (v_{ij})_{\substack{0 \leq t \leq \ell \\ 1 \leq j \leq S_t(\mathcal{C})}}.$$

We then have:

$$v^\top P_{\mathcal{C}} v = \sum_{t+u \leq \ell} \sum_{i,j} v_{ti} v_{uj} = \sum_{t+u \leq \ell} \left(\sum_i v_{ti} \right) \left(\sum_j v_{uj} \right).$$

By the Perron-Frobenius theorem, the coefficients of v are non-negative. For a given t , the coefficients v_{ti} are necessarily equal; otherwise, we could increase $\sum v_{ti}$ while leaving $\sum v_{ti}^2$ fixed, which leads to increasing $v^\top P_C^{(\ell)} v$ while keeping $\|v\|^2$ constant: this contradicts the definition of v .

Writing $v_{ti} = v_t$ for all $1 \leq i \leq S_t(\mathcal{C})$; we get:

$$v^\top P_C v = \sum_{t+u \leq \ell} S_t(\mathcal{C}) S_u(\mathcal{C}) v_t v_u \quad \text{and} \quad \|v\|^2 = \sum_t S_t(\mathcal{C}) v_t^2. \quad (3.21)$$

Let w be the size ℓ vector defined by $w_t = \sqrt{S_t(\mathcal{C})} v_t$. Rewriting the above expression in terms of w yields

$$v^\top P_C v = \sum_{t+u \leq \ell} \sqrt{S_t S_u} w_t w_u \quad \text{and} \quad \|v\|^2 = \|w\|^2, \quad (3.22)$$

where we omit the dependency of S_t in \mathcal{C} .

As a result, the spectral radius of P_C is equal to that of the $\ell \times \ell$ matrix Q_C defined by:

$$Q_C = \begin{pmatrix} S_0 & \sqrt{S_0 S_1} & \cdots & \sqrt{S_0 S_{\ell-1}} & \sqrt{S_0 S_\ell} \\ \sqrt{S_0 S_1} & S_1 & \cdots & \sqrt{S_1 S_{\ell-1}} & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \sqrt{S_0 S_{\ell-1}} & \sqrt{S_1 S_{\ell-1}} & \cdots & 0 & 0 \\ \sqrt{S_0 S_\ell} & 0 & \cdots & 0 & 0 \end{pmatrix}.$$

We now finally use the row sum bound to get:

$$\rho(P_C^{(\ell)}) = \rho(Q_C) \leq \max_t \sum_{u \leq \ell-t} \sqrt{S_t S_u} \quad (3.23)$$

$$\leq \max_t \sum_{u \leq \ell-t} \log(n) \alpha^{\frac{t+u}{2}} \quad \text{via lemma 25} \quad (3.24)$$

$$= O(\log(n) \alpha^{\ell/2}). \quad (3.25)$$

Combining the above inequality with Lemma 16 and inequality (3.20) eventually leads to

$$\rho(\delta^{(\ell)}) = \tilde{O}(\alpha^{\ell/2}), \quad (3.26)$$

which completes the proof of Proposition 4.

3.3.2 Stability to graph perturbation

In this subsection, we sketch the proofs for Theorems 18 and 19.

A note about computational complexity In the original algorithm, the computation of $B^{(\ell)}$ in polynomial time relies on the almost tree-like, tangle-free structure of the random graph \mathcal{G} ; this structure may be lost when we add cliques, and increase the algorithm complexity. As we want to devise polynomial algorithms in every case, this may be a hindrance.

Conversely, the computation of the distance matrix $D^{(\ell)}$ can be done in polynomial time (for example breadth-first search of the ℓ -neighbourhood of each vertex in G yields an algorithm in $O(n^{1+\kappa}) = O(n^{13/12})$ in the case of SBM, $O(n^2)$ in general) for any graph, which makes it all the more adapted to the problem at hand.

In order to prove Theorem 18, we need a less restrictive version of Proposition 4; indeed, bounding the spectral radius of the perturbation by $O(\alpha^{\ell/2})$ not only preserves the highest eigenvalues, but also bounds the remaining eigenvalues of $D^{(\ell)}$ by $\sqrt{\lambda_1(D^{(\ell)})}$. This bound is commonly referred to as a Ramanujan-like property of G .

This property, although interesting on its own, is not specifically needed for the reconstruction algorithm to work; rather, we only need one eigenvector associated to the second highest eigenvalue μ_2 to remain unchanged.

We'll therefore only need the following proposition:

Proposition 5. *We consider the same setting as Theorem 18. Let $D^{(\ell)}$ be the distance matrix of G , and \tilde{G} and $\tilde{D}^{(\ell)}$ be the perturbed versions (after adding adversarial noise) of G and $D^{(\ell)}$, respectively.*

Then

$$\rho(\tilde{D}^{(\ell)} - D^{(\ell)}) = o(\mu_2^\ell). \quad (3.27)$$

The proof relies on a bound similar to the one in Theorem 17, replacing matrices $P_{\mathcal{C}}$ and $Q_{\mathcal{C}}$ by matrices $P_{\mathcal{K}}$ and $Q_{\mathcal{K}}$ also depending only on the distance to the perturbed vertex set \mathcal{K} . The details can be found in the appendix.

3.4 Proof of Propositions 2 and 3

3.4.1 Outline of the proof and similarities with [29]

The main arguments of the proof rely on the study of three quantities:

- (i) a multi-type branching process Z_t ,
- (ii) a similar process based on exploring the neighbourhood of a vertex v in G , named $Y_t(v)$,
- (iii) the actual vectors we're aiming to study, $B^{(\ell)}\chi_k$.

When the ℓ -neighbourhood of v is cycle-free, we have that $B^{(\ell)}\chi_k = \langle \phi_k, Y_t(v) \rangle$ for $k \in [r_0]$; and there is a coupling between the laws of Z_t and $Y_t(v)$ for almost every v , which allows us to translate results on Z_t to results on $B^{(\ell)}\chi_k$.

The proof in [29] studies the matrix B^ℓ , where B is the non-backtracking matrix; B_{ij}^ℓ therefore counts the number of non-backtracking walks between i and j . When the ℓ -neighbourhood of i is tree-like, $(B^{(\ell)}\chi_k)_i = (B^\ell\bar{\chi}_k)_i$, where $\bar{\chi}_k$ is a similarly defined vector; most of the results from [29] can therefore be applied to this setting without further work. We will simply lay out the main steps of the proof, highlighting the main differences with [29] when necessary.

3.4.2 Local structure of G

For an integer $t \geq 0$, we introduce the vector $Y_t(v) = (Y_t(v)(i))_{i \in [r]}$, where

$$Y_t(v)(i) = |\{w \in V \mid d(v, w) = t, \sigma(w) = i\}|.$$

The proof of our first proposition, although quite lengthy, is completely identical to its equivalent in [29]; we therefore omit it.

Proposition 6. *Let $\ell \sim \kappa \log_\alpha(n)$ with $\kappa < 1/8$; then, for all $\gamma < 1/2$:*

- (i) *for any $k \in [r_0]$, there exists $\rho_k > 0$ such that in probability,*

$$\frac{1}{n} \sum_{v \in V} \frac{\langle \phi_k, Y_\ell(v) \rangle^2}{\mu_k^{2\ell}} \rightarrow \rho_k.$$

(ii) for any $j \neq k \in [r]$,

$$\mathbf{E} \left| \frac{1}{n} \sum_{v \in V} \langle \phi_j, Y_\ell(v) \rangle \langle \phi_k, Y_\ell(v) \rangle \right| = O \left(\alpha^{5\ell/2} n^{-\gamma/2} (\log(n))^{5/2} \right).$$

(iii) for any $j \neq k \in [r]$,

$$\mathbf{E} \left| \frac{1}{n} \sum_{v \in V} \langle \phi_j, Y_{2\ell}(v) \rangle \langle \phi_k, Y_\ell(v) \rangle \right| = O \left(\alpha^{7\ell/2} n^{-\gamma/2} (\log(n))^{5/2} \right).$$

For $t \geq 0$, define $\mathcal{Y}_t(v) = \{w \in V \mid d(v, w) = t\}$; for $k \in [r]$, we set

$$P_{k,\ell}(v) = \sum_{t=0}^{\ell-1} \sum_{w \in \mathcal{Y}_t(v)} L_k(w),$$

where

$$L_k(w) = \sum_{(x,y) \in \mathcal{Y}_1(w) \setminus \mathcal{Y}_t(v), x \neq y} \langle \phi_k, \tilde{Y}_t(x) \rangle \tilde{S}_{\ell-t-1}(y),$$

$\tilde{Y}_t(x)$ is the equivalent of $Y_t(x)$ when all vertices in $(G, v)_t$ (i.e. vertices at distance at most t from v) are removed and $\tilde{S}_{\ell-t-1}(y) = \|\tilde{Y}_{\ell-t-1}(y)\|_1$.

It can be seen from [29] that when $(G, v)_{2\ell}$ is a tree, then

$$(B^{(\ell)} B^{(\ell)} \chi_k)_v = P_{k,\ell}(v) + \chi_k(v) S_\ell(v) + \langle \phi_k, Y_{2\ell}(v) \rangle.$$

One main difference with the proof in [29] is the presence of the last term in the above sum, as well as the fact that dealing with $B^{(\ell)} B^{(\ell)} \chi_k$ is a little more difficult. The next proposition is an adaptation of Proposition 38 from [29], with an identical – and thus omitted – proof:

Proposition 7. *Let $\ell \sim \kappa \log_\alpha(n)$ with $\kappa < 1/10$. Then, for all $\gamma < 1/2$:*

(i) for all $k \in [r_0]$, there exists ρ'_k such that w.h.p

$$\frac{1}{n} \sum_{v \in V} \frac{(P_{k,\ell}(v) + \langle \phi_k, Y_{2\ell}(v) \rangle)^2}{\mu_k^{4\ell}} \rightarrow \rho'_k.$$

(ii) for any $j \neq k \in [r]$, for some $c > 0$:

$$\frac{1}{n} \sum_{v \in V} P_{k,\ell}(v) \langle \phi_j, Y_\ell(v) \rangle = O \left(\alpha^{7\ell/2} n^{-\gamma/2} (\log(n))^c \right).$$

3.4.3 From local neighbourhoods to the matrix $B^{(\ell)}$

For ease of notation, we define $N_{k,\ell}(v) = \langle \phi_k, Y_\ell(v) \rangle$; using the same methods as in [29], we have the following estimates:

Proposition 8. *Let $\ell \sim \kappa \log_\alpha(n)$ with $\kappa < 1/4$. Then w.h.p:*

$$\|B^{(\ell)} \chi_k - N_{k,\ell}\| = o(\alpha^{\ell/2} \sqrt{n}) \text{ and } \|B^{(\ell)} B^{(\ell)} \chi_k - P_{k,\ell} - N_{k,2\ell}\| = O(\alpha^\ell \sqrt{n}).$$

It then remains to follow the proof of Proposition 19 from [29]; we simply highlight the proof for estimation (iii) of Proposition 2, since it is the only difference:

Proof. (Proposition 2-(iii)): We have by definition

$$\langle \varphi_j, \zeta_k \rangle = \frac{\langle B^{(\ell)} \chi_j, B^{(\ell)} B^{(\ell)} \chi_k \rangle}{\|B^{(\ell)} \chi_j\| \|B^{(\ell)} B^{(\ell)} \chi_k\|}.$$

But $\|B^{(\ell)} \chi_j\| = \Theta(\sqrt{n} \mu_k^\ell)$, $\|B^{(\ell)} B^{(\ell)} \chi_k\| = \|B^{(\ell)} \chi_k\| \theta_k = \Theta(\sqrt{n} \mu_k^{2\ell})$ and:

$$\begin{aligned} \left| \langle B^{(\ell)} \chi_j, B^{(\ell)} B^{(\ell)} \chi_k \rangle - \langle N_{j,\ell}, P_{k,\ell} + N_{k,2\ell} \rangle \right| &\leq \|N_{j,\ell}\| \|B^{(\ell)} B^{(\ell)} \chi_k - P_{k,\ell} - N_{k,2\ell}\| \\ &\quad + \|B^{(\ell)} B^{(\ell)} \chi_k\| \|B^{(\ell)} \chi_j - N_{j,\ell}\| \\ &= \tilde{O}(\alpha^{4\ell} \sqrt{n}). \end{aligned}$$

Furthermore, from Propositions 6 and 7, we get

$$\langle N_{j,\ell}, P_{k,\ell} + N_{k,2\ell} \rangle = \tilde{O}(\alpha^{7\ell/2} n^{1-\gamma/2}).$$

This gives the desired result. \square

3.4.4 Ramanujan property of $B^{(\ell)}$

In order to complete the proof of Theorem 14, we need a control on the other eigenvalues of $B^{(\ell)}$. This is covered by the following proposition:

Proposition 9. *Let $H = \langle \varphi_1, \dots, \varphi_{r_0} \rangle$, and $\ell \sim \kappa \log_\alpha(n)$ with $\kappa < 1/12$. Then with high probability*

$$\sup_{x \in H^\perp, \|x\|=1} \|B^{(\ell)} x\| = \tilde{O}(\alpha^{\ell/2}). \quad (3.28)$$

The proof of this result relies on the following decomposition of $B^{(\ell)}$, whose proof can be found in [91]:

Lemma 17. *Matrix $B^{(\ell)}$ verifies the identity*

$$B^{(\ell)} = \delta^{(\ell)} + \sum_{m=1}^{\ell} \delta^{(\ell-m)} \bar{A} B^{(m-1)} - \sum_{m=0}^{\ell} \Gamma^{\ell,m}, \quad (3.29)$$

for matrices $\delta^{(j)}, \Gamma^{\ell,m}$ such that for $\ell = O(\log(n))$ and with high probability, for all $\varepsilon > 0$,

$$\rho(\delta^{(j)}) = \tilde{O}(\alpha^j), \quad j = 1, \dots, \ell, \quad (3.30)$$

$$\rho(\Gamma^{\ell,m}) = n^{\varepsilon-1} \alpha^{(\ell+m)/2}, \quad m = 1, \dots, \ell. \quad (3.31)$$

Here, \bar{A} refers to the expected value of the adjacency matrix A of G .

The next step is therefore to control $B^{(m-1)} x$ for $x \in H^\perp$; in what follows γ will be any constant below $1/2$. We begin with the following proposition from [29]:

Proposition 10. *Let $\ell \sim \kappa \log_\alpha(n)$ with $\kappa < \gamma/2$. There exists a subset $\mathcal{B} \subset V$, constants C and c such that w.h.p the following holds:*

(i) for all $i \in V \setminus \mathcal{B}$, $0 \leq m \leq \ell$,

$$\begin{aligned} |(B^{(m)}\chi_k)_i - \mu_k^{t-\ell}(B^{(\ell)}\chi_k)_i| &\leq C \log(n)^c \alpha^{m/2} && \text{if } k \in [r_0], \\ |(B^{(m)}\chi_k)_i| &\leq C \log(n)^c \alpha^{m/2} && \text{if } k \in [r] \setminus [r_0]. \end{aligned}$$

(ii) for all $i \in \mathcal{B}$, $0 \leq m \leq \ell$ and $k \in [r]$,

$$|(B^{(\ell)}\chi_k)_i| \leq C \log(n)^c \alpha^m.$$

(iii) $|\mathcal{B}| = \tilde{O}(\alpha^\ell n^{1-\gamma})$.

From this, we get the following corollary:

Corollary 2. *Let $\ell \sim \kappa \log_\alpha(n)$ with $\kappa < \gamma/2$; then, with high probability, for $0 \leq m \leq \ell - 1$ and $k \in [r_0]$:*

$$\sup_{x \perp B^{(\ell)}\chi_k, \|x\|=1} \langle B^{(m)}\chi_k, x \rangle = \tilde{O}(\sqrt{n} \alpha^{m/2}).$$

Additionally, for $k \in [r] \setminus [r_0]$,

$$\|B^{(m)}\chi_k\| = \tilde{O}(\sqrt{n} \alpha^{m/2}).$$

Proof. Write

$$\langle B^{(m)}\chi_k, x \rangle = \sum_{i \in \mathcal{B}} x_i (B^{(m)}\chi_k)_i + \sum_{i \notin \mathcal{B}} x_i (B^{(m)}\chi_k)_i = s_1 + s_2.$$

Using the Cauchy-Schwarz inequality, the first sum is bounded by

$$|s_1| \leq \log(n)^c \alpha^m \sqrt{|\mathcal{B}|} \leq \log(n)^d \alpha^m \alpha^{\ell/2} n^{(1-\gamma)/2} = o(\sqrt{n} \alpha^{m/2}),$$

while the second can be treated using Proposition 10 and the fact that $\langle B^{(\ell)}\chi_k, x \rangle = 0$:

$$\begin{aligned} |s_2| &\leq \mu_k^{t-\ell} \sum_{i \in \mathcal{B}} |x_i| |(B^{(\ell)}\chi_k)_i| + \sum_{i \notin \mathcal{B}} |x_i| |(B^{(m)}\chi_k)_i - \mu_k^{t-\ell}(B^{(\ell)}\chi_k)_i| \\ &\leq \log(n)^c \alpha^{t-\ell} \alpha^\ell \alpha^{\ell/2} n^{(1-\gamma)/2} + \log(n)^c \sqrt{n} \alpha^{t/2} \\ &= \tilde{O}(\sqrt{n} \alpha^{m/2}), \end{aligned}$$

where we used again the Cauchy-Schwarz inequality as before.

Let now $k \in [r] \setminus [r_0]$; as before, we write

$$\begin{aligned} \|B^{(m)}\chi_k\|^2 &= \sum_{i \in \mathcal{B}} (B^{(m)}\chi_k)_i^2 + \sum_{i \notin \mathcal{B}} (B^{(m)}\chi_k)_i^2 \\ &\leq |\mathcal{B}| \log(n)^c \alpha^{2m} + n \log(n)^c \alpha^m \\ &= n \log(n)^c (\alpha^{l+2m} n^{-\gamma} + \alpha^m) \\ &= \tilde{O}(n \alpha^m), \end{aligned}$$

and the result follows. \square

We are now ready to prove Proposition 9:

Proof. Let $x \in H^\perp$ such that $\|x\| = 1$ and the supremum in (3.28) is reached; using the decomposition from Lemma 17, we have

$$\|B^{(\ell)}x\| \leq \rho(\delta^{(\ell)}) + \sum_{m=1}^{\ell} \rho(\delta^{(\ell-m)}) \|\bar{A}B^{(m-1)}x\| + \sum_{m=1}^{\ell} \rho(\Gamma^{\ell,m}).$$

The first and third terms are bounded by $\tilde{O}(\alpha^{\ell/2})$. For the second term, we notice that defining the matrix P by

$$P = \frac{1}{n} \sum_{k=1}^r \mu_k \chi_k \chi_k^\top,$$

we have $\bar{A} = P - \text{diag}(P)$ since $W = \sum \mu_k \phi_k \phi_k^\top$.

Therefore, for fixed $1 \leq m \leq \ell$, we have:

$$\begin{aligned} \|\bar{A}B^{(m-1)}x\| &= \left\| \sum_{k=1}^r \mu_k \chi_k \chi_k^\top B^{(m-1)}x - \text{diag}(P)B^{(m-1)}x \right\| \\ &\leq \frac{\sup_i W_{ii}}{n} \|B^{(m-1)}x\| + \sum_{k \in [r_0]} \frac{\mu_k}{n} \|\chi_k \chi_k^\top B^{(m-1)}x\| + \sum_{k \in [r] \setminus [r_0]} \frac{\mu_k}{n} \|\chi_k \chi_k^\top B^{(m-1)}x\| \\ &= I + J + K. \end{aligned}$$

Notice first that $B_{ij}^{(\ell)} \leq 2$ for all i, j by the tangle-free property, so $I = O(1)$. Now, for $k \in [r_0]$, we have

$$\begin{aligned} \|\chi_k \chi_k^\top B^{(m-1)}x\| &= \|\chi_k\| \langle B^{(m-1)}\chi_k, x \rangle \\ &\leq \tilde{O}(\sqrt{n} \times \sqrt{n} \alpha^{m/2}). \end{aligned}$$

Therefore, $J = \tilde{O}(\alpha^{m/2})$; finally, using the Cauchy-Schwarz inequality, we have for $k \in [r] \setminus [r_0]$

$$\begin{aligned} \|\chi_k \chi_k^\top B^{(m-1)}x\| &\leq \|\chi_k\| \|B^{(m-1)}\chi_k\| \|x\| \\ &= \tilde{O}(\sqrt{n} \times \sqrt{n} \alpha^{m/2} \times 1). \end{aligned}$$

Putting this all together, we find that for $1 \leq m \leq \ell$

$$\|\bar{A}B^{(m-1)}x\| = \tilde{O}(\alpha^{m/2}).$$

Since $\rho(\delta^{(\ell-m)}) = \tilde{O}(\alpha^{(\ell-m)/2})$, we get $\|B^{(\ell)}x\| = \tilde{O}(\alpha^{\ell/2})$, which proves the desired result. \square

3.4.5 Proof of Proposition 3

Using Proposition 9, we are now able to prove our last result. Note that if $\kappa < 1/12$, there exists a $\gamma < 1/2$ such that $\kappa < \gamma/6$.

Let z_k be the Gram-Schmidt orthonormalization of φ_k ; using Lemma 9 from [29], we know that

$$\|\varphi_k - z_k\| = \tilde{O}(\alpha^{3\ell/2} n^{-\gamma/2}),$$

and thus z_k is asymptotically parallel to φ_k .

We only need a final lemma to complete our proof:

Lemma 18. Assume that $\ell \sim \kappa \log_\alpha(n)$ with $\kappa < \gamma/6$. Then

$$\|\zeta_k - z_k\| = \tilde{O}(\theta_k^{-1} \alpha^{\ell/2}).$$

Proof. Write

$$\zeta_k = \sum_{j \in [r_0]} \langle \zeta_k, z_j \rangle z_j + x,$$

where $x \in H^\perp$.

We have, for $j \neq k$, $\langle \zeta_k, z_j \rangle = \tilde{O}(\alpha^{2\ell} n^{-\gamma/2})$ by the above bound of $\|\varphi_j - z_j\|$; furthermore,

$$\begin{aligned} \|x\|^2 &= \langle \zeta_k, x \rangle \\ &= \theta_k^{-1} \langle B^{(\ell)} \varphi_k, x \rangle \\ &\leq \theta_k^{-1} \|B^{(\ell)} x\| \\ &= \tilde{O}(\theta_k^{-1} \alpha^{\ell/2}) \times \|x\|. \end{aligned}$$

Therefore, we can write

$$\begin{aligned} 1 = \|\zeta_k\|^2 &= \langle \zeta_k, z_k \rangle^2 + \sum_{j \neq k} \langle \zeta_k, z_j \rangle^2 + \|x\|^2 \\ &= \langle \zeta_k, z_k \rangle^2 + \tilde{O}(\alpha^{2\ell} n^{-\gamma/2}) + \tilde{O}(\theta_k^{-2} \alpha^\ell) \\ &= \langle \zeta_k, z_k \rangle^2 + \tilde{O}(\theta_k^{-2} \alpha^\ell), \end{aligned}$$

since $\kappa < \gamma/6$.

Then,

$$\|z_k - \zeta_k\|^2 = 2(1 - \langle \zeta_k, z_k \rangle) = \tilde{O}(\theta_k^{-2} \alpha^\ell),$$

which yields the desired result. \square

Proof. (of Proposition 3): We first bound $\|B^{(\ell)} z_k - Dz_k\|$ for $k \in [r_0]$. Notice that $Dz_k = \theta_k z_k$; this gives

$$\begin{aligned} \|B^{(\ell)} z_k - Dz_k\| &\leq \|B^{(\ell)} z_k - B^{(\ell)} \varphi_k\| + \|B^{(\ell)} \varphi_k - \theta_k z_k\| \\ &\leq \rho(B^{(\ell)}) \|z_k - \varphi_k\| + \theta_k \|\zeta_k - z_k\| \\ &= O(\alpha^\ell) \times \tilde{O}(\alpha^{3\ell/2} n^{-\gamma/2}) + \tilde{O}(\alpha^{\ell/2}) \\ &= \tilde{O}(\alpha^{\ell/2}). \end{aligned}$$

Consider now $x \in \mathbb{R}^V$ such that $\|x\| = 1$. Decomposing x as $\sum x_k z_k + x'$ where $x' \in H^\perp$, we have:

$$\begin{aligned} \|B^{(\ell)} x - Dx\| &\leq \sum_{k \in [r_0]} x_k \|B^{(\ell)} z_k - Dz_k\| + \|B^{(\ell)} x' - Dx'\| \\ &\leq \tilde{O}(\alpha^{\ell/2}) + \|B^{(\ell)} x'\| \\ &= \tilde{O}(\alpha^{\ell/2}), \end{aligned}$$

which completes the proof. \square

3.5 Proofs for Theorem 16

3.5.1 Proof of Lemma 14

We first recall a result from Kesten and Stigum: consider a multitype Galton-Watson process, where the type of the root node is distributed according to arbitrary probability vector ν , and a particle of type $j \in [r]$ has a $\text{Poi}(M_{ij})$ number of children of type i . Let Z_t be the vector of population at time t , and \mathcal{F}_t the natural filtration associated to Z_t ; we have the following statement:

Lemma 19. *For each μ eigenvalue of M such that $\mu^2 > \alpha$, and each eigenvector ϕ associated to μ ,*

$$t \mapsto X(\phi, \nu, t) = \mu_k^{-t} \langle \phi_k, Z_t \rangle \quad (3.32)$$

is an \mathcal{F}_t -martingale converging a.s. and in L^2 to a random variable with finite variance and expected value $\langle \phi_k, \nu \rangle$.

Let $\mu \neq \alpha$ be an eigenvalue of M of multiplicity d such that $\mu^2 > \alpha$, and $\phi^{(1)}, \dots, \phi^{(d)}$ an orthonormal basis of eigenvectors associated to d . We define for all $i \in [d], j \in [r]$, $X_j^{(i)}$ the limit variable of martingale (3.32), applied to $\phi = \phi^{(i)}$ and $\nu = \delta_j$. Similarly to previous notations, let $\phi^{(i)}$ (resp. $X^{(i)}$) be the vector $(\phi_j^{(i)})_{j \in [r]}$ (resp. $(X_j^{(i)})_{j \in [r]}$), and ϕ (resp. \mathbf{X}) the (random) matrix whose columns are the $\phi^{(i)}$ (resp. the $X^{(i)}$). Recall that from Lemma 19, the expected value of $X_j^{(i)}$ is $\phi_j^{(i)}$ for all i, j .

Now, let ξ be an eigenvector of $B^{(\ell)}$, normalized so that $\|\xi\|^2 = n$, with associated eigenvalue $\Theta(\mu^\ell)$; as shown in the proof of Theorem 14, there exists a vector $u \in \mathbb{R}^d$ such that

$$\|\xi - (\langle \phi u, Y_\ell(v) \rangle)_{v \in V}\| = o(1).$$

We let

$$\phi^{(\xi)} = \phi u \quad \text{and} \quad X^{(\xi)} = \mathbf{X} u.$$

From Proposition 6, u has norm $\Theta(1)$, and since $\mu^{-t} \langle \phi, Z_t \rangle$ (with $\nu = \delta_j$) converges to $X_j^{(i)}$ in L^1 , $\mu^{-t} \langle \phi^{(\xi)}, Z_t \rangle$ converges to $X_j^{(\xi)}$ in L^1 independently of ξ .

Using proposition 36 from [29], we have the following:

Lemma 20. *For all $i \in [r]$, we have the following convergence in L^1 :*

$$\frac{1}{n} \sum_{v \in V} \mathbf{1}_{\sigma(v)=i} \xi(v) \mathbf{1}_{|\xi(v)| \leq K} \rightarrow \pi(i) \mathbf{E} \left[X^{(\xi)} \mathbf{1}_{|X^{(\xi)}| \leq K} \right],$$

for all K that is a continuity point of the distribution of X_i , and independently of ξ .

Proof. We first recall the aforementioned proposition from [29]: we say that a function τ that takes a graph and a distinguished vertex as an argument is ℓ -local if $\tau(G, v)$ depends only on the ℓ -neighbourhood of v in G . Denote by T the multitype Galton-Watson tree discussed earlier, rooted at o , where o has the distribution π .

Lemma 21. *Assume that τ, ψ are two ℓ -local functions such that $|\tau| \leq \psi$ and ψ is non-decreasing by the addition of edges. Then, if $\ell \sim \kappa \log_\alpha(n)$ with $\kappa < 1/2$, we have, for $\gamma < 1/2$:*

$$\begin{aligned} & \mathbf{E} \left[\left| \frac{1}{n} \sum_{v \in V} \tau(G, v) - \mathbf{E}[\tau(T, o)] \right| \right] \\ & \leq c \frac{\alpha^{\ell/2} \sqrt{\log(n)}}{n^{\gamma/2}} \left(\left(\mathbf{E}[\max_{v \in V} \psi^4(G, v)] \right)^{1/4} \vee \left(\mathbf{E}[\psi^2(T, o)] \right)^{1/2} \right) \end{aligned}$$

We now apply this lemma with $\tau(G, v) = \mathbf{1}_{\sigma(v)=i} \langle \phi^{(\xi)}, Y_\ell(v) \rangle \mathbf{1}_{|\langle \phi^{(\xi)}, Y_\ell(v) \rangle| \leq K}$ where Y_ℓ is defined in Proposition 6.

We can set $\psi(G, v) = K$, and by Lemma 19 and the subsequent analysis, we have

$$\mathbf{E}[\tau(T, o)] \rightarrow \pi(i) \mathbf{E} \left[X_i^{(\xi)} \mathbf{1}_{|X_i^{(\xi)}| \leq K} \right]$$

independently of ξ .

Now, by definition, we have $\|\xi - (\langle \phi^{(\xi)}, Y_\ell(v) \rangle)_{v \in V}\| = o(1)$ (again independently of ξ). By the Cauchy-Schwarz inequality, we deduce that

$$\frac{1}{n} \sum_{v \in V} \left| \xi(v) - \langle \phi^{(\xi)}, Y_\ell(v) \rangle \right| = o(1)$$

as well, and the lemma follows if K is a continuity point of $X_i^{(\xi)}$. \square

It now remains to prove the desired properties of the $X_i^{(\xi)}$; first, since $\mu \neq \alpha$, then $\phi^{(\xi)}$ is orthogonal to the all-one vector and as such

$$\sum_{i \in [r]} \mathbf{E}[X_i] = \sum_{i \in [r]} \phi_i^{(\xi)} = 0$$

Moreover, $\|\phi^{(\xi)}\|^2 = \|u\|^2 = \Theta(1)$, which proves the second assertion.

Finally, let $\eta > 0$; since the $X_j^{(i)}$ all have finite variance, there exists a constant $K > 0$ such that $\mathbf{P}(|X_j^{(i)}| \leq K) \geq 1 - \eta$ for all i, j and thus

$$\mathbf{P}(\|\mathbf{X}\|_\infty \leq K) \geq 1 - dr\eta.$$

Using the equivalence of norms, we find

$$\mathbf{P}(\forall i, \|X^{(i)}\|_2^2 \leq rK'^2) \geq 1 - dr\eta$$

which implies (since $X^{(\xi)} = \mathbf{X}u$)

$$\mathbf{P}(\|X^{(\xi)}\|_2^2 \leq r\|u\|_2^2 K'^2) \geq 1 - dr\eta.$$

Using again norm equivalence yields finally, for $K = \sqrt{r}\|u\|K'$,

$$\mathbf{P}(\|X^{(\xi)}\|_\infty \leq K) \geq 1 - dr\eta. \tag{3.33}$$

Now, we have, for all i ,

$$\begin{aligned} \left| \mathbf{E}[X_i^{(\xi)} \mathbf{1}_{|X_i^{(\xi)}| \leq K}] - \mathbf{E}[X_i^{(\xi)}] \right| &= \mathbf{E} \left[X_i^{(\xi)} \mathbf{1}_{|X_i^{(\xi)}| > K} \right] \\ &\leq \sqrt{\mathbf{E} \left[(X_i^{(\xi)})^2 \right] \mathbf{P} \left(|X_i^{(\xi)}| \geq K \right)} \\ &\leq \sqrt{\mathbf{E} \left[(X_i^{(\xi)})^2 \right]} \cdot \sqrt{dr\eta} \end{aligned}$$

But by Doob's Theorem, $\mathbf{E} \left[(X_i^{(\xi)})^2 \right]$ is finite so choosing η accordingly yields the last inequality of Lemma 14.

3.5.2 Proof of limit (3.15)

For each $v \in V$, we define I_v to be the random variable equal to 1 if v is assigned to I^+ , and 0 otherwise. Conditionnally to ξ , it is straightforward to see that

$$I_v \sim \text{Ber}(q_v) \quad \text{with} \quad q_v = \frac{1}{2} + \frac{1}{2K} \xi(v) \mathbf{1}_{|\xi(v)| \leq K}$$

Now, let $P_i = (|\{v \in I^+ \mid \sigma(v) = i\}|) / n$; by definition,

$$P_i = \frac{1}{n} \sum_{v \in V} \mathbf{1}_{\sigma(v)=i} I_v$$

We therefore have

$$\mathbf{E}[P_i \mid \xi] = \frac{1}{n} \sum_{v \in V} \mathbf{1}_{\sigma(v)=i} \left(\frac{1}{2} + \frac{1}{2K} \xi(v) \mathbf{1}_{|\xi(v)| \leq M} \right), \quad \text{Var}(P_i) \leq \frac{1}{4n}$$

and thus with high probability, independently of ξ ,

$$P_i = \mathbf{E}[P_i \mid \xi] + n^{-1/3} \tag{3.34}$$

$$\rightarrow \pi(i) \left(\frac{1}{2} + \frac{1}{2K} \mathbf{E}[X_i^{(\xi)} \mathbf{1}_{|X_i^{(\xi)}| \leq M}] \right) = \pi(i) \tilde{p}_i \tag{3.35}$$

where the convergence speed is independent from ξ .

3.5.3 Explicit bounds on K

In this section, the goal is to perform a more precise analysis of the limit variables $X_i^{(\xi)}$, and to leverage this analysis to obtain an explicit value for K in Algorithm 15. For simplicity, we will assume that $\pi \equiv 1/r$ throughout this section, although most of the results hold for general π . We begin with a small lemma:

Lemma 22. *Let ϕ be a normed eigenvector of M associated to an eigenvalue $\alpha > \mu > \sqrt{\alpha}$, and denote by $X_i^{(\phi)}$ the limit random variables of Lemma 19. Let $c^{(\phi),j}$ (resp. $m^{(\phi),j}$) the vector of the j -th cumulants (resp. moments) of the $X_i^{(\phi)}$. We then have the following relation, for all $j \in \mathbb{N}$:*

$$c^{(\phi),j} = \frac{M}{\mu^j} m^{(\phi),j}$$

By definition, we have $c^{(\phi),1} = m^{(\phi),1} = \phi$, and we have the following corollary for $c^{(\phi),2}$ and $m^{(\phi),2}$:

Corollary 3. *We denote by ϕ^2 the vector whose coordinates are the ϕ_i^2 . Then*

$$c^{(\phi),2} = \left(\mathbf{I} - \frac{M}{\mu^2} \right)^{-1} \frac{M}{\mu^2} \phi^2 \quad \text{and} \quad m^{(\phi),2} = \left(\mathbf{I} - \frac{M}{\mu^2} \right)^{-1} \phi^2.$$

As a result, we have

$$\sum_{i \in [r]} \text{Var}(X_i^{(\phi)}) = \frac{1}{\tau - 1} \quad \text{and} \quad \sum_{i \in [r]} \mathbf{E}[(X_i^{(\phi)})^2] = \frac{\tau}{\tau - 1},$$

where $\tau = \mu_2^2/\alpha$.

Proof. (of Corollary 3). The first part is an easy calculation, observing that $c^{(\phi),2} = m^{(\phi),2} - \phi^2$.

For the second part, since the all-one vector e is an eigenvector of M associated to the eigenvalue α , we have:

$$\begin{aligned} \sum_{i \in [r]} \text{Var}(X_i^{(\phi)}) &= \langle e, c^{(\phi),2} \rangle \\ &= e^\top \left(\mathbf{I} - \frac{M}{\mu^2} \right)^{-1} \frac{M}{\mu^2} \phi^2 \\ &= \frac{\alpha/\mu^2}{1 - \alpha/\mu^2} e^\top \phi^2 \\ &= \frac{1}{\tau - 1}, \end{aligned}$$

and a similar calculation yields the second identity. \square

It now remains to prove Lemma 22:

Proof. (of Lemma 22). Using the Galton-Watson tree definition (and going one step down into the tree), we have the following characterization for the variables $X_i^{(\phi)}$:

$$X_i^{(\phi)} = \frac{1}{\mu} \sum_{j \in [r]} \sum_{k=1}^{\text{Poi}(M_{ij})} X_{j,k}^{(\phi)},$$

where the $X_{j,k}^{(\phi)}$ are independent copies of $X_j^{(\phi)}$ for all k . Applying the Laplace transform (denoted by $\psi_i^{(\phi)}$) and taking the logarithm on both sides, we find that for all $t \in \mathbb{R}$,

$$\log(\psi_i^{(\phi)}(t)) = \sum_{j \in [r]} M_{ij} \left(\psi_j^{(\phi)} \left(\frac{t}{\mu} \right) - 1 \right)$$

Now, the k -th Taylor coefficient of the LHS is $c_i^{(\phi),k}/k!$, and the one on the RHS is

$$\sum_{j \in [r]} M_{ij} \frac{m_j^{(\phi),k}}{k! \mu^k} = \frac{1}{\mu^k k!} \left[M m^{(\phi),k} \right]_i,$$

which completes the proof. \square

We now can prove our first result on the vector u defined before:

Lemma 23. *Let μ , ξ and u be defined as in the proof of Lemma 14. Then we have*

$$\|u\|^2 = r(\tau - 1) + o(1)$$

Proof. From Lemma 6, we know that for each $i \in [d]$,

$$\left\| \left(\langle \phi^{(i)}, Y_\ell(v) \rangle \right)_{v \in V} \right\|^2 = n(\rho^{(i)} + o(1)) \quad \text{where} \quad \rho^{(i)} = \sum_{i \in [r]} \pi(i) \mathbf{E} \left[(X^{(i)})^2 \right]$$

But since $\pi \equiv 1/r$, we know from Corollary 3 that

$$\rho^{(i)} = \rho := \frac{1}{r(\tau - 1)}$$

But the vectors $(\langle \phi^{(i)}, Y_\ell(v) \rangle)_{v \in V}$ are asymptotically orthogonal, and thus

$$n = \|\xi\|^2 = (\|v\|^2 + o(1)) \cdot n \cdot (\rho + o(1)),$$

which yields the desired result. \square

Now, we are ready to prove some bounds for K ; the main step is the following Markov bound on $(X^{(i)})^2$:

Lemma 24. *Let $\eta > 0$; then, for all $i \in [d]$, $j \in [r]$,*

$$\mathbf{P} \left(|X_j^{(i)}| \leq \sqrt{\frac{\tau}{\eta(\tau - 1)}} \right) \geq 1 - \eta$$

Proof. For all $C > 0$, we have by Markov's inequality

$$\mathbf{P} \left(|X_j^{(i)}| \geq C \right) \leq \frac{\mathbf{E}[(X_j^{(i)})^2]}{C^2} \leq \frac{\tau}{C^2(\tau - 1)},$$

where we bounded $\mathbf{E}[(X_j^{(i)})^2]$ by the sum of all $\mathbf{E}[(X_k^{(i)})^2]$. The lemma then follows easily. \square

Now, we have to unravel the calculations done in the proof for Lemma 14; let $\varepsilon < 0$. By the same bound as above (as well as the fact that the $\phi^{(i)}$ are orthogonal), we have

$$\mathbf{E} \left[(X_i^{(\xi)})^2 \right] \leq \|u\|^2 \frac{\tau}{\tau - 1} = r\tau + o(1)$$

Therefore, an asymptotically good choice of η is

$$\eta = \frac{\varepsilon^2}{r^2 d \tau},$$

which yields a value for K' of

$$K' = \sqrt{\frac{\tau}{\eta(\tau - 1)}} = \frac{r \tau}{\varepsilon} \sqrt{\frac{d}{\tau - 1}}$$

Finally, the bound for K becomes

$$K = \sqrt{r} \|u\| K' = \frac{r \tau}{\varepsilon} \sqrt{r d \tau}$$

We thus need to find a sufficient value for ϵ ; recall that

$$\sum_{i \in [r]} \mathbf{E} \left[X_i^{(\xi)} \right]^2 = \|u\|^2 = r(\tau - 1) + o(1) \quad \text{and} \quad \sum_{i \in [r]} \mathbf{E} \left[X_i^{(\xi)} \right] = 0,$$

and therefore there exists some values for i, j such that

$$\left| \mathbf{E} \left[X_i^{(\xi)} \right]^2 - \mathbf{E} \left[X_j^{(\xi)} \right]^2 \right| \geq \sqrt{r(\tau - 1)} + o(1).$$

A sufficient choice of ϵ is thus $\sqrt{r(\tau - 1)}/4$, which yields an explicit value for K :

$$K = \frac{r\tau}{\epsilon} \sqrt{rd\tau} = r\tau \sqrt{d \frac{\tau}{\tau - 1}}$$

3.6 Proof of Lemma 15

We first recall some results about the neighbourhoods of vertices, whose proofs can be found in [29]:

Lemma 25. *For a vertex i , define $S_t(i)$ as the number of vertices at distance t of i .*

Then there exist constants C and $\epsilon > 0$ such that with probability $1 - O(n^{-\epsilon})$, for all $i \in \{1, \dots, n\}$ and $\ell = O(\log(n))$:

$$S_t(i) \leq C \cdot \log(n) \cdot \alpha^t, \quad t \in \{1, \dots, \ell\}. \quad (3.36)$$

On the other hand, with high probability, when $\ell = \kappa \log_\alpha(n)$ with $\kappa < 1/2$:

$$\sum_{i=1}^n S_\ell(i)^2 = \Theta(n\alpha^{2\ell}). \quad (3.37)$$

Additionally, a result about the almost tree-like structure of vertex neighbourhoods:

Lemma 26. *Assume $\ell = \kappa \log(n)$, with $\kappa \log(\alpha) < 1/4$. Then with high probability no node i has more than one edge cycle in its ℓ -neighbourhood; we say that G is ℓ -tangle-free.*

Using those results, we are now able to prove Lemma 15:

Proof. From Lemma 26, we can deduce that if $d(i, j) \leq \ell$, there are at most two distinct paths between i and j . Therefore, $B_{ij}^{(\ell)} \leq 2$ for all i, j .

Additionally, if $D_{ij}^{(\ell)} = 1$, then there is a self-avoiding path of length ℓ between i and j , and thus $B_{ij}^{(\ell)} = 1$, so $\delta_{ij}^{(\ell)} \geq 0$ for all i, j .

Finally, assume that there exists a pair i, j such that $D_{ij}^{(\ell)} = 0$ and $B_{ij}^{(\ell)} = 2$; then there are two paths of length ℓ between i and j and $d(i, j) < \ell$ so there is also a path of length less than ℓ . This contradicts Lemma 26.

Consider now two vertices i and j such that $\delta_{ij}^{(\ell)} = 1$, there are two possibilities:

- (i) $D_{ij}^{(\ell)} = 0$ and $B_{ij}^{(\ell)} > 0$: then $d(i, j) < \ell$ and there is a path of length $< \ell$ and at least a path of length ℓ between i and j .
- (ii) $D_{ij}^{(\ell)} = 1$ and $B_{ij}^{(\ell)} > 1$: then there are at least two paths of length ℓ between i and j .

In both cases, there are at least two paths of length at most ℓ connecting i and j , which implies the statement of the lemma. \square

3.7 Proof of Lemma 16

Proof. (i) is obvious since $d(i, \mathcal{C}) + d(j, \mathcal{C}) \leq \ell$ implies $d(i, \mathcal{C}) \leq \ell$.

For (ii), note first that $V_{\mathcal{C}}$ and $V_{\mathcal{C}'}$ are disjoint for $\mathcal{C} \neq \mathcal{C}'$: if $i \in V_{\mathcal{C}} \cap V_{\mathcal{C}'}$, then \mathcal{C} and \mathcal{C}' are in the ℓ -neighbourhood of i , which contradicts Lemma 26.

Let $\pi_{\mathcal{C}}$ be the projection on $V_{\mathcal{C}}$ for all \mathcal{C} ; the $\pi_{\mathcal{C}}$ are mutually orthogonal and for a vector v , we have:

$$v^\top P^{(\ell)} v = \sum_{\mathcal{C}} v^\top \pi_{\mathcal{C}} P^{(\ell)} \pi_{\mathcal{C}} v = \sum_{\mathcal{C}} (\pi_{\mathcal{C}} v)^\top P_{\mathcal{C}}^{(\ell)} (\pi_{\mathcal{C}} v) \quad (3.38)$$

$$\leq \sum_{\mathcal{C}} \rho(P_{\mathcal{C}}^{(\ell)}) \cdot \|\pi_{\mathcal{C}} v\|^2 \quad (3.39)$$

$$\leq \max_{\mathcal{C}} \rho(P_{\mathcal{C}}^{(\ell)}) \cdot \sum_{\mathcal{C}} \|\pi_{\mathcal{C}} v\|^2. \quad (3.40)$$

On the other hand,

$$\|v\|^2 \geq \sum_{\mathcal{C}} \|\pi_{\mathcal{C}} v\|^2. \quad (3.41)$$

Combining inequalities (3.40) and (3.41) yields $\rho(P^{(\ell)}) \leq \max_{\mathcal{C}} \rho(P_{\mathcal{C}}^{(\ell)})$; the reverse inequality comes from the decomposition $P^{(\ell)} = \sum P_{\mathcal{C}}^{(\ell)}$. \square

3.8 Proof of Proposition 5

In the same vein as Lemma 25, for a vertex set \mathcal{X} , define $S_t(\mathcal{X})$ as the number of vertices at distance t of \mathcal{X} . By taking the union on all vertices of \mathcal{X} , we easily get the following corollary:

Corollary 4. *For the same constants C and ε as above, with probability $1 - O(n^{-\varepsilon})$, we have for all vertex subsets $\mathcal{X} \in \mathcal{P}(\{1, \dots, n\})$ and $\ell = O(\log(n))$:*

$$S_t(\mathcal{X}) \leq C \cdot |\mathcal{X}| \log(n) \cdot \alpha^t, \quad t \in \{1, \dots, \ell\}.$$

We are now able to prove Proposition 5:

Proof. Let \mathcal{K} be the modified vertex set, and consider vertices i and j such that $D_{ij}^{(\ell)} \neq \tilde{D}_{ij}^{(\ell)}$. Then we have one of four possibilities:

- (i) $\tilde{d}(i, j) = \ell$ and $d(i, j) < \ell$
- (ii) $\tilde{d}(i, j) > \ell$ and $d(i, j) = \ell$
- (iii) $\tilde{d}(i, j) = \ell$ and $d(i, j) > \ell$
- (iv) $\tilde{d}(i, j) < \ell$ and $d(i, j) = \ell$

In cases (i) and (ii), there is a path between i and j in G through \mathcal{K} of length at most ℓ , and in cases (iii) and (iv) there is a path between i and j in \tilde{G} through \mathcal{K} . Therefore, in all cases, we have that

$$d(i, \mathcal{K}) + d(j, \mathcal{K}) \leq \ell.$$

Write $|\tilde{D}^{(\ell)} - D^{(\ell)}|$ for the matrix whose (i, j) coefficient is $|\tilde{D}_{ij}^{(\ell)} - D_{ij}^{(\ell)}|$, and $P_{\mathcal{K}}$ for the matrix such that $P_{\mathcal{K},ij} = \mathbf{1}(d(i, \mathcal{K}) + d(j, \mathcal{K}) \leq \ell)$; the previous analysis and the Perron-Frobenius theorem imply that

$$\rho(\tilde{D}^{(\ell)} - D^{(\ell)}) \leq \rho(|\tilde{D}^{(\ell)} - D^{(\ell)}|) \leq \rho(P_{\mathcal{K}}). \quad (3.42)$$

We can then perform the same analysis as in the proof of Proposition 4 to find that the spectral radius of $P_{\mathcal{K}}$ is the same as that of

$$Q_{\mathcal{K}} = \begin{pmatrix} S_0 & \sqrt{S_0 S_1} & \cdots & \sqrt{S_0 S_{\ell-1}} & \sqrt{S_0 S_{\ell}} \\ \sqrt{S_0 S_1} & S_1 & \cdots & \sqrt{S_1 S_{\ell-1}} & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \sqrt{S_0 S_{\ell-1}} & \sqrt{S_1 S_{\ell-1}} & \cdots & 0 & 0 \\ \sqrt{S_0 S_{\ell}} & 0 & \cdots & 0 & 0 \end{pmatrix},$$

where we write S_t instead of $S_t(\mathcal{K})$ for ease of notation.

Corollary 4 then gives $S_t(\mathcal{K}) = O(\alpha^t \log(n) |\mathcal{K}|) = o(\alpha^t \tau^{\ell/2})$, and the same calculation as in Proposition 4 yields:

$$\rho(Q_{\mathcal{K}}) = o(\alpha^{\ell/2} \tau^{\ell/2}) = o(\mu_2^{\ell}), \quad (3.43)$$

and the theorem follows. \square

3.9 Proof of Theorem 19

In order to prove Theorem 19, we need to show that the controls in the proof of Theorem 18 are actually sharp. We begin with the following lemma, which comes from the fact that ℓ -neighbourhoods of the vertices of G are roughly of the same size:

Lemma 27. *Assume that $\gamma = \Theta(\tau^{\ell/2})$. Then there exists a set of vertices \mathcal{K} of size γ such that:*

$$S_{\ell}(\mathcal{K}) = \Omega(\alpha^{\ell} \cdot \gamma). \quad (3.44)$$

Proof. Let $\varepsilon > 0$ to be determined later, S be the set consisting of the $n^{1-\varepsilon}$ vertices i with the largest values $S_{\ell}(i)$; we first show that, for all $i \in S$

$$S_{\ell}(i) = \Theta(\alpha^{\ell}). \quad (3.45)$$

Indeed, from Lemma 25, we have the following inequalities:

$$Kn\alpha^{2\ell} \leq \sum_{i=1}^n S_{\ell}(i)^2 \leq n \min_{i \in S} S_{\ell}(i)^2 + |S|(C \log(n)\alpha^{\ell})^2, \quad (3.46)$$

and the second term is negligible before the two others, which implies (3.45).

We then build a set \mathcal{K} of size γ as follows: begin with any member of S , and at each step add a vertex x such that $d(x, \mathcal{K}) > 2\ell$. This is possible as long as the 2ℓ -neighbourhood of \mathcal{K} does not cover S , i.e. as long as:

$$\gamma \cdot C \log(n)\alpha^{2\ell} < n^{1-\varepsilon}. \quad (3.47)$$

But the LHS of this inequality is bounded by $C \log(n)n^{3/4}$, so this condition is satisfied as long as $\varepsilon < 1/4$.

By this construction, the vertices of \mathcal{K} have ℓ -neighbourhoods that are pairwise disjoint, so by equation (27) we have:

$$S_{\ell}(\mathcal{K}) = \Omega(\alpha^{\ell} \times \gamma). \quad (3.48)$$

\square

Consider now the vector v such that:

$$v_i = \begin{cases} \gamma^{-1/2} & \text{if } i \in \mathcal{K} \\ S_\ell(\mathcal{K})^{-1/2} & \text{if } d(i, \mathcal{K}) = \ell \\ 0 & \text{otherwise .} \end{cases} \quad (3.49)$$

The aim is to show the following equalities:

$$\frac{v^\top D^{(\ell)} v}{\|v\|^2} = \Omega(\mu_2^\ell) \quad \text{and} \quad \langle v, B^{(\ell)} \chi_k \rangle = o(\|v\| \|B^{(\ell)} \chi_k\|) \quad \forall k \in [r_0] \quad (3.50)$$

Indeed, Theorem 19 will then follow from a simple application of Courant-Fisher's Theorem.

Proof. (of Eq. (3.50)). We notice that $\|v\|^2 = 2$; furthermore:

$$\begin{aligned} v^\top D^{(\ell)} v &= \sum_{i,j} v_i D_{ij}^{(\ell)} v_j \\ &\geq 2 \sum_{i \in S_\ell(\mathcal{K})} \sum_{j \in \mathcal{K}} v_i v_j \\ &= 2\gamma S_\ell(\mathcal{K}) \gamma^{-1/2} S_\ell(\mathcal{K})^{-1/2} \\ &= 2\sqrt{\gamma S_\ell(\mathcal{K})} \\ &= \Omega(\mu_2^\ell), \end{aligned}$$

which proves the first inequality.

It remains then to prove that v is asymptotically orthogonal to $B^{(\ell)} \chi_k$ for $k \in [r_0]$: noticing that $v_i \leq 1$ for all i and $\|v\|_0 = \gamma + S_\ell(\mathcal{K})$, we find, using Corollary 4:

$$\begin{aligned} \langle v, B^{(\ell)} \chi_k \rangle &\leq (\gamma + S_\ell(\mathcal{K})) \cdot \|B^{(\ell)} \chi_k\|_\infty \\ &\leq (\gamma + S_\ell(\mathcal{K})) \cdot \tilde{O}(\alpha^\ell) \\ &= \tilde{O}(\gamma \alpha^{2\ell}) \\ &= o(\sqrt{n} \mu_k^\ell) \quad \text{since } \kappa < 1/4 \\ &= o(\|v\| \|B^{(\ell)} \chi_k\|), \end{aligned}$$

where we used part (ii) of proposition 10 to bound $\|B^{(\ell)} \chi_k\|_\infty$. □

Chapter 4

Non-backtracking spectrum of inhomogeneous random graphs

This chapter is based on the paper [116].

4.1 Introduction

Let $P \in \mathcal{M}_n(\mathbb{R})$ be a symmetric $n \times n$ matrix with entries in $[0, 1]$, and W a (symmetric) weight matrix with independent random entries. We define the inhomogeneous undirected random graph $G = (V, E)$ associated with the couple (P, W) as follows: the vertex set is simply $V = [n]$, and each edge $\{u, v\}$ is present in E independently with probability P_{uv} , and holds weight W_{uv} .

The entrywise expected value and variance of the weighted adjacency matrix of G are

$$\mathbf{E}A = P \circ \mathbf{E}W \quad \text{and} \quad \text{Var}(A) := P \circ \mathbf{E}[W \circ W] - P \circ P \circ \mathbf{E}W \circ \mathbf{E}W, \quad (4.1)$$

where \circ denotes the Hadamard product. When the entries of P are small, the second term of $\text{Var}(A)$ is negligible and the variance can be well approximated by the entrywise second moment; we thus define

$$Q := P \circ \mathbf{E}W \quad \text{and} \quad K := P \circ \mathbf{E}[W \circ W]. \quad (4.2)$$

A natural question, arising from matrix perturbation theory, is then as follows:

What is the relationship between the eigendecomposition of A and the one of Q ?

Unfortunately, at least in the unweighted case, when the mean degree of G is low ($o(\log(n))$), it is known that the largest eigenvalues (and associated eigenvectors) of A are determined by the large degree vertices; see [19] for a complete description of this phenomenon. To extract meaningful information on the spectrum of Q , another matrix has shown better performance: the non-backtracking matrix, whose application to community detection has been studied in [80, 29].

Given a weighted graph G , we define its associated non-backtracking matrix B as follows: B is a $2|E| \times 2|E|$ matrix indexed by the oriented edges of G , whose coefficients are

$$B_{ef} = W_f \mathbf{1}\{e \rightarrow f\} = W_f \mathbf{1}\{e_2 = f_1\} \mathbf{1}\{e_1 \neq f_2\},$$

where $e = (e_1, e_2)$ and $f = (f_1, f_2)$. The above question rephrases in our setting as

What is the relationship between the eigendecomposition of B and the one of Q ?

and the main focus of this article is to provide an answer as precise as possible to this problem. To this end, let

$$Q = \sum_{i=1}^r \mu_i \varphi_i \varphi_i^\top \quad \text{with} \quad |\mu_1| \geq |\mu_2| \geq \cdots \geq |\mu_r|$$

be the eigendecomposition of Q , and $\rho = \rho(K)$ the largest eigenvalue (in absolute value) of K . Note that by definition, Q and K are symmetric and therefore all eigenvalues defined above are real.

We shall assume that there exists some deterministic bound L (possibly depending on n) such that $\max |W_{ij}| \leq L$. We can then state our main theorem, without detailing the needed hypotheses for now:

Theorem 21 (informal statement). *Assume the following conditions:*

- (i) $r = n^{o(1)}$,
- (ii) the graph G is sparse enough,
- (iii) the eigenvectors of Q are sufficiently delocalized.

Let r_0 be the number of eigenvalues of Q whose absolute value is larger than both $\sqrt{\rho}$ and L :

$$\mu_k > \sqrt{\rho} \vee L \text{ for all } k \in [r_0] \quad \text{and} \quad \mu_{r_0+1} \leq \sqrt{\rho} \vee L \quad (4.3)$$

Then, for $i \leq r_0$, the i -th largest eigenvalue of B is asymptotically (as n goes to infinity) equal to μ_i , and all the other eigenvalues of B are constrained in a circle of center 0 and radius $\max(\sqrt{\rho}, L)$. Further, if $i \leq r_0$ is such that μ_i is a sufficiently isolated eigenvalue of Q , then the eigenvector associated with the i -th eigenvalue of B is correlated to a lifted version of φ_i .

Next section consists in the detailed statement of this theorem (with precise hypotheses and bounds given).

4.2 Detailed setting and results

4.2.1 Notations

General notations: Throughout this paper, we use the following notations:

- for integer n , $[n]$ denotes the set $\{1, \dots, n\}$.
- for $x \in \mathbb{R}^n$, we shall denote by x_i or $x(i)$ the i -th coordinate of x , whichever is most convenient. $\|x\|$ is the 2-norm of x , and $\|x\|_\infty$ the infinity norm of x .
- the operator norm of a matrix M is noted $\|M\|$; it is the maximal singular value of M . Its Frobenius norm is noted $\|M\|_F$ and its infinity norm $\|M\|_\infty = \sup_{i,j} |M_{ij}|$.
- $\mathbf{1}$ denotes the all-one vector, and $\mathbf{1}\{\cdot\}$ is the indicator function of an event.
- the group of permutations on r elements is noted \mathfrak{S}_r .
- the max (resp. min) of two numbers a, b is noted $a \vee b$ (resp. $a \wedge b$).
- the letter c denotes any absolute constant, whose value should be assumed to be the maximum of any such constant encountered so far. To improve the readability of our computations, we use numbered constants c_i during proofs.

Graph theoretic notations: For a graph $g = (V, E)$, let \vec{E} be the set of oriented edges in E , and

$$\vec{E}(V) = \{(u, v) \mid u \neq v \in V\}$$

be the set of all directed edges of the *complete graph* on V . If t is an integer, $g = (V, E)$ is a graph and $x \in V$, then the *ball* $(g, x)_t$ is the subgraph induced by all edges at distance at most t from x , and $\partial(g, x)_t$ is the boundary of the ball, i.e. the set of vertices at distance exactly t from x . Finally, the set of all non-backtracking paths of length t starting with x will be denoted $\mathcal{P}_g(x, t)$.

Non-backtracking matrix: Since we are interested in the spectrum of the non-backtracking matrix B , we need to be able to translate “vertex” quantities such as the vectors φ_i into “edge” quantities. Recall that $V = [n]$, and identify \vec{E} with the set $[2m]$; we define the $2m \times n$ *start* and *terminal* matrices S and T as

$$\forall e \in \vec{E}, i \in [n], \quad S_{ei} = \mathbf{1}\{e_1 = i\} \quad \text{and} \quad T_{ei} = \mathbf{1}\{e_2 = i\}.$$

For a vector $\phi \in \mathbb{R}^n$, this implies that $[T\phi](e) = \phi(e_2)$ for every edge $e \in \vec{E}$. We then define the “lifted” eigenvectors $\chi_i = T\varphi_i$ for $i \in [r]$.

We also define the *reverse* operator J such that $Je = \bar{e} := (e_2, e_1)$, and the diagonal matrix D_W such that $D_W(e, e) = W_e$; from the definition of B and symmetry of W it is straightforward to see that $JD_W = D_WJ$ and for all $t \geq 0$

$$JD_W B^t = (B^*)^t D_W J, \tag{4.4}$$

which is known in mathematical physics as parity-time invariance. For any vector $x \in \mathbb{R}^{\vec{E}}$, we denote the vector Jx by \check{x} .

Building upon the sketch in the introduction, we now expand on the model definition. Recall that the *expectation* and *variance* matrices were defined as

$$Q = P \circ \mathbf{E}W \quad \text{and} \quad K = P \circ \mathbf{E}[W \circ W].$$

4.2.2 Defining the convergence parameters

In full generality, with no assumptions on P and W , we do not expect meaningful results to hold; however, we are still able to provide interesting properties on a large class of matrices. We define in the following the parameters that will govern the convergence behavior :

(i) the *rank*

$$r = \max \left(\text{rank}(Q), \sqrt{\text{rank}(K)} \right);$$

note that in most practical applications (such as the unweighted case), we shall have $r = \text{rank}(Q)$, but we also treat cases where $r \gg \text{rank}(Q)$.

(ii) the *sparsity* parameter

$$d = n \max_{i, j \in [n]} P_{ij};$$

(iii) the *eigenvector delocalization* parameter

$$b = \sqrt{n} \max_{i \leq \text{rank}(Q)} \|\varphi_i\|_\infty;$$

(iv) the signal-to-noise ratio

$$\tau = \max_{\mu_i^2 > \mu_1} \frac{\mu_1}{\mu_i^2};$$

(v) and finally the almost sure probability bound

$$\|W\|_\infty = L;$$

our results hold trivially whenever $L = +\infty$ so we shall restrict ourselves to the case where L is finite, and the W_{ij} are almost surely bounded. While Theorem 22 below requires an almost sure bound, techniques for dealing with high probability bounds are discussed in Theorem 26.

The average degree of a vertex i will be noted by

$$d_i = \sum_{j \in [n]} P_{ij} \leq d,$$

which corresponds to the entries of the vector $P\mathbf{1}$. To ensure that G is connected enough for spectral properties to hold, we make the (common) assumption that $d_i \geq 1$ for all $i \in [n]$. The entries of $K\mathbf{1}$ can be viewed as an extension of the average degrees in the weighted case (see [11] or [20] for examples), and for the same reason as above we require that $K\mathbf{1}$ is bounded away from zero by a constant.

4.2.3 Main theorem

In the following, $G = G(P, W)$ is the random graph defined in the introduction, B is the non-backtracking matrix associated with G , and $|\lambda_1| \geq \dots \geq |\lambda_{2m}|$ are its eigenvalues.

In its most general form, our main result is as follows:

Theorem 22. *Let $n \geq 0$ and (P, W) be a couple of $n \times n$ matrices defining a random graph G . Define $\rho = \rho(K)$, r_0 as in (4.3), and r, b, d, τ, L as in Subsection 4.2.2.*

Let

$$\ell = \frac{1 - \epsilon}{4} \frac{\log(n)}{\log(d^5(1 \vee \tilde{L})^2)},$$

for arbitrary $\epsilon > 0$, and where $\tilde{L} = L/\mu_1$. There exist numbers n_0 and C_0 , all depending on n and the convergence parameters, such that the following holds:

(i) C_0 is smaller than

$$c \left(\frac{rbd\tilde{L}\log(n)}{1 - \tau} \right)^{25},$$

and n_0 is smaller than

$$\exp \left(c \frac{\max \left\{ \log(r), \log(b), \log(d)^2, \log(\tilde{L})^2, \log(\log(n)) \right\}}{\log(\tau^{-1})} \right).$$

(ii) If $n \geq n_0$, define

$$\sigma := C_0 \mu_1 \tau^{\ell/2}. \tag{4.5}$$

Then the following holds with probability at least $1 - c/\log(n)$ there exists a permutation s of $[r_0]$ such that

$$\max_{i \in [r_0]} |\lambda_i - \mu_{s(i)}| \leq \sigma, \quad (4.6)$$

and all the remaining eigenvalues of B are less than $C_0^{1/\ell} (\sqrt{\rho} \vee L)$.

(iii) For any $i \in [r_0]$, if

$$\delta_i := \max_{j \neq s(i)} |\mu_{s(i)} - \mu_j| \geq 2\sigma, \quad (4.7)$$

then there exists a normed eigenvector ξ associated with λ_i such that

$$\langle \xi, \xi_i \rangle \geq \sqrt{1 - rd^2 \tilde{L}^2 \frac{\rho}{\mu_i^2}} + O\left(\frac{\sigma}{\delta_i - \sigma}\right) \quad \text{where} \quad \xi_i = \frac{T\varphi_i}{\|T\varphi_i\|}.$$

In order to get an applicable and useful result, we need $n \geq n_0$ when n is sufficiently large, and $C_0^{1/\ell}$ goes to 1 as n goes to infinity. Both conditions are verified in particular when

$$1 - \tau = \Omega(1), \quad r, b = n^{o(1)} \quad \text{and} \quad \log(d)^2 = o(\log(n)).$$

By definition of \tilde{L} , whenever $\tilde{L} > 1$ we have $\mu_1 < L$ and thus $r_0 = 0$. We can therefore safely assume $\tilde{L} \leq 1$ in applications and not focus on any bound for L .

The proof of this theorem follows the same method as in many spectral proofs, from [91] to more recent papers such as [28]. It consists of the following:

- show that the neighbourhood of any vertex v is close to a suitably defined random tree,
- study a family of graph functionals that give rise to approximate eigenvectors of the random tree,
- use a concentration argument to transpose those tree eigenvectors to pseudo-eigenvectors of the non-backtracking matrix,
- bound the remaining eigenvalues using a variant of the trace method in [58],
- conclude by a matrix perturbation argument.

A large portion of the remainder of this paper is dedicated to implementing this method; however, we first provide several applications of our result to the fields of random matrix theory and random graph theory.

4.3 Applications

4.3.1 Phase transition in random graphs

Matrix perturbation theory focuses on finding the eigenvalues and eigenvectors of matrices of the form $X + \delta$, where X is a known matrix and δ is a perturbation assumed “small” in a sense. Celebrated results in this field include the Bauer-Fike theorem [18] for asymmetric matrices, and the Weyl [125] and Davis-Kahan [126] theorems for symmetric ones; incidentally the present paper makes use of those results in its proofs. Finding sharp general theorems without additional assumptions is known to be hard, since the eigenvalues and eigenvectors depend on the interactions between the eigenspaces of X and δ .

In the last two decades, growing attention has been paid to problems of the following form: finding the eigenvectors of $X_n + P_n$ (or, in its multiplicative form, $X_n(I_n + P_n)$), where P_n

is an $n \times n$ matrix with low rank $r \ll n$ (usually fixed) and known eigenvalues, and X_n is a random matrix with known distribution. Examples of this setting are the spiked covariance model [13, 73] and additive perturbations of Wigner matrices [107, 56, 34]. A more systematic study has been performed in [21, 22] on orthogonally invariant random matrices.

A staple of those results is the existence of a so-called *BBP phase transition* (named after Baik-Ben Arous-Péché, from the seminal article [13]): in the limit $n \rightarrow \infty$, each eigenvalue of P_n that is above a certain threshold gets reflected (albeit perturbed) in the spectrum of $X_n + P_n$, with the associated eigenvector correlated to the one of P_n .

Phase transition for the adjacency matrix The adjacency matrix A of our random graph G can be viewed as a perturbation model by writing

$$A = \mathbf{E}A + (A - \mathbf{E}A) = Q - \text{diag}(Q) + (A - \mathbf{E}A).$$

The term $\text{diag}(Q)$ being negligible with respect to the others, we can see A as the sum of a deterministic low-rank matrix and a random noise matrix with i.i.d centered entries. Further, the entrywise variance of A is equal (up to a negligible term) to K , so the parameter ρ can be seen as an equivalent to the variance in the Wigner model. We thus expect, whenever $\sqrt{\rho} \gg L$ (so that $\sqrt{\rho}$ is the actual threshold in Theorem 22), to find a phase transition akin to the one in [21]; and indeed the following theorem holds:

Theorem 23. *Let (P, W) be a matrix couple of size $n \times n$ and r, b, d, τ, L as above. Assume further that:*

- (i) *the Perron-Frobenius eigenvector of K is $\mathbf{1}$; that is $K\mathbf{1} = \rho\mathbf{1}$,*
- (ii) *the above eigenvector equation concentrates, i.e. with high probability there exists $\varepsilon \leq 1/2$ such that for all $i \in [n]$,*

$$\left| \sum_{j \sim i} W_{ij}^2 - \rho \right| \leq \varepsilon \rho \quad (4.8)$$

Then, if $i \in [r_0]$ is such that $\mu_i^2 \geq 2L^2$, there exists an eigenvalue ν_i of A that verifies

$$\nu_i = \mu_i + \frac{\rho}{\mu_i} + \frac{\rho}{\mu_i} \cdot O\left(\frac{L}{\mu_i} + \frac{L^2}{\mu_i^2} + \varepsilon\right). \quad (4.9)$$

Further, if the mean degree d_j for all j is equal to $d_0 > 1$, and i is such that $\delta_i \geq 2\sigma$ (with σ and δ_i defined in (4.5) and (4.7)), then there exists a normed eigenvector ζ of A with corresponding eigenvalue ν_i such that

$$\langle \zeta, \varphi_i \rangle = \sqrt{1 - \frac{\rho}{\mu_i^2}} + O\left[\frac{1}{\delta_i - \sigma} \left(\frac{L\rho}{\mu_i^2} + \frac{L^2\rho}{\mu_i^3} + \varepsilon\frac{\rho}{\mu_i}\right)\right]. \quad (4.10)$$

Whenever $\rho \gg L^2$, and ε goes to zero as $n \rightarrow \infty$, then the condition $\mu_i^2 \geq 2L^2$ is always verified and the $O(\cdot)$ term in (4.9) vanishes, and the obtained expansion is therefore asymptotically correct. The presence of δ_i renders a similar result on the scalar product harder to obtain; however, assuming $\delta_i = \Theta(\sqrt{\rho})$ (that is, the eigenvalues of Q are somewhat regularly spaced) implies similarly that the $O(\cdot)$ term in (4.10) vanishes.

The obtained expression for ν_i , as well as the scalar product expansion, are identical to the ones in [21], for low-rank additive perturbations of Gaussian Wigner matrices. Our result is thus a direct extension of [21], for a larger class of matrices upon a sparsity and concentration

condition. Such an extension isn't unexpected, in view of results concerning the universality of the semicircle law for Bernoulli random matrices, such as [53].

An especially interesting particular case of Theorem 23 is the unweighted random graph setting, where $W_{ij} = 1$ for all i, j . In this case, we have $K = P$ so the eigenvector equation $K\mathbf{1} = \rho\mathbf{1}$ is equivalent to all the average degrees being equal, i.e. $d_i = d_0 = \rho$ for $i \in [n]$. It is a well known fact (see for example [55]) that for unweighted random graphs the degree concentration property holds with $\varepsilon = 2\sqrt{\log(n)/d_0}$. A slight modification of the proof of Theorem 23 further removes several error terms, and the following corollary ensues:

Corollary 5. *Let P be a $n \times n$ matrix and r, b, d, τ as above, with $W = \mathbf{1}^*\mathbf{1}$. Assume further that for all $i \in [n]$,*

$$\sum_{j \in [n]} P_{ij} = d_0 > 16 \log(n).$$

Then for all $i \in [r_0]$, there exists an eigenvalue ν_i of A that verifies

$$\nu_i = \mu_i + \frac{d_0}{\mu_i} + O\left(\sqrt{\frac{\log(n) d_0}{d_0 \mu_i}}\right),$$

and if i is such that $\delta_i > 2\sigma$, there exists a normed eigenvector of A such that

$$\langle \zeta, \varphi_i \rangle = \sqrt{1 - \frac{d_0}{\mu_i^2}} + O\left(\frac{1}{\delta_i - \sigma} \sqrt{\frac{\log(n) d_0}{d_0 \mu_i}}\right).$$

In particular we have

$$\nu_1 = d_0 + 1 + O\left(\sqrt{\frac{\log(n)}{d_0}}\right)$$

This is an improvement on the results of [20], which only give $\nu_i = \mu_i + O(\sqrt{d_0})$. The condition $d_0 > 16 \log(n)$ ensures that the degrees of G concentrate. Since our result is really only meaningful whenever $d_0 \gg \log(n)$, so that the error term is negligible before d_0/μ_i , we do not perform the same detailed analysis as in [11]. However, a more precise phase transition around $d_0 \asymp \log(n)$ is not excluded.

Theorem 23 is derived from Theorem 22 through an adaptation of the Ihara-Bass formula [17], obtained by expanding arguments from [19, 124]:

Proposition 11. *Let x be an eigenvector of the matrix B with associated eigenvalue λ , such that $\lambda^2 \neq W_{ij}^2$ for every i, j . Define the weighted adjacency matrix $\tilde{A}(\lambda)$ and the diagonal degree matrix $\tilde{D}(\lambda)$ by*

$$\tilde{A}(\lambda)_{ij} = \mathbf{1}\{i \sim j\} \frac{\lambda W_{ij}}{\lambda^2 - W_{ij}^2} \quad \text{and} \quad \tilde{D}(\lambda)_{ii} = \sum_{j \sim i} \frac{W_{ij}^2}{\lambda^2 - W_{ij}^2}$$

Then the vector $y = S^ D_W x$ is a null vector of the laplacian matrix*

$$\delta(\lambda) = \mathbf{I} - \tilde{A}(\lambda) + \tilde{D}(\lambda).$$

The details and computations are left to the appendix.

4.3.2 Community detection in random networks

Community detection is a clustering problem that aims to identify large subgroups (or communities) with similar characteristics inside a large population, with the only data available being the pairwise interactions between individuals. Starting from its introductory paper [69], the stochastic block model has been a popular generative model for algorithm design; it consists of a random graph G where vertices are partitioned randomly in communities, and edges are present independently with probability depending only on the community membership of their endpoints. Popular algorithms for recovering communities include semi-definite programming methods [97], belief propagation [7], and spectral methods [82, 91]; a comprehensive review of algorithms and results can be found in [1].

Unlabeled stochastic block model In a general form, we can define the stochastic block model $\text{SBM}(n, r, \theta, M)$, where $\theta \in [r]^n$ and $M \in [0, 1]^{r \times r}$ as follows:

- the vertex set is $V = [n]$,
- each vertex $i \in [n]$ has a community label θ_i in $[r]$,
- for any pair of vertices (i, j) , an edge is present between i and j independently from the others with probability $M_{\theta_i \theta_j}$.

It is common to assume $M = \frac{\alpha}{n} M_0$, where M_0 does not depend on n and α is a scaling parameter. It is easy to see that up to diagonal terms, the expected adjacency matrix has the form

$$P = \Theta M \Theta^*,$$

where Θ is a $n \times r$ matrix such that $\Theta_{ij} = 1$ if $\theta_i = j$, and 0 otherwise. We shall assume that for any $k \in [r]$,

$$\frac{\#\{i \in [n] \mid \theta_i = k\}}{n} = \pi_k > 0, \quad (4.11)$$

where π is a deterministic probability vector. Let $\mu_1 \geq \dots \geq |\mu_r|$ the eigenvalues of $\text{diag}(\pi)M_0$, with α chosen such that $|\mu_1| = 1$, and ϕ_1, \dots, ϕ_r the associated eigenvectors. Then the non-zero eigenvalues of P are easily found to be the $\alpha\mu_i$, with associated eigenvectors $\Theta\phi_i$.

A common assumption is that each vertex type has the same average degree, i.e.

$$P\mathbf{1} = \alpha\mathbf{1},$$

otherwise a simple clustering based on vertex degree correlates with the underlying communities. Making this additional assumption, the following theorem holds:

Theorem 24. *Assume that r is constant, and $\alpha = n^{o(1)}$. Let r_0 be defined as follows :*

- if $\alpha \geq 1$ is constant, r_0 is the only integer in $[r]$ such that

$$\alpha\mu_k^2 > 1 \quad \text{for } i \in r_0, \quad \alpha\mu_{r_0+1}^2 \leq 1.$$

- if $\alpha = \omega(1)$, $r_0 = r$.

Then, for any n larger than an absolute constant and all $i \in [r_0]$ one has

$$|\lambda_i - \mu_i| \leq c(\alpha \log(n))^a (\alpha\mu_{r_0})^{-\kappa \log_\alpha(n)} := \sigma$$

for some positive constants c, a, κ , and all other eigenvalues of B are confined in a circle with radius $(1 + o(1))\sqrt{\alpha}$. Further, if μ_i is an isolated eigenvalue of $\text{diag}(\pi)M_0$, then there exists an eigenvector ξ of the non-backtracking matrix B associated with λ_i such that

$$\langle \xi, \xi_i \rangle \geq \sqrt{1 - \frac{1}{\alpha\mu_i^2}} + O(\sigma') \quad \text{where} \quad \xi_i = \frac{T\Theta\phi_i}{\|T\Theta\phi_i\|}.$$

This theorem is essentially a corollary of Theorem 22, with some simplifications due to $Q = K = P$ and $P\mathbf{1} = \alpha\mathbf{1}$; the error bound σ is the same as in the main theorem. It is a direct generalization of Theorem 4 in [29], for a diverging degree sequence; further, the property $\langle \xi, \xi_i \rangle = 1 - o(1)$ as soon as $\alpha \gg 1$ suggests that a clustering algorithm such as k -means performed on the eigenvectors of B recovers all but a vanishing fraction of the community memberships in this regime, which would provide an alternative to the Sphere-comparison algorithm presented in [5].

Conjecture 3. *In the SBM defined as above, as soon as $\alpha = \omega(1)$, running an approximate k -means algorithm on the top r eigenvectors of B allows to recover the community memberships of every vertex but a vanishing fraction as $n \rightarrow \infty$.*

Proving this conjecture would require a more careful eigenspace analysis for eigenvalues with multiplicity more than one, such as the one performed in [117], as well as an error bound on the clustering step similar to the one in [82]. Note that in this setting the theoretical covariance matrices defined in (4.24) are diagonal, and the eigenvectors of B are therefore asymptotically orthogonal, which can greatly simplify the perturbation analysis of Section 4.4.

Labeled block models In real-world networks, pairwise interactions often carry more information than just a binary one. A popular variant of the stochastic block model is thus a model with added edge labels, as follows: let \mathcal{L} be a label space, and consider a SBM drawn under the model described above. We assign to an edge (i, j) a label $L_{ij} \in \mathcal{L}$, drawn independently from a distribution $\mathbb{P}_{\theta_i\theta_j}$. Such classes of models have been investigated in full generality in [67, 84], and a variant with the underlying graph being an Erdős-Rényi model in [113].

We shall focus here on the symmetric two-community SBM, with

$$\pi = \left(\frac{1}{2}, \frac{1}{2} \right), \quad M = \begin{pmatrix} a & b \\ b & a \end{pmatrix}, \quad \mathbb{P}_{11} = \mathbb{P}_{22} = \mathbb{P} \quad \text{and} \quad \mathbb{P}_{12} = \mathbb{P}_{21} = \mathbb{Q}, \quad (4.12)$$

and assume that both measures are absolutely continuous with respect to another measure m (note that we can take $m = \mathbb{P} + \mathbb{Q}$), with Radon-Nikodym derivatives f and g . Let $w : \mathcal{L} \rightarrow \mathbb{R}$ a bounded weight function, such that $w(\ell) \leq L$ for any $\ell \in \mathcal{L}$; and define the weight matrix $W_{ij} = w(L_{ij})$ and the associated weighted non-backtracking matrix B . Then, an application of Theorem 22 yields the following result:

Theorem 25. *Define the parameter τ by*

$$\tau = 2 \frac{(a\mathbf{E}_{\mathbb{P}}[w^2] + b\mathbf{E}_{\mathbb{Q}}[w^2]) \vee L}{(a\mathbf{E}_{\mathbb{P}}[w] - b\mathbf{E}_{\mathbb{Q}}[w])^2}$$

Then, whenever $\tau < 1$, let ξ be a normed eigenvector corresponding to the second eigenvalue of B . There exists a parameter

$$\sigma \leq (a \log(n))^{25} \tau^{\kappa \log_a(n)}$$

for some constant κ such that

$$\langle \xi, \xi_0 \rangle = \sqrt{1 - \tau} + O(\sigma) \quad \text{where} \quad \xi_0 = \frac{\Theta \begin{pmatrix} 1 \\ -1 \end{pmatrix}}{\sqrt{n}}$$

Whenever this result holds, a proof identical to the one in [91] implies that recovering a positive fraction of the community memberships is possible.

In order to maximize the region in which reconstruction is possible, we need to choose the weights $w(\ell)$ such that τ is minimized. This optimization step is performed in the appendix, and leads to the following:

Proposition 12. *Define the weight function w and signal-to-noise ratio β as*

$$w(\ell) = \frac{af(\ell) - bg(\ell)}{af(\ell) + bg(\ell)} \quad \text{and} \quad \beta = \frac{1}{2} \int \frac{(af - bg)^2}{af + bg} dm, \quad (4.13)$$

where a, f, b, g and m are defined in Equation (4.12) and below. Then, whenever $\beta > 1$, a spectral algorithm based on the matrix B is able to recover a positive fraction of the community memberships when $n \rightarrow \infty$.

This settles a conjecture of [67], generalizing the setting from finite to arbitrary label space. Whenever we allow for a higher number of communities, as well as arbitrary choices for the connectivity matrix Q and distributions \mathbb{P}_{ij} , the problem proves to be harder; an analog to Theorem 25 does hold, but the optimization problem required to minimize the ratio τ looks to be untractable. In the symmetric SBM case, where

$$\pi = \frac{\mathbf{1}}{k}, \quad M = a\mathbf{1}\{i = j\} + b\mathbf{1}\{i \neq j\} \quad \text{and} \quad \mathbb{P}_{ij} = \mathbb{P}\mathbf{1}\{i = j\} + \mathbb{Q}\mathbf{1}\{i \neq j\},$$

we make the following conjecture:

Conjecture 4. *In the labeled symmetric SBM, partial reconstruction is possible as soon as $\beta > 1$, where*

$$\beta = \frac{1}{k} \int \frac{(af - bg)^2}{af + (k-1)bg} dm,$$

and a spectral algorithm based on the non-backtracking matrix with weight function

$$w(\ell) = \frac{af(\ell) - bg(\ell)}{af(\ell) + (k-1)bg(\ell)}$$

recovers a positive fraction of the community memberships in polynomial time.

As with Theorem 24, whenever the mean degree α of the graph grows to infinity, we have $\langle \xi, \xi_0 \rangle = 1 - o(1)$, which brings us our second conjecture:

Conjecture 5. *If we have $a = \alpha a_0$, $b = \alpha b_0$ with $\alpha = \omega(1)$, a_0, b_0 fixed, then as $n \rightarrow \infty$ a clustering algorithm based on the second eigenvector of the weighted non-backtracking matrix B with the weight function defined in (4.13) recovers all but a vanishing fraction of the community memberships.*

As a final remark, note that the optimal weight function assumes perfect knowledge of all model parameters, especially the exact label distribution for each community pair. However, in some cases, this weight function is a rescaling of a more agnostic one; as an example, in the censored block model [2] we find that $w(\ell) = c\ell$ (with $\ell = \pm 1$), and thus the spectral algorithm mentioned here is the same as in [113].

4.3.3 Extension to gaussian weights

In the form presented in Theorem 22, our result is only meaningful with almost surely bounded random variables (i.e. with $L < \infty$). With a more careful analysis of the error bounds, this can be extended to

$$L = \sup_{i,j \in [n]} \sup_k \mathbf{E}[W_{ij}^k]^{1/k}; \quad (4.14)$$

however, we determined the class of distributions satisfying (4.14) was not different enough from the bounded case to warrant increasing the complexity of the proof.

To the contrary, the setting where the W_{ij} are gaussian random variables is of independent interest; it can be seen as a special case of noisy matrix completion as described in [33, 78]. In this case, the moment condition of (4.14) is far from satisfied, and at least at first glance our proof cannot be adapted readily. Still, we show the following:

Theorem 26. *Assume that the $W_{ij} \sim \mathcal{N}(m_{ij}, s_{ij}^2)$ are independent Gaussian random variables, and let*

$$m = \sum_{i,j} m_{ij} \quad \text{and} \quad s = \sup_{i,j} s_{ij}.$$

Then the conclusions of Theorem 22 apply with

$$L = m + 2s\sqrt{\log(n)}$$

The loss of a $\sqrt{\log(n)}$ factor comes from the use of a concentration bound for the W_{ij} ; details can be found in the appendix.

To the best of our knowledge, there isn't much literature to compare with on the topic of eigenvalue reconstruction for noisy matrix completion, the works cited above being focused on reconstructing the whole matrix. However, results on gaussian matrix perturbation such as [21] seem to indicate that the $\sqrt{\log(n)}$ factor is superfluous and can be improved upon with other methods.

4.4 A Bauer-Fike type bound for almost orthogonal diagonalization

One important tool in tying together the local analysis of G is a matrix perturbation theorem, derived from the Bauer-Fike theorem. It mostly consists in a simplification and adaptation of Theorem 8.2 in [28], tailored to our needs. We begin by recalling the original Bauer-Fike Theorem:

Theorem 27 (Bauer-Fike Theorem [18]). *Let D be a diagonalizable matrix, such that $D = V^{-1}\Lambda V$ for some invertible matrix V and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$. Let E be any matrix of size $n \times n$. Then, any eigenvalue μ of $D + E$ satisfies*

$$|\mu - \lambda_i| \leq \|E\| \kappa(V), \quad (4.15)$$

for some $i \in [n]$, where $\kappa(V) = \|V\| \|V^{-1}\|$ is the condition number of V .

Let R be the RHS of (4.15), and $C_i := \mathcal{B}(\lambda_i, R)$ the ball centered at λ_i with radius R (in \mathbb{C}). Let $\mathcal{I} \subseteq [n]$ be a set of indices such that

$$\left(\bigcup_{i \in \mathcal{I}} C_i \right) \cap \left(\bigcup_{i \notin \mathcal{I}} C_i \right) = \emptyset.$$

Then the number of eigenvalues of $D + E$ in $\bigcup_{i \in \mathcal{I}} C_i$ is exactly $|\mathcal{I}|$.

4.4.1 A custom perturbation lemma for almost diagonalizable matrices

Building on this theorem, we now expose this section's first result. Let $U = (u_1, \dots, u_r)$ and $V = (v_1, \dots, v_r)$ be $n \times r$ matrices; our nearly diagonalizable matrix shall be $S = U\Sigma V^*$ with $\Sigma = \text{diag}(\theta_1, \dots, \theta_r)$. We shall assume that the θ_i are in decreasing order of modulus:

$$|\theta_r| \leq |\theta_{r-1}| \leq \dots \leq |\theta_1| = 1.$$

Now, let A be a $n \times n$ matrix, not necessarily diagonalizable. The assumptions needed for our results are as follows:

- (i) For some small constant $\varepsilon > 0$,

$$\|A - S\| \leq \varepsilon.$$

- (ii) The matrices U and V are well-conditioned: both U^*U and V^*V are nonsingular, and there exist two constants $\alpha, \beta > 1$ such that

$$\begin{aligned} \|U^*U\| &\leq \alpha, & \|V^*V\| &\leq \alpha, \\ \|(U^*U)^{-1}\| &\leq \beta, & \|(V^*V)^{-1}\| &\leq \beta. \end{aligned}$$

- (iii) There exists another constant $0 < \delta < 1$ such that

$$\|U^*V - I_r\|_\infty \leq \delta.$$

- (iv) The θ_i are well-separated from 0, in the sense that

$$|\theta_r| > 2\sigma := 2 \times 84r^2\alpha^{7/2}\beta(\varepsilon + 5r\alpha^2\beta\delta). \quad (4.16)$$

This definition may seem obscure, but shall happen naturally in the proof of the theorem.

Then the following result, whose statement and proof (regarding the eigenvalue perturbation) are adapted from [28], holds:

Theorem 28. *Let A be a matrix satisfying assumptions (i)-(iv) above, and let $|\lambda_1| \geq |\lambda_2| \geq \dots \geq |\lambda_r|$ be the r eigenvalues of A with greater modulus. There exists a permutation π such that for all $i \in [r]$*

$$|\lambda_{\pi(i)} - \theta_i| \leq r \times \sigma = 84r^3\alpha^{7/2}\beta(\varepsilon + 5r\alpha^2\beta\delta),$$

and the other $n - r$ eigenvalues of A all have modulus at most σ . Additionally, if i is such that

$$B(\theta_i, \sigma) \cap \left(\bigcup_{j \neq i} B(\theta_j, \sigma) \right) = \emptyset, \quad (4.17)$$

then there exists a normed eigenvector ξ associated with $\lambda_{\pi(i)}$ such that

$$\left\| \xi - \frac{u_i}{\|u_i\|} \right\| \leq \frac{3\sigma}{\delta_i - \sigma},$$

where δ_i is the minimum distance from θ_i to another eigenvalue:

$$\delta_i = \min_{j \neq i} |\theta_j - \theta_i| \geq 2\sigma.$$

Proof. We begin with defining an alternative matrix \bar{U} such that $\bar{U}^*V = I_r$. Let H_i be the subspace of \mathbb{R}^n such that

$$H_i = \text{vect}(v_j \mid j \neq i),$$

and consider the vectors \tilde{u}_i and \bar{u}_i defined as

$$\tilde{u}_i = u_i - P_{H_i}(u_i) \quad \text{and} \quad \bar{u}_i = \frac{\tilde{u}_i}{\langle \tilde{u}_i, v_i \rangle}$$

with P_{H_i} the projection on H_i , and \tilde{U}, \bar{U} the associated $n \times r$ matrices. Then it is straightforward to see that

$$\langle \bar{u}_i, v_i \rangle = 1 \quad \text{and} \quad \langle \bar{u}_i, v_j \rangle = 0,$$

for all $j \neq i$, which shows that $\bar{U}^*V = I_r$. Now, if we let V_i be the matrix V with the i -th column and line deleted,

$$P_{H_i} = V_i(V_i^*V_i)^{-1}V_i^*,$$

and

$$\|V_i^*u_i\|^2 = \sum_{j \neq i} \langle v_j, u_i \rangle^2 \leq r\delta^2,$$

hence we can compute $\|u_i - \tilde{u}_i\|$:

$$\|u_i - \tilde{u}_i\| = \|P_{H_i}(u_i)\| \leq \|V_i\| \|(V_i^*V_i)^{-1}\| \|V_i^*u_i\|,$$

and by the interlacing theorem $\|V_i\| \leq \sqrt{\alpha}$ and $\|(V_i^*V_i)^{-1}\| \leq \beta$ since V_i is a principal submatrix of V . Using the fact that $\|M\| \leq \|M\|_F$ for any matrix M , we find

$$\|U - \tilde{U}\| \leq r^2\sqrt{\alpha}\beta\delta.$$

For the second part, note that by the Cauchy-Schwarz inequality,

$$\begin{aligned} |\langle \tilde{u}_i, v_i \rangle - 1| &\leq |\langle u_i, v_i \rangle - 1| + \|u_i - \tilde{u}_i\| \cdot \|v_i\| \\ &\leq \delta(1 + r\alpha\beta), \end{aligned}$$

with the (generous) inequality $\|v_i\| \leq \|V\|$ used in the last line. Whenever δ is small enough, we can use the inequality $\left| (1-t)^{-1} - 1 \right| \leq 2t$ which is valid for $t \leq 1/2$:

$$\left| \frac{1}{\langle \tilde{u}_i, v_i \rangle} - 1 \right| \leq 2\delta(1 + r\alpha\beta).$$

As a result,

$$\begin{aligned} \|\bar{u}_i - \tilde{u}_i\| &= \|\tilde{u}_i\| \left| \frac{1}{\langle \tilde{u}_i, v_i \rangle} - 1 \right| \\ &\leq 2\delta\sqrt{\alpha}(1 + r\alpha\beta) \\ &\leq 4r\alpha^{3/2}\beta\delta. \end{aligned}$$

Using again the norm equivalence bound and the triangular inequality,

$$\|\bar{U} - U\| \leq 5r^2\alpha^{3/2}\beta\delta, \tag{4.18}$$

which ends the preliminary part of the proof.

We now set accordingly $\bar{S} = \bar{U}\Sigma V^*$, and claim that S is now a truly diagonalizable matrix. Indeed, any \bar{u}_i is an eigenvector of \bar{S} with associated eigenvalue θ_i , and a basis of $\text{im}(V)^\perp$ provides a family of eigenvectors of Σ with eigenvalue 0. We consequently set

$$\Pi = (\bar{U} \quad Y),$$

where Y is an orthonormal basis of $\text{im}(V)^\perp$; Π is the matrix of an eigenvector basis for S . Further, we have

$$\|\bar{S} - S\| \leq \|U - \bar{U}\| \|\Sigma\| \|V\| \leq 5r^2 \alpha^2 \beta \delta := \varepsilon'.$$

The above bound implies that the matrices A and \bar{S} are still close:

$$\|A - \bar{S}\| \leq \|A - S\| + \|S - \bar{S}\| \leq \varepsilon + \varepsilon', \quad (4.19)$$

and we can apply the Bauer-Fike theorem to A and \bar{S} ; the eigenvalues of A are contained in the union of the balls $B(\theta_i, \varepsilon'')$ and $B(0, \varepsilon'')$, where

$$\varepsilon'' = (\varepsilon + \varepsilon') \kappa(\Pi).$$

The computation of $\kappa(\Pi)$ being cumbersome, we defer the following lemma to later:

Lemma 28. *Let X be a $n \times r$ matrix with rank r , and X' such that $X^*X' = I_r$. Let Y be a matrix for an orthonormal basis of $\text{im}(X')^\perp = \ker((X')^*)$, and $P = (X, Y)$. Then, if $\|X\| \geq 1$ and $\|X'\| \geq 1$,*

$$\|\Pi\| \leq \sqrt{2}\|X\| \quad \text{and} \quad \|\Pi^{-1}\| \leq \sqrt{2}(1 + \|X\|\|X'\|)$$

Applying this to $X = \bar{U}$ and $X' = V$ gives the bound

$$\kappa(\Pi) \leq 2(\|\bar{U}\| + \|\bar{U}\|^2\|V\|),$$

and we use the triangular inequality to bound $\|\bar{U}\|$:

$$\|\bar{U}\| \leq \|U\| + \|\bar{U} - U\| \leq 6r^2 \alpha^{3/2} \beta,$$

a very loose but sufficient bound, that entails

$$\kappa(\Pi) \leq 84r^2 \alpha^{7/2} \beta.$$

The corresponding bound on ε'' reads

$$\varepsilon'' \leq 84r^2 \alpha^{7/2} \beta (\varepsilon + 5r \alpha^2 \beta \delta) = \sigma,$$

enlightening the definition in (4.16). Going back to the Bauer-Fike application, the separation condition (4.16) implies that $B(0, \sigma)$ is disjoint from $B(\theta_i, \sigma)$ for $i \in [r]$ and we can apply the second part of the theorem: there are exactly r eigenvalues of A inside the region

$$\Omega = \bigcup_{i \in [r]} B(\theta_i, \sigma),$$

and all other eigenvalues of A have modulus less than σ . Further, all connected components of Ω have the same number of eigenvalues of A and B . As a result, there exists a permutation π such that for all $i \in [r]$, we have

$$|\lambda_{\pi(i)} - \theta_i| \leq \sup_{\Omega' \subseteq \Omega} \text{diam}(\Omega') \leq 2r\sigma,$$

where the supremum is taken over all connected subsets of Ω .

We now move on to the eigenvector perturbation bound; let ξ be a normed eigenvector of A associated with the eigenvalue $\lambda_{\pi(i)}$. We write $\xi = \Pi x$ with Π the matrix defined before, and use (4.15):

$$\left\| \lambda_{\pi(i)} \Pi x - \sum_{j=1}^r \theta_j x_j \bar{u}_j \right\| = \|(A - \bar{S}) x\| \leq \varepsilon + \varepsilon',$$

which we rewrite as

$$\left\| \Pi \left(\lambda_{\pi(i)} x - \sum_{j \in [r]} \theta_j x_j e_j \right) \right\| \leq \varepsilon + \varepsilon',$$

with (e_1, \dots, e_n) the usual orthonormal basis of \mathbb{R}^n . Using the inequality $\|v\| \leq \|P^{-1}\| \|Pv\|$ holding for any vector v ,

$$\left\| \lambda_{\pi(i)} x - \sum_{j \in [r]} \theta_j x_j e_j \right\| \leq \|\Pi^{-1}\| (\varepsilon + \varepsilon').$$

We introduce the notation $\theta_{r+1} = \dots = \theta_n = 0$; whenever the ball $B(\theta_i, \sigma)$ is disjoint from all other such balls, we have $|\lambda_{\pi(i)} - \theta_i| \leq \sigma$, and thus for $j \neq i$

$$|\lambda_{\pi(i)} - \theta_j| \geq |\theta_j - \theta_i| - |\lambda_{\pi(i)} - \theta_i| \geq \delta_i - \sigma,$$

so that

$$\|x - x_i e_i\| = \left\| \sum_{j \neq i} x_j e_j \right\| \leq \frac{1}{\delta_j - \sigma} \left\| \sum_{j \neq i} (\lambda_{\pi(i)} - \theta_j) x_j e_j \right\| \leq \frac{\|\Pi^{-1}\| (\varepsilon + \varepsilon')}{\delta_j - \sigma}.$$

We now apply Π inside the norm the LHS, and use the fact that $\kappa(\Pi)(\varepsilon + \varepsilon') \leq \sigma$:

$$\|\xi - x_i \bar{u}_i\| \leq \frac{\sigma}{\delta_i - \sigma}.$$

Now, for any vectors $w, w' \in \mathbb{R}^n$, we have

$$\left\| \frac{w}{\|w\|} - \frac{w'}{\|w'\|} \right\| \leq \frac{2\|w - w'\|}{\|w\|}, \quad (4.20)$$

and all that remains is to write

$$\begin{aligned} \left\| \xi - \frac{u_i}{\|u_i\|} \right\| &\leq \left\| \xi - \frac{\bar{u}_i}{\|\bar{u}_i\|} \right\| + \left\| \frac{u_i}{\|u_i\|} - \frac{\bar{u}_i}{\|\bar{u}_i\|} \right\| \\ &\leq \frac{2\sigma}{\delta_i - \sigma} + 2\|u_i - \bar{u}_i\| \\ &\leq \frac{3\sigma}{\delta_i - \sigma}, \end{aligned}$$

having used (4.20) twice and $2\|u_i - \bar{u}_i\| \leq \sigma$. This ends the proof. \square

As announced, we now prove the aforementioned Lemma 28 on the condition number of P :

Proof. Let $z \in \mathbb{R}^n$ be a unit vector, and write $z = \begin{pmatrix} x \\ y \end{pmatrix}$ with x of size r and y of size $n - r$. Then, using that $\|Y\| = 1$,

$$\begin{aligned} \|\Pi z\| &= \|Xx + Yy\| \leq \|X\| \cdot \|x\| + \|Y\| \cdot \|y\| \\ &\leq (1 \vee \|X\|) (\|x\| + \|y\|) \\ &\leq \sqrt{2}\|X\|, \end{aligned}$$

which proves the first inequality. The second one relies on the following explicit formula for Π^{-1} :

$$\Pi^{-1} = \begin{pmatrix} (X')^* \\ -Y^*X(X')^* + Y^* \end{pmatrix}.$$

Indeed, using the relations $Y^*Y = I_{n-r}$ and $(X')^*Y = 0$:

$$\begin{aligned} \begin{pmatrix} (X')^* \\ -Y^*X(X')^* + Y^* \end{pmatrix} P &= \begin{pmatrix} (X')^* \\ -Y^*X(X')^* + Y^* \end{pmatrix} (X \ Y) \\ &= \begin{pmatrix} (X')^*X & (X')^*Y \\ -Y^*X(X')^*X + Y^*X & -Y^*X(X')^*Y + Y^*Y \end{pmatrix} \\ &= \begin{pmatrix} I_r & 0 \\ -Y^*X + Y^*X & Y^*Y \end{pmatrix} \\ &= \begin{pmatrix} I_r & 0 \\ 0 & I_{n-r} \end{pmatrix} \\ &= I_n. \end{aligned}$$

Furthermore, we have

$$\|-Y^*X(X')^* + Y^*\| \leq \|Y\| \|I_n - X(X')^*\| \leq 1 + \|X\| \|X'\|,$$

and the exact same argument as in the first inequality yields

$$\|P^{-1}\| \leq \sqrt{2} (1 + \|X\| \|X'\|)$$

□

4.4.2 Matrix power perturbation and phase perturbation control

We aim in the following section to apply Theorem 28 to powers of the matrix B ; however, such a process introduces uncertainty on the phase of the eigenvalues of B . The next theorem, adapted from [28] and [29], develops a method to control such uncertainty. As before, let $\Sigma = \text{diag}(\theta_1, \dots, \theta_r)$ with

$$1 = |\theta_1| \geq \dots \geq |\theta_r|,$$

and U, U', V, V' four $n \times r$ matrices. We set

$$S = U\Sigma^\ell V^* \quad \text{and} \quad S' = U'\Sigma^{\ell'} (V')^*,$$

for two integers ℓ, ℓ' .

Theorem 29. *Assume the following:*

- (i) *the integers ℓ, ℓ' are relatively prime,*
- (ii) *the matrices U, U', V, V' are well-conditioned:*

- they all are of rank r ,
- for some $\alpha, \beta \geq 1$, for X in $\{U, V, U', V'\}$,

$$\|X^*X\| \leq \alpha \quad \text{and} \quad \|(X^*X)^{-1}\| \leq \beta,$$

- for some small $\delta < 1$,

$$\|U^*V - \mathbf{I}_r\| \leq \delta \quad \text{and} \quad \|(U')^*V' - \mathbf{I}_r\| \leq \delta,$$

(iii) there exists a small constant $\varepsilon > 0$ such that

$$\|A^\ell - S\| \leq \varepsilon \quad \text{and} \quad \|A^{\ell'} - S'\| \leq \varepsilon,$$

(iv) if we let

$$\sigma_0 := 84r^3\alpha^{7/2}\beta(\varepsilon + 5r\alpha^2\beta\delta),$$

then

$$\sigma_0 < \ell |\theta_r|^\ell \quad \text{and} \quad \sigma_0 < \ell' |\theta_r|^{\ell'}. \quad (4.21)$$

Assume without loss of generality that ℓ is odd, and let

$$\sigma := \frac{\sigma_0}{\ell |\theta_r|^\ell}.$$

Then, the r largest eigenvalues of A are close to the θ_i in the following sense: there exists a permutation π of $[r]$ such that for $i \in [r]$,

$$|\lambda_{\pi(i)} - \theta_i| \leq 4\sigma,$$

and all other eigenvalues of A are less than $\sigma_0^{1/\ell}$. Additionally, if i is such that

$$B(\theta_i, \sigma) \cap \left(\bigcup_{j \neq i} B(\theta_j, \sigma) \right) = \emptyset, \quad (4.22)$$

then there exists a normed eigenvector ξ associated to $\lambda_{\pi(i)}$ such that

$$\left\| \xi - \frac{u_i}{\|u_i\|} \right\| \leq \frac{3\sigma}{\delta_i - \sigma},$$

with δ_i defined as in Theorem 28.

Proof. We apply Theorem 28 to A^ℓ, S and $A^{\ell'}, S'$; for any $i \in [r]$,

$$\left| \lambda_{\pi(i)}^\ell - \theta_i \right| \leq \sigma_0 \quad \text{and} \quad \left| \lambda_{\pi'(i)}^{\ell'} - \theta_i \right| \leq \sigma_0. \quad (4.23)$$

Examining the proof of Theorem 28, we notice that we can take $\pi = \pi'$ since taking the ℓ -th power does not change the ordering. We fix $i \in [r]$ and let $\lambda = \lambda_{\pi(i)} = |\lambda|e^{i\omega}$ and $\theta = \theta_i$ for now; then

$$\left| \frac{\lambda^\ell}{\theta^\ell} - 1 \right| \leq \nu := \frac{\sigma_0}{|\theta|^\ell}.$$

The argument of $(\lambda/\theta)^\ell$ is thus between $-\xi$ and ξ , with

$$\xi = |2 \arcsin(\nu/2)| \leq \pi/2\nu,$$

and the same holds for ℓ' (with ν' defined accordingly). Thus, there exists two integers p, p' and two numbers s, s' with absolute value less than $\pi/2\nu$ (resp. $\pi/2\nu'$), such that

$$\ell\omega = p\pi + s \quad \text{and} \quad \ell'\omega = p'\pi + s'.$$

This implies

$$p\ell' - p'\ell = \frac{s'\ell - s\ell'}{\pi}$$

The LHS of this inequality is an integer, and using condition (4.21) both terms in the RHS have a magnitude strictly lower than $1/2$, so both sides are 0. As ℓ and ℓ' are relatively prime, ℓ divides p and ℓ' divides p' , so that

$$\omega = k\pi + \frac{s}{\ell}.$$

Whenever θ_i is positive, k is even and we can take $\omega = s/\ell$, and when k is odd we choose $\omega = \pi + s/\ell$.

We now come back to (4.23), and write

$$\lambda_i^\ell = \theta_i^\ell(1 + z)$$

with $|z| \leq \nu$. Taking the modulus on both sides we find $|\lambda_i| = |\theta_i||1 + z|^{\frac{1}{\ell}}$ and we use the inequality $||1 + z|^{\frac{1}{\ell}} - 1| \leq 2|z|/\ell$ (valid for $|z| \leq 1/2$) to find

$$||\lambda_i| - |\theta_i|| \leq \frac{2\sigma_0}{\ell|\theta_i|^\ell}.$$

We can now prove the lemma: whether θ is positive or negative, a case analysis yields

$$\begin{aligned} |\lambda_i - \theta_i| &\leq ||\lambda_i| - |\theta_i|| + |\theta_i| \left| e^{is/\ell} - 1 \right| \\ &\leq \frac{2\sigma_0}{\ell|\theta_i|^\ell} + |\theta_i| \frac{|s|}{\ell} \\ &\leq \frac{4\sigma_0}{\ell|\theta_i|^\ell}, \end{aligned}$$

the desired bound. Now, assuming that ℓ is odd, we have by the mean value theorem

$$|\theta_i^\ell - \theta_j^\ell| \geq \ell(|\theta_i| \wedge |\theta_j|)^{\ell-1} |\theta_i - \theta_j| \geq \ell|\theta_r|^\ell |\theta_i - \theta_j|,$$

so that condition (4.22) implies the separation condition (4.17) applied to A^ℓ . We can then apply the same proof as in Theorem 28 and get

$$\left\| \xi - \frac{u_i}{\|u_i\|} \right\| \leq \frac{3\sigma_0}{\ell|\theta_r|^\ell \delta_i - \sigma_0},$$

which is equivalent to the theorem bound. □

4.5 Proof of Theorem 22

We prove in this section the main result on the spectral properties of B . We shall use the same notations as in Theorem 22; since the statement of the theorem is invariant upon multiplying the entries of W by a common constant, we shall assume in the rest of the paper that $\mu_1 = 1$.

Our candidates for the singular vectors of B^ℓ are the vectors (u_1, \dots, u_{r_0}) and (v_1, \dots, v_{r_0}) , where for $i \in [r_0]$

$$u_i = \frac{B^\ell \chi_i}{\mu_i^\ell} \quad \text{and} \quad v_i = \frac{(B^*)^\ell D_W \tilde{\chi}_i}{\mu_i^{\ell+1}},$$

with associated eigenvalue μ_i^ℓ . We let U (resp. V) be the $n \times r$ matrix whose columns are the u_i (resp v_i), and $D = \text{diag}(\mu_1, \dots, \mu_{r_0})$. Finally, we'll need an approximation of the Gram matrix of the vectors u (and v); we define for every $t \geq 0$ the *covariance* matrices $\Gamma_U^{(t)}$ and $\Gamma_V^{(t)}$ such that for $i, j \in [r_0]$,

$$\Gamma_{U,ij}^{(\ell)} = \sum_{s=0}^{\ell} \frac{\langle P\mathbf{1}, K^s \varphi^{i,j} \rangle}{(\mu_i \mu_j)^s} \quad \text{and} \quad \Gamma_{V,ij}^{(\ell)} = \sum_{s=0}^{\ell} \frac{\langle K\mathbf{1}, K^s \varphi^{i,j} \rangle}{(\mu_i \mu_j)^{s+1}}, \quad (4.24)$$

where $\varphi^{i,j} = \varphi_i \circ \varphi_j$.

4.5.1 Structure of the matrices U and V

Following from the subsequent local analysis of G , as well as a trace bound argument, we gather the following relations between matrices B^ℓ , D^ℓ and U . We define the following parameter

$$\omega = d^5 (1 \vee L)^2,$$

which is an upper bound on the exponential scaling of our error terms.

Theorem 30. *Let r, d, b, τ, L be parameters as above, such that $a \leq n^{1/4}$, and (P, W) be any matrices in $\mathcal{C}(r, d, b, \tau, L)$. Let ℓ be any integer such that*

$$\ell \leq \frac{1 - \epsilon \log(n)}{4 \log(\omega)}, \quad (4.25)$$

for some $\epsilon > 0$, where ω is the parameter defined above. Then there exists an event with probability at least $1 - c/\log(n)$ and a parameter $N_0 \leq a^{12} L^6$ such that if $n \geq N_0$

$$\|U^*U - \Gamma_U^{(\ell)}\| \leq C \times \eta, \quad (4.26)$$

$$\|V^*V - \Gamma_V^{(\ell)}\| \leq C \times \eta, \quad (4.27)$$

$$\|U^*V - \mathbf{I}_{r_0}\|_\infty \leq C \times \eta, \quad (4.28)$$

$$\|B^\ell U - UD^\ell\| \leq C'(\sqrt{\rho} \vee L)^\ell, \quad (4.29)$$

$$\|B^\ell P_{H^\perp}\| \leq C'(\sqrt{\rho} \vee L)^\ell, \quad (4.30)$$

where η , C and C' satisfy

$$C \leq crd^4 b^2 L, \quad C' \leq cr^2 d^6 b^2 L^2 \log(n)^{20} \quad \text{and} \quad \eta \leq n^{-1/4} \wedge \sqrt{\rho}^\ell \wedge L^\ell.$$

Furthermore, on this same event, we have the following bound:

$$\|B^\ell\| \leq c \log(n) n^{1/4} L^\ell. \quad (4.31)$$

The proof of this theorem will occupy the next few pages of this article; we first show how it implies the statement of Theorem 22.

4.5.2 Proof of the perturbation bounds

The goal here is to apply Theorem 29 to B^ℓ , U and V : we choose ℓ equal to the upper bound in (4.25) (with arbitrary ϵ , say 0.01) and $\ell' = \ell + 1$, and let

$$S = UD^\ell V^* \quad \text{and} \quad S' = U'D^{\ell'} V'^*,$$

where U', V' are defined identically to U and V replacing ℓ by ℓ' . We now check all the conditions of Theorem 29:

Condition (i) Since $\ell' = \ell + 1$, ℓ and ℓ' are relatively prime.

Condition (ii) We shall need a small lemma on the spectral properties of the covariance matrices, which will be proven in a subsequent section:

Lemma 29. *For all $t \geq 1$; the matrix $\Gamma_U^{(t)}$ (resp. $\Gamma_V^{(t)}$) is a positive definite matrix, with all its eigenvalues greater than 1 (resp. c_0^{-1}) and such that*

$$1 \leq \|\Gamma_U^{(t)}\| \leq \frac{r^2 d^3 L^2}{1 - \tau} \quad \text{and} \quad c_0^{-1} \leq \|\Gamma_V^{(t)}\| \leq \frac{r^2 d^2 L^2}{1 - \tau}.$$

Then, the minimum eigenvalue of V^*V is at least $c_0^{-1} - Cn^{-1/4}$, which is more than $c_0^{-1}/2$ as soon as

$$n \geq c_1 r^4 a^4 d^{16} b^8 L^4,$$

and we can take $\beta = 2c_0$ whenever this holds. On the other hand,

$$\|V^*V\| \leq \frac{rd^2L^2}{1 - \tau} + \frac{rb^2d^4L}{n^{1/4}} \leq \frac{2rb^2d^4L^2}{1 - \tau}.$$

Performing the same computations on U^*U leads us to the choice

$$\alpha = \frac{2rb^2d^4L^2}{1 - \tau}.$$

Finally, equation (4.28) allows us to take

$$\delta = C\eta.$$

Condition (iii) This condition requires some additional computations. Define as before the orthogonal projection P_H on $H = \text{im}(V)$, and $P_{H^\perp} = I_n - P_H$; we have the formula

$$P_H = V(V^*V)^{-1}V^*.$$

Noticing that $SP_H = S$, we can bound $\|B^\ell - S\|$ as follows:

$$\begin{aligned} \|B^\ell - S\| &\leq \|B^\ell P_H - SP_H\| + \|SP_{H^\perp}\| + \|B^\ell P_{H^\perp}\| \\ &\leq \|B^\ell P_H - S\| + \|B^\ell P_{H^\perp}\| \\ &\leq \|B^\ell V(V^*V)^{-1} - UD\| \|V^*\| + \|B^\ell P_{H^\perp}\|. \end{aligned}$$

To apply (4.29), we let

$$U = P_H U + P_{H^\perp} U = V(V^*V)^{-1} + \tilde{U} + P_{H^\perp} U.$$

The second term is equal to $V(V^*V)^{-1}(V^*U - \mathbf{I}_{r_0})$, and we can thus use (4.28):

$$\|\tilde{U}\| \leq \|V\| \|(V^*V)^{-1}\| \|V^*U - \mathbf{I}_{r_0}\| \leq r\sqrt{\alpha}\beta\delta.$$

Going back to the above inequality, we find

$$\|B^\ell - S\| \leq \|B^\ell U - UD^\ell\| + \|B^\ell\| \|\tilde{U}\| + \|B^\ell P_{H^\perp}\| \|U\| + \|B^\ell P_{H^\perp}\|,$$

and the bounds in Theorem 30 readily imply that all terms in the above inequality are bounded above by $\varepsilon := C''(\sqrt{\rho} \vee L)^\ell$, with

$$C'' \leq \frac{c_3 r^3 a d^8 b^3 L^3 \log(n)^{20}}{1 - \tau}.$$

Condition (iv) Using all the bounds proven in the above computations, we find that

$$\sigma_0 \leq C_0(\sqrt{\rho} \vee L)^\ell \quad \text{with} \quad C_0 \leq \frac{c_4 a^2 r^{11} d^{25} b^{13} L^{12} \log(n)^{20}}{(1 - \tau)^6}.$$

The bound we have to check is therefore

$$C_0(\sqrt{\rho} \vee L)^\ell \leq \ell |\mu_{r_0}|^\ell \iff C_0 \tau^\ell \leq \ell,$$

which happens as soon as

$$\log(n) \geq \frac{5 \log(C_0) \log(\omega)}{\log(\tau^{-1})}.$$

The same proof holds for ℓ' , with the same constants.

Having checked all assumptions of Theorem 29, we can now apply it to B^ℓ ; this implies the existence of a permutation $\pi \in \mathfrak{S}_{r_0}$ (possibly depending on n) such that for $i \in [r_0]$,

$$|\lambda_i - \mu_{\pi(i)}| \leq \sigma := C_0 \tau^\ell,$$

and all the other eigenvalues of B satisfy

$$|\lambda| \leq C_0^{\frac{1}{2}}(\sqrt{\rho} \vee L).$$

Now, assume that for some $i \in [r_0]$, $\delta_i \geq 2\sigma$. Then, applying the last part of Theorem 29, there exists an eigenvector of B associated with λ_i such that

$$\left\| \xi - \frac{u_i}{\|u_i\|} \right\| \leq \frac{3\sigma}{\delta_i - \sigma}.$$

We define in the following

$$\gamma_i = \langle P\mathbf{1}, (\mathbf{I}_n - \mu_i^{-2}K)^{-1} \varphi^{i,i} \rangle.$$

If we rewrite the definition of $\Gamma_{U,ii}^{(t)}$ as

$$\Gamma_{U,ii}^{(t)} = \left\langle P\mathbf{1}, \sum_{s=0}^t (\mu_i)^{-2s} K^s \varphi^{i,j} \right\rangle,$$

the matrix sum converges as $t \rightarrow \infty$ since $\rho(K) < \mu_i^2$, and using Lemma 30 below we have

$$\begin{aligned} \left| \Gamma_{U,ii}^{(\ell)} - \gamma_i \right| &= \sum_{t=\ell+1}^{\infty} \mu_i^{-2t} \langle P\mathbf{1}, K^t \varphi^{i,j} \rangle \\ &\leq \sum_{t=\ell+1}^{\infty} rd^3 L^2 \rho^t \mu_i^{-2t} \\ &\leq \sigma, \end{aligned}$$

and combined with (4.26) yields

$$\left| \|u_i\|^2 - \gamma_i \right| \leq 2\sigma.$$

On the other hand, we shall prove the following inequality in the following sections (equation (4.58)): for all $t \leq 2\ell$,

$$\left| \langle B^t \chi_i, \chi_i \rangle - \mu_i^t \langle P\mathbf{1}, \varphi^{i,i} \rangle \right| \leq \sigma \mu_i^t.$$

Setting $t = 0$ and $t = \ell$ in this inequality yields at the same time

$$\left| \|\chi_i\|^2 - \langle P\mathbf{1}, \varphi^{i,i} \rangle \right| \leq \sigma \quad \text{and} \quad \left| \langle u_i, \chi_i \rangle - \langle P\mathbf{1}, \varphi^{i,i} \rangle \right| \leq \sigma.$$

We now have, using the Cauchy-Schwarz inequality,

$$\begin{aligned} \left| \langle \xi, \xi_i \rangle - \sqrt{\frac{\langle P\mathbf{1}, \varphi^{i,i} \rangle}{\gamma_i}} \right| &\leq \left\| \xi - \frac{u_i}{\|u_i\|} \right\| + \left| \left\langle \frac{u_i}{\|u_i\|}, \frac{\chi_i}{\|\chi_i\|} \right\rangle - \sqrt{\frac{\langle P\mathbf{1}, \varphi^{i,i} \rangle}{\gamma_i}} \right| \\ &\leq \frac{3\sigma}{\delta_i - \sigma} + c_5 \sigma \\ &\leq \frac{c_6 \sigma}{\delta_i - \sigma}. \end{aligned}$$

Finally, notice that

$$\gamma_i = \langle P\mathbf{1}, \varphi^{i,i} \rangle + \left\langle P\mathbf{1}, \sum_{s=1}^{\infty} (\mu_i)^{-2s} K^s \varphi^{i,j} \right\rangle \leq \langle P\mathbf{1}, \varphi^{i,i} \rangle + rd^2 L^2 \frac{\rho/\mu_i^2}{1 - \rho/\mu_i^2}.$$

Using that $rd^2 L^2 \geq 1$ and $\langle P\mathbf{1}, \varphi^{i,i} \rangle \geq 1$, we find

$$\frac{\langle P\mathbf{1}, \varphi^{i,i} \rangle}{\gamma_i} \geq 1 - rd^2 L^2 \frac{\rho}{\mu_i^2}.$$

4.6 Preliminary computations

We begin the proof of Theorem 30 with some elementary computations on the entries of K and $\Gamma^{(t)}$, which will be of use in the later parts of the proof. Most of the results from this section are adapted from [28], although sometimes improved and adapted to our setting.

Bounding ρ and L from below We begin with a simple bound on $\rho = \rho(K)$; by the Courant-Fisher theorem, $\rho \geq \langle w, Kw \rangle$ for every unit vector w , and applying it to $w = \mathbf{1}/\sqrt{n}$

yields

$$\begin{aligned}
\rho &\geq \frac{\langle w, Kw \rangle}{n} \\
&= \frac{1}{n} \sum_{i,j \in [n]} P_{ij} \mathbf{E}[W_{ij}^2] \\
&= \frac{1}{d} \sum_{i,j \in [n]} P_{ij}^2 \mathbf{E}[W_{ij}]^2 \\
&= \frac{\|Q\|_F^2}{d},
\end{aligned}$$

where we used that $P_{ij} \leq d/n$ and the Jensen inequality. The Frobenius norm of Q is then greater than $\mu_1^2 = 1$, which in turns implies

$$\rho \geq \frac{1}{d}, \quad (4.32)$$

so that ρ is bounded away from zero. In order to prove a similar bound on L , we write for $x \in [n]$

$$\varphi_1(x) = \sum_{y \in [n]} Q_{xy} \varphi_1(y) \leq \sqrt{\sum_y Q_{xy}^2} \leq \frac{dL}{\sqrt{n}}.$$

Squaring and summing those inequalities over x gives

$$1 = \|\varphi_1\|^2 \leq d^2 L^2,$$

so that as with ρ ,

$$L \geq \frac{1}{d}. \quad (4.33)$$

A scalar product lemma Our second step is an important lemma for the following proof, leveraging the entrywise bounds on W :

Lemma 30. *Let $\varphi, \varphi' \in \mathbb{R}^n$ be any unit vectors. Then, for any $t \geq 0$,*

$$\langle \mathbf{1}, K^t \varphi \circ \varphi' \rangle \leq rd^2 L^2 \rho^t$$

Proof. We write the eigendecomposition of K as

$$K = \sum_{k=1}^s \nu_k \psi_k \psi_k^*,$$

with $\nu_1 = \rho$ the Perron-Frobenius eigenvalue of K and $s \leq r^2$ its rank. Then, for all $i \in [n]$,

$$\begin{aligned}
\sum_{k=1}^s \nu_k^2 \psi_k(i)^2 &= (K^2)_{ii} = \sum_{j \in [n]} K_{ij}^2 \\
&= \sum_{j \in [n]} P_{ij}^2 \mathbf{E}[W_{ij}^2]^2 \\
&\leq \sum_{j \in [n]} \left(\frac{d}{n}\right)^2 L^4 \\
&\leq \frac{d^2 L^4}{n}.
\end{aligned}$$

This is akin to a delocalization property on the eigenvectors of K .

We can now prove the above lemma:

$$\begin{aligned}
\langle \mathbf{1}, K^t \varphi \circ \varphi' \rangle &= \sum_{k=1}^s \nu_k^t \langle \mathbf{1}, \psi_k \rangle \langle \psi_k, \varphi \circ \varphi' \rangle \\
&\leq \rho^{t-1} \sum_{k=1}^s \|\psi_k\| \|\mathbf{1}\| \cdot |\nu_k| |\langle \psi_k, \varphi \circ \varphi' \rangle| \\
&\leq \rho^{t-1} \sqrt{n} \sum_{i \in [n]} |\varphi(i)| |\varphi'(i)| \sum_{k=1}^s |\nu_k| |\psi_k(i)| \\
&\leq \rho^t d \sqrt{n} \sum_{i \in [n]} |\varphi(i)| |\varphi'(i)| \sqrt{s} \sqrt{\sum_{k=1}^s \nu_k^2 \psi_k(i)^2} \\
&\leq \rho^t a \sqrt{n} \sqrt{s} \frac{dL^2}{\sqrt{n}} \sum_i |\varphi(i)| |\varphi'(i)| \\
&\leq r d^2 L^2 \rho^t,
\end{aligned}$$

where we extensively used the Cauchy-Schwarz inequality, as well as the bound $\rho^{-1} \leq d$ from (4.32). \square

Entrywise bounds for K^t For a more precise estimation of entrywise bounds, we define the scale-invariant *delocalization* parameter

$$\Psi = \frac{dL^2}{\rho}.$$

Using the same proof technique as in (4.33), as well as (4.32), we have

$$1 \leq \Psi \leq d^2 L^2$$

for any $i, j \in [n]$. Recall that, as shown in the proof of Lemma 30, for all $i \in [n]$

$$(K^2)_{ii} \leq \frac{d^2 L^4}{n} = \frac{\Psi^2}{n} \rho^2.$$

Now, for $t \geq 0$ and $i, j \in [n]$,

$$\begin{aligned}
(K^t)_{ij} &= \sum_{k \in [s]} \nu_k^t \psi_k(i) \psi_k(j) \\
&\leq \rho^{t-2} \sum_k \nu_k^2 |\psi_k(i)| |\psi_k(j)| \\
&\leq \rho^{t-2} \sqrt{(K^2)_{ii} (K^2)_{jj}},
\end{aligned}$$

where we again used the Cauchy-Schwarz inequality at the last line. This yields

$$(K^t)_{ij} \leq \frac{\Psi^2}{n} \rho^t \tag{4.34}$$

for any $t \geq 1$ and $i, j \in [n]$.

The covariance matrices Recall that we defined the matrices $\Gamma_U^{(t)}$ and $\Gamma_V^{(t)}$ as

$$\Gamma_{U,ij}^{(\ell)} = \sum_{s=0}^{\ell} \frac{\langle P\mathbf{1}, K^s \varphi^{i,j} \rangle}{(\mu_i \mu_j)^s} \quad \text{and} \quad \Gamma_{V,ij}^{(\ell)} = \sum_{s=0}^{\ell} \frac{\langle K\mathbf{1}, K^s \varphi^{i,j} \rangle}{(\mu_i \mu_j)^{s+1}}$$

for $i, j \in [r_0]$. Our aim is to prove the following lemma:

Lemma 31. *For all $t \geq 1$; the matrix $\Gamma_U^{(t)}$ (resp. $\Gamma_V^{(t)}$) is a positive definite matrix, with all its eigenvalues greater than 1 (resp. c^{-1}) and such that*

$$1 \leq \|\Gamma_U^{(t)}\| \leq \frac{r^2 d^3 L^2}{1 - \tau} \quad \text{and} \quad c^{-1} \leq \|\Gamma_V^{(t)}\| \leq \frac{r^2 d^2 L^2}{1 - \tau}.$$

Proof. We first prove the bounds for $\Gamma_V^{(t)}$. Let $C^{(s)}$ be the $r_0 \times r_0$ matrix with

$$C_{ij}^{(s)} = \frac{\langle K\mathbf{1}, K^s \varphi^{i,j} \rangle}{(\mu_i \mu_j)^{s+1}}.$$

Then for every $w \in \mathbb{R}^{r_0}$ we have

$$\begin{aligned} w^* C^{(s)} w &= \sum_{i,j \in [r_0]} \frac{w_i w_j}{(\mu_i \mu_j)^{s+1}} \sum_{x \in [n]} [K^{s+1} \mathbf{1}](x) \varphi_i(x) \varphi_j(x) \\ &= \sum_{x \in [n]} [K^{s+1} \mathbf{1}](x) \left(\sum_{i \in [r_0]} \frac{w_i \varphi_i(x)}{\mu_i^{s+1/2}} \right)^2 \\ &\geq 0, \end{aligned}$$

hence every matrix $C^{(s)}$ is positive semi-definite. Further, we have

$$C^{(0)} = D^{-1} \Phi^* \text{diag}(K\mathbf{1}) \Phi D^{-1},$$

where Φ is the $n \times r$ matrix whose columns are the φ_i . Using $\mu_i \leq 1$ for any $i \in [r_0]$, the eigenvalues of $C^{(0)}$ are all greater than $\min_x [K\mathbf{1}](x) \geq c^{-1}$ by our initial assumptions. This settles the positive definite property, as well as the minimum eigenvalue of $\Gamma_V^{(t)}$.

Now, applying Lemma 30 to φ_i and φ_j , for all $i, j \in [r_0]$ one has

$$\begin{aligned} \Gamma_{V,ij}^{(t)} &\leq \sum_{s=0}^t \frac{r d^2 L^2 \rho^{s+1}}{(\mu_i \mu_j)^{s+1}} \\ &\leq r d^2 L^2 \sum_{s=0}^{\infty} \left(\frac{\rho}{\mu_i \mu_j} \right)^s. \end{aligned}$$

By definition of τ , the summand above is less than τ^s , whose sum converges since $\tau < 1$. As a result,

$$\|\Gamma_V^{(t)}\|_{\infty} \leq \frac{r d^2 L^2}{1 - \tau},$$

and the classic bound $\|\Gamma_V^{(t)}\| \leq r_0 \|\Gamma_V^{(t)}\|_{\infty}$ implies the upper bound.

The proof for $\Gamma_U^{(t)}$ is very similar; the upper bound simply ensues from the fact that $d_x \leq d$ for any $x \in [n]$. For the lower bound, if we let as above

$$C'_{ij}{}^{(s)} = \frac{\langle P\mathbf{1}, K^s \varphi^{i,j} \rangle}{(\mu_i \mu_j)^s},$$

then

$$C'^{(0)} = \Phi^* \text{diag}(P\mathbf{1})\Phi,$$

and the minimum of $P\mathbf{1}$ is at least 1. This implies that the eigenvalues of $C'^{(0)}$ are larger than one, and we conclude as before. \square

4.7 Local study of G

It is a well-known fact (see for example [29]) that when the mean degree is low enough ($d = n^{o(1)}$), the graph G is locally tree-like — that is, vertex neighbourhoods behave almost like random trees. The goal of this section is to establish rigorously this result, as well as provide bounds on neighbourhood sizes.

4.7.1 Setting and definitions

Labeled rooted graphs A labeled rooted graph is a triplet $g_* = (g, o, \iota)$ consisting of a graph $g = (V, E)$, a root $o \in V$, and a mark function $\iota : V \rightarrow \mathbb{N}$ with finite support. We shall denote by \mathcal{G}_* the set of labeled rooted graphs with $V = \mathbb{N}$, and will often write $g_* = (g, o)$ for an element of \mathcal{G}_* , dropping the mark function. Notions of subgraphs, induced subgraphs and distance extend naturally from regular graphs to this setting.

Labeling trees and graphs We recall that G is the inhomogeneous random graph defined earlier. For each vertex $x \in V$, we can define the associated element of \mathcal{G}_* as follows: the root is set to x , each vertex $y \in [n]$ is given a mark $\iota(y) = y$, and we let $\iota(z) = 0$ for all $z \in \mathbb{N} \setminus [n]$. The resulting triple (G, x, ι) is a random element of \mathcal{G}_* .

Now, let $o \in [n]$; we define the inhomogeneous random tree as follows: first, the root is given a mark $\iota(o) = o$. Then, for each vertex x already labeled, we draw the number of children of x according to $\text{Poi}(d_{\iota(x)})$, where we recall that

$$d_{\iota(x)} = \sum_j P_{\iota(x),j} \leq d.$$

Each child y of x receives a label drawn independently at random from the distribution

$$\pi_{\iota(x)} = \left(\frac{P_{\iota(x),1}}{d_{\iota(x)}}, \dots, \frac{P_{\iota(x),n}}{d_{\iota(x)}} \right), \quad (4.35)$$

which sums to 1 by definition. The resulting tree is a random element of \mathcal{G}_* , denoted by (T, o) .

4.7.2 Growth properties of trees and graphs

A number of growth properties for neighbourhoods in T and G are needed to ensure the successful couplings below. By definition of d , G (resp. (T, o)) is dominated by an Erdős-Rényi graph $\mathcal{G}(n, d/n)$ (resp. a Galton-Watson tree with offspring distribution $\text{Poi}(d)$); we are thus able to directly lift properties from [29], Sections 8 and 9.

Lemma 32. *Let v be an arbitrary vertex in G ; then, there exist absolute constants $c_0, c_1 > 0$ such that for every $s > 0$, we have*

$$\mathbf{P}(\forall t \geq 1, |\partial(G, v)_t| \leq sd^t) \geq 1 - c_0 e^{-c_1 s}. \quad (4.36)$$

The same result holds when replacing (G, v) with the tree (T, o) defined above.

Taking $s = c_1^{-1} \log(c_0 n^2)$ in the above inequality, one gets

$$\mathbf{P}(\forall t \geq 1, \forall v \in V, |\partial(G, v)_t| \leq c_3 \log(n) d^t) \geq 1 - \frac{1}{n}, \quad (4.37)$$

for any $n \geq 3$. Summing these inequalities for $1 \leq t \leq \ell$ yields a similar bound for the whole ball: with probability at least $1 - \frac{1}{n}$, we have

$$|(G, v)_t| \leq c_4 \log(n) d^t \quad (4.38)$$

for all $v \in V$ and $t \geq 1$. In particular, this implies the following useful bound: for any $v \in V$,

$$\deg(v) \leq c_4 d \log(n).$$

Another consequence of (4.36) is the following useful lemma:

Lemma 33. *For every $p \geq 2$, there is a constant c_p such that*

$$\mathbf{E} \left[\max_{v \in V} \sup_{t \geq 1} \left(\frac{|\partial(G, v)_t|}{d^t} \right)^p \right] \leq c_p \log(n)^p \quad (4.39)$$

Similarly to the proof of (4.38), we have

$$\max_{v \in V} |(G, v)_t|^p \leq d^{tp} t^p \max_{x \in V} \sup_{s \leq t} \frac{|\partial(G, v)_s|^p}{d^{sp}},$$

which yields

$$\mathbf{E} \left[\max_{v \in V} |(G, v)_t|^p \right] \leq c_p t^p \log(n)^p d^{tp} \quad (4.40)$$

An important note is that the above results apply to any collection of n random variables satisfying an inequality like (4.36); in particular, it also applies to an i.i.d collection of inhomogeneous random trees of size n .

4.7.3 Local tree-like structure

We first check that the random graph G is tree-like. We say that a graph g is ℓ -tangle-free if there is at most one cycle in the ℓ -neighbourhood of every vertex in the graph. As mentioned before, the random graph G is dominated by an Erdős-Rényi graph $\mathcal{G}(n, d/n)$; we can therefore lift the desired properties from [29].

Lemma 34. *Let $\ell \leq n$ be any integer parameter.*

(i) *the random graph G is ℓ -tangle-free with probability at least $1 - ca^2 d^{4\ell}/n$*

(ii) *the probability that a given vertex v has a cycle in its ℓ -neighbourhood is at most $cad^{2\ell}/n$.*

We shall assume in the following that the 2ℓ -tangle-free property happens with probability at least $1 - cn^{-\epsilon}$ for some $\epsilon > 0$, which happens whenever

$$\ell \leq \frac{1 - \epsilon}{10} \log_d(n) \leq c_3 \log(n). \quad (4.41)$$

We now gather all the result of the current section into one proposition, for ease of reading. The bound $\ell \leq c \log(n)$ assumed above is used to simplify the inequalities below.

Proposition 13. *Let G be an inhomogeneous random graph, and $(T_x, x)_{x \in [n]}$ a family of random trees as defined above. Let ℓ be small enough so that (4.41) holds. Then there exists an event \mathcal{E} with probability at least $1 - \frac{1}{\log(n)}$, under which:*

- (i) *the graph G is 2ℓ -tangle-free,*
- (ii) *for all $v \in G$, $t \leq 2\ell$, we have*

$$|(G, x)_t| \leq c \log(n) d^t, \quad (4.42)$$

- (iii) *for any $t \leq 2\ell$, the number of vertices in G whose t -neighbourhood contains a cycle is at most $c \log(n)^2 d^{t+1}$*

Furthermore, for any $t \leq 2\ell$ and $p \geq 1$, we have

$$\mathbf{E} \left[\max_{v \in V} |(G, v)_t|^p \right]^{\frac{1}{p}} \leq c \log(n)^2 d^t, \quad (4.43)$$

and the same holds for the family $(T_x, x)_{x \in [n]}$.

4.7.4 Coupling between rooted graphs and trees

We now turn onto the main argument of this proof: we bound the variation distance between the neighbourhoods of (G, x) and (T, x) up to size ℓ .

First, recall some definitions: if $\mathbb{P}_1, \mathbb{P}_2$ are two probability measures on the space (Ω, \mathcal{F}) , their *total variation distance* is defined as

$$d_{\text{TV}}(\mathbb{P}_1, \mathbb{P}_2) = \sup_{\mathcal{A} \in \mathcal{F}} |\mathbb{P}_1(\mathcal{A}) - \mathbb{P}_2(\mathcal{A})|.$$

The following two characterizations of the total variation distance shall be useful: first, whenever Ω is countable, we have

$$d_{\text{TV}}(\mathbb{P}_1, \mathbb{P}_2) = \frac{1}{2} \|\mathbb{P}_1 - \mathbb{P}_2\|_1 = \frac{1}{2} \sum_{\omega \in \Omega} |\mathbb{P}_1(\omega) - \mathbb{P}_2(\omega)|. \quad (4.44)$$

Additionally,

$$d_{\text{TV}}(\mathbb{P}_1, \mathbb{P}_2) = \min_{\mathbb{P} \in \pi(X_1, X_2)} \mathbb{P}(X_1 \neq X_2), \quad (4.45)$$

where $\pi(X_1, X_2)$ denotes the set of all *couplings* between \mathbb{P}_1 and \mathbb{P}_2 , i.e. probability measures on $(\Omega^2, \mathcal{F} \otimes \mathcal{F})$ such that the marginal distributions are \mathbb{P}_1 and \mathbb{P}_2 .

Denoting by $\mathcal{L}(X)$ the probability distribution of a variable X , the aim of this section is to prove the following:

Proposition 14. *Let $\ell \leq c_0 \log(n)$ for some constant $c_0 > 0$. Then, for every vertex $v \in V$,*

$$d_{\text{TV}}(\mathcal{L}((G, v)_\ell), \mathcal{L}((T, v)_\ell)) \leq \frac{c \log(n)^2 d^{2\ell+2}}{n}. \quad (4.46)$$

A total variation distance lemma for sampling processes

For an integer n , denote by $\mathcal{S}(n)$ the set of all multisets with elements in $[n]$, and by $\mathcal{P}(n) \subset \mathcal{S}(n)$ the powerset of $[n]$. Let $p_1, \dots, p_n \in [0, 1/2]$, with $\sum p_i = \lambda$ and $\sum p_i^2 = \alpha$, and consider the two probability laws on $\mathcal{S}(n)$:

- \mathbb{P}_1 : each element i of $[n]$ is picked with probability p_i ,
- \mathbb{P}_2 : the size of the multiset S is drawn according to a $\text{Poi}(\lambda)$ distribution, and each element of S has an i.i.d label with distribution $(p_1/\lambda, \dots, p_n/\lambda)$.

Note that \mathbb{P}_1 is actually supported on $\mathcal{P}(n)$.

Proposition 15. *Let $\mathbb{P}_1, \mathbb{P}_2$ be defined as above. Then*

$$d_{\text{TV}}(\mathbb{P}_1, \mathbb{P}_2) \leq \alpha + \frac{e^{2\alpha} - 1}{2}.$$

Proof. Using characterization (4.44), we have

$$2 d_{\text{TV}}(\mathbb{P}_1, \mathbb{P}_2) = \sum_{S \in \mathcal{P}(n)} |\mathbb{P}_1(S) - \mathbb{P}_2(S)| + \mathbb{P}_2(S \notin \mathcal{P}(n)). \quad (4.47)$$

We shall treat those two terms separately. First, notice that for $S \in \mathcal{P}(n)$, we have

$$\mathbb{P}_1(S) = \prod_{i \in S} p_i \prod_{i \notin S} (1 - p_i) \quad (4.48)$$

$$\begin{aligned} \mathbb{P}_2(S) &= \frac{e^{-\lambda} \lambda^{|S|}}{|S|!} \times |S|! \prod_{i \in S} \frac{p_i}{\lambda} \\ &= e^{-\lambda} \prod_{i \in S} p_i, \end{aligned} \quad (4.49)$$

and thus by summing over all sets S ,

$$\mathbb{P}_2(S \in \mathcal{P}(n)) = e^{-\lambda} \prod_{i=1}^n (1 + p_i).$$

Using the classical inequality $\log(1+x) \geq x - x^2/2$, we can bound the second member of (4.47) as follows:

$$\begin{aligned} \mathbb{P}_2(S \notin \mathcal{P}(n)) &= 1 - e^{-\lambda} \prod_{i=1}^n (1 + p_i) \\ &\leq 1 - e^{-\lambda} e^{\lambda - \alpha/2} \\ &\leq \alpha/2. \end{aligned}$$

On the other hand, using again (4.48) and (4.49), the first term reduces to

$$\begin{aligned} \sum_{S \in \mathcal{P}(n)} |\mathbb{P}_1(S) - \mathbb{P}_2(S)| &= \sum_{S \in \mathcal{P}(n)} \prod_{i \in S} p_i \left| \prod_{i \notin S} (1 - p_i) - e^{-\lambda} \right| \\ &\leq \sum_{S \in \mathcal{P}(n)} \prod_{i \in S} p_i \left(\left| e^{-\lambda} - \prod_{i=1}^n (1 - p_i) \right| + \left| \prod_{i \notin S} (1 - p_i) - \prod_{i=1}^n (1 - p_i) \right| \right). \end{aligned}$$

Both absolute values above can be removed since the expressions inside are nonnegative; further, for $0 \leq p \leq 1/2$, we have $\log(1-x) \geq -x - x^2$. Combining all those estimates, we find

$$\begin{aligned}
& \sum_{S \in \mathcal{P}(n)} |\mathbb{P}_1(S) - \mathbb{P}_2(S)| \\
& \leq e^{-\lambda}(1 - e^{-\alpha}) \sum_{S \in \mathcal{P}(n)} \prod_{i \in S} p_i + \prod_{i=1}^n (1 - p_i) \sum_{S \in \mathcal{P}(n)} \prod_{i \in S} p_i \left(\prod_{i \in S} \frac{1}{1 - p_i} - 1 \right) \\
& \leq \alpha e^{-\lambda} \prod_{i=1}^n (1 + p_i) + e^{-\lambda} \left(\prod_{i=1}^n \left(1 + \frac{p_i}{1 - p_i} \right) - \prod_{i=1}^n (1 + p_i) \right) \\
& \leq \alpha + e^{-\lambda} \exp \left(\sum_{i=1}^n \frac{p_i}{1 - p_i} \right) - e^{-\frac{\alpha}{2}},
\end{aligned}$$

where we again used the logarithm inequalities extensively. Finally, for $0 \leq p \leq 1/2$, we have $p/(1-p) \leq p + 2p^2$, which allows us to finish the computation:

$$\sum_{S \in \mathcal{P}(n)} |\mathbb{P}_1(S) - \mathbb{P}_2(S)| \leq \frac{3}{2}\alpha + e^{2\alpha} - 1. \quad (4.50)$$

Combining (4.50) with (4.47) easily implies the lemma. \square

We introduce now a family of probability laws on $\mathcal{S}(n)$; for a subset $S \subseteq [n]$, let \mathbb{P}_S be the measure corresponding to picking each element i of S with probability p_i .

The variation distance between those laws and $\mathbb{P}_1 = \mathbb{P}_{[n]}$ is then easier to bound:

Lemma 35. *For any $S \subseteq [n]$, we have:*

$$d_{\text{TV}}(\mathbb{P}_1, \mathbb{P}_S) \leq \sum_{i \notin S} p_i.$$

Proof. Consider the following coupling: we take a realization X of \mathbb{P}_1 , and set $Y = X \cap S$. Then, $Y \sim \mathbb{P}_S$, and we find

$$\mathbb{P}(X \neq Y) = \mathbb{P}_1(X \cap S^c \neq \emptyset) \leq \mathbf{E}[|X \cap S^c|] = \sum_{i \notin S} p_i$$

This ends the proof, since (4.45) ensures that $d_{\text{TV}}(\mathbb{P}_1, \mathbb{P}_S) \leq \mathbb{P}(X \neq Y)$. \square

Proof of Proposition 14

Gathering all the previous results, we are now ready to prove Proposition 14:

Proof. Define the classical breadth-first exploration process on the neighbourhood of a vertex v as follows : start with $A_0 = \{v\}$ and at stage $t \geq 0$, if A_t is not empty, take a vertex $v_t \in A_t$ at minimal distance from v , reveal its neighbours N_t in $V \setminus A_t$, and update $A_{t+1} = (A_t \cup N_t) \setminus \{v_t\}$. We denote by $(\mathcal{F}_t)_{t \geq 0}$ the filtration generated by the $(A_t)_{t \geq 0}$, and by $D_t = \bigcup_{s < t} A_s$ the set of vertices already visited at time t , and τ the first time at which all vertices in $(G, v)_\ell$ have been revealed.

We perform the same exploration process in parallel on (T, v) , which corresponds to a breadth-first search of the tree. At step t , we denote by \mathbb{P}_t the distribution of N_t given \mathcal{F}_t , and \mathbb{Q}_t the distribution of the offspring of v_t in T (no conditioning is needed there).

Let E_ℓ denote the event that $(G, v)_\ell$ is a tree and contains no more than $c_1 \log(n)d^\ell$ vertices; from (4.38) and Lemma 34, we can choose c_1 such that E_ℓ has probability at least $1 - c_2 d^{2\ell+1}/n$ for some absolute constant c_2 . By iteration, it suffices to show that if E_ℓ holds, there exists a constant $c_3 > 0$ such that

$$d_{\text{TV}}(\mathbb{P}_t, \mathbb{Q}_t) \leq \frac{c_3 \log(n)d^{\ell+2}}{n} \quad \text{for all } t \leq \tau. \quad (4.51)$$

Given \mathcal{F}_t , the probability measure \mathbb{P}_t is as follows: each element i of $V \setminus A_t$ is selected with probability $p_i = P_{v_t i}$. Let \mathbb{P}'_t denote the same probability measure, but where the selection is made over all of V . Using Lemma 35, we first find that

$$d_{\text{TV}}(\mathbb{P}_t, \mathbb{P}'_t) \leq \sum_{i \in A_t} P_{v_t i} \leq c_1 \log(n)d^\ell \cdot \frac{d}{n}.$$

On the other hand, Proposition 15 yields

$$d_{\text{TV}}(\mathbb{P}'_t, \mathbb{Q}_t) \leq c_4 \sum_{i=1}^n P_{v_t i}^2 \leq c_5 \frac{d^2}{n}.$$

Equation (4.51) then results from a straightforward application of the triangle inequality. \square

4.8 Near eigenvectors of G

4.8.1 Functionals on (T, o)

Vertex functionals on trees

Similarly to [29], quantities of interest in the study of B will be tied to functionals on the random inhomogeneous tree defined above. Define a functional $f_{\varphi, t}$ on the set of labeled rooted trees $\mathcal{T}_* \subset \mathcal{G}_*$ by

$$f_{\varphi, t}(T, o) = \sum_{x_t \in \partial(T, o)_t} W_{\iota(o), \iota(x_1)} \cdots W_{\iota(x_{t-1}), \iota(x_t)} \varphi(\iota(x_t)),$$

where (o, x_1, \dots, x_t) is the unique path of length t between o and x_t . Then the following proposition holds:

Proposition 16. *Let $t \geq 0$ be an integer. For any $i, j \in [r]$, the following identities are true:*

$$\mathbf{E}[f_{\varphi_i, t}(T, x)] = \mu_i^t \varphi_i(x), \quad (4.52)$$

$$\mathbf{E}[f_{\varphi_i, t}(T, x) f_{\varphi_j, t}(T, x)] = (\mu_i \mu_j)^t \sum_{s=0}^t \frac{[K^s \varphi^{i, j}](x)}{(\mu_i \mu_j)^s}, \quad (4.53)$$

$$\mathbf{E}\left[(f_{\varphi_i, t+1}(T, x) - \mu_i f_{\varphi_i, t}(T, x))^2\right] = [K^{t+1} \varphi^{i, i}](x). \quad (4.54)$$

where we recall that $\varphi^{i, j} = \varphi_i \odot \varphi_j$.

Adapting functionals to non-backtracking paths

The matrix B considered here acts on (directed) edges, whereas the functionals considered so far are defined on vertices. Consequently, we define the following transformation: for a function $f : \mathcal{G}_* \rightarrow \mathbb{R}$, and a random vector $w \in \mathbb{R}^V$ with expected value \bar{w} , let

$$\vec{\partial}_w f(g, o) = \sum_{e: e_2=o} w_{e_1} f(g_e, o),$$

where g_e denotes the graph g with the edge e_1, e_2 removed.

The expectations from Proposition 16 are then adapted as follows:

Proposition 17. *Let $t \geq 0$ be an integer. For any $i, j \in [r]$, and $\phi \in \ker(P)$, the following identities are true:*

$$\mathbf{E}\left[\vec{\partial}_w f_{\varphi_i,t}(T_x, x)\right] = [P\bar{w}](x) \cdot \mathbf{E}[f_{\varphi_i,t}(T_x, x)], \quad (4.55)$$

$$\mathbf{E}\left[\vec{\partial}_w (f_{\varphi_i,t} \cdot f_{\varphi_j,t})(T_x, x)\right] = [P\bar{w}](x) \cdot \mathbf{E}[f_{\varphi_i,t}(T_x, x) f_{\varphi_j,t}(T_x, x)], \quad (4.56)$$

$$\begin{aligned} \mathbf{E}\left[\vec{\partial}_w [(f_{\varphi_i,t+1} - \mu_i f_{\varphi_i,t})^2](T_x, x)\right] \\ = [P\bar{w}](x) \cdot \mathbf{E}\left[(f_{\varphi_i,t+1}(T_x, x) - \mu_i f_{\varphi_i,t}(T_x, x))^2\right]. \end{aligned} \quad (4.57)$$

The proof for those results makes use of properties specific to moments of Poisson random variables; as with the preceding results, it is deferred to a later section.

4.8.2 Spatial averaging of graph functionals

In this section, we leverage the coupling obtained above to provide bounds on quantities of the form $\frac{1}{n} \sum_{x \in V} f(G, x)$, for local functions f . The tools and results used in this section are essentially identical to those in [29], with a few improvements and clarifications added when necessary.

We begin with a result that encodes the fact that the t -neighbourhoods in G are approximately independent. We say that a function f from \mathcal{G}_* to \mathbb{R} is t -local if $f(g, o)$ is only function of $(g, o)_t$.

Proposition 18. *Let $t \leq c_0 \log(n)$ for some constant $c_0 > 0$. Let $f, \psi : \mathcal{G}_* \rightarrow \mathbb{R}$ be two t -local functions such that $|f(g, o)| \leq \psi(g, o)$ for all $(g, o) \in \mathcal{G}_*$ and ψ is non decreasing by the addition of edges. Then*

$$\text{Var}\left(\sum_{o \in V} f(G, o)\right) \leq c \log(n)^4 n d^{2t} \cdot \sqrt{\mathbf{E}\left[\max_{o \in V} \psi(G, o)^4\right]}.$$

Proof. For $x \in V$, denote by E_x the set $\{\{u, x\} \in E \mid u \leq x\}$; the vector (E_1, \dots, E_n) is an independent vector, and we have

$$Y := \sum_{v \in V} f(G, v) = F(E_1, \dots, E_n).$$

for some measurable function F .

Define now G_x the graph with vertex set V and edge set $\bigcup_{y \neq x} E_y$, and set

$$Y_x = \sum_{v \in V} f(G_x, v).$$

The random variable Y_x is $\bigcup_{y \neq x} E_y$ -measurable, so the Efron-Stein inequality applies:

$$\text{Var}(Y) \leq \sum_{x \in [n]} \mathbf{E}\left[(Y - Y_x)^2\right].$$

For a given $x \in V$, the difference $f(G, o) - f(G_x, o)$ is always zero except if $x \in (G, o)_t$, due to the locality property; consequently,

$$\begin{aligned} |Y - Y_x| &\leq \sum_{o \in V} |f(G, o) - f(G_x, o)| \\ &\leq \sum_{o \in (G, x)_t} \psi(G, o) + \psi(G_x, o) \\ &\leq 2 \max_{x \in [n]} |(G, x)_t| \cdot \max_{o \in V} \psi(G, o), \end{aligned}$$

where we used the non-decreasing property of ψ in the last line. By the Cauchy-Schwarz inequality and equation (4.40), we can write

$$\begin{aligned} \mathbf{E}[(Y - Y_x)^2] &\leq 4 \sqrt{\mathbf{E} \left[\left| \max_{x \in [n]} (G, x)_t \right|^4 \right]} \cdot \sqrt{\mathbf{E} \left[\max_{o \in V} \psi(G, o)^4 \right]} \\ &\leq c_1 t^2 \log(n)^2 d^{2t} \cdot \sqrt{\mathbf{E} \left[\max_{o \in V} \psi(G, o)^4 \right]}. \end{aligned}$$

Using that $t \leq c_0 \log(n)$, and the linearity of expectation, yields the desired bound. \square

We now use our previous coupling results to provide a concentration bound between a functional on graphs and its expectation on trees:

Proposition 19. *Let $t \in \mathbb{N}$ and $f, \psi : \mathcal{G}_* \rightarrow \mathbb{R}$ be as in the previous proposition. Then, with probability at least $1 - \frac{1}{r^2 \log(n)^2}$, the following inequality holds:*

$$\left| \sum_{v \in V} f(G, v) - \mathbf{E} \left[\sum_{x \in [n]} f(T_x, x) \right] \right| \leq c r \log(n)^3 d^{t+1} \sqrt{n} \|\psi\|_*,$$

where $\|\psi\|_*$ is defined as

$$\|\psi\|_* = \left(\mathbf{E} \left[\max_{v \in V} \psi(G, v)^4 \right] \right)^{\frac{1}{4}} \vee \left(\max_{x \in [n]} \mathbf{E} \left[\psi(T_x, x)^2 \right] \right)^{\frac{1}{2}}.$$

Proof. Using the Chebyshev inequality and the variance bound from the preceding proposition, we have with probability at least $1 - \frac{1}{r^2 \log(n)^2}$

$$\left| \sum_{v \in V} f(G, v) - \mathbf{E} \left[\sum_{v \in V} f(G, v) \right] \right| \leq c_1 r \log(n)^3 d^t \sqrt{n} \|\psi\|_*.$$

It then remains to bound the difference between the expectation term and its counterpart on trees. For $x \in V$, let \mathcal{E}_x denote the event that the coupling between $(G, x)_t$ and $(T_x, x)_t$ fails; by

the locality property, $f(G, x) = f(T_x, x)$ on \mathcal{E}_x . Therefore, using the Cauchy-Schwarz inequality,

$$\begin{aligned}
& \left| \sum_{x \in [n]} \mathbf{E}[f(G, x) - f(T_x, x)] \right| \leq \sum_{x \in [n]} \mathbf{E}[|f(G, x)|1_{\mathcal{E}_x} + |f(T_x, x)|1_{\mathcal{E}_x}] \\
& \leq \sum_{x \in [n]} \sqrt{\mathbf{P}(\mathcal{E}_x)} \left(\sqrt{\mathbf{E}[\psi(G, x)^2]} + \sqrt{\mathbf{E}[\psi(T_x, x)^2]} \right) \\
& \leq \sqrt{\frac{c_2 \log(n)^2 d^{2t+2}}{n}} \cdot \sum_{x \in [n]} \left(\mathbf{E}[\psi(G, x)^4]^{\frac{1}{4}} + \sqrt{\mathbf{E}[\psi(T_x, x)^2]} \right) \\
& \leq c_3 \log(n) a d^{t+1} \sqrt{n} \|\psi\|_{\star}.
\end{aligned}$$

It is then straightforward to check that both obtained bounds are less than the RHS in the proposition, upon adjusting c . \square

4.8.3 Structure of near eigenvectors

In the following, the aim is to obtain bounds on the norms and scalar product of the near eigenvectors defined previously. Recall that for $i \in [r]$, we set u_i and v_i as

$$u_i = \frac{B^\ell \chi_i}{\mu_i^\ell} \quad \text{and} \quad v_i = \frac{(B^*)^\ell D_W \tilde{\chi}_i}{\mu_i^{\ell+1}}.$$

The main result of this section is as follows:

Proposition 20. *Let ℓ be small enough so that (4.41) holds. On an event with probability $1 - c_1/\log(n)$, the following inequalities hold for all $i, j \in [r]$, $t \leq 2\ell$ and some absolute constant $c > 0$:*

$$|\langle B^t \chi_i, \chi_j \rangle - \mu_i^t \langle \varphi_i, D_P \varphi_j \rangle| \leq \frac{c r b^2 d^2 \log(n)^6 d^{2t} L^t}{\sqrt{n}}, \quad (4.58)$$

$$|\langle B^t \chi_i, D_W \tilde{\chi}_j \rangle - \mu_i^{t+1} \delta_{ij}| \leq \frac{c r b^2 d^3 L \log(n)^6 d^{2t} L^t}{\sqrt{n}}, \quad (4.59)$$

$$|\langle B^t \chi_i, B^t \chi_j \rangle - \mu_i^t \mu_j^t \Gamma_{U,ij}^{(t)}| \leq \frac{c r b^2 d^2 \log(n)^7 d^{3t} L^{2t}}{\sqrt{n}}, \quad (4.60)$$

$$|\langle (B^*)^t D_W \tilde{\chi}_i, (B^*)^t D_W \tilde{\chi}_j \rangle - \mu_i^{t+1} \mu_j^{t+1} \Gamma_{ij}^{(t+1)}| \leq \frac{c r b^2 d^2 L^2 \log(n)^6 d^{3t} L^{2t}}{\sqrt{n}}, \quad (4.61)$$

$$\|B^{t+1} \chi_i - \mu_i B^t \chi_i\|^2 \leq r d^3 L^2 \rho^{t+1} + \frac{c r b^2 d^3 \log(n)^7 d^{3t} L^{2t}}{\sqrt{n}}. \quad (4.62)$$

Proof. The proof of those inequalities relies on careful applications of Proposition 19 to previously considered functionals. We aim to prove that each of those inequalities hold with probability $1 - c_2/r \log(n)$; we fix in the following an integer $t \leq 2\ell$ and $i, j \in [r]$. Let \mathcal{V}_t be the set of vertices such that $(G, v)_t$ is not a tree; we place ourselves in the event described in Proposition 13 and as a consequence

$$|\mathcal{V}_t| \leq c_3 \log(n)^2 d^{t+1}.$$

We first prove (4.58); let

$$f(g, o) = \mathbf{1}_{(g,o)_t \text{ has no cycles}} \varphi_j(o) \vec{\partial}_1 f_{\varphi_i, t}(g, o).$$

The function f is clearly t -local, and

$$\begin{aligned} |f(g, o)| &\leq \|\varphi_i\|_\infty \|\varphi_j\|_\infty \deg(o) |\partial(g, o)_t| L^t \\ &\leq \frac{b^2}{n} \deg(o) |(g, o)_t| L^t := \psi(g, o). \end{aligned}$$

The function ψ thus defined is non-decreasing by the addition of edges. When $v \notin \mathcal{V}_t$, we notice that

$$f(G, v) = \varphi_j(v) \cdot [T^* B^t \chi_i](v),$$

hence,

$$\left| \langle B^t \chi_i, \chi_j \rangle - \sum_{v \in V} f(G, v) \right| = \left| \sum_{v \in \mathcal{V}_t} \varphi_i(v) T^* B^t \chi_j \right| \leq 2|\mathcal{V}_t| \max_v \psi(G, v),$$

since by the tangle-free property there are at most two paths from v to any vertex in $(G, v)_t$. Furthermore, using the results in [Subsection 4.7.2](#), we find that with probability at least $1 - 1/n$

$$\max_v \psi(G, v) \leq \frac{c_4 b^2 \log(n)^2 d^{t+1} L^t}{n} \quad \text{and} \quad \|\psi\|_\star \leq \frac{c_4 b^2 \log(n)^3 d^{t+1} L^t}{n}.$$

Finally, a direct computation shows that

$$\sum_{x \in [x]} \mathbf{E}[f(T_x, x)] = \sum_{x \in [n]} \varphi_j(x) \cdot d_x \mu_i^t \varphi_i(x) = \mu_i^t \langle \varphi_j, D_P \varphi_i \rangle.$$

Applying [Proposition 19](#) to f and ψ , and using the triangle inequality:

$$\begin{aligned} |\langle B^t \chi_i, \chi_j \rangle - \mu_i^t \langle \varphi_j, D_P \varphi_i \rangle| &\leq \frac{c_5 b^2 \log(n)^4 d^{2t+2} L^t}{n} + \frac{c_6 r b^2 \log(n)^6 d^{2t+2} L^t}{\sqrt{n}} \\ &\leq \frac{c_7 r b^2 d^2 \log(n)^6 d^{2t} L^t}{\sqrt{n}}. \end{aligned}$$

The proof of the other inequalities is very similar, applying [Proposition 19](#) to other functionals from [Subsection 4.8.1](#). To avoid clutter, it is deferred to the appendix. \square

4.9 Proof of [Theorem 30](#)

Having shown [Proposition 20](#), all that remains is simply to gather the preceding bounds, and simplify them to get an easy-to-read summary. Bounds [\(4.26\)](#)-[\(4.28\)](#), as well as [\(4.31\)](#), being straightforward computations, they are deferred to the appendix.

4.9.1 A telescopic trick: proof of [\(4.29\)](#)

Notice that for a $r_0 \times r_0$ matrix M , we have

$$\|M\| \leq r_0 \max_i \|M_i\|. \tag{4.63}$$

where M_i are the columns (or lines) of M . To apply this inequality, we write

$$\|B^\ell u_i - \mu_i^\ell u_i\| \leq \sum_{t=0}^{\ell-1} \mu_i^{\ell-t-1} \|B^{t+1} u_i - \mu_i B^t u_i\|, \tag{4.64}$$

and (4.62) yields

$$\begin{aligned} \|B^{t+1}u_i - \mu_i B^t u_i\|^2 &\leq \mu_i^{-2\ell} \|B^{t+\ell+1}\chi_i - \mu_i B^{t+\ell}\chi_i\|^2 \\ &\leq \mu_i^{-2\ell} \left(rd^3 L^2 \rho^{t+\ell+1} + \frac{c_1 b^2 d^3 \log(n)^7 d^{3(t+\ell)} L^{2(t+\ell)}}{\sqrt{n}} \right). \end{aligned}$$

Since $i \leq r_0$, the bounds $\mu_i^2 \geq \rho \geq 1/d$ apply, so that

$$\|B^{t+1}u_i - \mu_i B^t u_i\|^2 \leq rd^3 L^2 \rho^{t+\ell+1} \mu_i^{-2\ell} + \frac{c_1 b^2 d^3 \log(n)^7 d^{3t+5\ell} L^{2(t+\ell)}}{\sqrt{n}}. \quad (4.65)$$

We now use the (very crude) inequality $\sqrt{x+y} \leq \sqrt{x} + \sqrt{y}$ inside (4.65):

$$\begin{aligned} \|B^\ell u_i - \mu_i^\ell u_i\| &\leq \sum_{t=0}^{\ell-1} \left[\mu_i^{\ell-t-1} \sqrt{rd^3/2} L \rho^{\frac{t+\ell+1}{2}} \mu_i^{-\ell} + \frac{c_1 b d^{3/2} \log(n)^{7/2} d^{\frac{3t+5\ell}{2}} L^{t+\ell}}{n^{1/4}} \right] \\ &\leq \sqrt{rd^3/2} L \rho^{\ell/2} \sum_{t=0}^{\ell-1} \left(\frac{\sqrt{\rho}}{\mu_i} \right)^{t+1} + c_2 b d^2 \log(n)^{9/2} \frac{(Ld^4)^\ell}{n^{1/4}} L^\ell. \end{aligned}$$

The terms in the sum are all less than 1 since $i \leq r_0$, and $\ell < c_3 \log(n)$ implies

$$\|B^\ell u_i - \mu_i^\ell u_i\| \leq c_3 \sqrt{rd^3/2} L \log(n) \rho^{\ell/2} + c_2 b d^2 \log(n)^{9/2} \frac{(aLd^3)^\ell}{n^{1/4}} L^\ell.$$

The bound $(Ld^4)^\ell \leq n^{1/4}$ holds by definition of ℓ , and (4.29) ensues via (4.63).

4.9.2 Bounding $\|B^\ell P_{H^\perp}\|$

Having established the candidates and error bounds for the upper eigenvalues of B^ℓ , it remains to bound the remaining eigenvalues (also called the *bulk*) of the matrix. This is done using a method first employed in [91], and leveraged again in a similar setting in [29, 28]. Our approach will be based on the latter two, adapting the non-backtracking method to the weighted case.

Our first preliminary step is the following lemma:

Lemma 36. *On an event with probability at least $1 - 1/\log(n)$, for any $t \leq \ell$, any unit vector $w \in H^\perp$ and $i \in [r_0]$, one has*

$$|\langle (B^*)^t D_W \tilde{\chi}_i, w \rangle| \leq \sqrt{rd^3/2} L^2 \rho^{t/2} + \frac{c_4 b d^{3/2} \log(n)^{9/2} d^{2\ell} L^\ell}{n^{1/4}}.$$

Proving this bound is done through the same telescopic sum trick as above, and is done in the appendix.

Tangle-free decomposition of B^ℓ

We adapt here the decomposition first used in [29] to our setting. Through the remainder of this section, we shall consider B as an operator on $\vec{E}(V)$ instead of \vec{E} , setting $B_{ef} = 0$ whenever $e \notin \vec{E}$ or $f \notin \vec{E}$. This yields a matrix with B as a principal submatrix and zeros everywhere else, thus the non-zero spectrum stays identical.

For $e, f \in \vec{E}(V)$, and $t \geq 0$, we define Γ_{ef}^k the set of non-backtracking paths of length k from e to f ; further, for an edge e we define X_e the indicator variable of $e \in \vec{E}$, and $A_e = X_e W_e$, so that A is the (weighted) adjacency matrix of G .

We then have that

$$(B^k)_{ef} = \sum_{\gamma \in \Gamma_{ef}^{k+1}} X_e \prod_{s=1}^k A_{\gamma_s \gamma_{s+1}}.$$

Define F_{ef}^k the set of ℓ -tangle-free paths (i.e. the set of paths γ such that the subgraph induced by γ is tangle-free). Then, whenever the graph G is tangle-free, for all $k \leq \ell$ the matrix B^k is equal to $B^{(k)}$, with

$$(B^{(k)})_{ef} = \sum_{\gamma \in F_{ef}^{k+1}} X_e \prod_{s=1}^k A_{\gamma_s \gamma_{s+1}}.$$

Define now the ‘‘centered’’ versions of the weighted and unweighted adjacency matrices \underline{A} and \underline{X} by

$$\underline{A}_{ij} = A_{ij} - Q_{ij} \quad \text{and} \quad \underline{X}_{ij} = X_{ij} - P_{ij}$$

for every $i \neq j$, and its centered non-backtracking counterpart as

$$(\delta^{(k)})_{ef} = \sum_{\gamma \in F_{ef}^{k+1}} \underline{X}_{ij} \prod_{s=1}^k \underline{A}_{\gamma_s \gamma_{s+1}},$$

with the convention that the product over an empty set is equal to 1.

Recall that for any two sets of real numbers $(x_i), (y_i)$, we have the following:

$$\prod_{s=0}^{\ell} x_s = \prod_{s=0}^{\ell} y_s + \sum_{t=0}^{\ell} \prod_{s=0}^{t-1} y_s (x_t - y_t) \prod_{s=t+1}^{\ell} x_s.$$

Applying this formula to the above definitions, and separating the case $t = 0$ in the sum yields

$$\begin{aligned} B_{ef}^{(\ell)} &= \delta_{ef}^{(\ell)} + \sum_{\gamma \in F_{ef}^{\ell+1}} Q_e \prod_{s=1}^{\ell} A_{\gamma_s \gamma_{s+1}} \\ &\quad + \sum_{t=1}^{\ell} \sum_{\gamma \in F_{ef}^{\ell+1}} \underline{X}_e \prod_{s=1}^{t-1} \underline{A}_{\gamma_s \gamma_{s+1}} Q_{\gamma_t \gamma_{t+1}} \prod_{s=t+1}^{\ell} A_{\gamma_s \gamma_{s+1}}. \end{aligned} \quad (4.66)$$

Define now $F_{t,ef}^{\ell+1} \subset \Gamma_{ef}^{\ell+1}$ the set of non-backtracking *tangled* paths γ such that $(\gamma_0, \dots, \gamma_t) \in F_{eg}^t$, $(\gamma_{t+1}, \dots, \gamma_{\ell+1}) \in F_{g'f}^{\ell-t}$ for some edges $g, g' \in \vec{E}(V)$. As an edge case, $F_{0,ef}^{\ell+1}$ is the set of tangled paths γ such that $(\gamma_0, \gamma_1) = e_1$ and $(\gamma_1, \dots, \gamma_{\ell+1}) \in F_{g'f}^{\ell}$ for some $g' \in \vec{E}(V)$ (note that necessarily $e_2 = g'_1$), and similarly for $F_{\ell,ef}$. Finally, we introduce the two matrices M and $M^{(2)}$ as

$$M_{ef} = \mathbf{1}\{e \rightarrow f\} Q_e \quad \text{and} \quad M_{ef}^{(2)} = \mathbf{1}(e \xrightarrow{2} f) Q_{e_2 f_1}$$

for $e, f \in \vec{E}(V)$, where $e \xrightarrow{2} f$ means that there exists a non-backtracking path of length two between e and f . Then, equation (4.66) can be rewritten as

$$B^{(\ell)} = \delta^{(\ell)} + M D_W B^{(\ell-1)} + \sum_{t=1}^{\ell-1} \delta^{(t-1)} M^{(2)} D_W B^{(\ell-t-1)} + \delta^{(\ell-1)} M - \sum_{t=0}^{\ell} R_t^{(\ell)}, \quad (4.67)$$

where

$$(R_t^{(\ell)})_{ef} = \sum_{\gamma \in F_{t,ef}^{\ell+1}} \underline{X}_e \prod_{s=1}^{t-1} A_{\gamma_s \gamma_{s+1}} Q_{\gamma_t \gamma_{t+1}} \prod_{s=t+1}^{\ell} A_{\gamma_s \gamma_{s+1}}$$

$$(R_0^{(\ell)})_{ef} = \sum_{\gamma \in F_{t,ef}^{\ell+1}} Q_e \prod_{s=1}^{\ell} A_{\gamma_s \gamma_{s+1}}.$$

Note that $M^{(2)}$ is pretty close to a modified version of Q ; more specifically, we make the decomposition

$$M^{(2)} = TQT^* + \tilde{M} = \sum_{k=1}^r \mu_k \chi_k \tilde{\chi}_k^* + \tilde{M}.$$

Then, the following decomposition holds:

$$B^{(\ell)} = \delta^{(\ell)} + MD_W B^{(\ell-1)} + \sum_{t=1}^{\ell-1} \sum_{k=1}^r \mu_k \delta^{(t-1)} \chi_k \tilde{\chi}_k^* D_W B^{(\ell-t-1)}$$

$$+ \sum_{t=1}^{\ell-1} \delta^{(t-1)} \tilde{M} B^{(\ell-t-1)} + \delta^{(\ell-1)} M - \sum_{t=0}^{\ell} R_t^{(\ell)}.$$

Noticing that $\|M\| \leq d$ and $\|\chi_k\| \leq d \log(n)$, the following lemma ensues:

Lemma 37. *On an event with probability at least $1 - 1/\log(n)$, the following inequality holds for any normed vector $x \in \mathbb{R}^{\vec{E}(V)}$:*

$$\|B^\ell x\| \leq \|\delta^{(\ell)}\| + L \|MB^{\ell-1}\| + d \log(n) \sum_{t=1}^{\ell-1} \|\delta^{(t-1)}\| \sum_{k=1}^r \left| \langle D_W \tilde{\chi}_k, B^{\ell-t-1} x \rangle \right|$$

$$+ \sum_{t=1}^{\ell-1} \|\delta^{(t-1)} \tilde{M} B^{\ell-t-1}\| + d \|\delta^{(\ell-1)}\| - \sum_{t=0}^{\ell} \|R_t^{(\ell)}\|.$$

Norm bounds

It then remains to bound the different quantities in the lemma above; this is done in another section, using a trace bound method. The results are as follows:

Proposition 21. *On an event with probability $1 - c_0/\log(n)$, for any $k \leq c_1 \log(n)$, the following bounds hold with probability at least $1 - 1/\ln(n)^2$:*

$$\|\delta^{(k-1)}\| \leq cd^3 \log(n)^{17} (\sqrt{\rho} \vee L)^k, \quad (4.68)$$

$$\|MB^{k-1}\| \leq \frac{cd^{7/2} L \log(n)^7 d^k L^k}{\sqrt{n}}, \quad (4.69)$$

$$\|\delta^{(t-1)} \tilde{M} B^{k-t-1}\| \leq \frac{cd^{13/2} L \log(n)^{24} d^k (\sqrt{\rho} \vee L)^k}{\sqrt{n}}, \quad (4.70)$$

$$\|R_t^{(k)}\| \leq \frac{cd^2 \log(n)^{22} d^k L^k}{n}. \quad (4.71)$$

Using these bounds, we are now finally able to prove (4.30):

Proof. By definition of ℓ , $d^\ell \leq n^{1/4}$ so most of the summands in Lemma 37 are negligible with respect to the others. More precisely, we have

$$\|B^\ell x\| \leq c_1 \left(\|\delta^\ell\| + d \log(n) \sum_{t=1}^{\ell-1} \|\delta^{(t-1)}\| \sum_{k=1}^r \left| \langle D_W \check{\chi}_k, B^{\ell-t-1} x \rangle \right| \right). \quad (4.72)$$

When $k \in [r_0]$, Lemma 36 implies that

$$\left| \langle D_W \check{\chi}_k, B^{\ell-t-1} x \rangle \right| \leq \sqrt{r} d^{3/2} L^2 \rho^{t/2} + \frac{c_4 b d^{3/2} \log(n)^{9/2} d^{2\ell} L^\ell}{n^{1/4}},$$

and by definition of ℓ , $d^{2\ell} L^\ell \leq (1 \wedge \sqrt{\rho})^\ell$ so the second term is bounded above by the first. On the other hand, for $k \in [r] \setminus [r_0]$, we can use equation (4.61) as follows:

$$\|(B^*)^t D_W \check{\chi}_i\|^2 \leq \mu_i^{2t+2} \Gamma_{V,ii}^{(t+1)} + \frac{c r b^2 d^4 L^2 \log(n)^6 d^{3t} L^{2t}}{\sqrt{n}}.$$

We now apply Lemma 30:

$$\Gamma_{V,ii}^{(t+1)} \leq \sum_{s=0}^{t+1} \frac{r d^2 L^2 \rho^s}{\mu_i^{2s}} \leq c r d^2 \log(n) L^2 \rho^{t+1} \mu_i^{-2t-2},$$

since $\mu_i^2 < \rho$; the second term being negligible before the first,

$$\left| \langle D_W \check{\chi}_k, B^{\ell-t-1} x \rangle \right| \leq \left\| (B^*)^{\ell-t-1} D_W \check{\chi}_i \right\| \leq c r d \log(n) L \rho^{\frac{\ell-t}{2}}.$$

We can now apply the above bounds on the scalar product as well as those of Proposition 21 to equation (4.72), and we get

$$\begin{aligned} \|B^\ell x\| &\leq c_2 d^{5/2} L \log(n)^{17} (\sqrt{\rho} \vee L)^\ell + c_3 r^2 d^6 L^2 \log(n)^{20} (\sqrt{\rho} \vee L)^\ell \\ &\quad + c_4 d^4 \log(n)^{17} (\sqrt{\rho} \vee L)^\ell \\ &\leq c r^2 d^6 L^2 \log(n)^{20} (\sqrt{\rho} \vee L)^\ell, \end{aligned}$$

which ends the proof of (4.30). \square

4.10 Trace method: proof of Proposition 21

The aim of this section is to prove the bounds in Proposition 21; we leverage here the powerful trace method introduced by Füredi and Komlós [58], and already used with success in [29] and [28]. We only prove (4.68) in this section, all other bounds being proven in the appendix.

Let m be a parameter to be fixed later. We start with the classical bound

$$\begin{aligned} \|\delta^{(k-1)}\|^{2m} &= \|\delta^{(k-1)} \delta^{(k-1)*}\|^m \\ &= \|(\delta^{(k-1)} \delta^{(k-1)*})^m\| \\ &\leq \text{Tr}((\delta^{(k-1)} \delta^{(k-1)*})^m). \end{aligned}$$

Expanding the trace above gives

$$\begin{aligned}
\|\delta^{(k-1)}\|^{2m} &\leq \sum_{(e_1, \dots, e_{2m})} \prod_{i=1}^m (\delta^{(k-1)})_{e_{2i-1}, e_{2i}} (\delta^{(k-1)})_{e_{2i+1}, e_{2i}} \\
&= \sum_{\gamma \in W_{k,m}} \prod_{i=1}^{2m} \underline{X}_{\gamma_{i,0}\gamma_{i,1}} \prod_{s=2}^k \underline{A}_{\gamma_{i,s-1}\gamma_{i,s}},
\end{aligned} \tag{4.73}$$

where $W_{k,m}$ is the set of sequences of paths $(\gamma_1, \dots, \gamma_{2m})$ such that $\gamma_i = (\gamma_{i,0}, \dots, \gamma_{i,k})$ is non-backtracking tangle-free of length k , and with boundary conditions that for all $i \in [m]$,

$$(\gamma_{2i,k-1}, \gamma_{2i,k}) = (\gamma_{2i-1,k-1}, \gamma_{2i-1,k}) \quad \text{and} \quad (\gamma_{2i+1,0}, \gamma_{2i+1,1}) = (\gamma_{2i,0}, \gamma_{2i,1}), \tag{4.74}$$

with the convention $\gamma_{2m+1} = \gamma_1$. All the random variables in the expression above are centered and independent as soon as they are supported by distinct edges, so the expectation of each term in the sum is zero except when each (unoriented) edge is visited at least twice. We let $W'_{k,m}$ be the set of all such sequences of paths. To $\gamma \in W'_{k,m}$, we associate the graph $G_\gamma = (V_\gamma, E_\gamma)$ of visited vertices and edges, and let

$$v_\gamma = |V_\gamma| \quad \text{and} \quad e_\gamma = |E_\gamma|.$$

For an unoriented edge $e \in E_\gamma$, we define its multiplicity m_e as the number of times e is visited in γ ; we also let S_γ be the set of starting edges in γ , that is

$$S_\gamma = \{(\gamma_{i,0}, \gamma_{i,1}) \mid i \in [2m]\}.$$

Using these definitions, we can bound the expectation as follows:

$$\mathbf{E}\left[\|\delta^{(k-1)}\|^{2m}\right] \leq \sum_{\gamma \in W'_{k,m}} \prod_{e \in S_\gamma} \mathbf{E}[|\underline{X}_e| \cdot |\underline{A}_e|^{m_e-1}] \prod_{e \notin S_\gamma} \mathbf{E}[|\underline{A}_e|^{m_e}].$$

We now bound the two terms in the products above: let e be an edge, and $p \geq 2$ be any multiplicity. Then conditioning on X_e ,

$$\begin{aligned}
\mathbf{E}[|\underline{A}_e|^{m_e}] &= P_e \mathbf{E}[|W_e - P_e \mathbf{E}[W_e]|^p] + (1 - P_e) P_e^p \mathbf{E}[W_e]^p \\
&\leq P_e L^{p-2} \left(1 + \frac{d}{n}\right)^{p-2} \mathbf{E}[(W_e - P_e \mathbf{E}[W_e])^2] + \left(\frac{dL}{n}\right)^{p-2} \frac{dP_e}{n} \mathbf{E}[W_e]^2 \\
&\leq P_e L^{p-2} \left(1 + \frac{d}{n}\right)^{p-2} \mathbf{E}[W_e^2] + P_e L^{p-2} \mathbf{E}[W_e^2] \left(\frac{d}{n}\right)^{p-2} \\
&\leq K_e L^{p-2} \left(1 + \frac{d}{n}\right)^p.
\end{aligned}$$

The other product is trickier; whenever $p \geq 3$, a similar computation yields

$$\mathbf{E}[|\underline{X}_e| \cdot |\underline{A}_e|^{p-1}] \leq K_e L^{p-3} \left(1 + \frac{d}{n}\right)^p.$$

On the other hand if $p = 2$,

$$\mathbf{E}[|\underline{X}_e| \cdot |\underline{A}_e|] \leq \frac{d}{n} L \left(1 + \frac{d}{n}\right)^2.$$

As a consequence, for $\gamma \in W'_{k,m}$, we define $S'_\gamma \subseteq S_\gamma$ the set of starting edges with multiplicity 2. Then

$$\mathbf{E} \left[\|\delta^{(k-1)}\|^{2m} \right] \leq \sum_{\gamma \in W'_{k,m}} \left(1 + \frac{d}{n}\right)^{2km} \left(\frac{d}{n}\right)^{|S'_\gamma|} d^{2m} L^{2km-2e_\gamma} \prod_{e \notin S'_\gamma} K_e,$$

where we used $L^{-1} \leq d$ and $S_\gamma = 2m$.

We now partition the paths in $W'_{k,m}$ as follows: we say that $\gamma \sim \gamma'$ if there exists a permutation $\sigma \in \mathfrak{S}_n$ such that $\gamma_{i,t} = \sigma(\gamma'_{i,t})$ for all $i, t \in [2m] \times [k]$. Clearly, all parameters such as v_γ , e_γ and $|S'_\gamma|$ are constant on any equivalence class; therefore it makes sense to define $\mathcal{W}_{k,m}(v, e)$ the set of equivalence classes of $W'(k, m)$ such that $v_\gamma = v$ and $e_\gamma = e$. Then, a path counting argument performed in [29] yields the following estimation:

Lemma 38. *Let v, e be integers such that $e - v + 1 \geq 0$. Then*

$$\mathcal{W}_{k,m}(v, e) \leq k^{2m} (2km)^{6m(e-v+1)}. \quad (4.75)$$

All that remains to bound the sum above is to control the contribution of a single equivalence class; this is done through this lemma:

Lemma 39. *Let $\gamma \in W'_{k,m}$ such that $v_\gamma = v$, $e_\gamma = e$ and $|S'_\gamma| = s$. We have*

$$\sum_{\gamma' \sim \gamma} \prod_{f \notin S'_{\gamma'}} K_f \leq d^{2m} n^{v-e+s} \rho^e (\Psi^2)^{3(e-v)+8m}. \quad (4.76)$$

Proof. For a sequence of paths $\gamma \in W'_{k,m}$, denote by E'_γ the set $E_\gamma \setminus S'_\gamma$. Then, due to the boundary conditions in (4.74), the graph G'_γ induced by E'_γ is connected. We let v_j (resp. $v_{\geq j}$) be the number of vertices with degree j (resp. at least j) in G'_γ . Again, by (4.74), removing an edge in S'_γ does not create a vertex of degree 1; therefore we have

$$v_1 \leq 4m,$$

since a vertex of G_γ can only be of degree 1 if it is an endpoint of γ_i for some $i \in [2m]$. Additionally, edge and vertex counting yields

$$v_1 + v_2 + v_{\geq 3} \geq v - s \quad \text{and} \quad v_1 + 2v_2 + 3v_{\geq 3} \leq 2(e - s),$$

since removing an edge in S'_γ removes at most one vertex from G_γ . Combining those inequalities gives

$$v_{\geq 3} + v_1 \leq 2(e - s) - 2(v - s) + 2v_1 \leq 2(e - v) + 8m; \quad (4.77)$$

this inequality encodes the fact that in a union of paths most vertices are of degree 2. We now reduce G'_γ into a multigraph $\hat{G}_\gamma = (\hat{V}_\gamma, \hat{E}_\gamma)$ as follows: \hat{V}_γ is the set of vertices in G'_γ with degree different from 2, and we add an edge between two vertices x_1 and x_2 of \hat{V}_γ for each path between x_1 and x_2 in G'_γ . For $\hat{f} \in \hat{E}_\gamma$, we annotate \hat{f} with the length $q_{\hat{f}}$ of its corresponding path in G'_γ .

We let \hat{v} and \hat{e} be the number of vertices and edges of \hat{G}_γ ; a sequence $\gamma' \sim \gamma$ is uniquely determined by an embedding of \hat{V}_γ in $[n]$ and for each edge $\hat{f} \in \hat{E}_\gamma$, an embedding of \hat{f} as a

path of length $q_{\hat{f}}$. As a result, we have

$$\begin{aligned} \sum_{\gamma' \sim \gamma} \prod_{f \notin S'_{\gamma'}} Q_f &\leq \sum_{y_1, \dots, y_{\hat{v}} \in [n]^{\hat{v}}} \prod_{\hat{f}=(y_i, y_j) \in \hat{E}_{\gamma}} \sum_{x_1, \dots, x_{q_{\hat{f}}-1} \in [n]} \prod_{t=1}^{q_{\hat{f}}} K_{x_{t-1}, x_t} \\ &= \sum_{y_1, \dots, y_{\hat{v}} \in [n]^{\hat{v}}} \prod_{\hat{f}=(y_i, y_j) \in \hat{E}_{\gamma}} (K^{q_{\hat{f}}})_{y_i, y_j} \\ &\leq \sum_{y_1, \dots, y_{\hat{v}} \in [n]^{\hat{v}}} \prod_{\hat{f} \in \hat{E}_{\gamma}} \left(\frac{\Psi^2}{n} \rho^{q_{\hat{f}}} \right), \end{aligned}$$

using (4.34) and recalling that $\Psi = L^2/\rho$. Now, notice that

$$\sum_{\hat{f} \in \hat{E}_{\gamma}} q_{\hat{f}} = |E'_{\gamma}| = e - s \quad \text{and} \quad \hat{e} - \hat{v} = |E'_{\gamma}| - |V'_{\gamma}| \geq e - v - s;$$

further $\hat{e} \leq \hat{v} + e - v - s \leq 3(e - v) + 8m - s$ using (4.77) and the inequality above. We finally find

$$\begin{aligned} \sum_{\gamma' \sim \gamma} \prod_{f \notin S'_{\gamma'}} Q_f &\leq n^{\hat{v} - \hat{e}} (\Psi^2)^{\hat{e}} \rho^{e-s} \\ &\leq n^{v-e+s} \rho^{e-s} (\Psi^2)^{3(e-v)+8m-s}, \end{aligned}$$

which ends the proof of Lemma 39, since $\Psi^2 \geq 1$ and $\rho^{-1} \leq a$. \square

We now are able to conclude; the contribution of one equivalence class in $\mathcal{W}_{k,m}(v, e)$ is less than

$$\begin{aligned} C_{\gamma} &= \left(1 + \frac{d}{n}\right)^{2km} \left(\frac{d}{n}\right)^{|S'_{\gamma}|} d^{2m} L^{2km-2e} \sum_{\gamma' \sim \gamma} \prod_{e \notin S'_{\gamma'}} K_e \\ &\leq c_1^{2m} d^{6m} n^{-|S'_{\gamma}|} L^{2km-2e} n^{v-e+|S'_{\gamma}|} \rho^e (\Psi^2)^{3(e-v)+8m} \\ &\leq c_1^{2m} d^{6m} n^{v-e} (\rho\Psi)^{km-e} \rho^e (\Psi^2)^{3(e-v)+8m} \\ &\leq c_1^{2m} d^{6m} \rho^{km} \left(\frac{\Psi}{d}\right)^{km-e} n^{1-g} (\Psi^2)^{3g+8m}, \end{aligned}$$

with $g = e - v + 1$ and we used that $L = \sqrt{\rho\Psi/d}$ and the bound

$$\left(1 + \frac{d}{n}\right)^k \leq \exp\left(\frac{dk}{n}\right) \leq c_1.$$

Summing over all equivalence classes now gives

$$\begin{aligned} \mathbf{E}\left[\|\delta^{(k-1)}\|^{2m}\right] &\leq \sum_{e=1}^{km} \sum_{v=1}^{e+1} |\mathcal{W}_{k,m}(v, e)| \max_{[\gamma] \in \mathcal{W}_{k,m}(v, e)} C_{\gamma} \\ &\leq \sum_{e=1}^{km} \sum_{v=1}^{e+1} k^{2m} (2km)^{6m(e-v+1)} c_1^{2m} d^{6m} \rho^{km} \left(\frac{\Psi}{d}\right)^{km-e} n^{1-g} (\Psi^2)^{3g+8m} \\ &\leq n(c_1 d^3 k)^{2m} \rho^{km} \sum_{e=1}^{km} \left(\frac{\Psi}{d}\right)^{km-e} \sum_{g=0}^{\infty} \left(\frac{\Psi^6 (2km)^{6m}}{n}\right)^g. \end{aligned} \tag{4.78}$$

We set the parameter m to

$$m = \left\lceil \frac{\log\left(\frac{n}{\Psi^6}\right)}{12 \log(\log(n))} \right\rceil;$$

when $n \geq c_2 \Psi^6$ for some absolute constant c_2 , we have

$$\frac{\Psi^6 (2km)^{6m}}{n} < \frac{1}{2} \quad \text{and} \quad n^{\frac{1}{2m}} \leq \log(n)^{12}.$$

The infinite sum inside (4.78) thus converges, and

$$\mathbf{E} \left[\|\delta^{(k-1)}\|^{2m} \right]^{\frac{1}{2m}} \leq c_3 d^3 \log(n)^{14} \sqrt{\rho}^k \left(1 \vee \sqrt{\frac{\Psi}{d}} \right)^k.$$

Finally, from the definition of Ψ , $\sqrt{\rho} \left(1 \vee \sqrt{\Psi/d} \right) = \sqrt{\rho} \vee L$, hence (4.68) by a Markov bound.

4.11 Applications of Theorem 22

4.11.1 Proof of Proposition 11

Let x be an eigenvector of B associated with the eigenvalue λ ; the eigenvalue equation for x reads

$$\lambda x_e = \sum_{e \rightarrow f} W_f x_f. \quad (4.79)$$

On the other hand, the definition $y = S^* D_W x$ expands to

$$y_i = \sum_{e: e_1=i} W_e x_e.$$

Applying equation (4.79) to e and e^{-1} yields

$$\lambda x_e = y_{e_2} - W_e x_{e^{-1}} \quad \text{and} \quad \lambda x_{e^{-1}} = y_{e_1} - W_e x_e,$$

and as a result

$$\lambda^2 x_e = \lambda y_{e_2} - \lambda W_e x_{e^{-1}} = \lambda y_{e_2} - W_e (y_{e_1} - W_e x_e).$$

Rearranging the terms, we find an expression for x_e :

$$x_e = \frac{\lambda y_{e_2} - W_e y_{e_1}}{\lambda^2 - W_e^2}; \quad (4.80)$$

in particular $y \neq 0$ if $x \neq 0$. Plugging (4.80) into the eigenvalue equation (4.79), we get for $i, j \in [n]$

$$\frac{\lambda^2 y_i - \lambda W_{ij} y_j}{\lambda^2 - W_{ij}^2} = \sum_{\substack{k \sim i \\ k \neq j}} W_{ik} \frac{\lambda y_k - W_{ik} y_i}{\lambda^2 - W_{ik}^2},$$

and we rearrange to find

$$\frac{\lambda^2 y_i}{\lambda^2 - W_{ij}^2} - \frac{W_{ij}^2 y_i}{\lambda^2 - W_{ij}^2} = \sum_{k \sim i} \frac{\lambda W_{ik}}{\lambda^2 - W_{ik}^2} y_k - \sum_{k \sim i} \frac{W_{ik}^2}{\lambda^2 - W_{ik}^2} y_i.$$

The fraction on the LHS cancels out, and writing the RHS as a matrix product

$$y = \tilde{A}(\lambda) y - \tilde{D}(\lambda) y,$$

the desired result.

4.11.2 Proof of Theorem 23

Our first step is to show that the matrices involved in Proposition 11 approximate the matrices A and ρI . If $\lambda^2 \geq 2L^2$, we have

$$\left| \lambda \tilde{A}_{ij}(\lambda) - A_{ij} \right| = \mathbf{1}\{i \sim j\} \left| \frac{W_{ij}}{1 - \frac{W_{ij}^2}{\lambda^2}} - W_{ij} \right| \leq \mathbf{1}(i \sim j) \frac{2LW_{ij}^2}{\lambda^2},$$

which implies using the Gershgorin circle theorem

$$\left\| \lambda \tilde{A}_{ij}(\lambda) - A_{ij} \right\| \leq \frac{2L}{\lambda^2} \max_i \sum_{j \sim i} W_{ij}^2 \leq \frac{4L\rho}{\lambda^2}. \quad (4.81)$$

Similarly,

$$\begin{aligned} \left| \lambda^2 \tilde{D}_{ii}(\lambda) - \rho \right| &\leq \frac{2L^2}{\lambda^2} \sum_{j \sim i} W_{ij}^2 + \left| \sum_{j \sim i} W_{ij}^2 - \rho \right| \\ &\leq \left(\frac{4L^2}{\lambda^2} + \varepsilon \right) \cdot \rho. \end{aligned} \quad (4.82)$$

We now take $\lambda = \lambda_i$ with $i \in [r_0]$; then there is a vector y that is a singular value of

$$-\lambda_i \delta(\lambda_i) = A - \left(\lambda_i + \frac{\rho}{\lambda_i} \right) \mathbf{I} + (\lambda \tilde{A}(\lambda_i) - A) - \lambda_i^{-1} (\lambda^2 \tilde{D}(\lambda) - \rho \mathbf{I}).$$

We can thus apply Weyl's inequality [125] to find that there exists an eigenvalue ν_i of A such that

$$\left| \nu_i - \left(\lambda_i + \frac{\rho}{\lambda_i} \right) \right| \leq \frac{4L\rho}{\lambda_i^2} + \left(\frac{4L^2}{\lambda_i^2} + \varepsilon \right) \cdot \frac{\rho}{\lambda_i}.$$

Now, we use Theorem 22 to find that $|\lambda_i - \mu_i| \leq \sigma$, and we have $\sigma = o(\rho)$ whenever n is large enough by virtue of (4.32). Since

$$|\lambda_i - \mu_i| \leq \sigma \quad \text{and} \quad \left| \frac{\rho}{\lambda_i} - \frac{\rho}{\mu_i} \right| \leq \frac{\rho}{\lambda_i \mu_i} \sigma \leq c_0 \sigma,$$

equation (4.9) ensues by noticing that $\lambda_i > c_1 \mu_i$ for some constant c_1 and σ is negligible before the other error terms.

Assume now that $\delta_i \geq 2\sigma$; examining the proof of Theorem 22, we have the existence of an eigenvector ξ of B associated with λ_i such that

$$\|\xi - u_i\| \leq \frac{3\sigma \|u_i\|}{\delta_i - \sigma}.$$

Proposition 11 implies that the vector $y = S^* D_W \xi$ is a null vector of the deformed laplacian $\delta(\lambda) = \mathbf{I} - \tilde{A}(\lambda) + \tilde{D}(\lambda)$. Notice that the matrix $S^* D_W^2 S$ is a diagonal matrix such that

$$[S^* D_W^2 S]_{ii} = \sum_{j \sim i} W_{ij}^2 \leq 2\rho,$$

from which we have

$$\|y - S^* D_W u_i\| \leq \frac{6\sigma \sqrt{\rho} \|u_i\|}{\delta_i - \sigma}.$$

We now follow the line of proof of Theorem 22; we first find

$$\langle S^* D_W u_i, \varphi_i \rangle = \mu_i^{-\ell} \langle B^\ell \chi_i, D_W \check{\chi}_i \rangle,$$

and combine it with (4.59) to obtain

$$|\langle S^* D_W u_i, \varphi_i \rangle - \mu_i| \leq \sigma. \quad (4.83)$$

Computing $\|S^* D_W u_i\|$ is trickier; we find

$$\begin{aligned} \langle S^* D_W u_i, S^* D_W u_i \rangle &= \mu_i^{-2\ell} \langle S^* D_W B^\ell \chi_i, S^* D_W B^\ell \chi_i \rangle \\ &= \mu_i^{-2\ell} \langle S^* D_W B^\ell \chi_i, T^* J D_W B^\ell \chi_i \rangle \\ &= \mu_i^{-2\ell} \langle T S^* D_W B^\ell \chi_i, (B^*)^\ell D_W \check{\chi}_i \rangle. \end{aligned}$$

Writing the coefficients of TS^*D_W explicitly, we have

$$[TS^*D_W]_{ef} = W_f \sum_{i \in [n]} \mathbf{1}\{e_2 = i\} \mathbf{1}\{f_1 = i\} = B_{ef} + [JD_W]_{ef},$$

which yields

$$\langle S^*D_W u_i, S^*D_W u_i \rangle = \mu_i^{-2\ell} \left(\langle B^{2\ell+1} \chi_i, D_W \check{\chi}_i \rangle + \langle B^\ell D_W \check{\chi}_i, B^\ell D_W \check{\chi}_i \rangle \right).$$

Those scalar products correspond to equations (4.59) and (4.61), respectively, and we thus get

$$\left| \|S^*D_W u_i\|^2 - \mu_i^2(1 + \Gamma_{V,ii}^{(\ell)}) \right| \leq 2\sigma.$$

The hypothesis $K\mathbf{1} = \rho\mathbf{1}$ allows us to approximate $\Gamma_{V,ii}^{(\ell)}$ efficiently:

$$\Gamma_{V,ii}^{(\ell)} = \sum_{t=0}^{\ell} \frac{\langle \mathbf{1}, K^{t+1} \varphi^{i,i} \rangle}{\mu_i^{2t+2}} = \sum_{t=0}^{\ell} \left(\frac{\rho}{\mu_i^2} \right)^{t+1}$$

since $\|\varphi_i\| = 1$, and we have as in the proof of Theorem 22

$$\left| \Gamma_{V,ii}^{(\ell)} - \frac{\rho/\mu_i^2}{1 - \rho/\mu_i^2} \right| \leq \sigma.$$

Gathering the previous bounds, we eventually arrive at

$$\left| \|S^*D_W u_i\|^2 - \frac{\mu_i^2}{1 - \rho/\mu_i^2} \right| \leq 3\sigma. \quad (4.84)$$

The exact same computations imply that

$$\|u_i\|^2 \leq \frac{d}{1 - \rho/\mu_i^2} + c_5\sigma,$$

and thus noticing that $\mu_i \geq \sqrt{\rho}$

$$\left\| \frac{y}{\|y\|} - \frac{S^*D_W u_i}{\|S^*D_W u_i\|} \right\| \leq \frac{c_6 \sigma \sqrt{d}}{\delta_i - \sigma}.$$

Combining this error bound with (4.83) and (4.84), we find the following result:

$$\left| \left\langle \frac{y}{\|y\|}, \varphi_i \right\rangle - \sqrt{1 - \frac{\rho}{\mu_i^2}} \right| \leq \frac{c_7 \sigma \sqrt{d}}{\delta_i - \sigma}.$$

The final step is to use the Davis-Kahan theorem [126] as follows: there exists an eigenvector ζ of A with associated eigenvalue ν_i , and such that

$$\left\| \zeta - \frac{y}{\|y\|} \right\| \leq \frac{c_8 \left(\frac{4L\rho}{\lambda_i^2} + \left(\frac{4L^2}{\lambda_i^2} + \varepsilon \right) \cdot \frac{\rho}{\lambda_i} \right)}{\delta_i}.$$

This error term dominates all the other ones found above, hence the bound in Theorem 23.

The proof of Corollary 5 follows along the same lines; however, we have directly

$$\tilde{A}(\lambda) = \frac{\lambda A}{\lambda^2 - 1} \quad \text{and} \quad \tilde{D}(\lambda) = \frac{d_0}{\lambda^2 - 1} \mathbf{I},$$

and thus the approximation bounds (4.81) and (4.82) become superfluous.

4.11.3 Proof of Theorem 24

We first link the SBM setting to the one of Theorem 22. In the unweighted case, we have $Q = K = P$, and the eigenvector equation $P\mathbf{1} = \alpha\mathbf{1}$ yields $\rho = \alpha$. It is easy to check that whenever n is large enough, the r_0 defined in Theorem 24 satisfies the assumptions of Theorem 22, with $\tau = 1/(\alpha\mu_{r_0}^2) < 1$. Equation (4.11) ensures that $\|\varphi_i\|_\infty \leq c/\sqrt{n}$ for some absolute constant $c > 0$, therefore $b = O(1)$. Finally, since $\tau^{-1} = \alpha\mu_{r_0}$, we have

$$C_0 \leq c\alpha \log(n)^{25} \quad \text{and} \quad n_0 \leq \exp(c \log(d) \log(\log(n))).$$

An application of Theorem 22 thus directly yields the bound on the eigenvalues of B ; regarding the eigenvectors, notice that as in the proof of Theorem 23

$$\|u_i\|^2 = \frac{\alpha}{1 - 1/(\alpha\mu_i^2)} + O(\sigma) \quad \text{and} \quad \|T\varphi_i\| = \alpha + O(\sigma),$$

which gives

$$\langle \xi, \xi_i \rangle = \sqrt{1 - \frac{1}{\alpha\mu_i^2}} + O(\sigma).$$

4.11.4 Proof of Theorem 25 and Proposition 12

Letting again Θ be the $n \times 2$ group membership matrix, we find as in the proof of Theorem 24 that we have $Q = \Theta\tilde{Q}\Theta^*$ and $K = \Theta\tilde{K}\Theta^*$, with

$$\tilde{Q} = \frac{1}{2} \begin{pmatrix} a\mathbf{E}_{\mathbb{P}}[w] & b\mathbf{E}_{\mathbb{Q}}[w] \\ b\mathbf{E}_{\mathbb{Q}}[w] & a\mathbf{E}_{\mathbb{P}}[w] \end{pmatrix} \quad \text{and} \quad \tilde{K} = \frac{1}{2} \begin{pmatrix} a\mathbf{E}_{\mathbb{P}}[w^2] & b\mathbf{E}_{\mathbb{Q}}[w^2] \\ b\mathbf{E}_{\mathbb{Q}}[w^2] & a\mathbf{E}_{\mathbb{P}}[w^2] \end{pmatrix}.$$

This implies first that

$$\rho = \frac{a\mathbf{E}_{\mathbb{P}}[w^2] + b\mathbf{E}_{\mathbb{Q}}[w^2]}{2},$$

and that the vector $\Theta \begin{pmatrix} 1 \\ -1 \end{pmatrix}$ is an eigenvector of Q associated with the eigenvalue

$$\mu_2 = \frac{a\mathbf{E}_{\mathbb{P}}[w] - b\mathbf{E}_{\mathbb{Q}}[w]}{2}.$$

All other hypotheses of Theorem 22 are easy to check, and we find that the announced results hold as soon as $\mu_2^2 > \rho \vee L$, or

$$\frac{(a\mathbf{E}_{\mathbb{P}}[w^2] + b\mathbf{E}_{\mathbb{Q}}[w^2]) \vee L}{(a\mathbf{E}_{\mathbb{P}}[w] - b\mathbf{E}_{\mathbb{Q}}[w])^2} < 1.$$

Now, let us disregard for a moment the condition on L , and compute ρ :

$$\rho = \frac{1}{2} \int_{\mathcal{L}} (af(\ell) + bg(\ell))w(\ell)^2 dm(\ell)$$

Define a scalar product on $\ell^\infty(\mathcal{L})$, the set of all bounded functions from \mathcal{L} to \mathbb{R} , as

$$\langle h_1, h_2 \rangle_{\mathcal{L}} = \int_{\mathcal{L}} (af(\ell) + bg(\ell))h_1(\ell)h_2(\ell) dm(\ell);$$

then $\rho = \|w\|_{\mathcal{L}}^2$, and applying the Cauchy-Schwarz theorem

$$\rho \cdot \left\| \frac{af - bg}{af + bg} \right\|_{\mathcal{L}}^2 \geq \left\langle w, \frac{af - bg}{af + bg} \right\rangle_{\mathcal{L}}^2 = \mu_2^2.$$

This implies that the signal-to-noise ratio μ_2^2/ρ is maximized whenever

$$w(\ell) = \frac{af(\ell) - bg(\ell)}{af(\ell) + bg(\ell)},$$

and in this case

$$\beta = \frac{\mu_2^2}{\rho} = \frac{1}{2} \int_{\mathcal{L}} \frac{(af(\ell) - bg(\ell))^2}{af(\ell) + bg(\ell)} dm(\ell)$$

In particular, we have $\mu_2 = \rho = \beta$, so $\beta > 1$ implies $\mu_2 > 1$. It remains to notice that $w(\ell) \leq 1$ for any ℓ , so the condition $\mu_2 \geq L$ is redundant as assumed.

4.11.5 Proof of Theorem 26

For $i, j \in [n]$, we note $W_{ij} = m_{ij} + s_{ij}Z_{ij}$ with $Z \sim \mathcal{N}(0, 1)$ a standard gaussian random variable. Let $\tilde{L} = 2\sqrt{\log(n)}$; a well known tail bound for gaussians reads

$$\mathbf{P}(|Z_{ij}| \geq \tilde{L}) \leq \frac{2}{\tilde{L}} e^{-\tilde{L}^2/2} \leq \frac{1}{n^2 \sqrt{\log(n)}}. \quad (4.85)$$

We now define the modified matrix \tilde{W} with

$$\tilde{W}_{ij} = m_{ij} + s_{ij}Z_{ij}\mathbf{1}\{|Z_{ij}| \leq \tilde{L}\},$$

with \tilde{Q} and \tilde{K} the associated expected and variance matrices. It is readily seen that $\tilde{Q} = Q$, and that the variables \tilde{W}_{ij} are bounded by

$$L = \sup_{i,j} |m_{ij}| + \tilde{L} \sup_{i,j} s_{ij}.$$

By a union bound, we have

$$\mathbf{P}(\tilde{W} \neq W) = \mathbf{P}(Z_{ij} > \tilde{L} \text{ for some } i \in [n]) \leq \binom{n}{2} \frac{1}{n^2 \sqrt{\log(n)}} \leq \frac{1}{2\sqrt{\log(n)}},$$

and whenever $\tilde{W} = W$, then the modified non-backtracking matrix coincides with the original one. Finally, notice that for $i, j \in [n]$

$$\text{Var}(Z_{ij}\mathbf{1}\{|Z_{ij}| \leq \tilde{L}\}) \leq 1,$$

which implies using the Perron-Frobenius theorem that $\rho(\tilde{K}) \leq \rho(K)$. Theorem 22 then applies to the modified couple (P, \tilde{W}) and the announced result follows.

4.12 Computing functionals on trees

We prove in this section the martingale estimates of Proposition 16 and Proposition 17.

4.12.1 Study of compound Poisson processes

Many proofs in this section rely on computations of Poisson compound processes, i.e. Poisson sums of random variables. For convenience, we gather them all in the following lemma:

Lemma 40. *Let N be a $\text{Poi}(d)$ random variable, and (X_i) , (Y_i) , (Z_i) three iid sequences of random variables, independent from N , such that X_i and Y_j (resp. Y_i and Z_j , or Z_i and X_j) are independent whenever $i \neq j$. Denote by A, B the random variables*

$$A = \sum_{i=1}^N X_i \quad \text{and} \quad B = \sum_{i=1}^N Y_i,$$

Then the following identities hold:

$$\mathbf{E}[A] = d\mathbf{E}[X], \quad \mathbf{E}[B] = d\mathbf{E}[Y], \tag{4.86}$$

$$\mathbf{E}[AB] = d\mathbf{E}[XY] + d^2\mathbf{E}[X]\mathbf{E}[Y] = d\mathbf{E}[XY] + \mathbf{E}[A]\mathbf{E}[B], \tag{4.87}$$

$$\mathbf{E}\left[\sum_{i=1}^N Z_i \left(\sum_{j \neq i} X_j\right)\right] = d\mathbf{E}[A]\mathbf{E}[Z], \tag{4.88}$$

$$\mathbf{E}\left[\sum_{i=1}^N Z_i \left(\sum_{j \neq i} X_j\right) \left(\sum_{k \neq i} Y_k\right)\right] = d\mathbf{E}[AB]\mathbf{E}[Z]. \tag{4.89}$$

Although the first two identities are well-known, we provide a full proof of this lemma:

Proof. The sequence (X_i) being independent from N , we immediately find that

$$\mathbf{E}[A | N] = N\mathbf{E}[X],$$

from which eq. (4.86) is derived. We then write

$$AB = \left(\sum_{i=1}^N X_i\right) \left(\sum_{i=1}^N Y_i\right) = \sum_{i=1}^N X_i Y_i + \sum_{i \neq j} X_i Y_j,$$

and using the independence property of $(X_i)_i$ and $(Y_i)_i$ yields

$$\mathbf{E}[AB | N] = N\mathbf{E}[XY] + N(N-1)\mathbf{E}[X]\mathbf{E}[Y].$$

Since N is a Poisson random variable, $\mathbf{E}[N(N-1)] = d^2$, hence (4.87).

We now move onto the third equation; rearranging terms gives

$$\sum_{i=1}^N Z_i \left(\sum_{j \neq i} X_j\right) = \sum_{i \neq j} Z_i X_j,$$

and therefore the conditional expectation given N is $N(N-1)\mathbf{E}[X]\mathbf{E}[Z]$. Using again that $\mathbf{E}[N(N-1)] = d^2$ brings (4.88).

Similarly, we can rearrange

$$\sum_{i=1}^N Z_i \left(\sum_{j \neq i} X_j\right) \left(\sum_{k \neq i} Y_k\right) = \sum_{j \neq i} X_j Y_j Z_i + \sum_{i \neq j \neq k} X_i Y_j Z_j,$$

and take conditional expectations on both sides to arrive at

$$\begin{aligned} \mathbf{E} \left[\sum_{i=1}^N Z_i \left(\sum_{j \neq i} X_j \right) \left(\sum_{k \neq i} Y_k \right) \middle| N \right] \\ = N(N-1)\mathbf{E}[XY]\mathbf{E}[Z] + N(N-1)(N-2)\mathbf{E}[X]\mathbf{E}[Y]\mathbf{E}[Z]. \end{aligned}$$

Again, the expected value of $N(N-1)(N-2)$ is d^3 , and we finally find

$$\begin{aligned} \mathbf{E} \left[\sum_{i=1}^N Z_i \left(\sum_{j \neq i} X_j \right) \left(\sum_{k \neq i} Y_k \right) \right] &= d^2\mathbf{E}[XY]\mathbf{E}[Z] + d^3\mathbf{E}[X]\mathbf{E}[Y]\mathbf{E}[Z] \\ &= d\mathbf{E}[AB]\mathbf{E}[Z], \end{aligned}$$

which ends the proof. \square

4.12.2 Decomposing the tree functionals

We now fix $t \geq 1$, $x \in [n]$ and two vectors $\varphi, \varphi' \in \mathbb{R}^n$ for the rest of the section. Let N be the number of children of the root of T , and $(T_k, \mathbf{I}_k)_{k \leq N}$ the subtrees at depth 1. We further introduce the following first moment notations:

$$g_\varphi(t, x) = \mathbf{E}[f_{\varphi, t}(T_x, x)] \quad \text{and} \quad h_{\varphi, \varphi'}(t, x) = \mathbf{E}[f_{\varphi, t}(T_x, x)f_{\varphi', t}(T_x, x)].$$

We begin by a small elementary computation: let $\phi \in \mathbb{R}^n$ be any vector. Then,

$$\mathbf{E}[W_{x\mathbf{I}_k}\phi(\mathbf{I}_k)] = \sum_{y \in [n]} \frac{P_{xy}}{d_x} \mathbf{E}[W_{xy}]\phi(y) = \frac{[Q\phi](x)}{d_x}. \quad (4.90)$$

Now, by linearity, we have

$$f_{\varphi, t}(T_x, x) = \sum_{k=1}^N W_{x\mathbf{I}_k} f_{\varphi, t-1}(T_k, \mathbf{I}_k). \quad (4.91)$$

By definition of the Galton-Watson tree, the random variables $X_k = W_{x\mathbf{I}_k} f_{\varphi, t-1}(T_k, \mathbf{I}_k)$ and $Y_k = W_{x\mathbf{I}_k} f_{\varphi', t-1}(T_k, \mathbf{I}_k)$ satisfy the assumptions of Lemma 40. Furthermore, conditioning on the value of \mathbf{I}_k , we can compute $\mathbf{E}X_k$:

$$\begin{aligned} \mathbf{E}[W_{x\mathbf{I}_k} f_{\varphi, t-1}(T_k, \mathbf{I}_k)] &= \mathbf{E}[W_{x\mathbf{I}_k} g_\varphi(t-1, \mathbf{I}_k)] \\ &= \frac{[Qg_\varphi(t-1, \cdot)](x)}{d_x}. \end{aligned}$$

Applying (4.86), and from the definition of g_φ , we come to the following recurrence relation:

$$g_\varphi(t, x) = [Qg_\varphi(t-1, \cdot)](x).$$

Solving this recurrence is straightforward, and we find

$$g_\varphi(t, \cdot) = Q^t g_\varphi(0, \cdot) = Q^t \varphi,$$

which implies (4.52).

Using now equation (4.87) from Lemma 40, we derive

$$\begin{aligned} h_{\varphi,\varphi'}(t,x) &= d_x \mathbf{E}[W_{x\mathbf{I}_k}^2 h_{\varphi,\varphi'}(t-1,\mathbf{I}_k)] + g_\varphi(t,x)g_{\varphi'}(t,x) \\ &= [Kh_{\varphi,\varphi'}(t,\cdot)](x) + g_\varphi(t,x)g_{\varphi'}(t,x), \end{aligned} \quad (4.92)$$

from which we can solve for $h_{\varphi,\varphi'}$:

$$\begin{aligned} h_{\varphi,\varphi'}(t,\cdot) &= Kh_{\varphi,\varphi'}(t-1,\cdot) + (Q^t\varphi) \odot (Q^t\varphi') \\ &= \sum_{s=0}^t K^s [(Q^{t-s}\varphi) \odot (Q^{t-s}\varphi')]. \end{aligned}$$

The eigenvector equations for φ_i and φ_j then imply (4.53).

Consider now the function $F_{i,t}(T_x, x) = f_{\varphi_i,t+1}(T_x, x) - \mu_i f_{\varphi_i,t}(T_x, x)$, and its associated first moment functions

$$G(t, x) = \mathbf{E}[F_{i,t}(T_x, x)] \quad \text{and} \quad H(t, x) = \mathbf{E}[F_{i,t}(T_x, x)^2].$$

The linearity of $f_{\varphi_i,t}$ implies that $F_{i,t}$ also verifies equation (4.91), and therefore

$$G(t, \cdot) = Q^t G(0, \cdot) = 0$$

for all $t \geq 0$. Equation (4.92) thus reduces to

$$\begin{aligned} H(t, \cdot) &= K^t H(0, \cdot) \\ &= K^t (h_{\varphi_i,\varphi_i}(1, \cdot) - \mu_i^2 \varphi^{i,i}) \\ &= K^t (\mu_i^2 \varphi^{i,i} + K \varphi^{i,i} - \mu_i^2 \varphi^{i,i}) \\ &= K^{t+1} \varphi^{i,i}, \end{aligned}$$

which ends the proof.

4.12.3 Edge functionals

Most of the handiwork needed to prove Proposition 17 was done in Lemma 40; indeed, in the tree (T_x, x) , the edge transformation on $f_{\varphi,t}$ can be written as

$$\vec{\partial}_w f_{\varphi,t}(T_x, x) = \sum_{j=1}^N w_{\mathbf{I}_j} \sum_{k \neq j} W_{x\mathbf{I}_k} f_{\varphi,t-1}(T(k), \mathbf{I}_k).$$

We define accordingly the random variables

$$X_k = W_{x\mathbf{I}_k} f_{\varphi,t-1}(T(k), \mathbf{I}_k), \quad Y_k = W_{x\mathbf{I}_k} f_{\varphi',t-1}(T(k), \mathbf{I}_k) \quad \text{and} \quad Z_k = w_{\mathbf{I}_k},$$

that verify the assumptions of Lemma 40. Computing $\mathbf{E}[Z]$ is straightforward:

$$\mathbf{E}[Z] = \sum_{y \in [n]} \frac{P_{xy}}{d_x} \mathbf{E}[w_y] = [P\bar{w}](x).$$

Hence, we can apply equation (4.88) to those variables, to deduce (4.55). Similarly, the product transformation has the form

$$\vec{\partial}_w(f_{\varphi,t} \cdot f_{\varphi',t})(T_x, x) = \sum_{j=1}^N w_{I_j} \left(\sum_{k \neq j} W_{xI_k} f_{\varphi,t-1}(T(k), I_k) \right) \times \left(\sum_{k \neq j} W_{xI_k} f_{\varphi',t-1}(T(k), I_k) \right),$$

which using (4.89) implies (4.56). Finally, equation (4.57) is proved with the exact same technique, considering $F_{i,t}(T_x, x)$ instead of $f_{\varphi,t}(T_x, x)$.

4.13 Near eigenvectors: computations

We finish here the proof of Proposition 20. First, let

$$f(g, o) = \mathbf{1}_{(g,o)_{t+1} \text{ has no cycles}} \varphi_j(o) f_{\varphi,t+1}(g, o).$$

Then f is $(t+1)$ -local, and we have

$$\begin{aligned} |f(g, o)| &\leq \|\varphi_i\|_{\infty} \|\varphi_j\|_{\infty} |\partial(g, o)_{t+1}| L^{t+1} \\ &\leq \frac{b^2}{n} |(g, o)_{t+1}| L^{t+1} := \psi(g, o). \end{aligned}$$

On the other hand, the scalar product $\langle B^t \chi_i, D_W \check{\chi}_j \rangle$ can be written as

$$\begin{aligned} \langle B^t \chi_i, D_W \check{\chi}_j \rangle &= \sum_{e \in \vec{E}} W_e \varphi_j(e_1) \sum_{\gamma} \prod_{s=1}^t W_{\gamma_s \gamma_{s+1}} \varphi_i(\gamma_{t+1}) \\ &= \sum_{e \in \vec{E}} \varphi_j(e_1) \sum_{\gamma} \prod_{s=0}^t W_{\gamma_s \gamma_{s+1}} \varphi_i(\gamma_{t+1}), \end{aligned}$$

where the sum ranges over all non-backtracking paths $\gamma = (\gamma_0, \dots, \gamma_{t+1})$ such that $(\gamma_0, \gamma_1) = e$. It follows that

$$\begin{aligned} \left| \langle B^t \chi_i, D_W \check{\chi}_j \rangle - \sum_{v \in V} f(G, v) \right| &\leq \left| \sum_{e: e_1 \notin \mathcal{V}_{t+1}} [B^t \chi_i](e) [D_W \check{\chi}_j](e) \right| \\ &\leq 2 |\mathcal{V}_{t+1}| \max_v \psi(G, v), \end{aligned}$$

using the tangle-free property as before. This time, the results from Subsection 4.7.2 yield

$$\max_v \psi(G, v) \leq \frac{c_1 b^2 \log(n)^2 d^{t+1} L^{t+1}}{n} \quad \text{and} \quad \|\psi\|_{\star} \leq \frac{c_1 b^2 \log(n)^3 d^{t+1} L^{t+1}}{n},$$

and the expected value on the tree is

$$\sum_{x \in [n]} \mathbf{E}[f(T_x, x)] = \sum_x \varphi_j(x) \mu_i^{t+1} \varphi_i(x) = \mu_i^{t+1} \delta_{ij}.$$

Concluding,

$$\begin{aligned} |\langle B^t \chi_i, D_W \check{\chi}_j \rangle - \mu_i^{t+1} \delta_{ij}| &\leq \frac{c_2 b^2 \log(n)^4 d^{2t+3} L^{t+1}}{n} + \frac{c_3 r b^2 \log(n)^6 d^{2t+3} L^{t+1}}{\sqrt{n}} \\ &\leq \frac{c_4 r b^2 d^3 L \log(n)^6 d^{2t} L^t}{\sqrt{n}}, \end{aligned}$$

which proves (4.59).

Now, let

$$f(g, o) = \mathbf{1}_{(g,o)_t \text{ has no cycles}} \vec{\partial}_1[f_{\varphi_i,t} \cdot f_{\varphi_j,t}](g, o).$$

Again, f is t -local, and we have

$$\begin{aligned} |f(g, o)| &\leq \|\varphi_i\|_\infty \|\varphi_j\|_\infty \deg(o) |\partial(g, o)_t|^2 L^{2t} \\ &\leq \frac{b^2}{n} \deg(o) |(g, o)_t|^2 L^{2t} := \psi(g, o) \end{aligned}$$

By definition of the $\vec{\partial}$ operator, we have, for $v \in V$,

$$f(g, v) = \sum_{e: e_2=v} [B^t \chi_i](e) [B^t \chi_j](e).$$

Hence,

$$\begin{aligned} \left| \langle B^t \chi_i, B^t \chi_j \rangle - \sum_{v \in V} f(G, v) \right| &= \left| \sum_{e: e_2 \notin \mathcal{V}_t} [B^t \chi_i](e) [B^t \chi_j](e) \right| \\ &\leq 2 |\mathcal{V}_t| \max_v \psi(G, v), \end{aligned}$$

using the tangle-free property as before. This time, the results from [Subsection 4.7.2](#) yield

$$\max_v \psi(G, v) \leq \frac{c_5 b^2 \log(n)^3 d^{2t+1} L^{2t}}{n} \quad \text{and} \quad \|\psi\|_\star \leq \frac{c_5 b^2 \log(n)^4 d^{2t+1} L^{2t}}{n},$$

and we can compute the expected value on the tree:

$$\begin{aligned} \sum_{x \in [n]} \mathbf{E}[f(T_x, x)] &= \sum_x [P\mathbf{1}](x) \mu_i^t \mu_j^t \sum_{s=0}^t \frac{[K^s \varphi^{i,j}](x)}{(\mu_i \mu_j)^s} \\ &= (\mu_i \mu_j)^t \sum_{s=0}^t \frac{\langle P\mathbf{1}, K^s \varphi^{i,j} \rangle}{(\mu_i \mu_j)^s} \\ &= (\mu_i \mu_j)^t \Gamma_{U,ij}^{(t)}. \end{aligned}$$

Gathering those estimates, we find

$$\begin{aligned} \left| \langle B^t \chi_i, B^t \chi_j \rangle - (\mu_i \mu_j)^t \Gamma_{U,ij}^{(t)} \right| &\leq \frac{c_6 b^2 \log(n)^5 d^{3t+2} L^{2t}}{n} + \frac{c_7 r b^2 \log(n)^7 d^{3t+2} L^{2t}}{\sqrt{n}} \\ &\leq \frac{c_8 r b^2 d^2 \log(n)^7 d^{3t} L^{2t}}{\sqrt{n}}, \end{aligned}$$

which proves (4.60).

Next is (4.61); we first notice that the parity-time equation (4.4) implies that

$$\langle (B^*)^t D_W \check{\chi}_i, (B^*)^t D_W \check{\chi}_j \rangle = \langle D_W B^t \chi_i, D_W B^t \chi_j \rangle.$$

Similarly to the previous computation, we therefore let $w_o = (W_{1o}^2, \dots, W_{no}^2)$, and

$$f(g, o) = \mathbf{1}_{(g,o)_t \text{ has no cycles}} \vec{\partial}_{w_o} [f_{\varphi_i, t} f_{\varphi_j, t}](g, o).$$

We have similarly

$$|f(g, o)| \leq \frac{b^2}{n} |(g, o)_t|^2 L^{2t+2} := \psi(g, o),$$

Now,

$$\begin{aligned} \max_v \psi(G, v) &\leq \frac{c_9 b^2 \log(n)^2 d^{2t} L^{2t+2}}{n}, \\ \|\psi\|_\star &\leq \frac{c_{10} b^2 \log(n)^3 d^{2t} L^{2t+2}}{n}, \end{aligned}$$

and as above

$$\begin{aligned} \sum_{x \in [n]} \mathbf{E}[f(T_x, x)] &= \sum_x [Pw_x](x) \mu_i^t \mu_j^t \sum_{s=0}^t \frac{[K^s \varphi^{i,j}](x)}{(\mu_i \mu_j)^s} \\ &= \sum_x [K\mathbf{1}](x) \mu_i^t \mu_j^t \sum_{s=0}^t \frac{[K^s \varphi^{i,j}](x)}{(\mu_i \mu_j)^s} \\ &= \Gamma_{V,ij}^{(t)}. \end{aligned}$$

Equation (4.61) is then derived as we did earlier.

Our final inequality to prove is (4.62); we consider now the function

$$F_t(g, o) = \mathbf{1}_{(g,o)_{t+1} \text{ has no cycles}} \vec{\partial}_1 [F_{i,t}^2](g, o)$$

For all $t \geq 0$, the function F_t is $t+1$ -local, and

$$\begin{aligned} |F_t(g, o)| &\leq \deg(o) (2\|\varphi_i\|_\infty |(g, o)_{t+1}|)^2 L^{2t} \\ &\leq 4 \deg(o) \frac{b^2}{n} |(g, o)_{t+1}|^2 L^{2t} := \psi_t(g, o). \end{aligned}$$

Whenever $v \notin \mathcal{V}_t$,

$$F_t(G, v) = \sum_{e: e_2=v} ([B^{t+1} \chi_i](v) - \mu_i [B^t \chi_i](v))^2$$

The same computations as in the other equations then imply that

$$\left| \|B^{t+1} \chi_i - \mu_i B^t \chi_i\|^2 - \sum_{x \in [x]} F_t(G, v) \right| \leq 2|\mathcal{V}_t| \max_v \psi(G, v),$$

and

$$\max_v \psi(G, v) \leq \frac{c_5 b^2 \log(n)^3 d^{2t+1} L^{2t}}{n} \quad \text{and} \quad \|\psi\|_\star \leq \frac{c_5 b^2 \log(n)^4 d^{2t+1} L^{2t}}{n}.$$

Furthermore,

$$\sum_{x \in [n]} \mathbf{E}[F_t(T_x, x)] = \sum_{x \in [n]} [P\mathbf{1}](x)[K^{t+1}\varphi^{i,i}](x) = \langle P\mathbf{1}, K^{t+1}\varphi^{i,i} \rangle,$$

and we can apply Lemma 30 to find

$$\left| \sum_{x \in [n]} \mathbf{E}[F_t(T_x, x)] \right| \leq rd^3 L^2 \rho^{t+1}.$$

Concluding as above,

$$\begin{aligned} \|B^{t+1}\chi_i - \mu_i B^t \chi_i\|^2 &\leq rd^3 L^2 \rho^{t+1} + \frac{c_{11} b^2 \log(n)^5 d^{3t+3} L^{2t}}{n} \\ &\quad + \frac{c_{12} r b^2 \log(n)^7 d^{3t+3} L^{2t}}{\sqrt{n}} \\ &\leq rd^3 L^2 \rho^{t+1} + \frac{c_{13} r b^2 d^3 \log(n)^7 d^{3t} L^{2t}}{\sqrt{n}}. \end{aligned}$$

4.14 Proofs for Theorem 30

4.14.1 Proof of (4.26)-(4.28)

We shall make use of the following classical bound: for a $r_0 \times r_0$ matrix M , we have

$$\|M\| \leq r_0 \|M\|_\infty. \quad (4.93)$$

First, the (i, j) entry of matrix U^*U is $\langle u_i, u_j \rangle$, and using (4.60) we find

$$|\langle u_i, u_j \rangle - \Gamma_{U,ij}^{(\ell)}| \leq \frac{c r b^2 d^2 \log(n)^7 d^{3\ell} L^{2\ell}}{(\mu_i \mu_j)^\ell \sqrt{n}}.$$

Since $i, j \leq r_0$, we have $\mu_i \mu_j \geq \rho \geq 1/d$ by definition of τ , and inequality (4.32). Plugging this inequality into the one above yields

$$|\langle u_i, u_j \rangle - \Gamma_{ij}^{(\ell)}| \leq \frac{c r b^2 d^2 \log(n)^7 (d^4 L^2)^\ell}{\sqrt{n}}.$$

By definition of ω and ℓ , the three following inequalities hold simultaneously:

$$(d^4 L^2)^\ell \leq n^{1/4}, \quad (d^4 L)^\ell \leq \sqrt{n} \quad \text{and} \quad \left(\frac{d^4 L^2}{\rho} \right)^\ell \leq (d^5 L^2)^\ell \leq \sqrt{n}.$$

Via (4.93), these imply, respectively, that $\|U^*U - \Gamma_U^{(\ell)}\|$ is less than $Cn^{-1/4}$, CL^ℓ and $C\rho^\ell$, the desired result. The derivation of (4.27) is identical, the bound from Proposition 20 being essentially the same for both cases.

We now move onto the proof of (4.28); we write the scalar product $\langle B^\ell \chi_i, (B^*)^\ell D_W \check{\chi}_j \rangle$ as $\langle B^{2\ell} \chi_i, D_W \check{\chi}_j \rangle$ and use (4.59) to find

$$\begin{aligned} |\langle u_i, v_j \rangle - \delta_{ij}| &\leq \frac{c r b^2 d^3 L \log(n)^6 d^{4\ell} L^{2\ell}}{\mu_i^{2\ell+1} \sqrt{n}} \\ &\leq \frac{c r b^2 d^{7/2} L \log(n)^6 d^{5\ell} L^{2\ell}}{\sqrt{n}}. \end{aligned}$$

The bounds we now need are

$$(d^5 L^2)^\ell \leq n^{1/4}, \quad (d^5 L) \leq \sqrt{n} \quad \text{and} \quad (d^6 L^2)^\ell \leq \sqrt{n},$$

which are true by choice of ω and ℓ , and we conclude as above.

4.14.2 Bounding $\|B^\ell\|$: proof of (4.31)

Let w be any unit vector in $R^{\vec{E}}$, and assume that we are in the event described in Proposition 13. Then

$$\begin{aligned} \|B^t w\|^2 &= \sum_{e \in \vec{E}} \left(\sum_{(e_0, \dots, e_t) \in \mathcal{P}(e, t)} \prod_{i=0}^{t-1} W_{e_i e_{i+1}} w(e_t) \right)^2 \\ &\leq L^{2\ell} \sum_{e \in \vec{E}} |\mathcal{P}(e, t)| \sum_{(e_0, \dots, e_t) \in \mathcal{P}(e, t)} w(e_t)^2. \end{aligned}$$

by the Cauchy-Schwarz inequality. Under the good event from Proposition 13, we have

$$|\mathcal{P}(e, t)| \leq 2|(G, e)_t| \leq c_1 \log(n) d^\ell.$$

Additionally, note that the factor $w(e_t)^2$ appears for each path of length t ending at e_t , or equivalently (reversing edge orientation) for each path in $\mathcal{P}(e_t^{-1}, t)$. Hence,

$$\begin{aligned} \|B^t w\|^2 &\leq c_1 \log(n) d^\ell L^{2\ell} \sum_{e \in \vec{E}} w(e)^2 |\mathcal{P}(e^{-1}, t)| \\ &\leq c_2 \log(n)^2 d^{2\ell} L^{2\ell}, \end{aligned}$$

and the definition of ℓ ensures (generously) that $d^{2\ell} < \sqrt{n}$.

4.14.3 Proof of Lemma 36

Note first that for all $t \geq 0$, the parity-time equation (4.4) allows the simplification

$$\langle (B^*)^t D_W \chi_i, w \rangle = \langle B^t \chi_i, D_W J w \rangle.$$

and we have $\|D_W J w\| \leq L$. Further, the assumption $w \in H^\perp$ implies

$$\mu_i^{-t} \langle (B^*)^t D_W \chi_i, w \rangle = \mu_i^{-t} \langle (B^*)^t D_W \chi_i, w \rangle - \mu_i^{-\ell} \langle (B^*)^\ell D_W \chi_i, w \rangle;$$

combining the two above arguments and using a telescopic sum as in the proof of (4.29) gives

$$\begin{aligned} \left| \mu_i^{-t} \langle (B^*)^t D_W \chi_i, w \rangle \right| &= \left| \sum_{s=t}^{\ell-1} \mu_i^{-s} \langle B^s \chi_i, D_W J w \rangle - \mu_i^{-(s+1)} \langle B^{s+1} \chi_i, D_W J w \rangle \right| \\ &\leq \sum_{s=t}^{\ell-1} \mu_i^{-(s+1)} \left| \langle B^{s+1} \chi_i, D_W J w \rangle - \mu_i \langle B^s \chi_i, D_W J w \rangle \right| \\ &\leq L \sum_{s=t}^{\ell-1} \mu_i^{-(s+1)} \|B^{s+1} \chi_i - \mu_i B^s \chi_i\|, \end{aligned}$$

where we used the Cauchy-Schwarz inequality at the last line. Now, we can apply equation (4.62):

$$\|B^{s+1}\chi_i - \mu_i B^s \chi_i\|^2 \leq rd^3 L^2 \rho^{s+1} + \frac{crb^2 d^3 \log(n)^7 d^{3s} L^{2s}}{\sqrt{n}},$$

and still following the proof of (4.29) we find

$$\|B^{s+1}\chi_i - \mu_i B^s \chi_i\| \leq \sqrt{r} d^{3/2} L \rho^{\frac{s+1}{2}} + \frac{c_2 b d^{3/2} \log(n)^{7/2} d^{3s/2} L^s}{n^{1/4}}.$$

Summing these inequalities (and using $\ell \leq c_3 \log(n)$) yields

$$|\langle (B^*)^t D_W \check{\chi}_i, w \rangle| \leq \sqrt{r} d^{3/2} L^2 \mu_i^t \sum_{s=t}^{\ell-1} \left(\frac{\sqrt{\rho}}{\mu_i} \right)^{s+1} + \mu_i^{-\ell} \frac{c_4 b d^{3/2} \log(n)^{9/2} d^{3\ell/2} L^\ell}{n^{1/4}}.$$

Since $i \leq r_0$, we have $\mu_i > \sqrt{\rho}$. As a result, all terms in the sum are bounded by the one for $s = t - 1$, and $\mu_i^{-\ell} \leq d^{\ell/2}$. We finally get

$$|\langle (B^*)^t D_W \check{\chi}_i, w \rangle| \leq \sqrt{r} d^{3/2} L^2 \rho^{t/2} + \frac{c_4 b d^{3/2} \log(n)^{9/2} d^{2\ell} L^\ell}{n^{1/4}},$$

as desired.

4.15 Norm bounds: additional proofs

4.15.1 Bound (4.69) on $\|MB^k\|$

Since $\|M\|$ is of order 1, we notice that (4.69) improves by a factor of \sqrt{n} on the crude bound $\|KB^k\| \leq \|K\| \|B^k\|$. We use the same trace method as above; we have

$$\begin{aligned} \|MB^{k-2}\|^{2m} &\leq \text{Tr} \left[\left(MB^{k-2} (B^*)^{k-2} M^* \right)^m \right] \\ &\leq \left(\frac{d}{n} \right)^{2m} \sum_{\gamma \in W_{k,m}} \prod_{i=1}^m X_{\gamma_{2i-1,1} \gamma_{2i-1,2}} \prod_{s=3}^k A_{\gamma_{2i-1,s-1} \gamma_{2i-1,s}} \\ &\quad \times \prod_{s=1}^{k-2} A_{\gamma_{2i,s-1} \gamma_{2i,s}} X_{\gamma_{2i,k-2} \gamma_{2i,k-1}} \end{aligned}$$

where $W_{k,m}$ is the set of sequences of paths defined just below equation (4.73). The set of edges of the form $(\gamma_{2i-1,0}, \gamma_{2i-1,1})$ or $(\gamma_{2i,k-1}, \gamma_{2i,k})$, which support no random variable, has cardinality at most m by the boundary conditions, hence the bound for any $\gamma \in W_{k,m}$:

$$\begin{aligned} \prod_{i=1}^m X_{\gamma_{2i-1,1} \gamma_{2i-1,2}} \prod_{s=3}^k A_{\gamma_{2i-1,s-1} \gamma_{2i-1,s}} \prod_{s=1}^{k-2} A_{\gamma_{2i,s-1} \gamma_{2i,s}} X_{\gamma_{2i,k-2} \gamma_{2i,k-1}} \\ \leq \left(\frac{d}{n} \right)^{e_\gamma - m} L^{2(k-2)m}. \end{aligned}$$

Using bound (4.75) on $\mathcal{W}_{k,m}(v, e)$ and the fact that each equivalence class contains at most n^v elements, we get

$$\begin{aligned}
\|MB^{k-2}\|^{2m} &\leq \left(\frac{d}{n}\right)^m \sum_{e=1}^{2km} \sum_{v=1}^{e+1} k^{2m} (2km)^{6m(e-v+1)} n^v \left(\frac{d}{n}\right)^e L^{2(k-2)m} \\
&\leq n^{-m} d^{5m} L^{2km} k^{2m} \sum_{e=1}^{2km} \sum_{v=1}^{e+1} (2km)^{6m(e-v+1)} d^e n^{v-e} \\
&\leq n^{-m+1} d^{5m} L^{2km} k^{2m} (2km) d^{2km} \sum_{g=0}^{\infty} \left(\frac{(2km)^{6m}}{n}\right)^g. \tag{4.94}
\end{aligned}$$

The choice of parameter

$$m = \left\lceil \frac{\log(n)}{12 \log(\log(n))} \right\rceil$$

ensures that the infinite sum in (4.94) converges for n larger than an absolute constant, which yields (4.69).

4.15.2 Bound (4.70) on $\|\Delta^{(t-1)} \tilde{M} B^{k-t-1}\|$

First, notice that $M_{ef}^{(2)}$ is equal to $(TQT^*)_{ef}$ except when $\mathbf{1}\{e \xrightarrow{2} f\} = 0$, which happens only when $e = f$, $e \rightarrow f$, $e \rightarrow f^{-1}$ of $f^{-1} \rightarrow e$. Therefore, we can write

$$|L_{ef}| \leq \frac{d}{n} (\tilde{M}_1 + \tilde{M}_2 + \tilde{M}_3 + \tilde{M}_4),$$

where each entry of the matrix M_i is one whenever the i -th condition mentioned above is true. Then, for each i we can write

$$\|\delta^{(t-1)} \tilde{M}_i B^{k-t-1}\| \leq \|\delta^{(t-1)}\| \|\tilde{M}_i B^{k-t-1}\|,$$

and a straightforward adaptation of the proof of bound (4.69) gives

$$\frac{d}{n} \|\tilde{M}_i B^{k-t-1}\| \leq \frac{cd^{7/2} L \ln(n)^7 d^{k-t} L^{k-t}}{\sqrt{n}}.$$

Combining the above bound with (4.68) easily implies (4.70).

4.15.3 Bound (4.71) on $R_t^{(\ell)}$

The proof of (4.71) is very similar to those above, as well as the one in [29]; we only highlight the main differences. Let $t \geq 1$ (the case $t = 0$ is almost identical), and $k \leq \log(n)$. The same trace argument gives

$$\begin{aligned}
\|R_t^{(k-1)}\|^{2m} &\leq \text{Tr} \left[\left(R_t^{(k-1)} R_t^{(k-1)*} \right)^m \right] \\
&= \sum_{\gamma \in T_{k,m,t}} \prod_{i=1}^{2m} \underline{X}_{\gamma_{i,0}\gamma_{i,1}} \prod_{s=2}^t \underline{A}_{\gamma_{i,s-1}\gamma_{i,s}} Q_{\gamma_{i,t}\gamma_{i,t+1}} \prod_{s=t+2}^k A_{\gamma_{i,s-1}\gamma_{i,s}},
\end{aligned}$$

where $T_{k,m,t}$ is the set of sequences of paths $(\gamma_1, \dots, \gamma_{2m})$ such that for all i , $\gamma_i^1 = (\gamma_{i,0}, \dots, \gamma_{i,t})$ and $\gamma_i^2 = (\gamma_{i,t+1}, \dots, \gamma_{i,k})$ are tangle-free and γ_i is tangled, with similar boundary conditions as in (4.74).

We define G_γ as the union of the $G_{\gamma_i^z}$ for $z \in [2m], j \in \{1, 2\}$. Since we remove an edge to each path, G_γ need not be connected; however, since γ_i is tangled, each connected component in G_{γ_i} contains a cycle, and the same holds for G . It follows that

$$v_\gamma \leq e_\gamma$$

for all $\gamma \in \mathcal{T}_{k,m,t}$. As before, we define the equivalence relation \sim and $\mathcal{T}_{k,m,t}(v, e)$ the set of equivalence classes with $v_\gamma = v$ and $e_\gamma = e$. Then, the following lemma from [29] holds:

Lemma 41. *Let v, e be any integers such that $v \leq e$. Then*

$$|\mathcal{T}_{k,m,t}(v, e)| \leq (4km)^{12m(e-v+1)+8m}.$$

As for bounding the contribution of a single path, the computations already performed in bounding (4.68) work similarly:

$$\begin{aligned} \prod_{i=1}^{2m} X_{\gamma_{i,0}\gamma_{i,1}} \prod_{s=2}^t A_{\gamma_{i,s-1}\gamma_{i,s}} Q_{\gamma_{i,t}\gamma_{i,t+1}} \prod_{s=t+2}^k A_{\gamma_{i,s-1}\gamma_{i,s}} \\ \leq \left(\frac{a}{n}\right)^{e_\gamma+2m} \left(1 + \frac{d}{n}\right)^{2km} d^{2m} L^{2km}, \end{aligned}$$

using $Q_{ij} \leq dL/n$ for all i, j . Finally, for $[\gamma] \in \mathcal{T}_{k,m,t}(v, e)$, there are at most n^v sequences γ' such that $\gamma' \sim \gamma$. This yields

$$\begin{aligned} \mathbf{E} \left[\|R_t^{(k-1)}\|^{2m} \right] &\leq c_1^{2m} d^{2m} L^{2km} \left(\frac{d}{n}\right)^{2m} (4km)^{20m} \sum_{e=1}^{2km} (4km)^{12m(e-v)} \sum_{v=1}^e d^e n^{v-e} \\ &\leq c_2^{2m} d^{4m} L^{2km} n^{-2m} \log(n)^{40m} (2km) d^{2km} \sum_{g=0}^{\infty} \left(\frac{(4km)^{12m}}{n}\right)^g, \end{aligned}$$

using preemptively the bound $m \leq \log(n)$ and the change of variables $g = e - v$. This time, choosing

$$m = \left\lceil \frac{\log(n)}{24 \log(\log(n))} \right\rceil$$

yields a convergent sum, and (4.71) follows.

Chapter 5

Spectral clustering in directed networks

This chapter is based on the paper [42], a joint work with S. Coste.

5.1 Introduction

The proverbial effectiveness of spectral clustering (“good, bad, and spectral”, said [74]), observed for long by practitioners, begins to be well-understood from a theoretical point of view. More and more problem-specific spectral algorithms are periodically designed, with better computational performance and accuracy. Most of the theoretical studies used to be concentrated on symmetric interactions and undirected networks, but over the last decade a flurry of works has tackled the directed case: the 2013 survey [89] gave an account of the richness of directed network clustering, but since then the field expanded in various directions.

In directed networks, the directionality of interactions is taken into account. This is more realistic from a modeling point of view, but at the cost of intricate technicalities in the analysis. In this context, the aim of this paper is to take a step back at spectral algorithms and convince the readers of a simple, yet largely unnoticed truth: *when clustering directed networks, directly using eigenvalues and eigenvectors (as opposed to SVD) of the untouched adjacency matrix (as opposed to symmetrized and/or regularized versions), works very well, especially in harder and sparser regimes.* This statement was hinted in some works, but remained essentially ignored and not backed by theoretical results. In this paper, our goal is to rigorously prove and illustrate this statement.

Spectral clustering of directed networks: overview and related work As well-summarized in [122], spectral algorithms often consist of three steps: (1) a matrix representation of the data, (2) a spectral truncation procedure, and (3) a geometric clustering method on the eigen/singular vectors.

The matrix representation depends on the nature of the data and interaction measurements. Early works were focused on symmetric interactions: the interaction $A_{x,y}$ between two nodes u, v was considered undirected, ie $A_{x,y} = A_{y,x}$. The matrix of interactions A is then hermitian. But in many applications, the interaction are intrinsically *directional*, with $A_{x,y}$ and $A_{y,x}$ not necessarily equal, or, equivalently, the network is *directed*. This covers a wider range of models: buyer/seller networks, ecologic systems with predator-prey interactions, citations in scientific papers, etc.

Forerunners in digraph clustering avoided using the interaction matrix A ; one reason for that is the lack of an orthogonal decomposition for non-normal matrices. Instead, they represented their data with naive symmetrizations of A , such as $A + A^*$ or the so-called *co-citation and*

bibliometric symmetrizations AA^* and A^*A (which reduces to studying the SVD of A), see [114] for an overview. Variants of the graph Laplacian were then introduced ([37]); they were hermitian but incorporated in some way the directionality of edges and were provable relaxations of normalized-cut problems (see [83] and references in [89, §4.2-4.3]); others used random-walks approaches similar to PageRank ([35], [109]). It is quite striking that in the survey [89], the authors classify the clustering methods (Chap. 4 therein) *without mentioning the use of the adjacency matrix of directed networks*. Implicit in these early works was the belief that directed networks *need* to be transformed or symmetrized. More recently, [44] and [81] cleverly introduced \mathbb{C} -valued Hermitian matrices.

The use of adjacency matrices was advocated later, notably in [86], and several more theoretical works, among them [90, 128, 120]. These works can roughly be split in two categories. On one side, the authors of many papers seemed reluctant to use eigenvalues of non-symmetric matrices, and favored the SVD instead, whose theoretical analysis is tractable in some cases. This is notably the case for [118, 90, 128]. However, as we'll see later, the SVD for non-hermitian matrices suffers the same problem as the eigendecomposition of hermitian matrices: it is sensitive to heterogeneity, and thus less powerful in sparse regimes. On the other side, [86, 120, 36] are closer in spirit to our work. They explicitly advocate the use of eigenvalues of non-symmetric matrices as a better option for inference problems. The theoretical analysis performed in [86, 120] allows them to guarantee performance in very specific cases, where the underlying graphs have a strong and dense structure (upper-triangular, cyclic or purely acyclic). In a different context (matrix completion), the paper [36] was one of the first to prove the efficiency of eigenvalues of non-hermitian matrices in high-dimensional problems.

Spectral decompositions of network matrices are generally known to reflect some of the underlying interaction structure between the nodes; this non-rigorous statement has now been mathematically understood in a variety of ways, many of them based on a mathematical model for community networks called the stochastic block-model ([69, 1]). Any clustering algorithm can be tested on synthetic data from the SBM to evaluate the reconstruction accuracy, that is, the number of nodes which have been correctly assigned to their community by the algorithm. In this work, we deal with a vast generalization of SBMs, the weighted, inhomogeneous, directed Erdős-Rényi random graph.

Definition 6. *Let P, W be two real $n \times n$ matrices, with P having entries in $[0, 1]$. A random weighted graph is defined as follows: the edge set is $V = [n]$; each one of the n^2 potential edges $(x, y) \subset V \times V$ is present in the graph with probability $P_{x,y}$ and independently of the others; if present, its weight is $W_{x,y}$. The resulting directed graph will be noted $G = (V, E)$ and its weighted adjacency matrix A is defined by $A_{x,y} := W_{x,y} \mathbf{1}_{(x,y) \in E}$.*

This allows for virtually any structure: classical block-models (assortative or disassortative), cyclic structures ([120]), path-wise structures ([81]), overlapping communities ([50]), bipartite clustering when both sides have the same size ([128, 129]), contextual information on the edges...

In SBMs, a key parameter is the density d , the number of edges divided by the size n . For inference problems, a lower density means a sparser information. Analyzing the performance of spectral clustering methods can be done using classical perturbation results in regimes where d is large, often of order n (the 'dense' regime), see [111] for instance. However, many real-world networks lie in sparser regimes $d \ll n$, like $d \lesssim \log(n)$, $d \rightarrow \infty$ (the 'semi-sparse regime') or even $d = O(1)$ (the 'sparse regime'), a radically difficult regime in which node degrees are extremely heterogeneous and the graph is not even connected. This behaviour has an impact on spectral quantities when they satisfy Fisher-Courant-Weyl inequalities, like eigenvalues of normal matrices or SVD of non-normal ones, deeply reducing their performance, see [19]. This

is why most theoretical works (for both directed and undirected models) were concentrated on $d \sim \log(n)$ regimes ([4] among others). In the sparse undirected regime, the celebrated Kesten-Stigum threshold ([29]) gives a spectral condition for the emergence of a second eigenvalue λ_2 , beyond the Perron one, in the spectrum of the non-backtracking matrix. The entries of its eigenvector are correlated with the block structure when there are two blocks. In this paper, we describe a general theory of directed Kesten-Stigum-like thresholds for every directed sparse SBM, irregardless of the number of blocks, their size, etc.

Contributions. We prove a Master Theorem describing the sharp asymptotics of eigenvalues and eigenvectors of sparse non-symmetric matrices with independent entries, like adjacency matrices of inhomogeneous directed Erdős-Rényi graphs. This is of independent interest in the field of random matrices. We show how to apply this theorem to the directed SBM and we introduce an elementary community-detection algorithm based on the adjacency matrix. We show why using both left and right eigenvectors is mandatory in sparse regimes. We give numerical and heuristic evidence for why Gaussian Mixture clustering is much more adapted than the popular k-means algorithm. Finally, we illustrate the strength of our method on synthetic data¹.

Notations. We use the standard Landau notations $o(\cdot), O(\cdot), \sim$; for k integer, $[k]$ stands for $\{1, \dots, k\}$. The letters u, v, w will be kept for vectors, the letters x, y, z will be kept for elements of $[n]$ (nodes). We see vectors in \mathbb{C}^n as functions from $[n]$ to \mathbb{C} , that is $u = (u(1), \dots, u(n))$. The notation $|u|_p$ stands for the p -norm of a vector, $|u|_p^p = |u(1)|^p + \dots + |u(n)|^p$. We drop the index p iff $p = 2$ (euclidean norm). If M is a matrix, $\|M\| = \sup_{|u|=1} |Mx|$ and $\|M\|_\infty = \max M$. The Hadamard product of two matrices A, B with the same shape is defined as the entrywise product: $(A \odot B)_{x,y} = A_{x,y}B_{x,y}$. The Frobenius norm is $\|M\|_F = \sqrt{\sum_{x,y} |M_{x,y}|^2}$.

5.2 The Master Theorem

Let P, W be two real $n \times n$ matrices, with P having entries in $[0, 1]$. The weighted, inhomogeneous directed Erdős-Rényi model was defined in Definition 6. The weighted adjacency matrix will be noted A . We focus on the $n \rightarrow \infty$ limit and we suppose in the assumptions thereafter that the graph is sparse, the weights are bounded and the spectral decomposition of $\mathbf{E}[A]$ is not degenerate. Let $Q = \mathbf{E}[A]$ and $K = \mathbf{E}[A \odot A]$ be the first and second entrywise-moments of A , given by

$$Q = P \odot W \quad \text{and} \quad K = P \odot W \odot W, \quad (5.1)$$

in other words $Q_{x,y} = P_{x,y}W_{x,y}$ and $K_{x,y} = P_{x,y}|W_{x,y}|^2$. Our assumptions are as follows:

- (i) $\|P\|_\infty = O(1/n)$ and $\|W\|_\infty = O(1)$.
- (ii) The matrix Q has rank $r = O(1)$, is real diagonalizable, and its r eigenvalues μ_i are well-separated in the sense that there is a constant $c > 0$ such that $|\mu_i - \mu_j| > c, |\mu_j| > c$.
- (iii) the right (resp. left) unit eigenvectors φ_i and ξ_i associated with μ_i are delocalized, in that

$$|\varphi_i|_\infty, |\xi_i|_\infty = O\left(\frac{1}{\sqrt{n}}\right) \quad (5.2)$$

¹The Python software used for the numerical experiments in this paper will be available on a public repository.

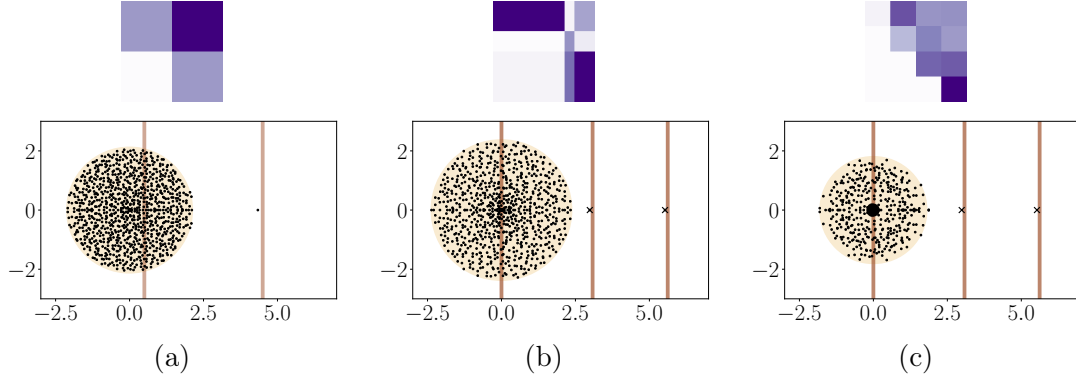


Figure 5.1: Spectra of inhomogeneous Erdős-Rényi graphs, with no weights, $n = 1000$ nodes. The underlying connectivity matrix P is a block-matrix; the inset of each picture shows a colorplot of P , with darker colors indicating higher values. The points are the eigenvalues of A . The brown lines indicate the non-zero eigenvalues μ_i of Q . The beige circle behind the eigenvalues has radius ϑ . The outliers close to μ_i for $i \in [r_0]$ are visible for each picture.

These assumptions will be commented later. We note $\rho = \|K\|$. The *detection threshold* is

$$\vartheta = \max(\sqrt{\rho}, \|W\|_\infty)$$

and we note r_0 the number of eigenvalues of Q with modulus strictly greater than ϑ .

Definition 7. Let μ_i be an eigenvalue of Q with left and right unit eigenvectors φ_i, ξ_i . If $|\mu_i| > \vartheta$, the (left and right) **eigendefects** of μ_i are defined by

$$R_i = |(K - \mu_i^2 \mathbf{I})^{-1} \varphi_i^2|_1 \quad L_i = |(K^* - \mu_i^2 \mathbf{I})^{-1} \xi_i^2|_1 \quad (5.3)$$

where φ_i^2, ξ_i^2 are the entrywise squares of φ_i, ξ_i .

These novel quantities will play a role of paramount importance in all this paper and will be commented later. Let us first state our main result, after which we will give some intuition on (5.3).

Theorem 31 (Master Theorem). *Under the above hypotheses, the following holds with probability going to 1 when $n \rightarrow \infty$. The r_0 eigenvalues of A with highest modulus, $\lambda_1, \dots, \lambda_{r_0}$, are asymptotically equal to the r_0 eigenvalues of Q with highest modulus: $|\lambda_i - \mu_i| = o(1)$. All the other $n - r_0$ eigenvalues of A are asymptotically smaller than ϑ . Moreover, if u_i, v_i is a left/right pair of unit eigenvectors of A associated with λ_i , and if φ_i, ξ_i is a left/right pair of unit eigenvectors of Q associated with μ_i , then*

$$\left| |\langle u_i, \varphi_j \rangle| - \frac{|\langle \varphi_i, \varphi_j \rangle|}{|\mu_i| \sqrt{R_i}} \right| = o(1) \quad \text{and} \quad \left| |\langle v_i, \xi_j \rangle| - \frac{|\langle \xi_i, \xi_j \rangle|}{|\mu_i| \sqrt{L_i}} \right| = o(1). \quad (5.4)$$

The eigendefects in Definition 7 measure how much the eigenequations of Q can be ‘entrywise squared’. For instance, let μ be an eigenvalue of Q with eigenvector φ . Then, $(Q - \mu \mathbf{I})\varphi = 0$. But is φ^2 an eigenvector of K with eigenvalue μ^2 ? If $i \in [r_0]$, the answer is obviously no since $\mu^2 > \|K\|$; then, the quantity $1/|(K - \mu^2 \mathbf{I})^{-1} \varphi^2|_1$ appearing in the theorem above is a measure of how far φ^2 is from being a μ^2 -eigenvector of K . The theorem says that when μ_i gives rise to an outlier λ_i in the spectrum of A , then the overlap $\langle u_i, \varphi_i \rangle$ between the real eigenvector and the sample eigenvector is higher when φ_i^2 is far from being an eigenvector of K . At a high level, this surprising and new phenomenon comes from an elementary formula regarding the covariance of Poisson sums (see Lemma 47 in Appendix 5.7), and we conjecture that a similar phenomenon will hold for every random matrix model which is asymptotically Poisson.

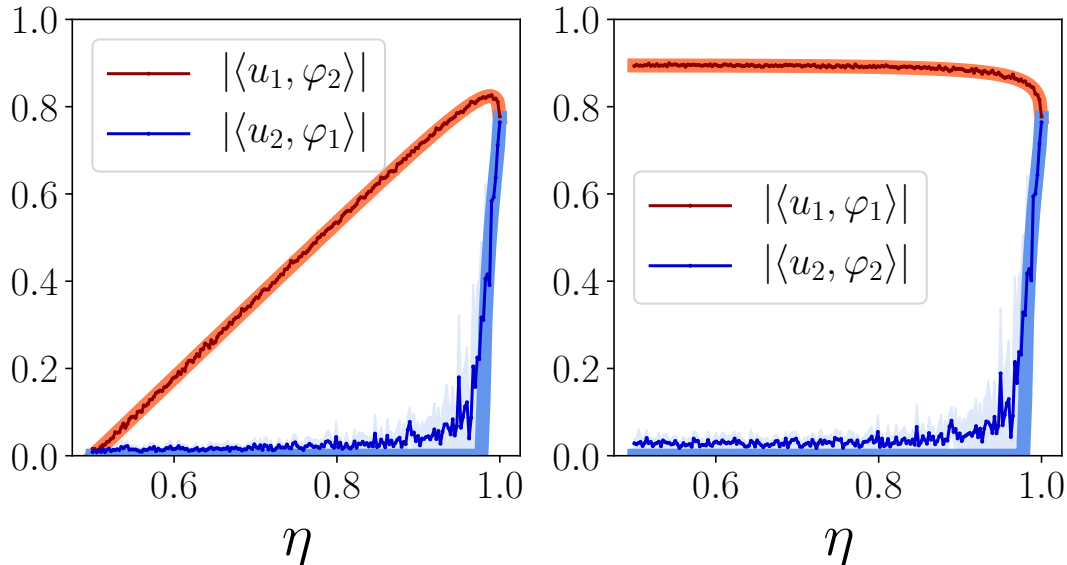


Figure 5.2: **Bottom:** An illustration of the ‘right eigenvector part’ of the Master Theorem, for the two-block SBM described in Theorem 32 with parameters $s = 10$ and η ranging from 0.5 to 1. For each η the results are averaged over 50 samples (coloured zone is for standard errors). The thin darker lines are $\langle u_i, \varphi_j \rangle$ for $i, j \in \{1, 2\}$. The thick lighter lines are our theoretical predictions $\langle \varphi_i, \varphi_j \rangle / \mu_i \sqrt{R_i}$. The second eigenvector begins to be informative as soon as the second eigenvalue reaches ϑ , which happens at around 0.979 in agreement with our predictions.

Remark (Comments on the hypotheses.). Under **Hypothesis 1**, if $\|P\|_\infty \leq C/n$, then the expected degree of $x \in [n]$ is $d_x = P_{x,1} + \dots + P_{x,n} \leq C$ so the average density d is smaller than C (sparse regime). Our proof holds in the semi-sparse regime $d \rightarrow \infty, d = n^{o(1)}$, but it is not our primary motivation. In this semi-sparse regime it can be proved that $|\mu_i| \sqrt{R_i}$ goes to 1, resulting in perfect alignment $\langle u_i, \varphi_i \rangle \rightarrow 1$.

Real diagonalizability in **Hypothesis 2** is here to simplify the proof but the Master Theorem will hold for complex eigendecompositions. The low-rank assumption is standard in the literature; it can be relaxed by replacing the rank with the effective rank, as in [28]. The separation assumption is not necessary for eigenvalue asymptotics, but strong separation is necessary for eigenvector overlaps.

Note that every bound in the hypotheses (such as the bound on W , the rank or the delocalization) can be extended to $n^{o(1)}$ at virtually no cost.

Theorem 31 follows the line of research initiated in [91, 29] and continued by different works, among which [116, 28, 106]. We give an overview of the proof of Theorem 31 in Appendix 5.7.

5.3 Spectral embeddings of directed SBM

5.3.1 Model definition

A powerful aspect of our Master Theorem lies in its application to clustering in directed stochastic blockmodels. In the following, we fix a number of clusters r , and a number of vertices n , understood to be large. Let $\sigma_g, \sigma_d : [n] \mapsto [r]$ be the *left* and *right* cluster assignments; that is, a vertex x is said to be in the i -th left (resp. right) cluster if $\sigma_g(x) = i$ (resp. $\sigma_d(x) = i$). Let F be an arbitrary $r \times r$ matrix with positive entries; the directed SBM is then a random graph

$G = ([n], E)$ with vertex set $[n]$ and such that for each directed edge (i, j) , we have

$$\mathbf{P}((i, j) \in E) = \frac{F_{\sigma_g(i), \sigma_d(j)}}{n}.$$

The aim is then to recover the left (or right) cluster memberships, given an observation of G . Note that, for ease of exposition, we did not include weights: W is thus the all-one matrix, and in this setting we have $Q = P = \mathbf{E}[A]$. It is relatively simple to compute the spectral decomposition of Q in this setting:

Proposition 22. *Let A be the adjacency matrix of G . The non-zero eigenvalues of $P = \mathbf{E}[A]$ are exactly those of the modularity matrix $F\Pi$, where*

$$\Pi_{ij} = \frac{|\sigma_g^{-1}(j) \cap \sigma_d^{-1}(i)|}{n}.$$

Additionally, the associated left eigenvectors of P are constant on the right clusters, while the right eigenvectors are constant on the left clusters.

Therefore, as per our main theorem, the left/right eigenvectors of the adjacency matrix of G are close to their expectation, which is constant on the right/left clusters. We thus expect a clustering algorithm on those eigenvectors to be able to recover at least a fraction of the community memberships. We refer to Appendix 5.9 for a more complete spectral analysis of the matrix P , as well as a formulation of Theorem 31 suited to the SBM setting.

Remark. *Whenever $\sigma_g = \sigma_d$, as is often the case, both left and right eigenvectors are constant on the clusters; this effectively doubles the signal to recover σ_d .*

An important question to ask is the following: how many eigenvectors of A do we need to be able to reconstruct the clusters? It is often assumed that r eigenvectors are needed to recover the memberships (see e.g. [122]). However, in our DSBM setting, we propose the following heuristic:

Partial cluster recovery is possible as soon as the first r_0 eigenvectors of $\mathbf{E}[A]$ are sufficient to recover the clusters.

Here, r_0 is the same as in Theorem 31, and denotes the number of eigenvalues of $\mathbf{E}[A]$ that get reflected in the spectrum of A . Since we showed that the eigenvectors of $\mathbf{E}[A]$ are constant on the clusters, this is equivalent to the function $k \mapsto (\psi_1(k), \dots, \psi_{r_0}(k))$ being injective, where the ψ_i are the right (resp. left) eigenvectors of $F\Pi$ (resp. ΠF). This can happen when $r_0 \ll r$, and even in some cases when $r_0 = 1$, which is a huge improvement on the threshold for reconstruction. Additional eigenvectors may of course increase the recovery accuracy; however, in some cases, the additional information they bring is nullified by the increase in dimensions for the clustering algorithms.

5.3.2 SBM with a pathwise structure

General case. We restrict to the classical SBM described earlier, with a specific shape known as *pathwise structure*, and notably studied in [81]. It is a good model for flow data. In this model, we have $\sigma_d = \sigma_g$, and the clusters partition $[n]$ in r parts of equal size. The underlying $r \times r$

connectivity F is given by

$$F = s \begin{pmatrix} 1/2 & \eta & & & \\ 1 - \eta & 1/2 & \eta & & \\ & \ddots & \ddots & \ddots & \\ & & & 1 - \eta & 1/2 & \eta \\ & & & & 1 - \eta & 1/2 \end{pmatrix} \quad (5.5)$$

where $s > 1$ is the density parameter and $\eta \in [1/2, 1]$. The modularity matrix is therefore given by F/r . The matrix F shown in (5.5) is a tridiagonal Toeplitz matrix; such matrices have been extensively studied and their eigendecomposition is known (see Appendix 5.10.3): as a result, cluster recovery is possible as soon as the top eigenvalue of F/r is at least one. This happens in particular whenever $s \geq 2r$.

Two blocks: explicit computations. In the case of two blocks $r = 2$ with the same size $n/2$, the connectivity matrix F is equal to

$$F = \begin{pmatrix} s/2 & s\eta \\ s(1 - \eta) & s/2 \end{pmatrix} \quad (5.6)$$

Define the parameter $\theta = 2\sqrt{\eta(1 - \eta)}$. The spectral structure of F is described in the following lemma:

Lemma 42. *The two eigenvalues of F are $v_1 = s\frac{1+\theta}{2}$ and $v_2 = s\frac{1-\theta}{2}$, with corresponding unit right eigenvectors f_i and unit left eigenvectors f_i given by*

$$\begin{aligned} f_1 &= (\sqrt{\eta}, \sqrt{1 - \eta}), & g_1 &= (\sqrt{1 - \eta}, \sqrt{\eta}), \\ f_2 &= (\sqrt{\eta}, -\sqrt{1 - \eta}), & g_2 &= (\sqrt{1 - \eta}, -\sqrt{\eta}). \end{aligned}$$

The eigenvectors of P thus verify

$$\varphi_i(x) \propto f_i(\sigma(x)) \quad \text{and} \quad \xi_i(x) \propto g_i(\sigma(x)),$$

and Theorem 31 applies in this setting:

Theorem 32. *Under the above assumptions, with high probability the following holds.*

1) *If $s < 4(1 + \theta)/(1 - \theta)^2$, then $r_0 = 1$. The Perron eigenvalue of A , namely λ_1 , is asymptotically equal to $v_1/2$, and all the other eigenvalues have modulus asymptotically smaller than $\sqrt{v_1/2}$. Moreover, if u_1, v_1 is a left/right pair of unit eigenvectors associated with λ_1 , then $\lim_{n \rightarrow \infty} |\langle u_1, \varphi_1 \rangle| = \lim_{n \rightarrow \infty} |\langle v_1, \xi_1 \rangle| = a_{1,1}$ where $a_{1,1}$ is a completely explicit function of s, η that satisfies*

$$a_{1,1} = 1 - \frac{2}{s} \cdot \frac{1 + \theta^2}{(1 + \theta)^2} + O\left(\frac{1}{s^2}\right)$$

2) *If instead $s < 4(1 + \theta)/(1 - \theta)^2$, then $r_0 = 2$. The Perron eigenvalue of A , namely λ_1 , is asymptotically equal to $sv_1/2$, the second eigenvalue λ_2 is asymptotically equal to $v_2/2$ and all the other eigenvalues have modulus asymptotically smaller than $\sqrt{v_1/2}$.*

Algorithm 1 Spectral clustering of n nodes, based on the adjacency matrix A .

- 1: **Data:** a $n \times n$ adjacency matrix A ; a number of clusters k ; a rank r_0 .
 - 2: Compute the r_0 largest eigenvalues of M and their unit left and right eigenvectors u_i, v_i .
 - 3: Define the spectral embedding $\mathcal{X}^A = \{\mathcal{X}_x^A : x \in [n]\}$ as in (5.7).
 - 4: Apply a GMM-clustering method on the cloud \mathcal{X}^M .
 - 5: **return** The partition of $[n]$ based on the output of GMM-clustering.
-

Moreover, if u_i, v_i is a left/right pair of unit eigenvectors associated with λ_i for $i = 1, 2$, then $\lim_{n \rightarrow \infty} |\langle u_1, \varphi_1 \rangle| = \lim_{n \rightarrow \infty} |\langle v_1, \xi_1 \rangle| = a_{1,1}$ as above, and additionally $\lim_{n \rightarrow \infty} |\langle u_2, \varphi_2 \rangle| = \lim_{n \rightarrow \infty} |\langle v_2, \xi_2 \rangle| = a_{2,2}$ with $a_{2,2}$ a completely explicit function of s, η that satisfies

$$a_{2,2} = 1 - \frac{2}{s} \cdot \frac{1 + \theta^2}{(1 - \theta)^2} + O\left(\frac{1}{s^2}\right).$$

The threshold $s > 4(1 + \theta)/(1 - \theta)^2$ can also be rewritten as $\eta > \eta(s)$, where η is an explicit function of s (see equation (5.48) in the Appendix). This will be the preferred formulation through the rest of the paper. Note that $\eta(s)$ decreases to $1/2$ quite slowly: as an example, we have $\eta(10) \approx 0.979$ and $\eta(50) \approx 0.885$, and it can be shown that $\eta(s) \sim cs^{-1/4}$.

5.4 Geometric clustering and community detection

Our Master Theorem describes the information given by the spectral embedding \mathcal{X} on the underlying model. Most spectral clustering pipelines then perform geometric clustering based on \mathcal{X} .

5.4.1 Algorithm and measure of performance

Our algorithm computes the left and right eigenvectors x_i, y_i associated with the r_0 largest eigenvalues of M , then defines an embedding of the nodes of $[n]$ in \mathbb{R}^{2k} by setting

$$\mathcal{X}^A(x) = (u_1(x), \dots, u_{r_0}(x), v_1(x), \dots, v_{r_0}(x)). \quad (5.7)$$

Then, we cluster these n points using the Gaussian Mixture Model for clustering ([94]). We insist on the fact that no data preprocessing is needed: no high-degree trimming, no pruning, no normalization. The complexity of our algorithm is similar to all the spectral clustering procedures: it needs the computation of at most $2r$ left/right eigenvectors where r is generally $O(\log(n)^c)$, and then doing a clustering method with at most r clusters on a $n \times 2r_0$ embedding.

Remark. *The number r_0 is a priori problem-dependent. However, since $r_0 \leq r$ and r is low in most problems, one can loop over r_0 at a minor cost. The Master Theorem allows for a more reasonable possibility, which is to directly estimate r_0 from the data as the number of isolated eigenvalues outside the bulk of the spectrum. This can easily be done either by visual inspection (see Figure 5.1) or by some ad hoc statistical rule and it does not require a priori knowledge of r — unlike many methods in the literature.*

In the stochastic block-model, we have a notion of ground-truth clustering $\sigma : V \rightarrow [k]$, where $\sigma(x) = i$ denotes the membership of node x to the i -th cluster. If our procedure outputs a clustering $\hat{\sigma}$, the performance of this clustering is measured through the overlap, also called Rand Index: it is the proportion of pairs of nodes on which σ and $\hat{\sigma}$ agree on membership, that is

$$\text{ov}(\sigma, \hat{\sigma}) = \frac{1}{\binom{n}{2}} \sum_{\{x,y\}} \mathbf{1}\{\sigma \text{ and } \hat{\sigma} \text{ agree on the edge } \{x,y\}\}, \quad (5.8)$$

where ‘agree’ means that either $\sigma(x) = \sigma(y)$ and $\hat{\sigma}(x) = \hat{\sigma}(y)$, or $\sigma(x) \neq \sigma(y)$ and $\hat{\sigma}(x) \neq \hat{\sigma}(y)$. Without any information on σ , a blind guess for $\sigma(x)$ is to randomly assign x to one of the k clusters. This is called a dummy label, $\hat{\sigma}_{\text{dummy}}$. The adjusted overlap is often preferred to the former:

$$\text{aov}(\sigma, \hat{\sigma}) = \frac{\text{ov}(\sigma, \hat{\sigma}) - \mathbf{E}[\text{ov}(\sigma, \hat{\sigma}_{\text{dummy}})]}{1 - \mathbf{E}[\text{ov}(\sigma, \hat{\sigma}_{\text{dummy}})]}. \quad (5.9)$$

An adjusted overlap of 1 indicates a perfect recovery of σ (up to permutation), while an overlap of 0 indicates that $\hat{\sigma}$ is not better than a dummy guess at recovering σ . In the literature this is often called Adjusted Rand Index. It will be our measure of performance in our numerical experiments.

Remark (notions of consistency). Strong consistency of a procedure corresponds to $\text{aov} = 1$, that is: all the labels are correctly guessed. Weak consistency is when $\lim \text{aov} = 1$ when $n \rightarrow \infty$. Partial consistency is when $\liminf \text{aov} > 0$, meaning that the algorithm does strictly better than random guess, a task called detection. In the undirected setting, strong consistency is feasible in the regime $d/\log(n) \rightarrow \infty$ ([4] and weak consistency as long as $d \rightarrow \infty$ ([60] and the survey section therein). Note that in the sparse regime with d constant (this paper), even weak consistency is not achievable because of a constant proportion of isolated nodes. Partial consistency in the undirected case was achieved in [29, 100] under a spectral condition on the modularity matrix. We expect similar results to the ones above to hold in the undirected case.

5.4.2 Choice of the clustering algorithm

When it comes to the last step of spectral clustering, i.e. geometric clustering on the spectral embedding, the overwhelming choice of algorithm is the k -means (see for example [120, 128]). It is very simple to implement, and its tractability allows for explicit performance bounds for spectral clustering [44, Theorem 2]. However, our numerical experiments (see Appendix 5.6) show an interesting phenomenon: the addition of a second informative eigenvector appears to decrease the algorithm’s performance ! Our assumption is that the increase in dimensions for clustering caused by the introduction of the second eigenvector nullifies the additional information it brings. To the contrary, our experiments showed a significative increase in performance when using GMMs, although very little is known in terms of their theoretical footing. A mix of theoretical and empirical results allow us to present a simple explanation: the eigenvectors of G are indeed close to a mixture of Gaussian distributions.

Theorem 33. Consider the SBM as described in Section 5.3, with $\sigma_d = \sigma_g$. Assume that μ_i is an isolated informative eigenvalue of P , and let u_i be a right eigenvector of the adjacency matrix of G such that $|u_i| = \sqrt{n}$. Then we have the following convergence in distribution: for all $j \in [r]$,

$$\frac{1}{\text{Card}(\sigma^{-1}(j))} \sum_{\sigma(x)=j} \delta_{u_i(x)} \xrightarrow[n \rightarrow +\infty]{d} \mathcal{L}(Z_{i,j}),$$

where $Z_{i,j}$ is a random variable with known mean μ_{ij} and variance σ_{ij}^2 (see equation (5.52) in the Appendix). Similar results hold for the right eigenvector v_i .

In the above theorem, δ denotes the Dirac delta; the LHS is therefore simply the discrete measure on the entries of u_i on the j -th community. The proof of this theorem, as well as an explicit derivation of μ_{ij} and σ_{ij} , can be found in the Appendix. It implies, in particular, that the distribution of u_i can be seen as a mixture of r different distributions.

We do not claim (and it is indeed false, see Remark 5.6) that $Z_{i,j} \sim \mathcal{N}(\mu_{ij}, \sigma_{ij}^2)$; however, numerical experiments performed in Figure 5.5 appear to show that, at least when the mean

degree of the graph is large, the distribution of u_i behaves as a mixture of Gaussian distributions. This brings us to the conjecture:

Conjecture 6. *When the top eigenvalue of P goes to infinity, the distribution of Z_{ij} approaches that of a normal random variable with same mean and variance.*

If proven true, this conjecture would give theoretical footing to the performance of GMMs, whose edge over other algorithms is only observed empirically for now.

5.4.3 Numerical validation of our results

Tested methods. We compare our Algorithm 1 with two other methods for digraph clustering. Both methods end with a k-means clustering on a spectral embedding. The first method uses the k left and k right top singular vectors of the adjacency matrix, where k is the number of blocks. The second one is SimpleHerm: we define a complex Hermitian matrix by $H_{x,y} = \omega A_{x,y} + \bar{\omega} A_{v,u}$ where ω is the $[2\pi k]$ -th root of unity. We then use the eigenvector of the smallest eigenvalue λ_1 of $L = I - D^{-1/2} H D^{-1/2}$ with D the diagonal degree matrix ($D_{x,x} = d_x^+ + d_x^-$); since its entries are complex, it is viewed as an embedding on \mathbb{R}^2 . This method was introduced in [81], and was convincingly shown to outperform other classical methods in semi-sparse regimes.

SVD and our method are agnostic to problem structure, but SimpleHerm is well-fitted to flow networks. The performance guarantees of our method relies on the probabilistic properties of the SBM, while SimpleHerm satisfies deterministic Cheeger-like inequalities (see [81]). It would be interesting to test these methods on more general models of directed networks.

Setting. For 50 values of η equally spread between 0.5 and 1 we sampled 20 directed SBMs with connectivity matrix F as in (5.5) and with $n = 2500$ nodes. The parameters are k , the number of blocks and equal to 2, 4, 6 (the blocks have the same size n/k). For the parameter s , we chose the unique $s(d, k)$ so that the the mean degree of our model with k blocks (given in the formula (5.42)) is equal to $d = 2, 3$ or 4, see Table 5.1 at page 160 and discussion therein. We insist on the fact that the mean degree in our model is extremely low, and in particular stays under the $\log(n) = \log(2500) \approx 7.82$ barrier. Our performance measure is the adjusted overlap (5.9), between the labelling output by the tested method, and the true labelling.

Results. The results of our experiments² are in Figure 5.3. With extremely low degrees ($d = 2$), our method (red curve) is the only one to catch a signal, the two other ones are unable to detect any community structure. For slightly higher $d = 3, 4$, our method seems to globally compare with SimpleHerm and be superior to SVD. When the asymmetry is closer to 1/2 (hard regime), our method performs better, see for instance the very neat advantage for $k = 6$ blocks, where our method reaches more than 20% overlap for $\eta \approx 1/2$, against no detection at all for the other methods. When η approaches 1, the performance of our method collapses back to low overlaps, while SimpleHerm has very high performances. We expect this phenomenon to be caused by the fact that when $\eta \rightarrow 1$, the eigenvectors of F all align with $(1, 0, \dots, 0)$. On a side note, we remark that our method has a better precision, with the standard error (coloured zones) being generally smaller.

5.5 Conclusion and future prospects

We rigorously described the behaviour of a simple spectral embeddings, using the eigenvectors of non-symmetric matrices, and we numerically show that our algorithm using Gaussian Mixture

²In a preliminary version of this paper, the method SimpleHerm was incorrectly implemented.

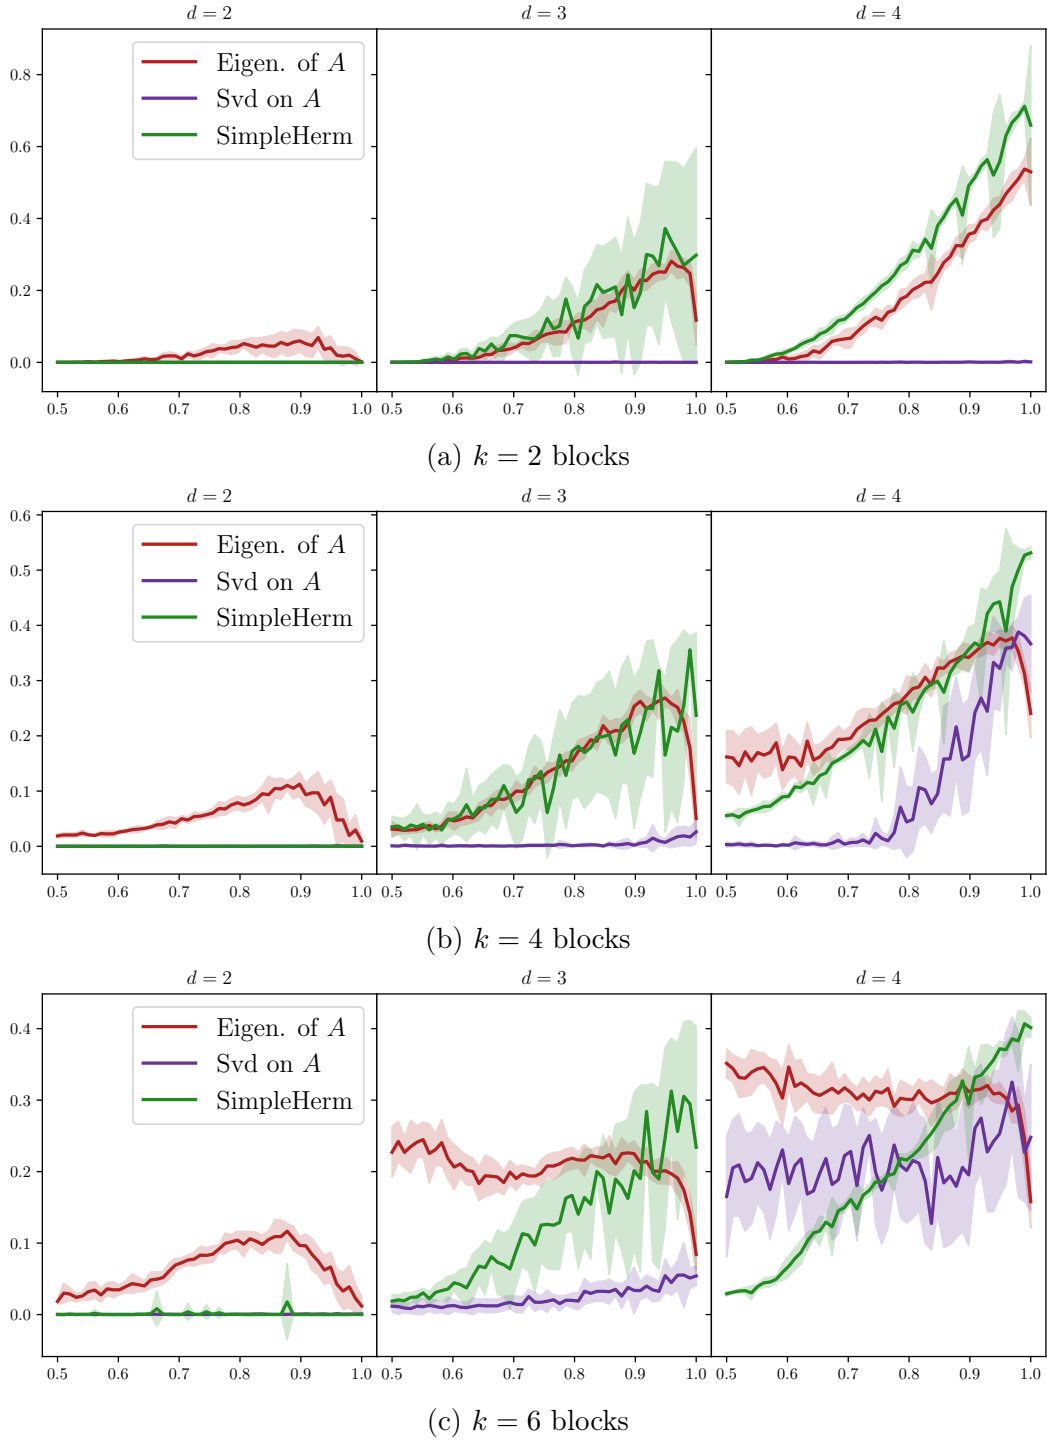


Figure 5.3: Averages of labels reconstruction scores (adjusted overlaps), averaged over 20 runs for 50 values of η ranging from 0.5 to 1, and for different mean degree d . The number of nodes is $n = 2500$. Coloured zones are for standard errors of the corresponding method. The parameter s corresponding to the various values d are given in Table 5.1 page 160).

clustering has surprisingly good results against state-of-the art methods in digraph clustering, especially in the difficult regime where the model density is $O(1)$. The main weakness of our theory is that it does not apply to rectangular matrices directly, but the randomsplit-squaring strategy as in [28] is directly applicable here. Since our theory is new, we chose to keep the exposition as simple and general as possible, but many new directions seem to be promising: among them is the possibility to use the distance-matrix $A_{x,y}^{(\ell)} = \mathbf{1}_{d+(x,y) \leq \ell}$ instead of A , which should result in a method which is more robust to adversarial perturbations, as in [3, 117]. Regarding the gaussianity of the model, we conjecture that the fluctuations of the eigenvalues are Gaussian in the sparse regime; as mentioned, the fluctuations of the eigenvector entries will not be Gaussian, but we now explore a proof of the convergence of these fluctuations when the density of the model increases.

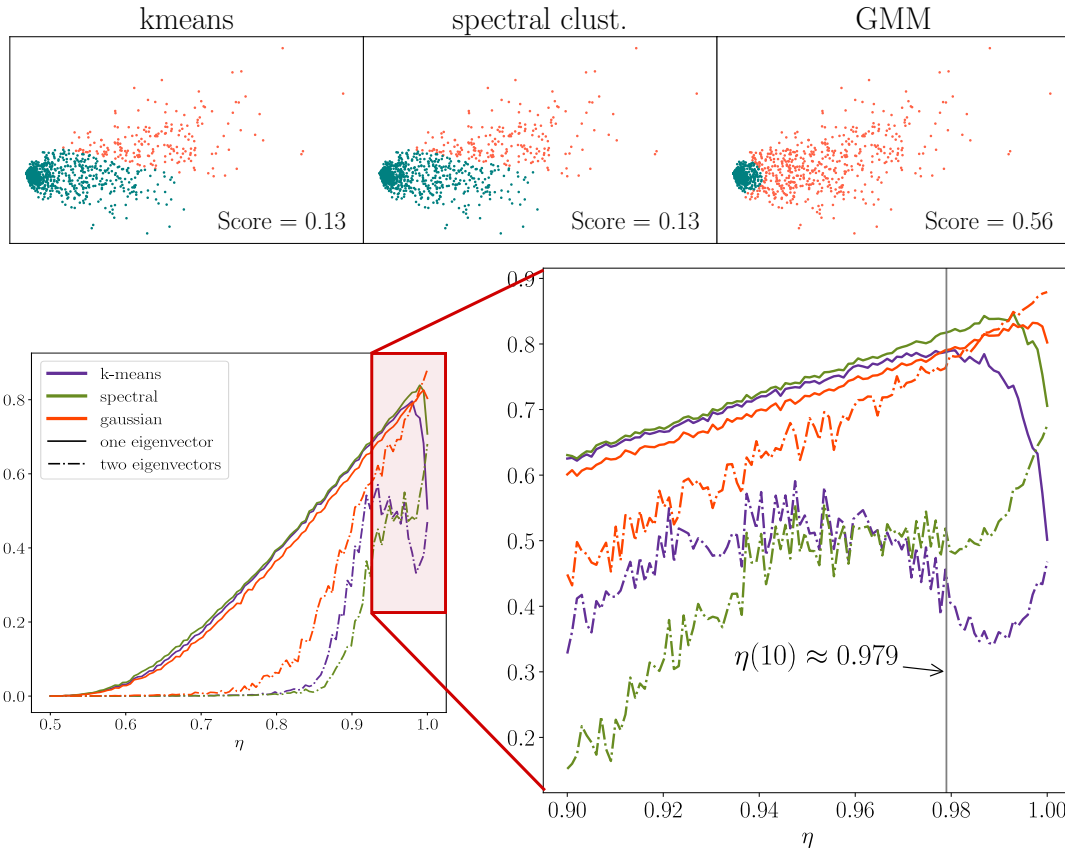


Figure 5.4: **Top.** Here $n = 2000$, $s = 10$ and $\eta = .99 > \eta(10) =$, thus $r_0 = 2$. The Score is the adjusted overlap in (5.9).

Middle. For the two-block model with $\eta \in [0.5, 1]$ we plotted the average Adjusted overlap over 100 runs of several clustering methods on spectral embeddings using either the embedding with the Perron vector $x \mapsto (u_1(x), 0)$ (solid lines) or the embedding with two dominant eigenvectors $x \mapsto (u_1(x), u_2(x))$ (dashdot lines). In the inset we see that the performance of GMM is not reduced by the addition of the second informative eigenvector at the critical point $\eta(10)$.

5.6 Gaussian Mixture clustering and Gaussian fluctuations

In Figure 5.4, we performed some experiments regarding which clustering method to use on the spectral embedding. We simply used three popular methods, implemented in Python’s Sklearn library ([108]):

- k-means, the most popular method in graph clustering,
- Spectral-Clustering, which solves a norm-cut problem on the singular vectors of a distance matrix, a method known to be powerful when the clusters are non-convex,
- Gaussian Mixture clustering, which fits the parameters of a mixture of gaussians to the data using the E-M algorithm.

The first panel in Figure 5.4 is only a visual illustration of what spectral embeddings on a two-block SBM looks like. Here, the parameters are $\eta = 0.99$ and $s = 10$; our theory shows that there are two outliers in the spectrum of A . Our spectral embedding \mathcal{X} in (5.7) has thus four dimensions (we use the left and right eigenvectors). For better visualization, we simply took the

two right eigenvectors. Each point in the figure is thus $(u_1(x), u_2(x))$ for some node x and the colors are the labels given by each clustering method.

The second panel in 5.4 shows the performance of these clustering methods, for η between 1/2 and 1. We also compared the use of only one eigenvector with the use of two eigenvectors, even when there is only one informative eigenvector ($r_0 = 1$).

- (i) When there is only one outlier ($\eta < \eta(s)$), clustering based on the Perron eigenvectors (solid lines) yields good results, while adding a second uninformative eigenvector (dashdot lines) deeply reduces the performance of any clustering method.
- (ii) When crossing $\eta(s)$, a second informative eigenvector appears; the performance of clustering methods based on the Perron eigenvector are reduced, in accordance with the correlation decrease of $|\langle u_1, \varphi_1 \rangle|$ predicted by Theorem 32 (see the golden line in Fig 5.1, first panel).
- (iii) But, when $\eta > \eta(s)$, the performance of kmeans and spectral-clustering based on the two informative eigenvectors u_1, u_2 first decreases, since these algorithms seem to struggle exploiting the extra information given by the second eigenvector (Figure 5.4, top panel). Only the Gaussian mixture model incorporates this extra information efficiently: it is the only method for which clustering based on two informative vectors is *better* than with only one (the two orange lines cross short after $\eta(s)$).

We did not try other clustering methods — these experiments are only indicative of a seemingly high performance for gaussian clustering. In Theorem 33, we showed that the spectral embeddings have a limiting distribution, accessible through the $Z_{i,j}$. If $Z_{i,j}$ was Gaussian, the performance of GMM would be completely understood, but as mentioned before Conjecture 6, in the sparse regime, the limiting distributions $Z_{i,j}$ (or equivalently, the spectral embeddings) *are not Gaussian*. The following remark explains why.

Remark. $Z_{i,j}$ has a positive atom at 0 (and possibly many other atoms): indeed, following the notations of the very last subsection, it is easily seen that the limit $Z_{i,j}$ is equal zero when the Galton-Watson tree \mathcal{T}_j is empty, which happens with strictly positive probability so $\mathbf{P}(Z_{i,j} > 0) > 0$; but clearly, the extinction probability of \mathcal{T}_j goes to zero when ν_1 , the highest modularity eigenvalue, goes to infinity. Note that the atom at zero is visible in Figure 5.5-(a).

The shapes of the random variables $Z_{i,j}$ are visible in Figure 5.5. In this figure, we plotted the histograms of the entries of u_i in two different 5000-nodes blockmodels, with connectivities F_1 (top) and F_2 (bottom) given by:

$$F_1 = \begin{pmatrix} 6 & 4 \\ 5 & 3 \end{pmatrix} \qquad F_2 = \begin{pmatrix} 48 & 6 \\ 12 & 24 \end{pmatrix} \qquad (5.10)$$

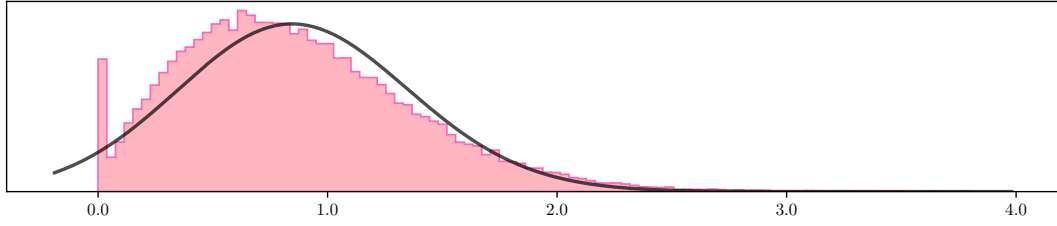
with two clusters of size $p_1 = 2/3, p_2 = 1/3$ — the same size on the left and on the right. In the first model, there is only one outlier; in the second, there are two.

Note that the Theorem implies that for $i \in [r_0]$, the densities of the discrete distributions

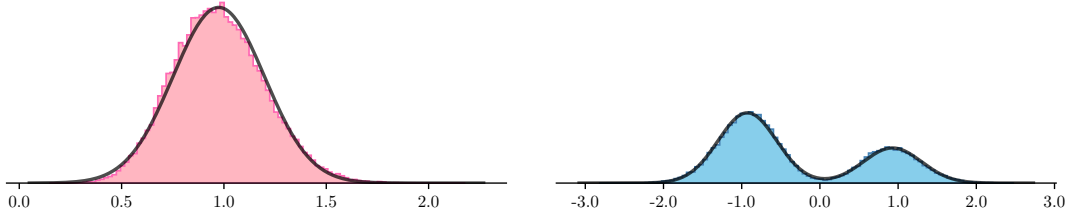
$$\frac{1}{n} \sum_{x \in [n]} \delta_{u_i(x)}$$

converge in distribution to $p_1 Z_{i,1} + p_2 Z_{i,2}$. The plots are the histograms of the $u_i(x)$ over 10 samples. In the figure, the grey lines are the densities of the *gaussian* mixtures

$$\mathcal{N}_i = p_1 \mathcal{N}(\mu_{i,1}, \sigma_{i,1}^2) + p_2 \mathcal{N}(\mu_{i,2}, \sigma_{i,2}^2), \qquad (5.11)$$



(a): $r_0 = 1$. Histogram of the entries of u_1 .



(b) $r_0 = 2$. Histogram of u_1 in pink, and u_2 in blue.

Figure 5.5: An illustration of Theorem 33, for the two-block SBM with connectivities as in (5.5) and clusters sizes $p_1 = 2/3, p_2 = 1/3$. Here, $n = 5000$ and there are 10 samples in each picture. The grey lines are the densities of the gaussian mixtures in (5.11).

where $i \in \{1, 2\}$ and $\mu_{i,j}, \sigma_{i,j}^2$ are the means and variances of the limiting random variables $Z_{i,j}$. It is clearly seen that in the first plot, the limit of u_1 is not Gaussian; in the second plot, the degrees of the graph are much higher (we are already on the semi-sparse regime), and the Gaussian approximations for both eigenvectors are strikingly convincing.

5.7 A bird's eye view on the proof of Theorem 31

The proof of Theorem 31 follows the celebrated high-trace method, introduced in [91, 29]. Considerable advances and simplifications have recently been made; on one hand, [116] performed this method on the *non-backtracking* matrix of weighted, *inhomogeneous undirected* graphs; on the other hand, [28] performed this method on the *adjacency* matrix of weighted, *homogeneous directed* graphs. In both of these papers, the underlying matrices P and W were Hermitian, which is no longer the case here. Our Master Theorem bridges the gap, and considers the *adjacency* matrix of weighted, *inhomogeneous directed* graphs with *general* P, W .

We hereby sketch the main ideas at a high level, and when needed we link our proof with the formerly cited papers. We emphasize that the proofs of theorems like Theorem 31 are often very technical. In this appendix, we tried to be as elementary as possible, to hide the technical details already written in other papers, and to give a short, accessible summary of the proof ideas — at the cost of completeness.

5.7.1 Warmup: notations

Probabilistic domination. For readability, we introduce notations regarding the asymptotic order of real random variables. Let

$$X = (X_n : n \in \mathbb{N}) \quad Y = (Y_n : n \in \mathbb{N})$$

be two families of real random variables. We write $X \preceq Y$ if there is a constant D such that for every constant $c > 0$, for n large enough,

$$\mathbf{P}(|X_n| > \log(n)^D |Y_n|) \leq \log(n)^{-c},$$

in other words, $|X_n|$ is smaller than $|Y_n|$ up to logarithmic terms, with probability smaller than every polylogarithm (typically, n^δ for small δ). Finally, we write $X \ll Y$ if for every constants $c, D > 0$, for n large enough,

$$\mathbf{P}(|X_n| > \log(n)^{-D} |Y_n|) \leq \log(n)^{-c}.$$

In other words, $|X_n|/|Y_n|$ goes to zero faster than every polylogarithm. With this handy device, it is easily seen that $X_n \preceq Y_n$ and $Y_n \ll Z_n$ imply $X_n \ll Z_n$. These notations are common in the field of random matrix theory, and they truly simplify the exposition compared with [116, 28].

Spectral decomposition. Before starting, we write the spectral decomposition of $Q = \mathbf{E}[A]$:

$$Q = \sum_{i=1}^r \mu_i \varphi_i \psi_i^*.$$

The μ_i are the r nonzero eigenvalues; we order them by decreasing modulus, $|\mu_1| > \dots > |\mu_r|$. The φ_i are the corresponding unit right eigenvectors; they do not always form an orthonormal basis, because Q has not been supposed normal. The ψ_i are the unique (up to a sign) left eigenvectors satisfying $\langle \varphi_i, \psi_j \rangle = \delta_{i,j}$, and in general they do not have unit-length. In the statement of the Master Theorem, the unit left eigenvectors ξ_i are thus $\psi_i/|\psi_i|$.

Thresholds, spectral gap. We recall that r_0 is the number of eigenvalues of Q with modulus greater than $\vartheta = \max(\|W\|_\infty, \sqrt{\rho})$ where $\rho = \sqrt{\|K\|}$ and we introduce $\Phi = (\varphi_1, \dots, \varphi_{r_0})$, $\Psi = (\psi_1, \dots, \psi_{r_0})$ and $\Sigma = \text{diag}(\mu_1, \dots, \mu_{r_0})$. The spectral gap of our model is defined as

$$\tau = \sqrt{\frac{\vartheta}{|\mu_{r_0}|}}.$$

It is very important to note that $\tau < 1$. The closer to 1, the harder the problem; the bounds of Theorem 31 are actually of the form $|\lambda_i - \mu_i| \preceq \tau^\ell$, where ℓ is a carefully chosen parameter that grows logarithmically with n .

Covariance functionals, eigendefects and cross-defects We will first need a notation for the Hadamard products of vectors in \mathbb{R}^n :

$$\varphi^{i,j} = \varphi_i \odot \varphi_j \quad \text{and} \quad \psi^{i,j} = \psi_i \odot \psi_j, \quad (5.12)$$

so that in the statement of the Master Theorem, φ_i^2 is equal to $\varphi^{i,i}$. In the proof we will only use the $\varphi^{i,i}$ notation.

We introduce two functions $\Gamma, \Gamma^* : \mathbb{C} \times \mathbb{C}^n \rightarrow \mathbb{C}$ by

$$\Gamma(z, \xi) = \sum_{t=0}^{\infty} \frac{\langle \mathbf{1}, K^t \xi \rangle}{z^t} \quad \text{and} \quad \Gamma^*(z, \xi) = \sum_{t=0}^{\infty} \frac{\langle \mathbf{1}, (K^*)^t \xi \rangle}{z^t}. \quad (5.13)$$

Let us recall that the norm of K is ρ ; consequently, the sums above are convergent when $|z|$ is strictly greater than ρ , and in particular when $z = \mu_i^2$ with $i \in [r_0]$. In this case, one has

$$\Gamma(\mu_i^2, \varphi^{i,i}) = \mu_i^2 \sqrt{R_i} \quad \text{and} \quad \Gamma^*(\mu_i^2, \psi^{i,i}) = \mu_i^2 \sqrt{L_i}. \quad (5.14)$$

Proof. Let $\xi = \varphi^{i,i}$. We have $\Gamma(\mu_i^2, \xi) = \sum_{t=0}^{\infty} \frac{\langle \mathbf{1}, K^t \xi \rangle}{\mu_i^{2t}} = \langle \mathbf{1}, (\sum_{t=0}^{\infty} (K/\mu_i^2)^t) \xi \rangle$. The Neumann summation formula shows that since $\|K/\mu_i^2\| < 1$, the matrix sum in this expression is nothing but $(\mathbf{I} - K/\mu_i^2)^{-1} = \mu_i^2 (\mu_i^2 \mathbf{I} - K)^{-1}$, and that this matrix has all entries nonnegative. Since ξ also has nonnegative entries, $|\mu_i^2 \mathbf{I} - K|^{-1} \xi|_1 = \langle \mathbf{1}, (\mu_i^2 \mathbf{I} - K)^{-1} \xi \rangle$, and we recognize the right eigendefects. \square

Since we'll also need 'cross defects' like $\Gamma(\mu_i \mu_j, \varphi^{i,j})$, we introduce the notations $\Gamma_{\text{left}}, \Gamma_{\text{right}}$ for two matrices of size $r_0 \times r_0$ defined by

$$(\Gamma_{\text{left}})_{i,j} = \Gamma(\mu_i \mu_j, \varphi^{i,j}) \quad \text{and} \quad (\Gamma_{\text{right}})_{i,j} = \Gamma^*(\mu_i \mu_j, \psi^{i,j}).$$

5.7.2 The Pseudo-Master Theorem: A is nearly diagonalized by pseudo-eigenvectors

The tools for studying matrix A are the *pseudo-eigenvectors*. Consider the eigenvalue μ_i of Q . We define two vectors $U_i = A^\ell \varphi_i / \mu_i^\ell$ and $V_i = (A^*)^\ell \psi_i / \mu_i^\ell$, where

$$\ell = \lfloor \kappa \log(n) \rfloor \quad (5.15)$$

and κ is a positive constant to be chosen later. To put it in matrix form, these vectors are the columns of the $n \times r_0$ matrices

$$U = A^\ell \Phi \Sigma^{-\ell} \quad \text{and} \quad V = (A^*)^\ell \Psi \Sigma^{-\ell}.$$

The key aspects of U and V are summarized in the following list of statements, which is a pseudo-version of the Master Theorem. We recall that the symbols \ll, \preceq were rigorously defined in the preceding subsection.

Theorem 34 (Pseudo-Master Theorem). *For a sufficiently small choice of κ , the following holds.*

(i) U and V are nearly inverses:

$$\|U^*V - \mathbf{I}\| \ll 1 \quad (5.16)$$

(ii) U and Ψ are nearly inverses, and V and Φ are nearly inverses:

$$\|\Psi^*U - \mathbf{I}\| \ll 1 \quad \|\Phi^*V - \mathbf{I}\| \ll 1 \quad (5.17)$$

(iii) V and U nearly diagonalize A^ℓ :

$$\|V^*A^\ell U - \Sigma^\ell\| \ll 1 \quad (5.18)$$

(iv) Γ_{left} is nearly the Gram matrix of U , and the same for Γ_{right} and V :

$$\|U^*U - \Gamma_{\text{left}}\| \ll 1 \quad \|V^*V - \Gamma_{\text{right}}\| \ll 1 \quad (5.19)$$

(v) A^ℓ is negligible outside of the vector spaces spanned by the pseudo-eigenvectors:

$$\|A^\ell \text{Proj}_{\text{im}(V)^\perp}\| \preceq \vartheta^\ell, \quad \|\text{Proj}_{\text{im}(U)^\perp} A^\ell\| \preceq \vartheta^\ell, \quad (5.20)$$

where Proj_{C^\perp} denotes the projection matrix on the orthocomplement of the subspace C .

A crucial point in this theorem is that the error scale (up to polylog factors) is ϑ : only the last bound, (5.20), is actually sharp. The other error terms are meant to be negligible.

Proving the Pseudo-Master Theorem (PMT) is really the core of the proof, and where lie most of the difficulties. From the PMT, it is only a matter of linear algebra and perturbation theory to prove the Master Theorem: we simply summarize the spirit in the next subsection, and we quickly jump to the proof of the PMT.

5.7.3 Master Theorem = Pseudo-Master Theorem + perturbation theory

Given the statements in the preceding subsection, the main theorem follows from a standard algebraic analysis, which builds on detailed quantitative variants of the Bauer-Fike theorem. We refer the reader to the comprehensive studies in [91, Section 4] or [28, Section 8], which are technical, but can be applied directly without any further modification. In this short paragraph, we simply explain the ideas.

The main trick is to define a matrix $S = U\Sigma^\ell V^*$. If we had $V^*U = \mathbf{I}$, this matrix would be exactly diagonalizable with eigenvalues μ_i^ℓ ; but V^*U is only close to \mathbf{I} . Fortunately, it is easily seen that if U, V are well-conditioned, then S is diagonalizable, with eigenvalues close to the μ_i^ℓ and eigenvectors close to the U_i, V_j . The fact that U, V are well-conditioned follows from (5.19): by continuity, their condition number is close to the condition number of $\Gamma_{\text{right}}, \Gamma_{\text{left}}$, who in turn are bounded:

Lemma 43. *The condition numbers of the matrices $\Gamma_{\text{right}}, \Gamma_{\text{left}}, U, V$ are all $\preceq 1$.*

Proof. Note $\pi_s(x) = \sqrt{(K^*)^s \mathbf{1}(x)}$ and $\Pi_s = \text{diag}(\pi_s)$. It is easily seen that $\langle \mathbf{1}, K^s \varphi^{i,j} \rangle = (\Phi^* \Pi_s^2 \Phi)_{i,j}$. The matrix Γ_{left} is thus a sum of semi-definite positive matrices, with first term \mathbf{I}_0 , so its smallest eigenvalue is ≥ 1 and its condition number is smaller than $\|\Gamma_{\text{left}}\|$. On the other hand, note that $|\langle \mathbf{1}, K^s \varphi^{i,j} \rangle| \leq |\mathbf{1}| |\varphi^{i,j}| \rho^s$. Thanks to (5.2), we have $|\varphi^{i,j}| \leq c\sqrt{1/n}$ for some c ; consequently, $(\Gamma_{\text{left}})_{i,j} \leq c \sum \rho^s / (\mu_i \mu_j)^s \leq c/(1-\tau) \preceq 1$. Finally, since the size r_0 is also ≤ 1 , we have $\|\Gamma_{\text{left}}\| \preceq 1$ and the condition number is $\preceq 1$. By the Weyl inequalities and (5.19), $|\text{cond}(U) - \text{cond}(\Gamma_{\text{left}})| \leq \|UU^* - \Gamma_{\text{left}}\| \ll 1$ and $\text{cond}(U) \preceq 1$. \square

The smallest eigenvalue ν_{r_0} of S is close to $\mu_{r_0}^\ell$, and since $\tau > 1$ independently of n and $\ell = \lfloor \kappa \log(n) \rfloor$, we get $|\mu_{r_0}^\ell| \ll \vartheta^\ell$. But we can write A^ℓ as a perturbation of S ,

$$A^\ell = S + (A^\ell - S),$$

and, crucially, statements (5.18)-(5.20) can be bootstrapped to show that $\|A^\ell - S\| \ll 1$. The Bauer-Fike theorem applies, and roughly says that the eigenvalues of A^ℓ are within distance $\ll \vartheta^\ell$ of the eigenvalues of S : consequently, A^ℓ has r_0 eigenvalues ϑ^ℓ -close to $\mu_1^\ell, \dots, \mu_{r_0}^\ell$, and the eigenvalue 0 of S gives rise to $n - r_0$ eigenvalues of A^ℓ with modulus $\preceq \vartheta^\ell$. A similar statement holds for the eigenvectors; extra work needs to be done for getting results on the eigenvalues of A , since there might be some phase effects.

The results on eigenvectors follow in the same way, with a Davis-Kahan-like custom theorem proved in [116]. To ensure performant bounds, we must assume that the eigenvalues μ_i are well-separated, which is the reason why we suppose that $|\mu_i - \mu_j| > c$ for some $c > 0$ in the Hypotheses before the Master Theorem. Since $\|A^\ell - S\| \ll 1$ and the condition numbers of the matrices of interest are all $\preceq 1$, the Davis-Kahan theorem yields a bound of the form $|u_i - U_i|/|U_i| \preceq \|A^\ell - S\|/c$, which is $\ll 1$:

$$\left| u_i - \frac{U_i}{|U_i|} \right| \ll 1. \quad (5.21)$$

The ‘eigenvector part’ of the Master Theorem then easily follows, by the continuity of $u \mapsto \langle u, \varphi \rangle$ and the limits in the Pseudo-Master Theorem.

5.8 Proof of the Pseudo-Master Theorem

This section is devoted to the proof of the Pseudo-Master Theorem.

The first two sections gather some results on local approximations of random graphs: to a large extent, they are classical and well-known. We state them for completeness and because they give a good intuition on the following parts, but these statements can be retrieved using routine methods from random graph theory. Subsection 5.8.3 states a powerful concentration result on random graphs, proved in [28].

We use these results as a computational toolbox: in Subsections 5.8.4 to 5.8.7, we perform all the necessary calculations on our pseudo-eigenvectors, and we rely on an elegant generalization of Kesten martingales on inhomogeneous Galton-Watson trees. These sections differ from previous works in the sense that we had to adapt the computations to the general setting of our Master Theorem, with full inhomogeneity and non-symmetry.

Finally, the main ideas of previous works on trace methods are summarized in Subsection 5.8.8, as they apply without modifications to our setting.

5.8.1 The graph has few short cycles, and small neighbourhood growth

Sparse random graphs, that is, random graphs where the mean degree of vertices is $O(1)$, have been known for long to be locally-tree like, in the sense they have very few short cycles. In our model, we supposed that $P_{x,y} \leq d/n$ for some d . As a consequence, the expected degree of a vertex x is $d_x = P_{x,1} + \dots + P_{x,n}$, and is smaller than d , and the graph is stochastically dominated by a directed homogeneous Erdős-Rényi random graph with connectivity d/n . In turn, many random variables (cycle counts, edge number) are stochastically dominated by the undirected Erdős-Rényi graph of degree $2d/n$. Most properties on the local structure directly follow from classical results: absolutely no problem-specific work is needed here. We simply gather the results we will tacitly use in the sequel.

We note $(G, x)_t$ the *forward* neighbourhood of radius t around x in G , that is: the subgraph spanned by vertices y , for which there is a directed path with length smaller than t from x to y . The crucial choice will be the depth ℓ at which we look into the graph. We recall

$$\ell = \lfloor \kappa \log(n) \rfloor$$

where κ is an explicit constant depending on d . With such a choice, we have the crucial property that shallow neighborhoods are nearly trees. We say that a graph is *t-tangle-free* if for every vertex x , the subgraph $(G, x)_t$ has no more than one directed cycle.

Lemma 44. *There is a constant $c = c(d)$ such that G is 2ℓ -tangle-free with probability $\geq 1 - n^{-c}$. Moreover, if N is the number of vertices x such that $(G, x)_{2\ell}$ contains a cycle, then $N \leq 1$.*

The proof Lemma 44 follow from the choice of the constant κ in (5.8.1). We refer to [29] for the details.

5.8.2 The graph is locally approximated by trees

Let x be a vertex in G , and $\mathcal{N}(x)$ the set of its neighbours (in the forward sense: y is a neighbor of x if $(x, y) \in E$, not necessarily when $(y, x) \in E$). Then $\mathcal{N}(x)$ has the following distribution : each vertex $y \neq x$ is present in $\mathcal{N}(x)$ with probability $P_{x,y}$, independently from all other vertices. It is a well-known fact that whenever the $P_{x,y}$ are small, the law of $|\mathcal{N}(x)|$ is well approximated by $\text{Poi}(d_x)$ (the so-called ‘rare events theorem’). Moreover, conditionnally on $\{\mathcal{N}(x) \neq \emptyset\}$, the distribution of a random element in $\mathcal{N}(x)$ is equal to π_x , where

$$(\pi_x)_z = P_{x,z}/d_x. \quad (5.22)$$

That being said, the distribution of k elements in $\mathcal{N}(x)$ is *not* the product distribution $(\pi_x)^{\otimes k}$ because the elements have dependencies (they cannot be chosen multiple times, for instance), but these dependencies are nearly nonexistent. In fact, let us introduce a new distribution \mathbb{Q}_x on random multi-sets.

- The number of elements of the random multiset $\mathcal{M}(x)$ under \mathbb{Q}_x is a Poisson random variable with mean d_x ;
- Conditionnally on $|\mathcal{M}(x)| = k$, each of these k element is sampled independently on $[n]$ with probability distribution π_x .

The following proposition is a rigorous formulation of the intuitions given above. We set $\sigma_x = P_{x,1}^2 + \dots + P_{x,n}^2$; in our regime, $\sigma_x = O(1/n) = o(1)$.

Proposition 23. [116, Lemma 8] *We have*

$$d_{\text{TV}}(\mathbf{P}_x, \mathbb{Q}_x) \leq 2\sigma_x. \quad (5.23)$$

Building from those remarks, we define a random tree T_x as follows :

- the root (at depth 0) is a single vertex labeled x ,
- for each vertex i at depth t with label x_i , the children of i at depth $t + 1$ with their labels have the joint distribution $\mathcal{M}(x_i)$, independently from all other vertices at depth $\leq t$.

With this definition, the tree T_x is formally undirected, but we can view it as a directed tree with directions flowing out of the root. Proposition 23 then implies that the neighbourhood distributions in T_x and G are similar, as summarized in the following result:

Lemma 45. *Whenever $\ell = \lfloor \kappa \log(n) \rfloor$ with κ small enough, we have*

$$d_{\text{TV}}((G, x)_{2\ell}, (T_x, x)_{2\ell}) \ll 1. \quad (5.24)$$

This lemma will not directly be used in our proof; it is only a step in the proof of Proposition 24 thereafter. We stated it anyways because it gives a rigorous meaning to the fact that G is well-approximated locally by trees, and it intuitively gives a justification for our tree computations in subsequent parts of the proofs.

Note that we formulated this section only with forward neighborhoods; the propositions are also true with backward neighborhoods, and in this case we only have to see T_x as a directed tree, with edges oriented towards the root.

5.8.3 Concentration of linear functionals

The graph G (or the matrix A) is a collection of n^2 independent random variables: it is thus natural that if $F(G, x)$ is a function which depends only on a small neighborhood of the graph, then its space average

$$\frac{1}{n} \sum_{x \in [n]} F(G, x) \quad (5.25)$$

should be concentrated around its mean. This is the case, and a stronger statement actually holds: we saw in (5.24) that (G, x) and (T_x, x) nearly have the same distribution, and it is actually known that the functional in (5.25) is indeed concentrated around the expectation of the same functional applied on the trees (T_x, x) .

To formalize this, we say that a function $f : \mathcal{G} \times \mathbb{N} \mapsto \mathbb{R}$, where \mathcal{G} denotes the set of all graphs, is ℓ -local if $f(H, x)$ only depends on the ℓ -neighbourhood of x in H . Let (T_x, x) be the family of independent random rooted trees as in (5.24). Then the following lemma is true as long as the constant κ in ℓ is small enough ($\kappa < 0.01/\log(2d)$ will be sufficient).

Proposition 24. *Let \mathcal{F} be a family of functions $f : \mathcal{G} \times \mathbb{N} \mapsto \mathbb{R}$, with less than ≤ 1 elements. We suppose that each $f \in \mathcal{F}$ is a 2ℓ -local function, that for all graphs H and node x one has $\sup_{f \in \mathcal{F}} f(H, x) \leq |(H, x)_{2\ell}|^2 \times (c_n/n)$, for some $c_n > 0$. Then*

$$\sup_{f \in \mathcal{F}} \left| \sum_{x \in G} f(G, x) - \sum_{x \in [n]} \mathbf{E}[f(T_x, x)] \right| \preceq \frac{c_n}{n^{0.1}}.$$

Proof. This is a restatement of Theorem 12.5 in [28], with $\beta = 2$ and $\alpha = c_n/n$. The bound therein is $c_n n^{-1/2+2\kappa}$. A sufficiently small κ leads to our statement. Note that if $f(H, x) \leq |(H, x)_{2\ell}|$, then f also satisfies $f(H, x) \leq |(H, x)_{2\ell}|^2$, which is the reason why we chose to keep the exponent 2. \square

A crucial point in the proof of the Pseudo-Master Theorem is that the pseudo-eigenvectors U_i and V_i are ℓ -local functions of G , and thus so are their scalar products with one another. As such, all computations in (5.16)-(5.18) reduce to computing expectations on the random trees defined in the preceding subsection. We can apply Lemma 24 to G and the family $(T_x)_x$; the problem was reduced to computing the expectations of our functionals on the trees T_x .

5.8.4 Pseudo-eigenvectors on the random tree

Let us take a quick look at the entries of U and V , and pick one vertex $x \in [n]$ and some index i . By definition,

$$\begin{aligned} U_i(x) &= \mu_i^{-\ell} (A^\ell \varphi_i)(x) \\ &= \mu_i^{-\ell} \sum_{x_1, \dots, x_\ell} W_{x, x_1} W_{x_1, x_2} \cdots W_{x_{\ell-1}, x_\ell} \varphi_i(x_\ell), \end{aligned} \quad (5.26)$$

where the sum runs over all the paths in the graph G , ie sequences of vertices such that the edges (x_s, x_{s+1}) are present in the graph.

This is, by definition, an ℓ -local function ; its counterpart on the tree can thus be defined as

$$\tilde{U}_i(x) = \mu_i^{-\ell} \sum_{x_1, \dots, x_\ell \in T_x} W_{x, \iota(x_1)} W_{\iota(x_1), \iota(x_2)} \cdots W_{\iota(x_{\ell-1}), \iota(x_\ell)} \varphi_i(\iota(x_\ell)), \quad (5.27)$$

where the sum ranges over the vertices x_ℓ at depth ℓ in T_x and the unique path x, \dots, x_ℓ connecting x to x_ℓ , and $\iota(x_i)$ is the label of x_i .

5.8.5 The martingale equation

Let $\mathbf{E}_{\ell-1}$ denote the conditional expectation with respect to the first $\ell - 1$ generations of T_x . Write

$$\mathbf{E}_{\ell-1}[\tilde{U}_i(x)] = \mu_i^{-\ell} \sum_{x_1, \dots, x_{\ell-1}} W_{x, \iota(x_1)} \cdots W_{\iota(x_{\ell-2}), \iota(x_{\ell-1})} \mathbf{E}_{\ell-1} \left[\sum_{x_\ell} W_{\iota(x_{\ell-1}), \iota(x_\ell)} \varphi_i(\iota(x_\ell)) \right].$$

Let us note $y = \iota(x_{\ell-1})$; the inner expectation reads

$$\mathbf{E}_{\ell-1} \left[\sum_{x_\ell} W_{\iota(x_{\ell-1}), \iota(x_\ell)} \varphi_i(\iota(x_\ell)) \right] = \mathbf{E} \left[\sum_{z \in \mathcal{M}(y)} W_{y, z} \varphi_i(z) \right]$$

This is a sum of a $\text{Poi}(d_y)$ number of independent random variables with distribution π_y , so

$$\begin{aligned} \mathbf{E} \left[\sum_{z \in \mathcal{M}(y)} W_{y, z} \varphi_i(z) \right] &= d_y \times \sum_{z \in [n]} \frac{P_{y, z}}{d_y} W_{y, z} \varphi_i(z) \\ &= \sum_{y \in [n]} P_{y, z} W_{y, z} \varphi_i(z) \\ &= \sum_{y \in [n]} Q_{y, z} \varphi_i(z) \\ &= (Q \varphi_i)(y) \\ &= \mu_i \varphi_i(y) \end{aligned}$$

where, in the last line, we used the fact the φ_i is an eigenvector of Q . Recalling that $y = \iota(x_{\ell-1})$, we have

$$\mathbf{E}_{\ell-1}[\tilde{U}_i(x)] = \mu_i^{-\ell+1} \sum_{x_1, \dots, x_{\ell-1}} W_{x, \iota(x_1)} \cdots W_{\iota(x_{\ell-2}), \iota(x_{\ell-1})} \varphi_i(\iota(x_{\ell-1})).$$

In fact, defining $\tilde{U}_i(x, t)$ by replacing ℓ by t in the definition of $\tilde{U}_i(x)$, we showed that

$$\text{The random process } t \mapsto \tilde{U}_i(x, t) \text{ is a martingale.} \quad (5.28)$$

The common expectation is easily seen to be $\mathbf{E}[\tilde{U}_i(x, 0)] = \varphi_i(x)$.

5.8.6 Proving (5.16)-(5.17)-(5.18)

We only give the argument for (5.16), the others are done in a similar fashion. The (i, j) coefficient of U^*V is $\langle U_i, V_j \rangle$, which can be rewritten as

$$\langle U_i, V_j \rangle = (\mu_i \mu_j)^{-\ell} \langle A^\ell \varphi_i, (A^*)^\ell \psi_j \rangle = (\mu_i \mu_j)^{-\ell} \langle A^{2\ell} \varphi_i, \psi_j \rangle$$

Let $f(x) = f_{i,j}(x) = \psi_j(x)[A^{2\ell} \varphi_i](x)$; it is easily seen that $\langle U_i, V_j \rangle = \sum_{x \in [n]} f(G, x)$, and we are ready to apply the concentration property in Proposition 24 to the family $\{f_{i,j}\}$.

Lemma 46 (correlations between pseudo-eigenvectors are concentrated).

$$\sup_{i,j} \left| \langle U_i, V_j \rangle - \mu_i^{-\ell} \mu_j^{-\ell} \sum_{x \in [n]} \psi_j(x) \mathbf{E}[\tilde{U}_i(x, 2\ell)] \right| \ll 1 \quad (5.29)$$

Proof. The delocalization properties of the φ_i (Hypothesis (5.2)), the tangle-free property (there is no more than one cycle in (G, x)) and the fact that $|W_{x,y}| \leq \|W\|_\infty$ all together imply that

$$f(G, x) \leq \frac{c \|W\|_\infty^\ell}{n} |(G, x)_{2\ell}|,$$

for some universal constant c . But since $\|W\|_\infty < c'$ for some c' , any choice for κ sufficiently small will give (for instance) $\|W\|_\infty^\ell \leq n^{0.01}$. Proposition 24 straightforwardly leads to a $\leq n^{-\varepsilon}$ error for some small ε . With our notations, this is $\ll 1$. □

We are now in a position to use property (5.28): $\mathbf{E}[\tilde{U}_i(x, 2\ell)] = \varphi_i(x)$, and the orthogonality property of the left and right eigenvectors yields $|\langle U_i, V_j \rangle - \delta_{ij}| \ll 1$. It is then straightforward to go from this elementwise bound to (5.16): for any $r_0 \times r_0$ matrix M , one has $\|M\| \leq r \|M\|_\infty$. Since we also have $r \leq 1$, we obtain $\|U^*V - I\| \leq \|U^*V - I\|_\infty \ll 1$.

The same proof works for V , with the subtle difference that these left-pseudo-eigenvectors V_i are backward-looking: $V_i(x)$ is a local function of $(G, x)_\ell^-$, where the $-$ superscript denotes the backward ball. But the proof is the same: Lemma 24 is true for backward functionals and the couplings in Subsection 5.8.2 are the same, but with the T_x oriented towards the root; everything works exactly the same.

5.8.7 Proving (5.19): martingale correlations

Statements in (5.19) are trickier: even if Lemma 24 still allows approximating $\langle U_i, U_j \rangle$ with a tree computation, there are some strong dependencies between $\tilde{U}_i(x, t)$ and $\tilde{U}_j(x, t)$ that we cannot neglect.

Rewriting the correlation term. Proceeding as before with $f(G, x) = [A^\ell \varphi_i](x) \cdot [A^\ell \varphi_j](x)$, we have

$$\sup_{i,j} \left| \langle U_i, U_j \rangle - \sum_{x \in [n]} \mathbf{E} \left[\tilde{U}_i(x) \tilde{U}_j(x) \right] \right| \ll 1. \quad (5.30)$$

We recognize a covariance term between two martingales; let us then introduce the increments

$$\delta_t = \mathbf{E}_{t-1} \left[(\tilde{U}_i(x, t) - \tilde{U}_i(x, t-1)) (\tilde{U}_j(x, t) - \tilde{U}_j(x, t-1)) \right]$$

A classical use of the martingale property implies that

$$\mathbf{E} \left[\tilde{U}_i(x) \tilde{U}_j(x) \right] = \varphi_i(x) \varphi_j(x) + \mathbf{E}[\delta_1 + \cdots + \delta_\ell]$$

The increment δ_t has an explicit expression:

$$\delta_t = \mu_i^{-t} \mu_j^{-t} \sum_{\substack{x_1, \dots, x_{t-1} \\ x'_1, \dots, x'_{t-1}}} \prod_{s=1}^{t-1} W_{\iota(x_{s-1}), \iota(x_s)} W_{\iota(x'_{s-1}), \iota(x'_s)} \times \\ \mathbf{E}_{t-1} \left[\sum_{\substack{x_{t-1} \rightarrow x_t \\ x'_{t-1} \rightarrow x'_t}} W_{\iota(x_{t-1}), \iota(x_t)} W_{\iota(x'_{t-1}), \iota(x'_t)} \varphi_i(\iota(x_t)) \varphi_j(\iota(x'_t)) - \mu_i \mu_j \varphi_i(\iota(x_{t-1})) \varphi_j(\iota(x'_{t-1})) \right]. \quad (5.31)$$

Covariance of Poisson sums. In the sum above, the only nonzero terms are when $x_{t-1} = x'_{t-1}$: otherwise, the inner sum becomes a product of two independent random variables of respective expectations $\mu_i \varphi_i(x_{t-1})$ and $\mu_j \varphi_j(x'_{t-1})$. Writing again $y = \iota(x_{t-1})$, the conditional expectation becomes

$$C_y := \mathbf{E} \left[\sum_{z, z' \in \mathcal{M}(y)} W_{y,z} W_{y,z'} \varphi_i(z) \varphi_j(z') - \mu_i \mu_j \varphi_i(y) \varphi_j(y) \right] \\ = \text{Cov} \left(\sum_{z \in \mathcal{M}(y)} W_{y,z} \varphi_i(z), \sum_{z \in \mathcal{M}(y)} W_{y,z} \varphi_j(z) \right). \quad (5.32)$$

We then make use of the following elementary lemma:

Lemma 47. *If N is a Poisson random variable and $(A_1, B_1), (A_2, B_2), \dots$ are iid copies of a couple of random variable (A, B) , then*

$$\text{Cov} \left(\sum_{k=1}^N A_k, \sum_{k=1}^N B_k \right) = \mathbf{E}[N] \mathbf{E}[AB]. \quad (5.33)$$

In (5.32), $N = |\mathcal{M}(y)|$ is a $\text{Poi}(d_y)$ random variable, and the couple (A, B) is simply

$$A = W_{y,Z} \varphi_i(Z) \quad \text{and} \quad B = W_{y,Z} \varphi_j(Z),$$

where Z is a random index on $[n]$ with distribution π_y . Computing C_y is now straightforward:

$$C_y = d_y \mathbf{E}[W_{y,Z} \varphi_i(Z) W_{y,Z} \varphi_j(Z)] \\ = d_y \sum_{z \in [n]} \frac{P_{yz}}{d_y} W_{yz}^2 \varphi_i(z) \varphi_j(z) \\ = (K \varphi^{i,j})(y).$$

Let us now come back to (5.31). We know that we can remove every term with $x_{t-1} \neq x'_{t-1}$, and the computations above further reduce it to

$$\delta_t = (\mu_i \mu_j)^{-t} \sum_{x_1, \dots, x_{t-1}} \prod_{s=1}^{t-1} W_{\iota(x_{s-1}), \iota(x_s)}^2 (K \varphi^{i,j})(\iota(x_{t-1}))$$

We recognize an expression similar to the definition of \tilde{U}_i in (5.27). Using the same methods, we are able to show that

$$\mathbf{E}[\delta_t] = \frac{[K^t \varphi^{i,j}(x)]}{(\mu_i \mu_j)^t}.$$

Finally, summing the increments, we get

$$\begin{aligned} \mathbf{E}[\tilde{U}_i(x)\tilde{U}_j(x)] &= \varphi_i(x)\varphi_j(x) + \sum_{t=1}^{\ell} K^t \varphi^{i,j}(x) \\ &= \sum_{t=0}^{\ell} \frac{K^t \varphi^{i,j}(x)}{\mu_i^t \mu_j^t}. \end{aligned} \quad (5.34)$$

The expression for $\langle U_i, U_j \rangle$ is obtained by summing (5.34) over all values of x and plugging this into (5.19), thus obtaining:

$$\sup_{i,j} \left| \langle U_i, U_j \rangle - \sum_{t=0}^{\ell} \frac{\langle \mathbf{1}, K^t \varphi^{i,j} \rangle}{(\mu_i \mu_j)^t} \right| \ll \vartheta^\ell. \quad (5.35)$$

Statement (5.35) is pretty close to (5.19), at the sole difference of the summation index, which is stopped at ℓ . Also, note that (5.35) is valid for every i, j in $[r]$, not just in $[r_0]$. However, whenever $i, j \in [r_0]$, we have by definition

$$\sqrt{\rho} \leq \sqrt{\tau} \mu_i \quad \text{and} \quad \sqrt{\rho} \leq \sqrt{\tau} \mu_j,$$

and the spectral radius of the matrix $K/(\mu_i \mu_j)$ is thus at most τ . Since the entries of $\varphi^{i,j}$ are of order $O(n^{-1})$, we have the bound

$$\sup_{i,j} \left| \sum_{t=\ell+1}^{\infty} \frac{\langle \mathbf{1}, K^t \varphi^{i,j} \rangle}{(\mu_i \mu_j)^t} \right| \leq C\tau^\ell \ll 1.$$

Combined with (5.35), this finally ends the proof of (5.19).

5.8.8 The trace method

All that remains now is to prove equation (5.20); that is, once we showed that the first r_0 eigenvalues of A^ℓ are close to the μ_i^ℓ , it remains to show that the $n - r_0$ eigenvalues are confined in a circle of radius ϑ^ℓ . This is done in three parts:

- (i) a *tangle-free* decomposition, expressing A^ℓ as a product involving its expectation Q , the powers A^t for $t \leq \ell$ and some additional random matrices $\underline{A}^{(t)}$, which can be understood as the centered versions of A^t .
- (ii) a trace method on the aforementioned centered matrices, inspired by [58], to bound their spectral radius;
- (iii) finally, a scalar product bound to control the whole sum whenever $x \in H^\perp$.

Tangle-free decomposition. We showed in Lemma 44 that with high probability the graph G is 2ℓ -tangle-free; as a result, for all vertices u, v and $t \leq \ell$ we have

$$A_{uv}^t = \sum_{\gamma \in F_{uv}^{(t)}} \prod_{s=1}^t A_{\gamma_{s-1}\gamma_s},$$

where the sum ranges over all tangle-free paths (i.e. paths whose induced graph is tangle-free) of length t in the complete graph K_n . The centered matrices $\underline{A}^{(t)}$ are thus similarly defined as

$$\underline{A}_{uv}^{(t)} = \sum_{\gamma \in F_{uv}^{(t)}} \prod_{s=1}^t \underline{A}_{\gamma_{s-1}\gamma_s},$$

where $\underline{A} = A - Q$ is the centered version of A .

To decompose A^t in terms of the latter matrices, we make use of the following equality, valid for any $(a_i), (b_i)$:

$$\prod_{s=1}^t a_s = \prod_{s=1}^t b_s + \sum_{k=1}^t \left(\prod_{s=1}^{k-1} b_s \right) (a_k - b_k) \left(\prod_{s=k+1}^t a_s \right)$$

Applying this to the two above equations yields

$$A_{uv}^\ell = \underline{A}_{uv}^{(\ell)} + \sum_{k=1}^{\ell} \sum_{\gamma \in F_{uv}^{(\ell)}} \left(\prod_{s=1}^{k-1} \underline{A}_{\gamma_{s-1}\gamma_s} \right) (Q_{\gamma_{k-1}\gamma_k}) \left(\prod_{s=k+1}^{\ell} A_{\gamma_{s-1}\gamma_s} \right)$$

Each term in the sum above is close to $[\underline{A}^{(k-1)} Q A^{\ell-k-1}]_{uv}$, with the following caveat : the concatenation of a path in $F_{uv}^{(k-1)}$ and one in $F_{wx}^{(\ell-k-1)}$ is not necessarily tangle-free ! Nevertheless, we write

$$[\underline{A}^{(k-1)} Q A^{\ell-k-1}]_{uv} = \sum_{\gamma \in F_{uv}^{(\ell)}} \left(\prod_{s=1}^{k-1} \underline{A}_{\gamma_{s-1}\gamma_s} \right) (Q_{\gamma_{k-1}\gamma_k}) \left(\prod_{s=k+1}^{\ell} A_{\gamma_{s-1}\gamma_s} \right) + [R_k^{(\ell)}]_{uv},$$

so that we finally get

$$A^\ell = \underline{A}^{(\ell)} + \sum_{k=1}^{\ell} \underline{A}^{(k-1)} Q A^{\ell-k-1} - \sum_{k=1}^{\ell} R_k^{(\ell)} \quad (5.36)$$

Bounding $\|\underline{A}^{(k)}\|$. The trace method gets its name from its leverage of the following inequality:

$$\|\underline{A}^{(k)}\| = \left\| \left(\underline{A}^{(k)} \underline{A}^{(k)*} \right)^m \right\|^{\frac{1}{2m}} \leq \text{tr} \left[\left(\underline{A}^{(k)} \underline{A}^{(k)*} \right)^m \right]^{\frac{1}{2m}}.$$

The above trace can be expanded as

$$\text{tr} \left[\left(\underline{A}^{(k)} \underline{A}^{(k)*} \right)^m \right] = \sum_{\gamma} \prod_{i=1}^{2m} \prod_{t=1}^k \underline{A}_{\gamma_{i,t-1}\gamma_{i,t}}, \quad (5.37)$$

where the sum ranges over all concatenations of $2m$ k -paths $\gamma = (\gamma_1, \dots, \gamma_{2m})$ such that γ_i is tangle-free for all i , and with adequate boundary conditions.

The goal is now to use a Markov bound, and thus to compute the expectation in (5.37); the key argument is the following:

Each term in the sum (5.37) has expectation zero unless γ visits each of its edges at least twice.

We now classify the subgraphs γ by their number of vertices $v(\gamma)$ and edges $e(\gamma)$, and we say that γ and γ' are equivalent if there exists a permutation $\sigma \in \mathfrak{S}_n$ such that $\sigma(\gamma_{i,t}) = \gamma'_{i,t}$ for all i, t . All that remains is to bound the number of such equivalence classes and their contributions to the overall expectation; this is done in [28, 116] and yields the following results:

Lemma 48. *The number $\mathcal{N}(v, e)$ of equivalence classes of subgraphs γ with v vertices and e edges such that each edge is visited at least twice satisfies*

$$\mathcal{N}(v, e) \leq (2km)^{6m(e-v+1)+2m}, \quad (5.38)$$

and for each $\tilde{\gamma} \in \mathcal{N}(v, e)$, the contribution $\mathcal{W}(\tilde{\gamma})$ of the equivalence class to the trace expectation is bounded above:

$$\mathcal{W}(\tilde{\gamma}) \leq \mathcal{W}(v, e) := n^{v-e} \rho^e \left(\frac{dL^2}{\rho} \right)^{3(e-v)+8m}$$

Now, all that remains is to sum the terms $\mathcal{N}(v, e)\mathcal{W}(v, e)$ over all possible choices of v and e , to find the following bound:

$$\left\| \underline{A}^{(k)} \right\| \preceq \vartheta^k.$$

The operator norm of $R_k^{(\ell)}$ is bounded using similar arguments.

A scalar product bound. Let $w \in \mathbb{R}^n$; with the previously established bounds, we have

$$\left\| A^\ell w \right\| \preceq \vartheta^\ell + \sum_{k=1}^{\ell} \vartheta^k \left\| QA^{\ell-k-1} w \right\|$$

It remains to bound the rightmost norm whenever w is orthogonal to the $A^\ell \psi_i$. First, we use the eigendecomposition of Q :

$$\begin{aligned} \left\| QA^{\ell-k-1} w \right\| &= \left\| \sum_{i \in [r]} \mu_i \varphi_i \psi_i^* A^{\ell-k-1} w \right\| \\ &\leq \mu_1 \sum_{i \in [r]} \langle (A^*)^{\ell-k-1} \psi_i, w \rangle \end{aligned}$$

Since $\langle w, (A^*)^\ell \psi_i \rangle = 0$ by assumption, we can use Cauchy-Schwarz and a telescopic sum to bound the scalar product:

$$\langle (A^*)^{\ell-k-1} \psi_i, w \rangle \leq |\mu_i|^{\ell-k-1} \sum_{t=\ell-k-1}^{\ell-1} |\mu_i|^{-t} \left\| (A^*)^t \psi_i - \mu_i^{-1} (A^*)^{t+1} \psi_i \right\|$$

The final bound thus stems from the following lemma [28]:

Lemma 49. *For every $t \leq \ell$ and $i \in [r_0]$ we have*

$$\left\| A^t \varphi_i - \mu_i^{-1} A^{t+1} \varphi_i \right\|^2 \preceq \vartheta^{2t},$$

and the same holds for A^* and ψ_i .

Indeed, adopting again the notations from 5.8.4, we have to compute the expectation on the tree of the $t + 1$ -local function

$$f(T, x) = (\tilde{U}_i(x, t) - \mu_i^{-1} \tilde{U}_i(x, t + 1))^2,$$

which can be understood as the increment variance of the martingale $\tilde{U}_i(x, t)$. Subsequently, we can use the same arguments as in 5.8.7 with $i = j$, which shows

$$\mathbf{E}[f(T, x)] = [K^{t+1} \varphi^{i,i}](x) \leq \frac{\vartheta^{2t}}{n}.$$

Summing this for $x \in [n]$ and applying Proposition 24, we are done.

5.9 Master Theorem for the stochastic block model

We recall the definition of P : given the cluster membership functions $\sigma_g, \sigma_d : [n] \rightarrow [r]$ and a connectivity matrix F of size $r \times r$, the entries of P are given by

$$P_{x,y} = \frac{F_{\sigma_g(x), \sigma_d(y)}}{n}.$$

We introduce the probability vectors p, q , which are equal to the relative clusters sizes, as well as the ‘cluster intersection’ matrix Π :

$$(p)_i = \frac{\text{Card}(\sigma_g^{-1}(i))}{n}, \quad (q)_i = \frac{\text{Card}(\sigma_d^{-1}(i))}{n} \quad \text{and} \quad \Pi_{i,j} = \frac{\text{Card}(\sigma_g^{-1}(j) \cap \sigma_d^{-1}(i))}{n}.$$

The cluster membership matrices Σ_g, Σ_d are matrices of size $n \times r$, defined by

$$(\Sigma_g)_{x,i} = \mathbf{1}_{\sigma_g(x)=i} \quad \text{and} \quad (\Sigma_d)_{x,i} = \mathbf{1}_{\sigma_d(x)=i}.$$

With these notations, the following identities hold:

$$\begin{aligned} p &= \frac{1}{n} \Sigma_g \mathbf{1}, & \text{diag}(p) &= \frac{1}{n} (\Sigma_g)^* \Sigma_g, & P &= \frac{1}{n} \Sigma_g F (\Sigma_d)^*; \\ q &= \frac{1}{n} \Sigma_d \mathbf{1}, & \text{diag}(q) &= \frac{1}{n} (\Sigma_d)^* \Sigma_d, & \Pi &= \frac{1}{n} (\Sigma_d)^* \Sigma_g \end{aligned} \tag{5.39}$$

5.9.1 Spectral decomposition of P

We prove in this section a slightly refined version of Proposition 22.

Proposition 25. *The non-zero eigenvalues of P are exactly those of the modularity matrix $F\Pi$, with the same multiplicities. Further, each right eigenvector of P is of the form $\Sigma_g f$, where f is a right eigenvector of $F\Pi$, while each left eigenvector of P is of the form $\Sigma_d g$ with g a left eigenvector of ΠF .*

The proof of this proposition relies on the following elementary lemma, a consequence of the Sylvester identity $\det(z - XY) = \det(z - YX)$.

Lemma 50. *Let X be a $n \times m$ matrix and Y a $m \times n$ matrix. Then the non-zero eigenvalues of XY are the same as those of YX , with identical multiplicities.*

of Proposition 25. We apply the above lemma to $X = \Sigma_g$ and $Y = \frac{1}{n}F(\Sigma_d)^*$; the identities in (5.39) show that $XY = P$ and $YX = F\Pi$, which directly gives the desired result. Now, let f be a right eigenvector of $F\Pi$, with associated eigenvalue λ , and define $\varphi = \Sigma_g f$. Then

$$P\varphi = \frac{1}{n}\Sigma_g F(\Sigma_d)^*\Sigma_g f = \Sigma_g F\Pi f = \lambda\Sigma_g f = \lambda\varphi.$$

Combined with the previous result, this implies that all right eigenvectors of P with non-zero eigenvalues are of the form $\Sigma_g f$ for an eigenvector f of $F\Pi$. In particular, they are constant on the left clusters. The result on left eigenvectors is proved similarly. \square

Let f_1, \dots, f_r (resp. g_1, \dots, g_r) be a basis of right (resp. left) eigenvectors of $F\Pi$ (resp. ΠF), not necessarily normalized. We define the entrywise products $f^{i,j}$ and $g^{i,j}$ as in equation (5.12). The following statement describes the unit eigenvectors of P in terms of f_i, g_j .

Lemma 51. *Let (φ_i) (resp. (ξ_i)) be a basis of normed right (resp. left) eigenvectors of P . Then*

$$\varphi_i = \frac{\Sigma_g f_i}{\sqrt{n\langle p, f^{i,i} \rangle}} \quad \text{and} \quad \xi_i = \frac{\Sigma_d g_i}{n\sqrt{n\langle q, g^{i,i} \rangle}}.$$

Proof. In light of Proposition 25, we only have to compute the norms of $\Sigma_g f_i$ and $\Sigma_d g_i$:

$$|\Sigma_g f_i|^2 = f_i^*(\Sigma_g)^*\Sigma_g f_i = n f_i^* \text{diag}(p) f_i = n\langle p, f^{i,i} \rangle.$$

The first equality follows, and the second is proved in identical fashion. \square

5.9.2 Master Theorem for SBM

We are now ready to prove the version of the Master Theorem, adapted to the directed SBM.

Theorem 35. *Let r_0 be the number of eigenvalues ν_i such that $\nu_i^2 > \nu_1$. Then, with high probability the following holds: the r_0 highest eigenvalues $\lambda_1, \dots, \lambda_{r_0}$ of A satisfy*

$$|\lambda_i - \nu_i| = o(1),$$

and all other eigenvalues of A are asymptotically smaller than $\sqrt{\nu_1}$. Further, if v_i, u_i are a pair of left/right unit eigenvectors of A associated with λ_i , then

$$|\langle u_i, \varphi_j \rangle| = a_{i,j} + o(1) \quad \text{and} \quad |\langle v_i, \xi_j \rangle| = b_{i,j} + o(1),$$

where $a_{i,j}$ and $b_{i,j}$ are defined as

$$a_{i,j} = \frac{|\langle p, f^{i,j} \rangle|}{\sqrt{\langle p, f^{j,j} \rangle \langle p, (\mathbf{I} - \nu_i^{-2} F\Pi)^{-1} f^{i,i} \rangle}} \quad (5.40)$$

$$b_{i,j} = \frac{|\langle q, g^{i,j} \rangle|}{\sqrt{\langle q, g^{j,j} \rangle \langle q, (\mathbf{I} - \nu_i^{-2} (\Pi F)^*)^{-1} g^{i,i} \rangle}} \quad (5.41)$$

The first part of this theorem is an application of Theorem 31, by means of Proposition 25. It remains to compute the $a_{i,j}$ and $b_{i,j}$ as a function of the SBM parameters. We recall that the definition of Γ and Γ^* are in (5.13).

Lemma 52. *Let $z \in \mathbb{C}$, and $h \in \mathbb{R}^r$. Then*

$$\Gamma(z, \Sigma_g h) = n \langle p, (\mathbf{I} - z^{-1} F\Pi)^{-1} h \rangle \quad \text{and} \quad \Gamma^*(z, \Sigma_d h) = n \langle q, (\mathbf{I} - z^{-1} (\Pi F)^*)^{-1} h \rangle.$$

Proof. Since the graph is unweighted, we have $K = P$. By an immediate recursion, we have

$$P^t \Sigma_g h = \Sigma_g (F\Pi)^t h,$$

so that using the first identity of (5.39)

$$\langle \mathbf{1}, P^t \Sigma_g h \rangle = n \langle p, (F\Pi)^t h \rangle.$$

Summing over all t and using the Von Neumann summation implies the first equality, and the second is alike. \square

As a result,

$$\nu_i^2 R_i = \Gamma(\nu_i^2, \varphi^{i,i}) = \Gamma(\nu_i^2, \Sigma_g f^{i,i}) = n \langle p, (\mathbf{I} - \nu_i^{-2} F\Pi)^{-1} f^{i,i} \rangle,$$

where we used the previous lemma, and similarly

$$\nu_i^2 L_i = n \langle q, (\mathbf{I} - \nu_i^{-2} (\Pi F)^*)^{-1} g^{i,i} \rangle.$$

of Theorem 35. Using the expressions in Theorem 31, we have

$$a_{i,j} = \frac{|\langle \varphi_i, \varphi_j \rangle|}{|\nu_i| \sqrt{R_i}} \quad \text{and} \quad b_{i,j} = \frac{|\langle \xi_i, \xi_j \rangle|}{|\nu_i| \sqrt{L_i}}.$$

Computing the numerators is straightforward using Lemma 51:

$$\begin{aligned} \langle \varphi_i, \varphi_j \rangle &= \frac{\langle \Sigma_g f_i, \Sigma_g f_j \rangle}{n \sqrt{\langle p, f^{i,i} \rangle \langle p, f^{j,j} \rangle}} \\ &= \frac{\langle p, f^{i,j} \rangle}{\sqrt{\langle p, f^{i,i} \rangle \langle p, f^{j,j} \rangle}}. \end{aligned}$$

On the other hand, for the denominator, we have

$$\nu_i^2 R_i = \Gamma(\nu_i^2, \varphi^{i,i}) = \frac{\Gamma(\nu_i^2, \Sigma_g f^{i,i})}{n \langle p, f^{i,i} \rangle},$$

and using Lemma 52 we find

$$\nu_i^2 R_i = \frac{\langle p, (\mathbf{I} - \nu_i^{-2} F\Pi)^{-1} f^{i,i} \rangle}{\langle p, f^{i,i} \rangle}$$

It simply remains to simplify the expressions to prove the formula for $a_{i,j}$, and the exact same method works for $b_{i,j}$ as well. \square

5.10 Pathwise SBM

In this section, we derive the thresholds shown in Subsection 5.3.2. Let us first give some motivation on this model.

5.10.1 Motivation

Stochastic block-models with a pathwise structure as in (5.5) are well suited for modeling flow data: the intra-block connectivity is the same $s/2$ in any blocks; connections can only happen between adjacent blocks and the rate depends on the flow order: edges have a higher chance of appearing from one block V_i to the following V_{i+1} , than between one block V_i and the preceding one V_{i-1} (η versus $1 - \eta$). The model in [81] is a small variant of this one: in their model, undirected edges appear between adjacent blocks, and then one direction is chosen uniformly at random with probability η for edges between adjacent blocks, and with probability $1/2$ for edges inside the same block. Our model allows the appearance of a double edge $(x, y), (y, x)$, which is not the case in their model. However, in the sparse regime where s does not depend on n , the two models are contiguous and our results can be shown to hold for both.

Remark. *The works [119, 120] are close in spirit to ours, although they do not approach sparse regimes. We chose to perform the computations for the specific F above, but the same computations can be done for other models. In particular, it would be interesting to perform these computations for the matrix given in [119], p. 73 and to recover the shape observed by the author in Figs 2.23-24.*

5.10.2 Model density

Let us compute the mean degree d in this model, when there are k blocks and the asymmetry parameter is η . All the blocks have the same size n/k , hence they have $(n/k)^2$ entries; on the k diagonal blocks, the mean degree is $s/2$, on the $k - 1$ upper-diagonal blocks it is $s(1 - \eta)$ and on the $k - 1$ lower diagonals they are $s\eta$, so

$$d = \frac{1}{n^2} \left[k \frac{s n^2}{2 k^2} + (k - 1) s (1 - \eta) \frac{n^2}{k^2} + (k - 1) s \eta \frac{n^2}{k^2} \right] = \frac{s}{k} \left(\frac{3}{2} - \frac{1}{k^2} \right). \quad (5.42)$$

For each k , the unique parameter s such that the model has mean degree d is given by $s(k, d) = kd(3/2 - k^{-2})^{-1}$. Table 5.1 gives the values of s used in our simulations in Figure 5.3.

requested mean degree d	2	3	4
number of blocks $k = 2$	3.2	4.8	6.4
$k = 4$	5.5	8.3	11.1
$k = 6$	8.1	12.2	16.3

Table 5.1: Value of $s = s(k, d)$ such that, for the given number of blocks k , the mean degree of the model is equal to d .

5.10.3 Eigendecomposition of tridiagonal Toeplitz matrices.

Since the blocks on the right and on the left are identical and have the same size n/r , the formulas in Theorem 35 are really easy to use.

Let F be the tridiagonal $k \times k$ Toeplitz matrix defined in (5.5). We extract the following formulas from [104] (see (4) for eigenvalues and (7)-(8) for eigenvectors). The k eigenvalues are

$$\tilde{\nu}_k = \frac{s}{2} + 2s \cos \left(\frac{\pi k}{r+1} \right) \sqrt{\eta(1-\eta)} \quad (1 \leq k \leq r). \quad (5.43)$$

and the corresponding right-eigenvectors f_i and left-eigenvectors g_i are

$$f_i(j) \propto \left(\frac{1-\eta}{\eta} \right)^{j/2} \sin \left(\frac{ij\pi}{r+1} \right) \quad g_i(j) \propto \left(\frac{\eta}{1-\eta} \right)^{j/2} \sin \left(\frac{ij\pi}{r+1} \right) \quad (5.44)$$

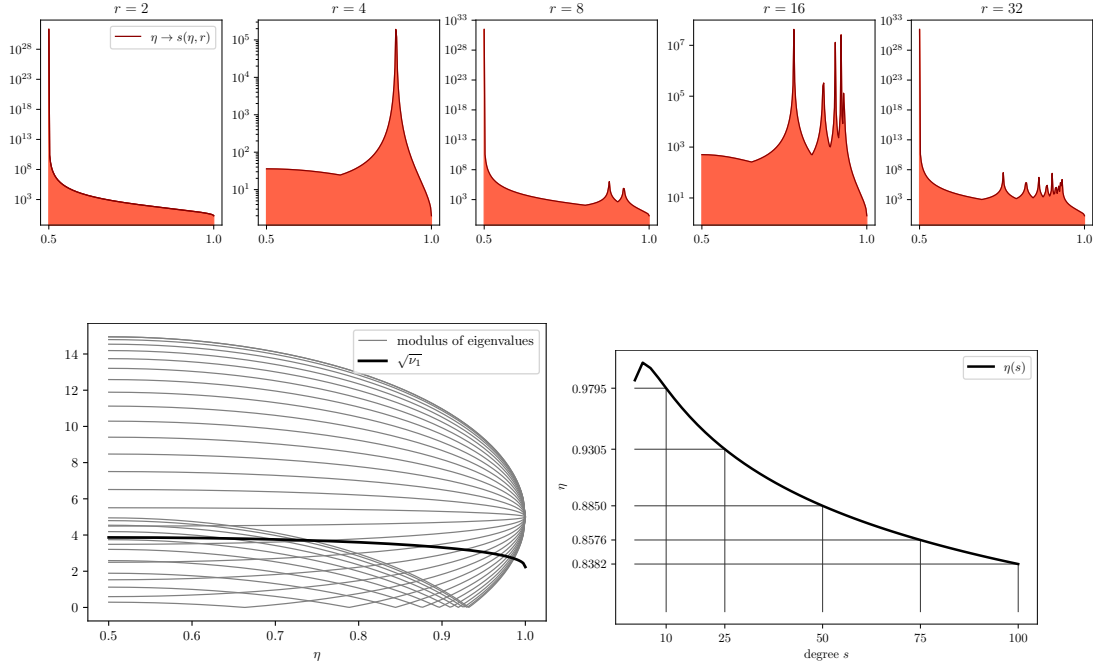


Figure 5.6: **Top:** A plot of the shape of the RHS of (5.45) as a function of η , for different values of r ranging from 2 to 2^5 , with a log-scale on the y -axis. The red zone represent the set of degree parameters s that lie below the threshold $s(\eta, r)$ in (5.45): they are the values for which our method does not yield full reconstruction guarantees, since at least one eigenvalue/eigenvector couple is ‘lost in the bulk’ of the spectrum of A .

Bottom left: Here $s = 10$ and $r = 30$. The absolute values of the eigenvalues $\nu_k = s(1/2 + c_k\theta)$ are plotted in grey, while the threshold $\rho = \sqrt{\nu_1}$ is in bold black.

Bottom right: the shape of $\eta \mapsto \eta(s)$.

5.10.4 Digression: the full threshold

We hereby state an auxiliary result that might be of potential interest, which simply consists in an application of Theorem 35.

Proposition 26. *In the pathwise SBM as in (5.5) with r blocks, degree parameter $s > 1$ and asymmetry parameter $1/2 \leq \eta \leq 1$, the r distinct eigenvectors can be detected if*

$$\frac{s}{r} > \frac{1/2 + c_1\theta}{\min_{k \in [r]} (1/2 + c_k\theta)^2} \quad (5.45)$$

where $\theta = 2\sqrt{\eta(1-\eta)}$ and $c_k = \cos(k\pi/(r+1))$.

We plotted the threshold for several values of r in Figure 5.6. It is interesting to note that for specific values of θ , the threshold for s is $+\infty$; this corresponds to cases where $\theta = -2c_k$, and one eigenvalue is zero. This is a good illustration of the principle discussed earlier : $r_0 = 1$ suffices to recover cluster information (since the top eigenvector is nonconstant), even though it’s completely impossible to recover as many informative eigenvectors as there are clusters.

5.10.5 Computations for the two-block case

We place ourselves in the setting of Theorem 35; recall that

$$F = \begin{bmatrix} s/2 & s\eta \\ s(1-\eta) & s/2 \end{bmatrix}, \quad F\Pi = \Pi F = \begin{bmatrix} s/4 & s\eta/2 \\ s(1-\eta)/2 & s/4 \end{bmatrix} =: M. \quad (5.46)$$

Define $\theta = 2\sqrt{\eta(1-\eta)}$; as in (5.43), the eigenvalues of M are

$$\nu_1 = \frac{s(1+\theta)}{4} \quad \text{and} \quad \nu_2 = \frac{s(1-\theta)}{4},$$

with associated right and left eigenvectors

$$\begin{aligned} f_1 &= (\sqrt{\eta}, \sqrt{1-\eta}) & f_2 &= (\sqrt{\eta}, -\sqrt{1-\eta}) \\ g_1 &= (\sqrt{1-\eta}, \sqrt{\eta}) & g_2 &= (\sqrt{1-\eta}, -\sqrt{\eta}). \end{aligned}$$

Applying Theorem 35, we have $r_0 = 2$ whenever $\nu_2^2 > \nu_1$, which simplifies to

$$s > \frac{4(1+\theta)}{(1-\theta)^2},$$

which settles the first part of Theorem 32. Now, simplifying (5.40) whenever $i = j$, we have

$$a_{i,i} = \sqrt{\frac{\langle p, f^{i,i} \rangle}{\langle p, (I - \nu_i^{-2}M)^{-1} f^{i,i} \rangle}},$$

and since $p = (1/2, 1/2)$ and the f_i have unit length, this simplifies further to

$$a_{i,i} = \frac{1}{\sqrt{\langle \mathbf{1}, (I - \nu_i^{-2}M)^{-1} f^{i,i} \rangle}}. \quad (5.47)$$

The standard adjoint formula yields, whenever $\alpha < 1/\nu_1$,

$$(I - \alpha M)^{-1} = \frac{1}{4(1 - \alpha\nu_1)(1 - \alpha\nu_2)} \begin{pmatrix} 4 - \alpha s & 2\alpha s \eta \\ 2\alpha s(1 - \eta) & 4 - \alpha s, \end{pmatrix}$$

and since $f^{1,1} = f^{2,2} = (\eta, 1 - \eta)$, we have

$$\begin{aligned} \langle \mathbf{1}, (I - \alpha M)^{-1} f^{i,i} \rangle^{-1} &= \frac{4(1 - \alpha\nu_1)(1 - \alpha\nu_2)}{4 - \alpha s + \alpha s \theta^2} \\ &= \frac{4 - 2\alpha s + \alpha^2 s^2 \frac{1-\theta^2}{4}}{4 - \alpha s + \alpha s \theta^2} \\ &=: \gamma(\alpha s). \end{aligned}$$

Since we will choose $\alpha = \nu_i^{-2}$, we have $\alpha s > 0$ and

$$\begin{aligned} \gamma(x) &= \left(1 - \frac{x}{2}\right) \left(1 + \frac{(1 - \theta^2)x}{4}\right) + O(x^2) \\ &= 1 - \frac{1 + \theta^2}{4}x + O(x^2). \end{aligned}$$

Substituting $\alpha = \nu_i^{-2}$ and taking the square root, we find

$$a_{i,i} = 1 - \frac{1 + \theta^2}{8} \frac{s}{\nu_i^2} + O\left(\frac{1}{s^2}\right) = 1 - \frac{2}{s} \frac{1 + \theta^2}{(1 \pm \theta)^2} + O\left(\frac{1}{s^2}\right),$$

which are the expressions in Theorem 32.

Note that it is possible to continue the computations and find explicit expressions for the $a_{i,i}$ in terms of η and s , but the resulting expressions are too complex to give any more insight than the asymptotic expressions.

5.10.6 Computations when there are two blocks

We defined $\eta(s)$ as the unique number in $[0, 1]$ such that

$$s > \frac{4(1 + \theta)}{(1 - \theta)^2} \iff \eta > \eta(s),$$

with $\theta(\eta) = 2\sqrt{\eta(1 - \eta)}$ (see also Figure 5.6). The inverse $\theta^{-1} : [0, 1] \rightarrow [1/2, 1]$ is given by

$$\theta^{-1}(t) = \frac{1 + \sqrt{1 - t^2}}{2}.$$

Now, the solution x of the equation

$$s = \frac{4(1 + x)}{(1 - x)^2}$$

is the solution of the quadratic $sx^2 - (2s + 4)x + s - 4 = 0$. The discriminant is $\delta(s) = 16(2s + 1)$ and the unique solution in $[0, 1]$ is

$$x(s) = 1 + \frac{2 - 2\sqrt{2s + 1}}{s}.$$

Finally, the smallest $\eta(s)$ for which (5.45) is satisfied is $\eta(s) = \theta^{-1}(x(s))$, that is,

$$\eta(s) = \frac{1 + \sqrt{1 - x(s)^2}}{2} = \frac{1 + \sqrt{1 - \left(1 + \frac{2 - 2\sqrt{2s + 1}}{s}\right)^2}}{2}. \quad (5.48)$$

It is possible to expand the term inside the square root, but with no meaningful gain. The function $\eta(s)$ has the series expansion

$$\eta(s) = \frac{1}{2} + \sqrt[4]{\frac{2}{s}} + O(s^{-3/4}),$$

but the convergence is very slow: the truncated RHS is less than one only whenever $s \geq 32$.

5.11 Convergence of eigenvectors

The goal of this section is to prove Theorem 33. We place ourselves in the stochastic block model setting as in Section 5.3, with $\sigma_g = \sigma_d = \sigma$ and $q = p$. We assume that F and p are constant with n , so that the modularity matrix M doesn't depend on n .

5.11.1 Convergence of $\tilde{U}_i(x)$

We consider the multitype Galton-Watson trees (\mathcal{T}_j, o) as defined in [29]: the root o has *type* $\sigma(o) = j$, and afterwards, each vertex with type i has $\text{Poi}(M_{i,k})$ children of type k . Unlike our initial trees T_x , which are heavily n -dependent with node labels in $[n]$ and edge probabilities $P_{x,y}$, the tree \mathcal{T}_j does not depend on n , and its labels are in $[r]$. We define on those trees the random processes

$$\mathcal{U}_i(j, t) = \nu_i^{-t} \sum_{k=1}^r N_k(\mathcal{T}_j, t) f_i(k),$$

where $N_k(\mathcal{T}_j, t)$ counts the number of vertices of type k at depth t in \mathcal{T}_j . Assume that we have chosen the f_i such that $\phi_i = \sum_d f_i$ has unit norm; we recall that the processes $\tilde{U}_i(x, t)$ were defined in (5.27)-(5.28).

Lemma 53. *Let $x \in [n]$, and define $j = \sigma(x)$. Then the processes $\tilde{U}_i(x, t)$ and $\mathcal{U}_i(j, t)$ have the same distribution.*

Proof. Let $\sigma(T_x)$ be the tree where a vertex with label y is mapped to a vertex with label $\sigma(y)$. Then the root of $\sigma(T_x)$ has label $\sigma(x) = j$. Take a vertex in T_x with label y , and let $\sigma(y) = i$; the number of children of y with type k has distribution $\text{Poi}(\tilde{M}_{y,k})$, with

$$\tilde{M}_{y,k} = d_y \sum_{\sigma(z)=k} \frac{P_{yz}}{d_y}.$$

Using the definition of P for the SBM, we have

$$\tilde{M}_{y,k} = \sum_{\sigma(z)=k} \frac{F_{\sigma(y),\sigma(z)}}{n} = F_{i,k} p_k = M_{i,k}.$$

Therefore, the laws of $\sigma(T_x)$ and \mathcal{T}_j coincide. Now, since there are no weights the product in (5.27) is equal to 1 and we have

$$\tilde{U}_i(x, t) = \sum_{x_t} \varphi_i(\iota(x_t));$$

replacing φ by its definition in terms of f_i ,

$$\tilde{U}_i(x, t) = \nu_i^{-t} \sum_{x_t} f_i(\sigma(\iota(x_t))) = \nu_i^{-t} \sum_{k=1}^r N_k(\sigma(T_x), t) f_i(k),$$

which ends the proof. \square

This lemma allows us to translate results back and forth between the T_x and the $\mathcal{T}_{\sigma(x)}$; we thus know from the expectation/correlation computations in Subsections 5.8.5-5.8.7 that $\mathcal{U}_i(j, t)$ (with $j = \sigma(x)$) is a martingale with

$$\mathbf{E}[\mathcal{U}_i(j, t)] = f_i(j) \quad \text{and} \quad \mathbf{E}[\mathcal{U}_i(j, t)^2] \leq \Gamma_{\text{right}}(i, i).$$

By the Doob martingale convergence theorem, this implies that $\mathcal{U}_i(j, t)$ converges in L^2 as $t \rightarrow \infty$ towards a random variable $\mathcal{Z}_{i,j}$. Since we have convergence in L^2 , it entails

$$\begin{aligned} \mathbf{E}[\mathcal{Z}_{i,j}] &= \lim_{t \rightarrow +\infty} \mathbf{E}[\mathcal{U}_i(j, t)] = f_i(j) \\ \mathbf{E}[\mathcal{Z}_{i,j}^2] &= \lim_{t \rightarrow +\infty} \mathbf{E}[\mathcal{U}_i(j, t)^2] = [(I - \nu_i^{-2} M)^{-1} f^{i,i}](j). \end{aligned}$$

In the last line, we used computations from Appendix 5.9 to determine the limit.

Another important fact is that the law of $\mathcal{U}_i(j, t)$ does not depend on n whatsoever. This implies that $\mathcal{U}_i(j, \ell)$ converges to $\mathcal{Z}_{i,j}$ as $n \rightarrow +\infty$, which in turn yields the following proposition.

Proposition 27. *Let $\mathcal{Z}_{i,j}$ be the limit of the random process $\mathcal{U}_i(j, t)$. Then,*

$$\tilde{U}_i(x) \xrightarrow[n \rightarrow \infty]{L^2} \mathcal{Z}_{i,\sigma(x)}. \quad (5.49)$$

5.11.2 Convergence of the pseudo-eigenvectors

Now that we showed convergence on the random tree, we shall use the concentration proposition 24 to translate it on the graph. Let $h : \mathbb{R} \rightarrow \mathbb{R}$ be a bounded continuous function, and define

$$f(G, x) = \frac{1}{p_j n} \mathbf{1}_{\sigma(x)=j} h(U_i(x)).$$

The function f satisfies the hypotheses of Proposition 24 since h is bounded, and we have

$$f(T_x, x) = \frac{1}{p_j n} \mathbf{1}_{\sigma(x)=j} h(\tilde{U}_i(x))$$

Using Proposition 27, and the fact that L^2 convergence implies convergence in distribution,

$$\mathbf{E}[\mathbf{1}_{\sigma(x)=j} h(\tilde{U}_i(x))] \rightarrow \mathbf{1}_{\sigma(x)=j} \mathbf{E}[h(\mathcal{Z}_{i,j})]$$

uniformly in x , so that summing over all vertices

$$\lim_{n \rightarrow +\infty} \frac{1}{p_j n} \sum_{\sigma(x)=j} h(U_i(x)) = \mathbf{E}[h(\mathcal{Z}_{i,j})]. \quad (5.50)$$

The above equation implies immediately that the discrete distribution of the $U_i(x)$ with $\sigma(x) = j$ converges weakly to $\mathcal{Z}_{i,j}$; in other words,

$$\frac{1}{p_j n} \sum_{\sigma(x)=j} \delta_{U_i(x)} \xrightarrow[n \rightarrow \infty]{d} \mathcal{Z}_{i,j}. \quad (5.51)$$

5.11.3 Convergence of the eigenvector

Define the normalized pseudo eigenvectors:

$$\bar{U}_i = \frac{\sqrt{n} U_i}{|U_i|},$$

then thanks to (5.19) we have $|U_i| - \sqrt{n}\gamma \ll 1$, where

$$\gamma = \sqrt{\langle p, (\mathbf{I} - \nu_i^{-2} M)^{-1} f^{i,i} \rangle}$$

from the computations in Appendix 5.9. A consequence of our notation \ll is that $\sqrt{n}/|U_i|$ converges in probability towards $\gamma > 0$. Thanks to (5.51) and Slutsky's lemma,

$$\frac{1}{p_j n} \sum_{\sigma(x)=j} \delta_{\bar{U}_i(x)} \xrightarrow[n \rightarrow \infty]{d} \mathcal{Z}_{i,j} := \frac{\mathcal{Z}_{i,j}}{\gamma},$$

and consequently $\mathcal{Z}_{i,j}$ has mean and variance

$$\mu_{i,j} = \frac{f_i(j)}{\sqrt{\langle p, (\mathbf{I} - \nu_i^{-2} M)^{-1} f^{i,i} \rangle}} \quad \text{and} \quad \sigma_{i,j}^2 = \frac{[(\mathbf{I} - \nu_i^{-2} M)^{-1} f^{i,i}](j)}{\langle p, (\mathbf{I} - \nu_i^{-2} M)^{-1} f^{i,i} \rangle}. \quad (5.52)$$

Now comes the last step of our proof; let \bar{u}_i be a right eigenvector of A such that $|\bar{u}_i| = \sqrt{n}$ (one can take $\bar{u}_i = \sqrt{n} u_i$). Then, (5.21) implies that

$$|\bar{u}_i - \bar{U}_i| \ll \sqrt{n}. \quad (5.53)$$

We want to show convergence for \bar{u}_i ; using the Portmanteau lemma, it suffices to show bounds like (5.50) for Lipschitz functions. Let h be a bounded function with Lipschitz constant C ; we write

$$\begin{aligned} \left| \frac{1}{p_j n} \sum_{\sigma(x)=i} [h(\bar{u}_i(x)) - h(\bar{U}_i(x))] \right| &\leq \frac{1}{p_j n} \sum_{\sigma(x)=i} |h(\bar{u}_i(x)) - h(\bar{U}_i(x))| \\ &\leq \frac{C}{p_j n} \sum_{\sigma(x)=i} |\bar{u}_i(x) - \bar{U}_i(x)| \\ &\leq \frac{C}{p_j n} \sqrt{p_j n} |\bar{u}_i - \bar{U}_i|, \end{aligned}$$

using the Cauchy-Schwarz inequality. Using (5.53), we are done.

Bibliography

- [1] Emmanuel Abbe. Community Detection and Stochastic Block Models: Recent Developments. *Journal of Machine Learning Research*, 18(177):1–86, 2018.
- [2] Emmanuel Abbe, Afonso S. Bandeira, Annina Bracher, and Amit Singer. Decoding Binary Node Labels from Censored Edge Measurements: Phase Transition and Efficient Recovery. *IEEE Transactions on Network Science and Engineering*, 1(1):10–22, January 2014.
- [3] Emmanuel Abbe, Enric Boix-Adserà, Peter Ralli, and Colin Sandon. Graph Powering and Spectral Robustness. *SIAM Journal on Mathematics of Data Science*, 2(1):132–157, January 2020.
- [4] Emmanuel Abbe, Jianqing Fan, Kaizheng Wang, Yiqiao Zhong, et al. Entrywise eigenvector analysis of random matrices with low expected rank. *Annals of Statistics*, 48(3):1452–1474, 2020.
- [5] Emmanuel Abbe and Colin Sandon. Community Detection in General Stochastic Block Models: Fundamental Limits and Efficient Algorithms for Recovery. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 670–688, October 2015.
- [6] Emmanuel Abbe and Colin Sandon. Achieving the KS threshold in the general stochastic block model with linearized acyclic belief propagation. *Advances in Neural Information Processing Systems*, 29, 2016.
- [7] Emmanuel Abbe and Colin Sandon. Proof of the Achievability Conjectures for the General Stochastic Block Model. *Communications on Pure and Applied Mathematics*, 71(7):1334–1406, 2018.
- [8] Lada A. Adamic and Natalie Glance. The political blogosphere and the 2004 u.s. election: Divided they blog. In *Proceedings of the 3rd International Workshop on Link Discovery, LinkKDD '05*, page 36–43, New York, NY, USA, 2005. Association for Computing Machinery.
- [9] Edoardo M. Airoldi, David M. Blei, Stephen E. Fienberg, and Eric P. Xing. Mixed Membership Stochastic Blockmodels. *Journal of machine learning research : JMLR*, 9:1981–2014, September 2008.
- [10] Noga Alon, Michael Krivelevich, and Benny Sudakov. Finding a large hidden clique in a random graph. *Random Structures & Algorithms*, 13(3-4):457–466, 1998.
- [11] Johannes Alt, Raphaël Ducatez, and Antti Knowles. Extremal eigenvalues of critical Erdős–Rényi graphs. May 2020.
- [12] Vivek Bagaria, Jian Ding, David Tse, Yihong Wu, and Jiaming Xu. Hidden Hamiltonian Cycle Recovery via Linear Programming. *Operations Research*, 68(1):53–70, January 2020.

- [13] Jinho Baik, Gérard Ben Arous, and Sandrine Péché. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *Annals of Probability*, 33(5):1643–1697, September 2005.
- [14] J. Banks, C. Moore, R. Vershynin, N. Verzelen, and J. Xu. Information-Theoretic Bounds and Phase Transitions in Clustering, Sparse PCA, and Submatrix Localization. *IEEE Transactions on Information Theory*, 64(7):4872–4894, July 2018.
- [15] B. Barak, S. B. Hopkins, J. Kelner, P. Kothari, A. Moitra, and A. Potechin. A Nearly Tight Sum-of-Squares Lower Bound for the Planted Clique Problem. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 428–437, October 2016. ISSN: 0272-5428.
- [16] A. D. Barbour and Louis Hsiao Yun Chen. *An Introduction to Stein’s Method*. World Scientific, 2005. Google-Books-ID: Z8BoDQAAQBAJ.
- [17] Hyman Bass. The Ihara-Selberg Zeta Function of a Tree Lattice. *International Journal of Mathematics*, 03(06):717–797, December 1992.
- [18] F. L. Bauer and C. T. Fike. Norms and exclusion theorems. *Numerische Mathematik*, 2(1):137–141, December 1960.
- [19] Florent Benaych-Georges, Charles Bordenave, and Antti Knowles. Largest eigenvalues of sparse inhomogeneous Erdős–Rényi graphs. *The Annals of Probability*, 47(3):1653–1676, May 2019.
- [20] Florent Benaych-Georges, Charles Bordenave, and Antti Knowles. Spectral Radii of Sparse Random Matrices. *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, 56(3):2141–2161, August 2020.
- [21] Florent Benaych-Georges and Raj Rao Nadakuditi. The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices. *Advances in Mathematics*, 227(1):494–521, May 2011.
- [22] Florent Benaych-Georges and Raj Rao Nadakuditi. The singular values and vectors of low rank perturbations of large rectangular random matrices. *Journal of Multivariate Analysis*, 111:120–135, October 2012.
- [23] Itai Benjamini and Oded Schramm. Recurrence of Distributional Limits of Finite Planar Graphs. *Selected Works in Probability and Statistics*, pages 533–545. Springer, New York, NY, 2011.
- [24] Quentin Berthet and Philippe Rigollet. Optimal detection of sparse principal components in high dimension. *The Annals of Statistics*, 41(4):1780–1815, August 2013.
- [25] Rajendra Bhatia. *Matrix Analysis*. Graduate Texts in Mathematics. Springer-Verlag, New York, 1997.
- [26] Béla Bollobás. Threshold functions for small subgraphs. *Mathematical Proceedings of the Cambridge Philosophical Society*, 90(2):197–206, September 1981.
- [27] Béla Bollobás. *Random Graphs*. Cambridge University Press, August 2001. Google-Books-ID: o9WecWgilzYC.

- [28] Charles Bordenave, Simon Coste, and Raj Rao Nadakuditi. Detection thresholds in very sparse matrix completion. August 2020.
- [29] Charles Bordenave, Marc Lelarge, and Laurent Massoulié. Non-backtracking Spectrum of Random Graphs: Community Detection and Non-regular Ramanujan Graphs. *The Annals of Probability*, 46(1):1–71, January 2018.
- [30] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, Oxford, New York, February 2013.
- [31] Matthew Brennan, Guy Bresler, and Wasim Huleihel. Reducibility and Computational Lower Bounds for Problems with Planted Sparse Structure. In *Conference On Learning Theory*, pages 48–166. PMLR, July 2018.
- [32] Andrei Z. Broder, Alan M. Frieze, and Eli Shamir. Finding hidden hamiltonian cycles. In *Proceedings of the Twenty-Third Annual ACM Symposium on Theory of Computing*, STOC '91, page 182–189, New York, NY, USA, 1991. Association for Computing Machinery.
- [33] Emmanuel J. Candes and Yaniv Plan. Matrix Completion With Noise. *Proceedings of the IEEE*, 98(6):925–936, June 2010.
- [34] Mireille Capitaine, Catherine Donati-Martin, and Delphine Féral. The largest eigenvalues of finite rank deformation of large Wigner matrices: Convergence and nonuniversality of the fluctuations. *The Annals of Probability*, 37(1):1–47, January 2009.
- [35] Mo Chen, Qiong Yang, and Xiaoou Tang. Directed graph embedding. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, IJCAI'07, page 2707–2712, San Francisco, CA, USA, 2007. Morgan Kaufmann Publishers Inc.
- [36] Yuxin Chen, Chen Cheng, and Jianqing Fan. Asymmetry helps: Eigenvalue and eigenvector analyses of asymmetrically perturbed low-rank matrices. *The Annals of Statistics*, 49(1):435–458, February 2021.
- [37] Fan Chung. Laplacians and the Cheeger Inequality for Directed Graphs. *Annals of Combinatorics*, 9(1):1–19, April 2005.
- [38] Aaron Clauset, Cosma Rohilla Shalizi, and M. E. J. Newman. Power-Law Distributions in Empirical Data. *SIAM Review*, 51(4):661–703, November 2009.
- [39] Amin Coja-Oghlan. Graph Partitioning via Adaptive Spectral Techniques. *Combinatorics, Probability and Computing*, 19(2):227–284, March 2010.
- [40] Thomas H. Cormen, Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction To Algorithms*. MIT Press, 2001. Google-Books-ID: NLngYy-WF1_YC.
- [41] Simon Coste and Justin Salez. Emergence of extended states at zero in the spectrum of sparse random graphs. *arXiv:1809.07587 [math]*, September 2018. arXiv: 1809.07587.
- [42] Simon Coste and Ludovic Stephan. A simpler spectral approach for clustering in directed networks. *arXiv:2102.03188 [cs, math]*, February 2021. arXiv: 2102.03188.
- [43] Simon Coste and Yizhe Zhu. Eigenvalues of the non-backtracking operator detached from the bulk. *Random Matrices: Theory and Applications*, page 2150028, August 2020.

- [44] Mihai Cucuringu, Huan Li, He Sun, and Luca Zanetti. Hermitian matrices for clustering directed graphs: insights and applications. In Silvia Chiappa and Roberto Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 983–992. PMLR, 26–28 Aug 2020.
- [45] Lorenzo Dall' Amico, Romain Couillet, and Nicolas Tremblay. Revisiting the bethe-hessian: Improved community detection in sparse heterogeneous graphs. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [46] A. Dasgupta, J. E. Hopcroft, and F. McSherry. Spectral analysis of random graphs with skewed degree distributions. In *45th Annual IEEE Symposium on Foundations of Computer Science*, pages 602–610, October 2004. ISSN: 0272-5428.
- [47] Aurelien Decelle, Florent Krzakala, Cristopher Moore, and Lenka Zdeborová. Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Physical Review E*, 84(6):066106, December 2011.
- [48] Yael Dekel, Ori Gurel-Gurevich, and Yuval Peres. Finding Hidden Cliques in Linear Time with High Probability. *Proceedings*, pages 67–75. Society for Industrial and Applied Mathematics, January 2011.
- [49] Yash Deshpande and Andrea Montanari. Finding Hidden Cliques of Size $\sqrt{N/e}$ in Nearly Linear Time. *Foundations of Computational Mathematics*, 15(4):1069–1128, August 2015.
- [50] Zhuanlian Ding, Xingyi Zhang, Dengdi Sun, and Bin Luo. Overlapping community detection based on network decomposition. *Scientific reports*, 6:24115, 2016.
- [51] Ioana Dumitriu and Yizhe Zhu. Sparse general Wigner-type matrices: Local law and eigenvector delocalization. *Journal of Mathematical Physics*, 60(2):023301, February 2019.
- [52] Jack Edmonds. Paths, Trees, and Flowers. *Canadian Journal of Mathematics*, 17:449–467, 1965.
- [53] László Erdős, Antti Knowles, Horng-Tzer Yau, and Jun Yin. The local semicircle law for a general class of random matrices. *Electronic Journal of Probability*, 18, 2013.
- [54] William Evans, Claire Kenyon, Yuval Peres, and Leonard J. Schulman. Broadcasting on trees and the Ising model. *The Annals of Applied Probability*, 10(2):410–433, May 2000.
- [55] Uriel Feige and Eran Ofek. Spectral techniques applied to sparse random graphs. *Random Structures & Algorithms*, 27(2):251–275, 2005.
- [56] Delphine Féral and Sandrine Péché. The Largest Eigenvalue of Rank One Deformation of Large Wigner Matrices. *Communications in Mathematical Physics*, 272(1):185–228, May 2007.
- [57] Santo Fortunato. Community detection in graphs. *Physics Reports*, 486(3):75–174, February 2010.
- [58] Z. Füredi and J. Komlós. The eigenvalues of random symmetric matrices. *Combinatorica*, 1(3):233–241, September 1981.

- [59] Sainyam Galhotra, Arya Mazumdar, Soumyabrata Pal, and Barna Saha. The Geometric Block Model. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1), April 2018.
- [60] Chao Gao, Zongming Ma, Anderson Y Zhang, and Harrison H Zhou. Achieving optimal misclassification proportion in stochastic block models. *The Journal of Machine Learning Research*, 18(1):1980–2024, 2017.
- [61] Michael R. Garey and David S. Johnson. *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., USA, 1990.
- [62] Lennart Gulikers, Marc Lelarge, and Laurent Massoulié. Non-Backtracking Spectrum of Degree-Corrected Stochastic Block Models. In Christos H. Papadimitriou, editor, *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*, volume 67 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 44:1–44:27, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [63] Lennart Gulikers, Marc Lelarge, and Laurent Massoulié. A spectral method for community detection in moderately sparse degree-corrected stochastic block models. *Advances in Applied Probability*, 49(3):686–721, September 2017.
- [64] Olivier Guédon and Roman Vershynin. Community detection in sparse networks via Grothendieck’s inequality. *Probability Theory and Related Fields*, 165(3):1025–1049, August 2016.
- [65] Bruce Hajek, Yihong Wu, and Jiaming Xu. Semidefinite Programs for Exact Recovery of a Hidden Community. In *Conference on Learning Theory*, pages 1051–1095. PMLR, June 2016.
- [66] Ki-ichiro Hashimoto. Zeta Functions of Finite Graphs and Representations of p -Adic Groups. *Automorphic Forms and Geometry of Arithmetic Varieties*, pages 211–280, January 1989.
- [67] Simon Heimlicher, Marc Lelarge, and Laurent Massoulié. Community Detection in the Labelled Stochastic Block Model. September 2012.
- [68] Remco van der Hofstad. *Random Graphs and Complex Networks*, volume 1 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 2016.
- [69] Paul W. Holland, Kathryn Blackmond Laskey, and Samuel Leinhardt. Stochastic block-models: First steps. *Social Networks*, 5(2):109–137, June 1983.
- [70] Svante Janson, Tomasz Łuczak, and Andrzej Ruciński. *Random Graphs*. John Wiley & Sons, September 2011. Google-Books-ID: RjnqVoB4VmUC.
- [71] Adel Javanmard, Andrea Montanari, and Federico Ricci-Tersenghi. Phase transitions in semidefinite relaxations. *Proceedings of the National Academy of Sciences*, 113(16):E2218–E2223, April 2016.
- [72] Mark Jerrum. Large Cliques Elude the Metropolis Process. *Random Structures & Algorithms*, 3(4):347–359, 1992.
- [73] I. M. Johnstone and D. Paul. PCA in High Dimensions: An Orientation. *Proceedings of the IEEE*, 106(8):1277–1292, August 2018.

- [74] Ravi Kannan, Santosh Vempala, and Adrian Vetta. On clusterings: Good, bad and spectral. *J. ACM*, 51(3):497–515, May 2004.
- [75] Richard M. Karp. Reducibility among Combinatorial Problems. The IBM Research Symposia Series, pages 85–103. Springer US, Boston, MA, 1972.
- [76] Brian Karrer and M. E. J. Newman. Stochastic blockmodels and community structure in networks. *Physical Review E*, 83(1):016107, January 2011.
- [77] Tosio Kato. *Perturbation Theory for Linear Operators*. Classics in Mathematics. Springer-Verlag, Berlin Heidelberg, 2 edition, 1995.
- [78] Raghunandan Keshavan, Andrea Montanari, and Sewoong Oh. Matrix Completion from Noisy Entries. In Y. Bengio, D. Schuurmans, J. D. Lafferty, C. K. I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems 22*, pages 952–960. Curran Associates, Inc., 2009.
- [79] H. Kesten and B. P. Stigum. A Limit Theorem for Multidimensional Galton-Watson Processes. *The Annals of Mathematical Statistics*, 37(5):1211–1223, October 1966.
- [80] Florent Krzakala, Cristopher Moore, Elchanan Mossel, Joe Neeman, Allan Sly, Lenka Zdeborová, and Pan Zhang. Spectral redemption in clustering sparse networks. *Proceedings of the National Academy of Sciences*, 110(52):20935–20940, December 2013.
- [81] Steinar Laenen and He Sun. Higher-order spectral clustering of directed graphs. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 941–951. Curran Associates, Inc., 2020.
- [82] Jing Lei and Alessandro Rinaldo. Consistency of spectral clustering in stochastic block models. *The Annals of Statistics*, 43(1):215–237, February 2015.
- [83] E. A. Leicht and M. E. J. Newman. Community Structure in Directed Networks. *Physical Review Letters*, 100(11):118703, March 2008.
- [84] Marc Lelarge, Laurent Massoulié, and Jiaming Xu. Reconstruction in the Labelled Stochastic Block Model. *IEEE Transactions on Network Science and Engineering*, 2(4):152–163, October 2015.
- [85] Jure Leskovec and Andrej Krevl. SNAP Datasets: Stanford large network dataset collection. <http://snap.stanford.edu/data>, June 2014.
- [86] Y. Li, X. Wu, and A. Lu. Analysis of Spectral Space Properties of Directed Graphs Using Matrix Perturbation Theory with Application in Graph Partition. In *2015 IEEE International Conference on Data Mining*, pages 847–852, November 2015. ISSN: 1550-4786.
- [87] Alexander Lubotzky. Cayley graphs: eigenvalues, expanders and random walks. London Mathematical Society Lecture Note Series, pages 155–190. Cambridge University Press, Cambridge, 1995.
- [88] Konstantin Makarychev, Yury Makarychev, and Aravindan Vijayaraghavan. Learning Communities in the Presence of Errors. In *Conference on Learning Theory*, pages 1258–1291. PMLR, June 2016.

- [89] Fragkiskos D. Malliaros and Michalis Vazirgiannis. Clustering and community detection in directed networks: A survey. *Physics Reports*, 533(4):95–142, December 2013.
- [90] Mahendra Mariadassou, Stéphane Robin, Corinne Vacher, et al. Uncovering latent structure in valued graphs: a variational approach. *The Annals of Applied Statistics*, 4(2):715–742, 2010.
- [91] Laurent Massoulié. Community Detection Thresholds and the Weak Ramanujan Property. In *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing*, STOC '14, pages 694–703, New York, NY, USA, 2014. ACM.
- [92] Laurent Massoulié, Ludovic Stephan, and Don Towsley. Planting trees in graphs, and finding them back. In *Conference on Learning Theory*, pages 2341–2371. PMLR, June 2019.
- [93] David W. Matula. The employee party problem. In *Notices of the American Mathematical Society*, volume 19, pages A382–A382. AMER MATHEMATICAL SOC 201 CHARLES ST, PROVIDENCE, RI 02940-2213, 1972.
- [94] G. McLachlan and K. Basford. *Mixture Models: Inference and Applications to Clustering*, volume 38. January 1988.
- [95] Marc Mezard, Giorgio Parisi, and Miguel Angel Virasoro. *Spin Glass Theory And Beyond: An Introduction To The Replica Method And Its Applications*. World Scientific Publishing Company, November 1987. Google-Books-ID: DwY8DQAAQBAJ.
- [96] Ankur Moitra, William Perry, and Alexander S. Wein. How robust are reconstruction thresholds for community detection? In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '16, page 828–841, New York, NY, USA, 2016. Association for Computing Machinery.
- [97] Andrea Montanari and Subhabrata Sen. Semidefinite programs on sparse random graphs and their application to community detection. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '16, pages 814–827, Cambridge, MA, USA, June 2016. Association for Computing Machinery.
- [98] Cristopher Moore, Santa Fe Institute, and The Computational Complexity Column by Vikraman Arvind. The Computer Science and Physics of Community Detection: Landscapes, Phase Transitions, and Hardness. *Bulletin of EATCS*, 1(121), February 2017.
- [99] Cristopher Moore and Mertens. *The Nature of Computation*. Oxford University Press, Oxford, New York, August 2011.
- [100] Elchanan Mossel, Joe Neeman, and Allan Sly. Reconstruction and estimation in the planted partition model. *Probability Theory and Related Fields*, 162(3):431–461, August 2015.
- [101] Elchanan Mossel, Joe Neeman, and Allan Sly. Belief propagation, robust reconstruction and optimal recovery of block models. *The Annals of Applied Probability*, 26(4):2211–2256, August 2016.
- [102] Elchanan Mossel, Joe Neeman, and Allan Sly. A Proof of the Block Model Threshold Conjecture. *Combinatorica*, 38(3):665–708, June 2018.
- [103] Elchanan Mossel and Jiaming Xu. Local Algorithms for Block Models with Side Information. August 2015.

- [104] Silvia Noschese, Lionello Pasquini, and Lothar Reichel. Tridiagonal Toeplitz matrices: Properties and novel applications. *Numerical Linear Algebra with Applications*, 20, March 2013.
- [105] Richard Otter. The Number of Trees. *Annals of Mathematics*, 49(3):583–599, 1948.
- [106] Soumik Pal and Yizhe Zhu. Community Detection in the Sparse Hypergraph Stochastic Block Model. July 2020.
- [107] S. Péché. The largest eigenvalue of small rank perturbations of Hermitian random matrices. *Probability Theory and Related Fields*, 134(1):127–173, January 2006.
- [108] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [109] William Pentney and Marina Meila. Spectral Clustering of Biological Sequence Data. volume 20, pages 845–850, January 2005.
- [110] Boris Pittel, Joel Spencer, and Nicholas Wormald. Sudden emergence of a giant core in a random graph. *Journal of Combinatorial Theory, Series B*, 67(1):111–151, 1996.
- [111] Karl Rohe, Sourav Chatterjee, Bin Yu, et al. Spectral clustering and the high-dimensional stochastic blockmodel. *The Annals of Statistics*, 39(4):1878–1915, 2011.
- [112] Alaa Saade, Florent Krzakala, and Lenka Zdeborová. Spectral Clustering of graphs with the Bethe Hessian. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 406–414. Curran Associates, Inc., 2014.
- [113] Alaa Saade, Marc Lelarge, Florent Krzakala, and Lenka Zdeborová. Spectral detection in the censored block model. In *2015 IEEE International Symposium on Information Theory (ISIT)*, pages 1184–1188, June 2015.
- [114] Venu Satuluri and Srinivasan Parthasarathy. Symmetrizations for clustering directed graphs. In *Proceedings of the 14th International Conference on Extending Database Technology, EDBT/ICDT '11*, page 343–354, New York, NY, USA, 2011. Association for Computing Machinery.
- [115] Guilhem Semerjian, Gabriele Sicuro, and Lenka Zdeborová. Recovery thresholds in the sparse planted matching problem. *Physical Review. E*, 102(2-1):022304, August 2020.
- [116] Ludovic Stephan and Laurent Massoulié. Non-backtracking spectra of weighted inhomogeneous random graphs. April 2020.
- [117] Ludovic Stephan and Laurent Massoulié. Robustness of Spectral Methods for Community Detection. In *Conference on Learning Theory*, pages 2831–2860, Phoenix, USA, June 2019. PMLR.
- [118] Daniel Sussman, Minh Tang, Donniell Fishkind, and Carey Priebe. A Consistent Adjacency Spectral Embedding for Stochastic Blockmodel Graphs. *Journal of the American Statistical Association*, 107, August 2011.
- [119] Hadrien Van Lierde. *Spectral clustering algorithms for directed graphs*. PhD thesis, 2015.

- [120] Hadrien Van Lierde, Tommy WS Chow, and Jean-Charles Delvenne. Spectral clustering algorithms for the detection of clusters in block-cyclic and block-acyclic graphs. *Journal of Complex Networks*, 7(1):1–53, 2019.
- [121] Nicolas Verzelen and Ery Arias-Castro. Community detection in sparse random networks. *The Annals of Applied Probability*, 25(6):3465–3510, December 2015.
- [122] Ulrike von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, December 2007.
- [123] Yuchung J. Wang and George Y. Wong. Stochastic Blockmodels for Directed Graphs. *Journal of the American Statistical Association*, 82(397):8–19, 1987.
- [124] Yusuke Watanabe and Kenji Fukumizu. Graph Zeta Function in the Bethe Free Energy and Loopy Belief Propagation. In Y. Bengio, D. Schuurmans, J. D. Lafferty, C. K. I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems 22*, pages 2017–2025. Curran Associates, Inc., 2009.
- [125] Hermann Weyl. Das asymptotische verteilungsgesetz der eigenwerte linearer partieller differentialgleichungen (mit einer anwendung auf die theorie der hohlraumstrahlung). *Mathematische Annalen*, 71(4):441–479, December 1912.
- [126] Y. Yu, T. Wang, and R. J. Samworth. A useful variant of the Davis–Kahan theorem for statisticians. *Biometrika*, 102(2):315–323, June 2015.
- [127] Pan Zhang. Robust spectral detection of global structures in the data by learning a regularization. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- [128] Zhixin Zhou and Arash A. Amini. Analysis of spectral clustering algorithms for community detection: the general bipartite setting. *Journal of Machine Learning Research*, 20(47):1–47, 2019.
- [129] Zhixin Zhou and Arash A. Amini. Optimal Bipartite Network Clustering. *Journal of Machine Learning Research*, 21(40):1–68, 2020.