THEORY OF COMPUTING, Volume 13 (8), 2017, pp. 1–22
www.theoryofcomputing.org

SPECIAL ISSUE: APPROX-RANDOM 2015

The Minimum Bisection in the Planted Bisection Model

Amin Coja-Oghlan* Oliver Cooley† Mihyun Kang‡
Kathrin Skubch

Received December 19, 2015; Revised August 31, 2016; Published September 23, 2017

Abstract: In the planted bisection model a random graph $G(n, p_+, p_-)$ with n vertices is created by partitioning the vertices randomly into two classes of equal size (up to ± 1). Any two vertices that belong to the same class are linked by an edge with probability p_+ and any two that belong to different classes with probability $p_- < p_+$ independently. The planted bisection model has been used extensively to benchmark graph partitioning algorithms. If $p_\pm = 2d_\pm/n$ for numbers $0 \le d_- < d_+$ that remain fixed as $n \to \infty$, then w. h. p. the "planted" bisection (the one used to construct the graph) will not be a minimum bisection. In this paper we derive an asymptotic formula for the minimum bisection width under the assumption that $d_+ - d_- > c\sqrt{d_+ \ln d_+}$ for a certain constant $c > 0$.

ACM Classification: G.2.2

AMS Classification: 05C80, 05C15

Key words and phrases: random graphs, minimum bisection, planted bisection, belief propagation

A preliminary version of this paper appeared in the Proceedings of the 19th International Workshop on Randomization and Computation (RANDOM 2015) [11].

*The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 278857–PTCC.

†Supported by Austrian Science Fund (FWF): P26826 and W1230.

‡Supported by Austrian Science Fund (FWF): P26826 and W1230.

 DOI: 10.4086/toc.2017.v013a008

1 Introduction

1.1 Background and motivation

The minimum bisection problem is a well-known NP-hard problem [21], in which given a graph G one aims to find a partition of the vertex set of G into two classes of equal size (up to ± 1) so as to minimise the number of crossing edges between the two classes, called the bisection width. We denote by $\mathrm{bis}(G)$ the minimum bisection width of G.

Since the early days of computational complexity, graph partitioning problems, e. g., the minimum bisection problem, have played a central role in computer science [21, 28]. Over the years they have inspired some of the most important algorithmic techniques that we have at our disposal today, such as network flows or semidefinite programming [3, 18, 22, 29, 42].

In the context of the probabilistic analysis of algorithms, it is hard to think of a more intensely studied problem than the *planted bisection model*. In this model a random graph $G = G(n, p_{+1}, p_{-1})$ on vertex set $[n] = \{1, \ldots, n\}$ is created by choosing a map $\sigma : V \to \{-1, 1\}$ uniformly at random subject to $\big|\, |\sigma^{-1}(1)| - |\sigma^{-1}(-1)| \,\big| \leq 1$ and connecting any two vertices $v \neq w$ with probability $p_{\sigma(v)\sigma(w)}$ independently, where $0 \leq p_{-1} < p_{+1} \leq 1$. To ease notation, we often write p_+ for p_{+1} and p_- for p_{-1}, and handle subscripts similarly for other parameters.

Given the random graph G (but not the planted bisection σ), the task is to find a *minimum bisection* of G, i. e., to partition the vertices into two disjoint sets $S, \bar{S} = [n] \setminus S$ whose sizes satisfy $\big|\, |S| - |\bar{S}| \,\big| \leq 1$ such that the number of S-$\bar{S}$-edges is minimum. The planted bisection model has been employed to gauge algorithms based on spectral, semidefinite programming, flow and local search techniques, to name but a few [5, 6, 7, 8, 9, 12, 15, 16, 17, 26, 27, 31, 33, 35].

Remarkably, for a long time the algorithm with the widest range of $n, p_\pm$ for which a minimum bisection can be found efficiently was one of the earliest ones, namely Boppana's spectral algorithm [6].[1] It succeeds if

$$n(p_+ - p_-) \geq c\sqrt{np_+ \ln n}$$

for a certain constant $c > 0$. Under this assumption the planted bisection is minimum w. h. p. (with high probability, meaning with probability tending to one as $n \to \infty$). In fact, recently the critical value $c^* > 0$ for which this statement is true was identified explicitly [1, 38]. In particular, for

$$n(p_+ - p_-) > c^* \sqrt{np_+ \ln n}$$

the minimum bisection width simply equals $(1/4 + o(1))n^2 p_-$ w. h. p.

But if $n(p_+ - p_-) < c^* \sqrt{np_+ \ln n}$, then the minimum bisection width will be strictly smaller than the width of the planted bisection w. h. p. Yet there is another spectral algorithm [9] that finds a minimum bisection w. h. p. under the weaker assumption that

$$n(p_+ - p_-) \geq c\sqrt{np_+ \ln(np_+)}\,, \tag{1.1}$$

[1] Boppana's paper, which only appeared in conference proceedings, did not contain the proof of an important statement about the spectrum of the adjacency matrix. A similar result could, however, be proved along the lines of the work of Friedman, Kahn and Szemerédi [20]; see also [19, 44].

for a certain constant $c > 0$, and even certifies the optimality of its solution. However, [9] does not answer what is arguably the most immediate question: what is the *typical* value of the minimum bisection width $\mathrm{bis}(G)$?

In this paper we derive the value to which the (suitably scaled) minimum bisection width converges in probability (Theorem 1.1). We confine ourselves to the case that $(n/2)p_{\pm} = d_{\pm}$ remain fixed as $n \to \infty$. Hence, the random graph G has bounded average degree. This is arguably the most interesting case because the discrepancy between the planted and the minimum bisection gets larger as the graphs get sparser. In fact, it is easy to see that in the case of fixed $(n/2)p_{\pm} = d_{\pm}$ the difference between the planted and the minimum bisection width is $\Theta(n)$ as the planted bisection is not even locally optimal w. h. p.

Although we build upon some of the insights from [9], it seems difficult to prove our main result by tracing the fairly complicated algorithm from that paper. Instead, our main tool is an elegant message passing algorithm called *Warning Propagation* that plays an important role in the study of random constraint satisfaction problems via ideas from statistical physics [36]. Running Warning Propagation on G naturally corresponds to a fixed-point problem on the 2-simplex, and the minimum bisection width can be cast as a function of the fixed point.

1.2 The main result

To state the fixed-point problem, we consider the functions

$$\psi : \mathbb{R} \to \mathbb{R}, \quad x \mapsto \begin{cases} -1 & \text{if } x < -1, \\ x & \text{if } -1 \le x \le 1, \\ 1 & \text{if } x > 1, \end{cases}$$

and

$$\tilde{\psi} : \mathbb{R} \to \mathbb{R}, \quad x \mapsto \begin{cases} -1 & \text{if } x \le -1, \\ 1 & \text{if } x > -1. \end{cases}$$

Let $\mathcal{P}(\{-1,0,1\})$ be the set of probability measures on $\{-1,0,1\}$. Clearly, we can identify $\mathcal{P}(\{-1,0,1\})$ with the set of all maps $p : \{-1,0,1\} \to [0,1]$ such that $p(-1)+p(0)+p(1) = 1$, i. e., the 2-simplex. For a $\{-1,0,1\}$-valued random variable Z we denote by $\mathcal{L}(Z) \in \mathcal{P}(\{-1,0,1\})$ the distribution of Z. Given $p \in \mathcal{P}(\{-1,0,1\})$, let $(\eta_{p,i})_{i \ge 1}$ be a family of i. i. d. $\{-1,0,1\}$-valued random variables with distribution p. Moreover, let $\gamma_{\pm} = \mathrm{Po}(d_{\pm})$ be Poisson variables that are independent of each other and of the $\eta_{p,i}$. Let

$$Z_{p,d_+,d_-} := \sum_{i=1}^{\gamma_+} \eta_{p,i} - \sum_{i=\gamma_+ +1}^{\gamma_+ + \gamma_-} \eta_{p,i} \tag{1.2}$$

and define a map

$$\begin{aligned} \mathcal{T}_{d_+,d_-} : \mathcal{P}(\{-1,0,1\}) &\to \mathcal{P}(\{-1,0,1\}), \\ p &\mapsto \mathcal{L}(\psi(Z_{p,d_+,d_-})) \end{aligned} \tag{1.3}$$

that maps $p \in \mathcal{P}(\{-1,0,1\})$ to the distribution of $\psi(Z_{p,d_+,d_-})$. Further, with $(\eta_{p,i})_{i\geq 1}$ and $\gamma_\pm$ as before, let

$$\varphi_{d_+,d_-} : \mathcal{P}(\{-1,0,1\}) \to \mathbb{R},$$
$$p \mapsto \frac{1}{2}\mathbb{E}\left[\sum_{i=1}^{\gamma_+} 1\left\{\eta_{p,i} = -\tilde{\psi}(Z_{p,d_+,d_-})\right\} + \sum_{i=\gamma_++1}^{\gamma_++\gamma_-} 1\left\{\eta_{p,i} = \tilde{\psi}(Z_{p,d_+,d_-})\right\}\right].$$

Moreover, let us call $p \in \mathcal{P}(\{-1,0,1\})$ *skewed* if $p(1) \geq 1 - d_+^{-10}$.

Theorem 1.1. *There exists a constant $c > 0$ such that for any $d_\pm > 0$ satisfying*

$$d_+ \geq 2 \qquad \text{and} \qquad d_+ - d_- \geq c\sqrt{d_+ \ln d_+}$$

the map $\mathcal{T}_{d_+,d_-}$ has a unique skewed fixed point p^* and $n^{-1}\operatorname{bis}(G)$ converges in probability to $\varphi_{d_+,d_-}(p^*)$.

Note that $\mathcal{T}_{d_+,d_-}$ may have further fixed points besides p^* for example the Dirac measure $(0,1,0)$, but p^* is the only fixed point which is skewed. We also note that the condition $d_+ \geq 2$ is not optimised—any constant larger than 1 would do as a lower bound, but then in any case the condition $d_+ \geq 2$ follows from the lower bound on $d_+ - d_-$ for sufficiently large c.

In the following sections we will use that the assumptions of Theorem 1.1 allow us to assume that also d_+ is sufficiently large.

1.3 Further related work

Determining the minimum bisection width of a graph is NP-hard [21] and there is evidence that the problem does not even admit a PTAS [30]. On the positive side, it is possible to approximate the minimum bisection width within a factor of $O(\ln n)$ for graphs on n vertices in polynomial time [42].

The planted bisection model has been studied in statistics under the name "stochastic block model" [23]. However, in the context of statistical inference the aim is to recover the planted partition σ as closely as possible given G rather than to determine the minimum bisection width. Recently there has been a lot of progress, much of it inspired by non-rigorous work [13], on the statistical inference problem. The current status of the problem is that matching upper and lower bounds are known for the values of $d_\pm$ for which it is possible to obtain a partition that is non-trivially correlated with σ [34, 37, 40]. Furthermore, there are algorithms that recover a best possible approximation to σ under certain conditions on $d_\pm$ [1, 39, 38]. But since our objective is different, the methods employed in the present paper are somewhat different and, indeed, rather simpler.

Finally, there has been recent progress on determining the minimum bisection width on the Erdős-Rényi random graph. Although its precise asymptotics remain unknown in the case of bounded average degrees d, it was proved in [14] that the main correction term corresponds to the "Parisi formula" in the Sherrington-Kirkpartrick model [43]. Additionally, regarding the case of very sparse random graphs (i. e., with constant average degree), there is a sharp threshold (at $np = \ln 4$) for the minimum bisection width to be linear in n [32].

Generally speaking, the approach that we pursue is somewhat related to the notion of "local weak convergence" of graph sequences as it was used in [2]. More specifically, we are going to argue that the minimum bisection width of G is governed by the "limiting local structure" of the graph, which is a two-type Galton–Watson tree. The fixed-point problem in Theorem 1.1 mirrors the execution of a message passing algorithm on the Galton–Watson tree. The study of this fixed-point problem, for which we use the *contraction method* [41], is the key technical ingredient of our proof. We believe that this strategy provides an elegant framework for tackling many other problems in the theory of random graphs as well. In fact, in a recent paper [10] we combined Warning Propagation with a fixed point analysis on Galton–Watson trees to the k-core problem. A similar approach was used in [24] in the context of random constraint satisfaction problems. Further, in [4] Warning Propagation was applied to the random graph coloring problem.

2 Outline

From here on we keep the notation and the assumptions of Theorem 1.1. In particular, we assume that $d_+ - d_- \geq c\sqrt{d_+ \ln d_+}$ for a large enough constant $c > 0$ and that $d_\pm$ remain fixed as $n \to \infty$. Furthermore we assume that d_+ is bounded from below by a large enough constant. Throughout the paper all graphs will be locally finite and of countable size.

Three main insights enable the proof of Theorem 1.1. The first one, which we borrow from [9], is that w. h. p. G features a fairly large set $\mathcal{C}$ of vertices such that for any two optimal bisections τ_1, τ_2 of G (i. e., maps $\tau_1, \tau_2 : V(G) \to \{\pm 1\}$), we either have $\tau_1(v) = \tau_2(v)$ for all $v \in \mathcal{C}$ or $\tau_1(v) = -\tau_2(v)$ for all $v \in \mathcal{C}$ (Lemma 2.1). In the language of random constraint satisfaction problems, the vertices in $\mathcal{C}$ are "frozen." While there remain $\Omega(n)$ unfrozen vertices, the subgraph that they induce is subcritical, i. e., all components are of size $O(\ln n)$ and indeed most are of bounded size.

The second main ingredient is an efficient message passing algorithm called *Warning Propagation*, (cf. [36, Chapter 19]). We will show that a bounded number of Warning Propagation iterations suffice to arrange almost all of the unfrozen vertices optimally (i. e., to assign almost all of the vertices to two classes such that there is a minimum bisection respecting this assignment) and thus to obtain a very good approximation to the minimum bisection w. h. p. (Proposition 2.2). This insight reduces our task to tracing Warning Propagation for a bounded number of rounds.

This last problem can be solved by studying Warning Propagation on a suitable Galton–Watson tree, because G only contains a negligible number of short cycles w. h. p. (Lemma 2.3). Thus, the analysis of Warning Propagation on the random tree is the third main ingredient of the proof. This task will turn out to be equivalent to studying the fixed-point problem from Section 1.2 (Proposition 2.5). We proceed to outline the three main components of the proof.

2.1 The core

Given a vertex u of a graph G let $\partial_G u$ denote the neighbourhood of u in G. We sometimes omit the subscript G when the graph is clear from the context. More particularly, in the random graph G, let $\partial_\pm u$ denote the set of all neighbours w of u in G with $\sigma(w)\sigma(u) = \pm 1$. Following [9], we define $\mathcal{C}$ as the largest subset $U \subset [n]$ such that

$$\Big| |\partial_\pm u| - d_\pm \Big| \leq \frac{c}{4}\sqrt{d_+ \ln d_+} \qquad \text{and} \qquad |\partial u \setminus U| \leq 100 \quad \text{for all } u \in U\,. \tag{2.1}$$

Clearly, the set $\mathcal{C}$, which we call the *core*, is uniquely defined because any union of sets U that satisfy (2.1) also has the property. Let $\sigma_{\mathcal{C}} : \mathcal{C} \to \{\pm 1\}$, $v \mapsto \sigma(v)$ be the restriction of the "planted assignment" σ to $\mathcal{C}$.

Furthermore, for a graph G, a set $U \subset V(G)$ and a map $\sigma : U \to \{-1, 1\}$ we let

$$\mathrm{cut}(G,\sigma) := \min\Bigg\{ \sum_{\{v,w\}\in E(G)} \frac{1-\tau(v)\tau(w)}{2} \quad \Bigg| \quad \tau : V(G) \to \{\pm 1\} \text{ satisfies } \tau(v) = \sigma_U(v) \text{ for all } v \in U \Bigg\}.$$

Note that U does not appear explicitly in the notation $\mathrm{cut}(G,\sigma)$ despite being integral to the definition—it is however implicit in the notation since U is the domain of σ.

In words, $\mathrm{cut}(G,\sigma)$ is the smallest number of edges in a cut of G that separates the vertices in $U \cap \sigma^{-1}(-1)$ from those in $U \cap \sigma^{-1}(1)$. In particular, $\mathrm{cut}(G,\sigma_{\mathcal{C}})$ is the smallest cut of G that separates the vertices in the core $\mathcal{C}$ that are frozen to -1 from those that are frozen to 1. Finally, for any vertex v we define a set $\mathcal{C}_v = \mathcal{C}_v(G,\sigma)$ of vertices via the following process.

C1 Let $\mathcal{C}_v^{(0)} = \{v\} \cup \partial v$.

C2 Inductively, let $\mathcal{C}_v^{(t+1)} = \mathcal{C}_v^{(t)} \cup \bigcup_{u \in \mathcal{C}_v^{(t)} \setminus \mathcal{C}} \partial u$ and let $\mathcal{C}_v = \bigcup_{t \geq 0} \mathcal{C}_v^{(t)}$.

Lemma 2.1 ([9], Proposition 19 and Section 3.6)**.** *We have*

$$\mathrm{bis}(G) = \mathrm{cut}(G,\sigma_{\mathcal{C}}) + O(1) \qquad \textit{and} \qquad |\mathcal{C}| \geq n(1 - d_+^{-100})$$

w. h. p. Furthermore, for any $\varepsilon > 0$ there exists $\omega > 0$ such that w. h. p.

$$\sum_{v \in [n]} |\mathcal{C}_v| \cdot 1\{|\mathcal{C}_v| \geq \omega\} \leq \varepsilon n\,.$$

2.2 Warning Propagation

To calculate $\mathrm{cut}(G,\sigma_{\mathcal{C}})$ we adopt the *Warning Propagation* ("WP") message passing algorithm.[2] Let us first introduce WP for a generic graph $G = (V(G), E(G))$ and a map $\sigma : U \subset V(G) \to \{-1, 1\}$. At each time $t \geq 0$, WP sends a "message" $\mu_{v \to w}(t \mid G, \sigma) \in \{-1, 0, 1\}$ from v to w for any edge $\{v, w\} \in E(G)$. The messages are directed objects, i. e., $\mu_{v\to w}(t \mid G,\sigma)$ and $\mu_{w \to v}(t \mid G, \sigma)$ may differ. They are defined inductively by

$$\mu_{v\to w}(0 \mid G,\sigma) := \begin{cases} \sigma(v) & \text{if } v \in U, \\ 0 & \text{otherwise,} \end{cases}$$

[2]A discussion of Warning Propagation in the context of the "cavity method" from statistical physics can be found in [36].

and for $t \geq 0$

$$\mu_{v\to w}(t+1 \mid G,\sigma) := \psi\left(\sum_{u\in\partial v\setminus w} \mu_{u\to v}(t \mid G,\sigma)\right). \tag{2.2}$$

Again, U is implicit in the notation $\mu_{v\to w}(t \mid G,\sigma)$ since U is the domain of σ.

Thus, the WP messages are initialised according to $\sigma : U \to \{-1,1\}$. Subsequently, v sends message ± 1 to w if it receives more ± 1 than ∓ 1 messages from its neighbours $u \neq w$. If there is a tie, v sends out 0. Finally, for $t \geq 0$ define

$$\mu_v(t \mid G,\sigma) := \sum_{w\in\partial v} \mu_{w\to v}(t \mid G,\sigma).$$

The intuition is that the message $\mu_{v\to w}$ which v sends to w indicates which class v is most likely to be in based on the current local information it receives from its other neighbours. To minimise the cut, we would like to place v into the class in which most of its neighbours lie. The initialisation is given by the set U, which we will choose to be the core.

Proposition 2.2. *For any $\varepsilon > 0$ there exists $t_0 = t_0(\varepsilon, d_+, d_-)$ such that for all $t \geq t_0$ w. h. p.*

$$\left|\mathrm{cut}(G,\sigma_{\mathcal{C}}) - \frac{1}{2}\sum_{v\in[n]}\sum_{w\in\partial v} 1\{\mu_{w\to v}(t \mid G,\sigma) = -\tilde{\psi}(\mu_v(t \mid G,\sigma))\}\right| \leq \varepsilon n.$$

We defer the proof of Proposition 2.2 to Section 3.

2.3 The local structure

Proposition 2.2 shows that w. h. p. in order to approximate $\mathrm{cut}(G,\sigma_{\mathcal{C}})$ up to a small error of εn we merely need to run WP for a number t_0 of rounds that is bounded in terms of ε. The upshot is that the WP messages $\mu_{w\to v}(t \mid G,\sigma)$ that are required to figure out the minimum bisection width are determined by the *local* structure of G. We show that the local structure of G "converges to" a suitable Galton–Watson tree. For this purpose, for simplicity we always say that the number of potential neighbours of any vertex in each class is $n/2$. This ignores the fact that if n is odd the classes do not have quite this size and the fact that a vertex cannot be adjacent to itself. However, ignoring these difficulties will not affect our calculations in any significant way.

Our task boils down to studying WP on that Galton–Watson tree. Specifically, let $T = T_{d_+,d_-}$ be the Galton–Watson tree with two types $+1,-1$ and offspring matrix

$$\begin{pmatrix} \mathrm{Po}(d_+) & \mathrm{Po}(d_-) \\ \mathrm{Po}(d_-) & \mathrm{Po}(d_+) \end{pmatrix}. \tag{2.3}$$

Hence, a vertex of type ± 1 spawns $\mathrm{Po}(d_+)$ vertices of type ± 1 and independently $\mathrm{Po}(d_-)$ vertices of type ∓ 1. Moreover, the type of the root vertex r_T is chosen uniformly at random. Let $\tau = \tau_{d_+,d_-} : V(T) \to \{\pm 1\}$ assign each vertex of T its type.

The random graph (G,σ) "converges to" (T,τ) in the following sense. For two triples (G,r,σ), (G',r',σ') of graphs G,G', root vertices $r \in V(G)$, $r' \in V(G')$ and maps $\sigma : V(G) \to \{\pm 1\}$, $\sigma' : V(G') \to$

$\{\pm 1\}$ we write $(G,r,\sigma) \cong (G',r',\sigma')$ if there is a graph isomorphism $\varphi : G \to G'$ such that $\varphi(r) = r'$ and $\sigma = \sigma' \circ \varphi$. Further, we denote by $\partial^t(G,r,\sigma)$ the rooted graph obtained from (G,r,σ) by deleting all vertices at distance greater than t from r together with the restriction of σ to this subgraph. The following lemma characterises the local structure of (G,σ).

Lemma 2.3. *Let $t > 0$ be an integer and let T be any tree with root r and map $\tau : V(T) \to \{\pm 1\}$. Then*

$$\frac{1}{n}\sum_{v\in[n]} 1\left\{\partial^t(G,v,\sigma) \cong \partial^t(T,r,\tau)\right\} \quad \xrightarrow{n\to\infty} \quad \mathbb{P}\left[\partial^t(T,r_T,\tau) \cong \partial^t(T,r,\tau)\right] \quad \textit{in probability}.$$

Furthermore, w. h. p. G does not contain more than $\ln n$ vertices v such that $\partial^t(G,v,\sigma)$ contains a cycle.

Proof. Given a tree T with root r and map $\tau : V(T) \to \{\pm 1\}$, let

$$X_t = X_t(T,r,\tau) := \frac{1}{n}\sum_{v\in[n]} 1\left\{\partial^t(G,v,\sigma) \cong \partial^t(T,r,\tau)\right\}.$$

Similarly, for a set C of isomorphism classes of rooted $\{\pm 1\}$-marked trees let

$$q_t = q_t(C) := \mathbb{P}\left[\partial^t(T,r_T,\tau) \in C\right].$$

The proof proceeds by induction on t. If $t = 0$, pick a vertex $v \in [n]$ uniformly at random, then

$$X_0 = \mathbb{P}_v\left(\sigma(v) = \tau(r)\right) = 1/2 \qquad \text{and} \qquad p_0 = \mathbb{P}_T\left(\tau(r_T) = \tau(r)\right) = 1/2$$

for any $\tau(r) \in \{\pm 1\}$. To proceed from t to $t+1$, let d denote the number of children $v_1,\dots,v_d$ of r in T. For each $i = 1,\dots,d$, let T_i denote the tree rooted at v_i in the forest obtained from T by removing r and let $\tau_i : V(T_i) \to \{\pm 1\}$ denote the restriction of τ to the vertex set of T_i. Finally, let $C_1,\dots,C_{\tilde{d}}$ for some $\tilde{d} \le d$ denote the distinct isomorphism classes among $\{\partial^t(T_i,v_i,\tau_i) : i = 1,\dots,d\}$, and let $c_j = |\{i : \partial^t(T_i,v_i,\tau_i) \in C_j\}|$. Let $v \in [n]$ be an arbitrary vertex in G. Our aim is to determine the probability of the event $\{\partial^{t+1}(G,v,\sigma) \cong \partial^{t+1}(T,r,\tau)\}$. Therefore, we think of G as being created in three rounds. First, partition $[n]$ into two classes. Second, randomly insert edges between vertices in $[n] \setminus \{v\}$ according to their planted sign. Finally, reveal the neighbours of v. For the above event to happen, v must have d neighbours in G. Since $|\partial_\pm v|$ are independent binomially distributed random variables with parameters $n/2$ and $p_\pm$ and because $(n/2)p_\pm = d_\pm$, we may approximate $|\partial_\pm v|$ with a Poisson distribution, and v has degree d with probability

$$\frac{(d_+ + d_-)^d}{d!\exp(d_+ + d_-)} + \mathrm{o}(1),$$

where the error term is bounded by $(d_+ + d_-)/n$ by Le Cam's inequality. Conditioned on v having degree d, by induction v is adjacent to precisely c_j vertices with neighbourhood isomorphic to $\partial^t(T_i,v_i,\tau_i) \in C_j$ with probability

$$\binom{d}{c_1 \dots c_{\tilde{d}}} \prod_{j=1}^{\tilde{d}} q_t(C_j) + \mathrm{o}(1).$$

The number of cycles of length $\ell \leq 2t+3$ in G is stochastically bounded by the number of such cycles in $G(n, d_+/n)$ (the Erdős-Rényi graph). For each ℓ, this number tends in distribution to a Poisson variable with bounded mean (see, e. g., Theorem 3.19 in [25]) and so the total number of such cycles is bounded w. h. p. Thus all the pairwise distances (in $G - v$) between neighbours of v are at least $2t+1$ w. h. p. (and in particular this proves the second part of the lemma). Therefore

$$\mathbb{E}_G[X_{t+1}] = \frac{(d_+ + d_-)^d}{d! \exp(d_+ + d_-)} \binom{d}{c_1 \ldots c_{\tilde{d}}} \prod_{j=1}^{\tilde{d}} q_t(C_j) + \mathrm{o}(1).$$

By definition of T, we obtain $\mathbb{E}[X_{t+1}] = q_{t+1} + \mathrm{o}(1)$. To apply Chebyshev's inequality, it remains to determine $\mathbb{E}[X_{t+1}^2]$. Let $v, w \in [n]$ be two randomly choosen vertices. Then w. h. p. v and w have distance at least $2t+3$ in G, conditioned on which $\partial^{t+1}(G, v, \sigma)$ and $\partial^{t+1}(G, w, \sigma)$ are independent. Therefore we obtain

$$\begin{aligned}
&\mathbb{P}_{v,w}\left(\partial^{t+1}(G, v, \sigma) \cong \partial^{t+1}(T, r, \tau) \wedge \partial^{t+1}(G, w, \sigma) \cong \partial^{t+1}(T, r, \tau)\right) \\
&\quad = \mathbb{P}_v\left(\partial^{t+1}(G, v, \sigma) \cong \partial^{t+1}(T, r, \tau)\right) \mathbb{P}_w\left(\partial^{t+1}(G, w, \sigma) \cong \partial^{t+1}(T, r, \tau)\right) + \mathrm{o}(1),
\end{aligned}$$

and finally

$$\begin{aligned}
\mathbb{E}_G[X_{t+1}^2] =& \mathbb{E}_G\left[\mathbb{P}_v\left(\partial^{t+1}(G, v, \sigma) \cong \partial^{t+1}(T, r, \tau)\right) \mathbb{P}_w\left(\partial^{t+1}(G, w, \sigma) \cong \partial^{t+1}(T, r, \tau)\right)\right] \\
&+ \frac{1}{n}\mathbb{E}_G[X_{t+1}] + \mathrm{o}(1) \\
=& \mathbb{E}_G[X_{t+1}]^2 + \mathrm{o}(1).
\end{aligned}$$

The first assertion follows from Chebyshev's inequality. □

2.4 The fixed point

Let (T, r, τ) be a rooted tree together with a map $\tau : V(T) \to \{\pm 1\}$. Then for any pair v, w of adjacent vertices we have the WP messages $\mu_{v \to w}(t \mid T, \tau)$, $t \geq 0$, as defined in (2.2). Since we are going to be particularly interested in the messages directed towards the root, we introduce the following notation. Given the root r, any vertex $v \neq r$ of T has a unique parent vertex u (the neighbour of v on the unique path from v to r). Initially, let

$$\mu_{v\uparrow}(0 \mid T, r, \tau) = \tau(v) \tag{2.4}$$

and define

$$\mu_{v\uparrow}(t \mid T, r, \tau) = \mu_{v \to u}(t \mid T, \tau) \tag{2.5}$$

for $t > 0$. In addition, set $\mu_{r\uparrow}(0 \mid T, r, \tau) = \tau(r)$ and let

$$\mu_{r\uparrow}(t+1 \mid T, r, \tau) = \psi\left(\sum_{v \in \partial r} \mu_{v\uparrow}(t \mid T, r, \tau)\right) \qquad (t \geq 0) \tag{2.6}$$

be the message that r would send to its parent if there was one.

For $p = (p(-1), p(0), p(1)) \in \mathcal{P}(\{-1,0,1\})$ we let $\bar{p} = (p(1), p(0), p(-1))$. Recalling the map

$$\mathcal{T} = \mathcal{T}_{d_+,d_-} : \mathcal{P}(\{-1,0,1\}) \to \mathcal{P}(\{-1,0,1\}),$$
$$p \mapsto \mathcal{T}(p(-1), p(0), p(1)) = \mathcal{L}(\psi(Z_{p,d_+,d_-}))$$

defined by (1.3) in Section 1.2 and writing $\mathcal{T}^t$ for its t-fold iteration, we observe the following.

Lemma 2.4. *Let $p_t := \mathcal{T}^t(0,0,1)$.*

1. *Given that $\tau(r_T) = +1$, the message $\mu_{r_T\uparrow}(t \mid T, r_T, \tau)$ has distribution p_t.*
2. *Given that $\tau(r_T) = -1$, the message $\mu_{r_T\uparrow}(t \mid T, r_T, \tau)$ has distribution $\bar{p}_t$.*

Proof. The proof is by induction on t. In the case $t = 0$ the assertion holds because $\mu_{r_T\uparrow}(0 \mid T, r_T, \tau) = \tau(r_T)$. Now assume that the assertion holds for t. To prove it for $t+1$, let $C_\pm$ be the set of all children v of r_T with $\tau(r_T)\tau(v) = \pm 1$. By construction, $|C_\pm|$ has distribution $\mathrm{Po}(d_\pm)$. Furthermore, let (T_v, v, τ_v) signify the subtree pending on a child v of r_T. Because T is a Galton–Watson tree, the random subtrees T_v are mutually independent. Moreover, each T_v is distributed as a Galton–Watson tree with offspring matrix (2.3) and a root vertex of type $\pm\tau(r_T)$ for each $v \in C_\pm$. Therefore, by induction, the message $\mu_{v\uparrow}(t \mid T_v, v, \tau_v)$ has distribution p_t if $\tau(v) = 1$ and $\bar{p}_t$ if $\tau(v) = -1$. As a consequence,

$$\mu_{r_T\uparrow}(t+1 \mid T, r_T, \tau) := \psi\left(\sum_{v \in C_+} \mu_{v\uparrow}(t \mid T_v, v, \tau_v) + \sum_{v \in C_-} \mu_{v\uparrow}(t \mid T_v, v, \tau_v)\right)$$

has distribution p_{t+1} if $\tau(r_T) = 1$ and $\bar{p}_{t+1}$ otherwise. □

Lemma 2.4 shows that the operator $\mathcal{T}$ mimics WP on the Galton–Watson tree (T, r_T, τ). Hence, to understand the behaviour of WP after a large enough number of iterations we need to investigate the fixed point to which $\mathcal{T}^t(0,0,1)$ converges as $t \to \infty$. In Section 4 we will establish the following.

Proposition 2.5. *The operator $\mathcal{T}$ has a unique skewed fixed point p^* and $\lim_{t\to\infty} \mathcal{T}^t(0,0,1) = p^*$.*

Proof of Theorem 1.1. Consider the random variables

$$X_n := \frac{1}{n}\,\mathrm{bis}(G),$$
$$Y_n^{(t)} := \frac{1}{2}\frac{1}{n}\sum_{v \in [n]}\sum_{w \in \partial_G v} \mathbf{1}\{\mu_{w\to v}(t \mid G, \sigma) = -\tilde{\psi}(\mu_v(t \mid G, \sigma))\}.$$

Then Lemma 2.1 and Proposition 2.2 imply that for any $\varepsilon > 0$,

$$\lim_{t\to\infty}\lim_{n\to\infty} \mathbb{P}\left[\left|X_n - Y_n^{(t)}\right| > \varepsilon\right] = 0. \tag{2.7}$$

By Definition (2.2), $\mu_{w\to v}(t \mid G,\sigma)$ and $\mu_v(t \mid G,\sigma)$ are both determined by $\partial_G^t v$ and the initialisation $\mu_{u\to w}(0 \mid G,\sigma)$ for all $u,w \in \partial_G^t v$, $\{u,w\} \in E(G)$. Since (2.5) and (2.6) match the recursive definition (2.2) of $\mu_{w\to v}(t \mid G,\sigma)$ and $\mu_v(t \mid G,\sigma)$, Lemma 2.3 implies that for any fixed $t>0$ (as n tends to infinity),

$$Y_n^{(t)} \quad \overset{n\to\infty}{\to} \quad x^{(t)} := \frac{1}{2}\mathbb{E}\left[\sum_{w\in\partial_T r_T} 1\{\mu_{w\uparrow}(t \mid T, r_T, \tau) = -\tilde{\psi}(\mu_{r_T}(t \mid T, r_T, \tau))\}\right] \qquad \text{in probability.} \tag{2.8}$$

Now let p^* denote the unique skewed fixed point of $\mathcal{T}$ guaranteed by Proposition 2.5. Since each child of r_T can be considered a root of an independent instance of T to which we can apply Lemma 2.4, we obtain that given $(\tau(w))_{w\in\partial r_T}$ the sequence $(\mu_{w\uparrow}(t \mid T, r_T, \tau))_{w\in\partial r_T}$ converges to a sequence of independent random variables $(\eta_w)_{w\in\partial r_T}$ with distribution p^* (if $\tau(w)=1$) and $\bar{p}^*$ (if $\tau(w)=-1$). By definition $\mu_{r_T}(t \mid T, r_T, \tau)$ converges to

$$\sum_{w\in\partial r_T, \tau(w)=1} \eta_w + \sum_{w\in\partial r_T, \tau(w)=-1} \eta_w$$

Considering the offspring distributions of r_T in both cases, i. e., $\tau(r_T)=\pm 1$, we obtain from $\varphi_{d_+,d_-}(p) = \varphi_{d_+,d_-}(\bar{p})$ for all $p \in \mathcal{P}(\{-1,0,1\})$ that

$$\lim_{t\to\infty} x^{(t)} = \varphi_{d_+,d_-}(p^*). \tag{2.9}$$

Finally, combining (2.7)–(2.9) completes the proof. $\square$

3 Proof of Proposition 2.2

Lemma 3.1. *If $v \in \mathcal{C}$ and $w \in \partial_G v$, then $\mu_{v\to w}(t \mid G,\sigma) = \sigma(v) = \mu_{v\to w}(t \mid G,\sigma_{\mathcal{C}})$ for all $t \geq 0$.*

Proof. We proceed by induction on t. For $t=0$ the assertion is immediate from the initialisation of the messages. To go from t to $t+1$, consider $v \in \mathcal{C}$ and $w \in \partial_G v$. We may assume without loss of generality that $\sigma(v)=1$. By the definition of the WP message,

$$\mu_{v\to w}(t+1 \mid G,\sigma) = \psi\left(\sum_{u\in\partial_G v\setminus\{w\}} \mu_{u\to v}(t \mid G,\sigma)\right) = \psi(S_+ + S_- + S_0) \tag{3.1}$$

where

$$S_+ := \sum_{u\in\mathcal{C}\cap\sigma^{-1}(+1)\cap\partial_G v\setminus\{w\}} \mu_{u\to v}(t \mid G,\sigma),$$
$$S_- := \sum_{u\in\mathcal{C}\cap\sigma^{-1}(-1)\cap\partial_G v\setminus\{w\}} \mu_{u\to v}(t \mid G,\sigma),$$
$$S_0 := \sum_{u\in\partial_G v\setminus(\mathcal{C}\cup\{w\})} \mu_{u\to v}(t \mid G,\sigma).$$

Now, (2.1) ensures that

$$S_+ \geq d_+ - \frac{c}{4}\sqrt{d_+ \ln d_+}\,, \qquad S_- \geq -d_- - \frac{c}{4}\sqrt{d_+ \ln d_+}\,, \qquad |S_0| \leq 100 \leq \frac{c}{4}\sqrt{d_+ \ln d_+}\,, \tag{3.2}$$

provided that the constant $c > 0$ is chosen large enough. Combining (3.1) and (3.2), we see that $S_+ + S_- + S_0 \geq 1$ and thus $\mu_{v\to w}(t+1 \mid G, \sigma) = 1$. The same argument works for $\mu_{v\to w}(t+1 \mid G, \sigma_{\mathcal{C}}) = 1$. □

Let G_v denote the subgraph of G induced on $\mathcal{C}_v$. To prove Proposition 2.2, fix $s > 0$ large enough. Let $\mathcal{S} = \mathcal{S}(s)$ be the set of all vertices such that either $|\mathcal{C}_v| > \sqrt{s}$ or G_v is cyclic. Then Lemma 2.1 (with slightly smaller ε) and Lemma 2.3 imply that $|\mathcal{S}| \leq \varepsilon n$ w. h. p. For the rest of this section, let $v \notin \mathcal{S}$ be fixed.

For $w \in \mathcal{C}_v \setminus \{v\}$ we let $w_{\uparrow v}$ be the neighbour of w on the path from w to v. We define $G_{w\to v}$ as the component of w in the graph obtained from G_v by removing the edge $\{w, w_{\uparrow v}\}$. The vertex set of $G_{w\to v}$ will be denoted by $\mathcal{C}_{w\to v}$. Further, $h_{w\to v}$ is the maximum distance between w and any other vertex in $G_{w\to v}$. Additionally, h_v is the maximum distance between v and any other vertex in G_v. Finally, let $\sigma_v : \mathcal{C}_v \to \{\pm 1\}$, $w \mapsto \sigma(w)$ and let $\sigma_{\mathcal{C},v} : \mathcal{C}_v \cap \mathcal{C} \to \{\pm 1\}$, $w \mapsto \sigma_{\mathcal{C}}(w)$.

Lemma 3.2.

(1) For any $w \in \mathcal{C}_v \setminus \{v\}$ and any $t > h_{w\to v}$ we have

$$\mu_{w\to w_{\uparrow v}}(t \mid G, \sigma) = \mu_{w\to w_{\uparrow v}}(h_{w\to v} + 1 \mid G, \sigma) = \mu_{w\to w_{\uparrow v}}(t \mid G, \sigma_{\mathcal{C}})\,.$$

(2) For any $t \geq h_v$ we have $\mu_v(t \mid G, \sigma) = \mu_v(h_v + 1 \mid G, \sigma) = \mu_v(t \mid G, \sigma_{\mathcal{C}})$.

Proof. The proof of (1) proceeds by induction on $h_{w\to v}$. The construction **C1–C2** of $\mathcal{C}_v$ in Section 2.1 ensures that any $w \in \mathcal{C}_v$ with $h_{w\to v} = 0$ either belongs to $\mathcal{C}$ or has no neighbour besides $w_{\uparrow v}$. Hence for the first case the assumption follows from Lemma 3.1. If $\partial_G w \setminus \{w_{\uparrow v}\} = \emptyset$ we obtain that

$$\mu_{w\to w_{\uparrow v}}(t \mid G, \sigma) = \mu_{w\to w_{\uparrow v}}(t \mid G, \sigma_{\mathcal{C}}) = 0$$

for all $t \geq 1$ by the definition of the WP messages. Now, assume that $h_{w\to v} > 0$ and let $t > h_{w\to v}$. Then all neighbours $u \neq w_{\uparrow v}$ of w in $G_{w\to v}$ satisfy $h_{u\to v} < h_{w\to v}$. Thus, by induction

$$\begin{aligned}\mu_{w\to w_{\uparrow v}}(t \mid G, \sigma) &= \psi\left(\sum_{u\in\partial_G w\setminus\{w_{\uparrow v}\}} \mu_{u\to w}(t-1 \mid G, \sigma)\right)\\ &= \psi\left(\sum_{u\in\partial_G w\setminus\{w_{\uparrow v}\}} \mu_{u\to w}(h_{u\to v}+1 \mid G, \sigma)\right) = \mu_{w\to w_{\uparrow v}}(h_{w\to v}+1 \mid G, \sigma)\,.\end{aligned}$$

An analogous argument applies to $\mu_{w\to w_{\uparrow v}}(t \mid G, \sigma_{\mathcal{C}})$. The proof of (2) is similar. □

For each vertex $w \in \mathcal{C}_v$, $w \neq v$, let $\mu^*_{w\to v} = \mu_{w\to w_{\uparrow v}}(s \mid G, \sigma)$. Further, let $\mu^*_w = \mu_w(s \mid G, \sigma)$. In addition, for $z \in \{\pm 1\}$ let

$$\sigma^z_{w\to v} : \mathcal{C}_{w\to v} \cap (\{w\} \cup \mathcal{C}) \to \{\pm 1\}, \qquad u \mapsto \begin{cases} z & \text{if } u = w, \\ \sigma(u) & \text{otherwise.} \end{cases}$$

In words, $\sigma^z_{w\to v}$ freezes w to z and all other $u \in \mathcal{C}_{w\to v}$ that belong to the core to $\sigma(u)$. Analogously, let

$$\sigma^z_v : \mathcal{C}_v \cap (\{v\} \cup \mathcal{C}) \to \{\pm 1\}, \qquad u \mapsto \begin{cases} z & \text{if } u = v, \\ \sigma(u) & \text{otherwise.} \end{cases}$$

Lemma 3.3. *Suppose that $u \in \mathcal{C}_v \setminus \{v\}$ is such that $h_{u\to v} \geq 1$.*

(1) If $z = \mu^*_{u\to v} \in \{-1, 1\}$, then

$$\operatorname{cut}(G_{u\to v}, \sigma^z_{u\to v}) < \operatorname{cut}(G_{u\to v}, \sigma^{-z}_{u\to v}). \tag{3.3}$$

Similarly, if $z = \psi(\mu^*_v) \in \{-1, 1\}$, then

$$\operatorname{cut}(G_v, \sigma^z_v) < \operatorname{cut}(G_v, \sigma^{-z}_v). \tag{3.4}$$

(2) If $\mu^*_{u\to v} = 0$, then

$$\operatorname{cut}(G_{u\to v}, \sigma^{+1}_{u\to v}) = \operatorname{cut}(G_{u\to v}, \sigma^{-1}_{u\to v}). \tag{3.5}$$

Similarly, if $\mu^*_v = 0$, then

$$\operatorname{cut}(G_v, \sigma^{+1}_v) = \operatorname{cut}(G_v, \sigma^{-1}_v). \tag{3.6}$$

Proof. We prove (3.3) and (3.5) by induction on $h_{u\to v}$. If $h_{u\to v} = 1$ then we have that all neighbours $w \in \partial_{\mathcal{C}_{u\to v}} u$ of u with $\mu^*_{u\to v} \neq 0$ are in $\mathcal{C}$, i. e., fixed under $\sigma^z_{u\to v}$. Since $\mathcal{C}_{u\to v} = \partial_G u \setminus \{u_{\uparrow v}\} \cup \{u\}$, we obtain

$$\operatorname{cut}(\mathcal{C}_{u\to v}, \sigma^{-z}_{u\to v}) - \operatorname{cut}(\mathcal{C}_{u\to v}, \sigma^z_{u\to v}) = \left| \sum_{w \in \partial_G u \setminus \{u_{\uparrow v}\}} \mu^*_{w\to v} \right| \tag{3.7}$$

by definition of z. By the induction hypothesis and because $G_{u\to v}$ is a tree (as $v \notin \mathcal{S}$) we have that (3.7) holds for $h_{u\to v} > 1$ as well. A similar argument yields (3.4) and (3.6). □

Now, let $\mathcal{U}_v$ be the set of all $w \in \mathcal{C}_v$ such that $\mu^*_{w\to v} \neq 0$. Furthermore, let

$$\sigma_{\uparrow v} : \mathcal{U}_v \cup \{v\} \to \{-1, +1\}, \qquad w \mapsto \begin{cases} \tilde{\psi}(\mu^*_v) & \text{if } w = v, \\ \mu^*_{w\to v} & \text{otherwise.} \end{cases} \tag{3.8}$$

Thus, $\sigma_{\uparrow v}$ sets all $w \in \mathcal{C}_v \cap \mathcal{C} \setminus \{v\}$ to their planted sign and all $w \in \mathcal{U}_v \setminus \mathcal{C}$ to $\mu^*_{w\to v}$. Moreover, $\sigma_{\uparrow v}$ sets v to $\psi(\mu^*_v)$ if $\psi(\mu^*_v) \neq 0$ and to 1 if there is a tie.

Corollary 3.4. *We have* $\mathrm{cut}(G_v, \sigma_{\mathcal{C}}) = \mathrm{cut}(G_v, \sigma_{\uparrow v})$.

Proof. This is immediate from Lemma 3.3. □

Hence, in order to determine an optimal cut of G_v we merely need to figure out the assignment of the vertices in $\mathcal{C}_v \setminus (\{v\} \cup \mathcal{U}_v)$. Suppose that $\sigma^*_{\uparrow v} : \mathcal{C}_v \to \{\pm 1\}$ is an optimal extension of $\sigma_{\uparrow v}$ to a cut of G_v, i. e.,

$$\mathrm{cut}(G_v, \sigma_{\uparrow v}) = \sum_{\{u,w\} \in E(G_v)} \frac{1}{2}(1 - \sigma^*_{\uparrow v}(u)\sigma^*_{\uparrow v}(w)).$$

Corollary 3.5. *It holds that*

$$\sum_{w \in \partial_G v} \frac{1}{2}(1 - \sigma^*_{\uparrow v}(v)\sigma^*_{\uparrow v}(w)) = \sum_{w \in \partial_G v} 1\{\mu^*_{w \to v} = -\tilde{\psi}(\mu_v)\}.$$

Proof. Part (2) of Lemma 3.3 implies that $\sigma^*_{\uparrow v}(v)\sigma^*_{\uparrow v}(w) = 1$ for all $w \in \partial_G v$ such that $\mu^*_{w \to v} = 0$. The claim follows from the definition of $\sigma^*_{\uparrow v}$ and (3.8). □

Proof of Proposition 2.2. Given $\varepsilon > 0$ choose $\delta = \delta(\varepsilon, d_+, d_-)$ sufficiently small and $s = s(\varepsilon, \delta, d_+, d_-) > 0$ sufficiently large. In particular, pick s large enough so that

$$\mathbb{P}(|\mathcal{S}| \geq \delta n) < \varepsilon. \tag{3.9}$$

Provided that δ is sufficiently small, the Chernoff bound implies that for large n

$$\mathbb{P}\left(\frac{1}{2}\sum_{v \in \mathcal{S}} |\partial_G v| \geq \varepsilon n \;\middle|\; |\mathcal{S}| < \delta n\right) < \varepsilon. \tag{3.10}$$

Now, suppose that $\sigma^*_{\mathcal{C}}$ is an optimal extension of $\sigma_{\mathcal{C}}$ to a cut of G and let $v \notin \mathcal{S}$. Then using the definition of $\mathcal{C}_v$, Corollary 3.4 implies that

$$\sum_{w \in \partial_G v} (1 - \sigma^*_{\mathcal{C}}(v)\sigma^*_{\mathcal{C}}(w)) = \sum_{w \in \partial_G v} (1 - \sigma^*_{\uparrow v}(v)\sigma^*_{\uparrow v}(w)).$$

Therefore, we obtain

$$\mathbb{P}\left(\left|\mathrm{cut}(G, \sigma_{\mathcal{C}}) - \frac{1}{2}\sum_{v \notin \mathcal{S}} \sum_{w \in \partial_G v} (1 - \sigma^*_{\uparrow v}(v)\sigma^*_{\uparrow v}(w))\right| \geq \varepsilon n\right) \leq \mathbb{P}\left(\frac{1}{2}\sum_{v \in \mathcal{S}} |\partial_G v| \geq \varepsilon n\right) \leq 2\varepsilon.$$

The assertion follows from Lemma 3.2 for $t \geq s$. □

4 Proof of Proposition 2.5

We continue to denote the set of probability measures on $\mathcal{X} \subset \mathbb{R}^k$ by $\mathcal{P}(\mathcal{X})$. For an $\mathcal{X}$-valued random variable X we denote by $\mathcal{L}(X) \in \mathcal{P}(\mathcal{X})$ the distribution of X. Furthermore, if $p, q \in \mathcal{P}(\mathcal{X})$, then we write $\mathcal{P}_{p,q}(\mathcal{X})$ to denote the set of all probability measures μ on $\mathcal{X} \times \mathcal{X}$ such that the marginal distribution of the first component coincides with p and the marginal distribution of the second component coincides with q. The space $\mathcal{P}(\{-1,0,1\})$ is complete with respect to the L_1-Wasserstein metric that is defined by

$$\ell_1(p,q) := \inf\left\{\mathbb{E}|X-Y| : X, Y \text{ are random variables with } \mathcal{L}(X,Y) \in \mathcal{P}_{p,q}(\{-1,0,1\})\right\}.$$

In words, the infimum of $\mathbb{E}|X-Y|$ is over all couplings (X,Y) of the distributions p, q. Such a coupling (X,Y) is *optimal* if $\ell_1(p,q) = \mathbb{E}|X-Y|$. Finally, let $\mathcal{P}^*(\{-1,0,1\})$ be the set of all skewed probability measures on $\{-1,0,1\}$, that is, $\mathcal{P}^*(\{-1,0,1\})$ is the set of all $p = (p(-1), p(0), p(1)) \in \mathcal{P}(\{-1,0,1\})$ satisfying $p(1) \geq 1 - d_+^{-10}$. Being a closed subset of $\mathcal{P}(\{-1,0,1\})$, $\mathcal{P}^*(\{-1,0,1\})$ is complete with respect to $\ell_1(\cdot\,,\cdot)$.

As in the definition (1.2)-(1.3) of the operator $\mathcal{T} = \mathcal{T}_{d_+,d_-}$ for $p \in \mathcal{P}(\{-1,0,1\})$ we let $(\eta_{p,i})_{i\geq 1}$ be a family of independent random variables with distribution p. Further, let $\gamma_\pm = \mathrm{Po}(d_\pm)$ be independent of each other and of the $(\eta_{p,i})_{i\geq 1}$. We introduce the shorthands

$$Z_p := Z_{p,d_+,d_-}\,, \qquad Z_{p,+} := \sum_{i=1}^{\gamma_+} \eta_{p,i}\,, \qquad Z_{p,-} := \sum_{i=\gamma_+ +1}^{\gamma_+ + \gamma_-} \eta_{p,i}\,,$$

so that

$$Z_p = Z_{p,+} - Z_{p,-}\,.$$

Also set $\lambda = c\sqrt{d_+ \ln d_+}$ and recall that $c > 0$ is a constant that we assume to be sufficiently large.

Lemma 4.1. *The operator $\mathcal{T}$ maps $\mathcal{P}^*(\{-1,0,1\})$ into itself.*

Proof. Suppose that $p \in \mathcal{P}(\{-1,0,1\})$ is skewed. Then

$$\mathbb{P}\left(Z_p < 1\right) \leq \mathbb{P}\left(Z_{p,+} \leq d_+ - \frac{\lambda - 1}{2}\right) + \mathbb{P}\left(Z_{p,-} \geq d_- + \frac{\lambda - 1}{2}\right). \tag{4.1}$$

Since $|\eta_{p,i}| \leq 1$ for all i, we can bound the second summand from above by invoking the Chernoff bound on a binomial approximation of the Poisson distribution. Using that $d_- + c\sqrt{d_+ \ln d_+} \leq d_+$ we obtain

$$\mathbb{P}\left(\gamma_- \geq d_- + \frac{c}{2}\sqrt{d_+ \ln d_+} - \frac{1}{2}\right) \leq d_+^{-\frac{c^2}{2}} < \frac{1}{3} d_+^{-10} \tag{4.2}$$

provided c is large enough.

To bound the other summand from above we use that $(\eta_{p,i})_{i\geq 1}$ is a sequence of independent *skewed* random variables, whence by the Chernoff bound

$$\begin{aligned}
&\mathbb{P}\left(Z_{p,+} \leq d_+ - \frac{\lambda - 1}{2}\right) \\
&\leq \mathbb{P}\left(|\gamma_+ - d_+| > \lambda/8\right) + \mathbb{P}\left(Z_{p,-} \leq d_+ - \frac{\lambda - 1}{2} \,\middle|\, \gamma_+ \geq d_+ - \lambda/8\right) \\
&\leq \frac{1}{3} d_+^{-10} + \mathbb{P}\left[\mathrm{Bin}(d_+ - \lambda/8, 1 - d_+^{-10}) \leq d_+ - \lambda/7\right] < \frac{2}{3} d_+^{-10}\,,
\end{aligned} \tag{4.3}$$

provided that c is sufficiently large. Combining (4.1)–(4.3) completes the proof. □

Lemma 4.2. *The operator $\mathcal{T}$ is strictly ℓ_1-contracting on $\mathcal{P}^*(\{-1,0,1\})$.*

Proof. Let $p,q \in \mathcal{P}^*(\{-1,0,1\})$. We aim to show that $\ell_1(\mathcal{T}(p),\mathcal{T}(q)) \le (1/2)\,\ell_1(p,q)$. To this end, we let $(\eta_{p,i},\eta_{q,i})_{i\ge 1}$ be a family of random variables such that the $(\eta_{p,i})_{i\ge 1}$ are independent with distribution p and the $(\eta_{q,i})_{i\ge 1}$ are independent with distribution q but the pair $(\eta_{p,i},\eta_{q,i})$ is an optimal coupling for every i. Then by the definition of $\ell_1(\,\cdot\,,\,\cdot\,)$,

$$\ell_1(\mathcal{T}(p),\mathcal{T}(q)) \le \mathbb{E}\left|\psi(Z_p)-\psi(Z_q)\right|. \tag{4.4}$$

To estimate the r. h. s., let $\tilde\eta_{p,i} = 1\{\eta_{p,i}=1\}$ and $\tilde\eta_{q,i} = 1\{\eta_{q,i}=1\}$. Further, let $\mathfrak{F}_i$ be the σ-algebra generated by $\tilde\eta_{p,i},\tilde\eta_{q,i}$ and let $\mathfrak{F}$ be the σ-algebra generated by γ_+,γ_- and the random variables $(\tilde\eta_{p,i},\tilde\eta_{q,i})_{i\ge 1}$. Additionally, let $\gamma=\gamma_+ + \gamma_-$ and consider the three events

$$A_1 := \left\{\sum_{i=1}^{\gamma}\tilde\eta_{p,i}\tilde\eta_{q,i} \ge \gamma-10\right\}, \qquad A_2 := \{\gamma \ge 2d_+\}, \qquad A_3 := \{\gamma_+ - \gamma_- \le 20\}.$$

We are going to bound $|\psi(Z_p)-\psi(Z_q)|$ on $A := A_1 \setminus (A_2\cup A_3)$, $C := \overline{A_1\cup A_2\cup A_3}$, A_2 and $A_3\setminus A_2$ separately. The bound on the first event is immediate: if A occurs, then $\psi(Z_p)=\psi(Z_q)=1$ with certainty. Hence,

$$\mathbb{E}\left[|\psi(Z_p)-\psi(Z_q)|\;1_A\right] = 0. \tag{4.5}$$

Let us turn to the second event C. Because the pairs $(\eta_{p,i},\eta_{q,i})_{i\ge 1}$ are mutually independent, we find

$$\mathbb{E}\left[|\eta_{p,i}-\eta_{q,i}| \mid \mathfrak{F}\right] = \mathbb{E}\left[|\eta_{p,i}-\eta_{q,i}| \mid \mathfrak{F}_i\right] \qquad \text{for all } i\ge 1. \tag{4.6}$$

Clearly, if $\tilde\eta_{p,i}\tilde\eta_{q,i}=1$, then $\eta_{p,i}-\eta_{q,i}=0$. Consequently,

$$\mathbb{E}\left[|\eta_{p,i}-\eta_{q,i}| \mid \mathfrak{F}_i\right] \le \frac{\mathbb{E}|\eta_{p,i}-\eta_{q,i}|}{\mathbb{P}[\tilde\eta_{p,i}\tilde\eta_{q,i}=0]} = \frac{\mathbb{E}|\eta_{p,1}-\eta_{q,1}|}{\mathbb{P}[\tilde\eta_{p,1}\tilde\eta_{q,1}=0]}. \tag{4.7}$$

Since the events A_1,A_2,A_3 are $\mathfrak{F}$-measurable and because $\bar A_2$ ensures that $\gamma<2d_+$, (4.6) and (4.7) yield

$$\mathbb{E}[|\psi(Z_p)-\psi(Z_q)| \mid \mathfrak{F}]1_C \le \frac{2d_+\mathbb{E}|\eta_{p,1}-\eta_{q,1}|}{\mathbb{P}[\tilde\eta_{p,1}\tilde\eta_{q,1}=0]}\cdot 1_C. \tag{4.8}$$

Further, because the pairs $(\eta_{p,i},\eta_{q,i})_{i\ge 1}$ are independent and because p,q are skewed,

$$\mathbb{P}(C) \le \mathbb{P}\left(\gamma\le 2d_+, \sum_{i=1}^{\gamma}\tilde\eta_{p,i}\tilde\eta_{q,i} \le \gamma-10\right) \le \left(2d_+\,\mathbb{P}(\tilde\eta_{p,1}\tilde\eta_{q,1}=0)\right)^{10}. \tag{4.9}$$

Combining (4.8) and (4.9), we obtain

$$\mathbb{E}\left[\mathbb{E}\left[|\psi(Z_p)-\psi(Z_q)| \mid \mathfrak{F}\right]1_C\right] \le (2d_+)^{11}\,\mathbb{P}(\tilde\eta_{p,1}\tilde\eta_{q,1}=0)^{9}\,\mathbb{E}|\eta_{p,1}-\eta_{q,1}|. \tag{4.10}$$

Since p,q are skewed, we furthermore obtain $\mathbb{P}(\tilde{\eta}_{p,1}\tilde{\eta}_{q,1}=0)\leq 2d_+^{-10}$. Therefore

$$\begin{aligned}\mathbb{E}\left[|\psi(Z_p)-\psi(Z_q)|1_C\right] &= \mathbb{E}\left[\mathbb{E}\left[|\psi(Z_p)-\psi(Z_q)|\mid\mathfrak{F}\right]1_{\overline{A_1\cup A_2\cup A_3}}\right]\\ &\leq 2^{20}d_+^{-79}\mathbb{E}|\eta_{p,1}-\eta_{q,1}|\,.\end{aligned}$$

With respect to A_2, the triangle inequality yields

$$\mathbb{E}[|\psi(Z_p)-\psi(Z_q)|1_{A_2}]\leq 2\mathbb{E}|\eta_{p,1}-\eta_{q,1}|\cdot\mathbb{E}[\gamma 1_{A_2}]\,. \tag{4.11}$$

Further, since $\gamma=\mathrm{Po}(d_++d_-)$, the Chernoff bound entails that $\mathbb{E}[\gamma 1_{A_2}]\leq d_+^{-1}$ if the constant c is chosen large enough. Combining this estimate with (4.11), we get

$$\mathbb{E}[|\psi(Z_p)-\psi(Z_q)|1_{A_2}]\leq 2d_+^{-1}\mathbb{E}|\eta_{p,1}-\eta_{q,1}|\,. \tag{4.12}$$

Finally, on $A_3\setminus A_2$ we have

$$\mathbb{E}[|\psi(Z_p)-\psi(Z_q)|1_{A_3\setminus A_2}]\leq 4d_+\mathbb{E}|\eta_{p,1}-\eta_{q,1}|\,\mathbb{P}\left[\gamma_+-\gamma_-\leq 20\right]. \tag{4.13}$$

Since $\gamma_\pm=\mathrm{Po}(d_\pm)$ and $d_+-d_-\geq\lambda$, the Chernoff bound yields $\mathbb{P}\left[\gamma_+-\gamma_-\leq 20\right]\leq d_+^{-2}$, if c is large enough. Hence, (4.13) implies

$$\mathbb{E}[|\psi(Z_p)-\psi(Z_q)|1_{A_3\setminus A_2}]\leq 4d_+^{-1}\mathbb{E}|\eta_{p,1}-\eta_{q,1}|\,. \tag{4.14}$$

Finally, the assertion follows from (4.4), (4.5), (4.10), (4.12) and (4.14). □

Proof of Proposition 2.5. The assertion follows from Lemmas 4.1 and 4.2 and the Banach fixed point theorem. □

References

[1] EMMANUEL ABBE, AFONSO S. BANDEIRA, AND GEORGINA HALL: Exact recovery in the stochastic block model. *IEEE Trans. Inform. Theory*, 62(1):471–487, 2016. [doi:10.1109/TIT.2015.2490670, arXiv:1405.3267] 2, 4

[2] DAVID ALDOUS AND JOHN M. STEELE: The objective method: Probabilistic combinatorial optimization and local weak convergence. In *Probability on Discrete Structures*, pp. 1–72. Springer, 2004. [doi:10.1007/978-3-662-09444-0_1] 4

[3] SANJEEV ARORA, SATISH RAO, AND UMESH V. VAZIRANI: Expander flows, geometric embeddings and graph partitioning. *J. ACM*, 56(2):5:1–5:37, 2009. Preliminary version in STOC'04. [doi:10.1145/1502793.1502794] 2

[4] VICTOR BAPST, AMIN COJA-OGHLAN, SAMUEL HETTERICH, FELICIA RASSMANN, AND DAN VILENCHIK: The condensation phase transition in random graph coloring. *Communications in Mathematical Physics*, 341(2):543–606, 2016. Preliminary version in RANDOM'14. [doi:10.1007/s00220-015-2464-z, arXiv:1404.5513] 5

[5] BÉLA BOLLOBÁS AND ALEX D. SCOTT: Max cut for random graphs with a planted partition. *Combin. Probab. Comput.*, 13(4-5):451–474, 2004. [doi:10.1017/S0963548304006303] 2

[6] RAVI B. BOPPANA: Eigenvalues and graph bisection: An average-case analysis. In *Proc. 28th FOCS*, pp. 280–285. IEEE Comp. Soc. Press, 1987. [doi:10.1109/SFCS.1987.22] 2

[7] THANG NGUYEN BUI, SOMA CHAUDHURI, FRANK THOMSON LEIGHTON, AND MICHAEL SIPSER: Graph bisection algorithms with good average case behavior. *Combinatorica*, 7(2):171–191, 1987. Preliminary version in FOCS'84. [doi:10.1007/BF02579448] 2

[8] TED CARSON AND RUSSELL IMPAGLIAZZO: Hill-climbing finds random planted bisections. In *Proc. 12th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'01)*, pp. 903–909. ACM Press, 2001. ACM DL. 2

[9] AMIN COJA-OGHLAN: A spectral heuristic for bisecting random graphs. *Random Structures Algorithms*, 29(3):351–398, 2006. Preliminary version in SODA'05. [doi:10.1002/rsa.20116] 2, 3, 5, 6

[10] AMIN COJA-OGHLAN, OLIVER COOLEY, MIHYUN KANG, AND KATHRIN SKUBCH: How does the core sit inside the mantle? *Electronic Notes in Discrete Mathematics*, 49:489 – 496, 2015. [doi:10.1016/j.endm.2015.06.068, arXiv:1503.09030] 5

[11] AMIN COJA-OGHLAN, OLIVER COOLEY, MIHYUN KANG, AND KATHRIN SKUBCH: The minimum bisection in the planted bisection model. In *Proc. 19th Internat. Workshop on Randomization and Computation (RANDOM'15)*, pp. 710–725. Schloss Dagstuhl, 2015. [doi:10.4230/LIPIcs.APPROX-RANDOM.2015.710, arXiv:1505.02985] 1

[12] ANNE CONDON AND RICHARD M. KARP: Algorithms for graph partitioning on the planted partition model. *Random Structures Algorithms*, 18(2):116–140, 2001. Preliminary version in RANDOM'99. [doi:10.1002/1098-2418(200103)18:2<116::AID-RSA1001>3.0.CO;2-2] 2

[13] 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, 2011. [doi:10.1103/PhysRevE.84.066106, arXiv:1109.3041] 4

[14] AMIR DEMBO, ANDREA MONTANARI, AND SUBHABRATA SEN: Extremal cuts of sparse random graphs. *Ann. Probab.*, 45(2):1190–1217, 2017. [doi:10.1214/15-AOP1084, arXiv:1503.03923] 4

[15] TASSOS DIMITRIOU AND RUSSELL IMPAGLIAZZO: Go with the winners for graph bisection. In *Proc. 9th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'98)*, pp. 510–520. ACM Press, 1998. ACM DL. 2

[16] MARTIN E. DYER AND ALAN M. FRIEZE: The solution of some random NP-hard problems in polynomial expected time. *J. Algorithms*, 10(4):451–489, 1989. [doi:10.1016/0196-6774(89)90001-1] 2

[17] URIEL FEIGE AND JOE KILIAN: Heuristics for semirandom graph problems. *J. Comput. System Sci.*, 63(4):639–671, 2001. [doi:10.1006/jcss.2001.1773] 2

[18] URIEL FEIGE AND ROBERT KRAUTHGAMER: A polylogarithmic approximation of the minimum bisection. *SIAM J. Comput.*, 31(4):1090–1118, 2002. Preliminary version in FOCS'00. See also SIAM Rev. 2006. [doi:10.1137/S0097539701387660] 2

[19] URIEL FEIGE AND ERAN OFEK: Spectral techniques applied to sparse random graphs. *Random Structures Algorithms*, 27(2):251–275, 2005. [doi:10.1002/rsa.20089] 2

[20] JOEL FRIEDMAN, JEFF KAHN, AND ENDRE SZEMERÉDI: On the second eigenvalue of random regular graphs. In *Proc. 21st STOC*, pp. 587–598. ACM Press, 1989. [doi:10.1145/73007.73063] 2

[21] MICHAEL R. GAREY, DAVID S. JOHNSON, AND LARRY J. STOCKMEYER: Some simplified NP-complete graph problems. *Theoret. Comput. Sci.*, 1(3):237–267, 1976. Preliminary version in STOC'74. [doi:10.1016/0304-3975(76)90059-1] 2, 4

[22] MICHEL X. GOEMANS AND DAVID P. WILLIAMSON: Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42(6):1115–1145, 1995. Preliminary version in STOC'94. [doi:10.1145/227683.227684] 2

[23] PAUL W. HOLLAND, KATHRYN LASKEY, AND SAMUEL LEINHARDT: Stochastic blockmodels: First steps. *Social networks*, 5(2):109–137, 1983. [doi:10.1016/0378-8733(83)90021-7] 4

[24] MORTEZA IBRAHIMI, YASH KANORIA, MATT KRANING, AND ANDREA MONTANARI: The set of solutions of random XORSAT formulae. *Ann. Appl. Probab.*, 25(5):2743–2808, 2015. Preliminary version in SODA'12. [doi:10.1214/14-AAP1060, arXiv:1107.5377] 5

[25] SVANTE JANSON, TOMASZ ŁUCZAK, AND ANDRZEJ RUCIŃSKI: *Random Graphs*. John Wiley & Sons, 2000. [doi:10.1002/9781118032718] 9

[26] MARK JERRUM AND GREGORY B. SORKIN: The Metropolis algorithm for graph bisection. *Discrete Applied Mathematics*, 82(1-3):155–175, 1998. [doi:10.1016/S0166-218X(97)00133-9] 2

[27] ARI JUELS: *Topics in black-box combinatorial function optimization*. Ph. D. thesis, UC Berkeley, 1996. Available at Research Gate. 2

[28] RICHARD KARP: Reducibility among combinatorial problems. In *Complexity of Computer Computations*, pp. 85–103. Plenum Press, 1972. [doi:10.1007/978-1-4684-2001-2_9] 2

[29] MAREK KARPINSKI: Approximability of the minimum bisection problem: An algorithmic challenge. In *Proc. 27th Internat. Symp. Math. Found. Comput. Sci. (MFCS'02)*, volume 2420 of *LNCS*, pp. 59–67. Springer, 2002. [doi:10.1007/3-540-45687-2_4] 2

[30] SUBHASH KHOT: Ruling out PTAS for graph min-bisection, dense k-subgraph, and bipartite clique. *SIAM J. Comput.*, 36(4):1025–1071, 2006. Preliminary version in FOCS'04. [doi:10.1137/S0097539705447037] 4

[31] LUDĚK KUČERA: Expected complexity of graph partitioning problems. *Discr. Appl. Math.*, 57(2-3):193–212, 1995. [doi:10.1016/0166-218X(94)00103-K] 2

[32] MALWINA J. LUCZAK AND COLIN MCDIARMID: Bisecting sparse random graphs. *Random Structures Algorithms*, 18(1):31–38, 2001. [doi:10.1002/1098-2418(200101)18:1<31::AID-RSA3>3.0.CO;2-1] 4

[33] KONSTANTIN MAKARYCHEV, YURY MAKARYCHEV, AND ARAVINDAN VIJAYARAGHAVAN: Approximation algorithms for semi-random partitioning problems. In *Proc. 44th STOC*, pp. 367–384. ACM Press, 2012. [doi:10.1145/2213977.2214013, arXiv:1205.2234] 2

[34] LAURENT MASSOULIÉ: Community detection thresholds and the weak Ramanujan property. In *Proc. 46th STOC*, pp. 694–703. ACM Press, 2014. [doi:10.1145/2591796.2591857, arXiv:1311.3085] 4

[35] FRANK MCSHERRY: Spectral partitioning of random graphs. In *Proc. 42nd FOCS*, pp. 529–537. IEEE Comp. Soc. Press, 2001. [doi:10.1109/SFCS.2001.959929] 2

[36] MARC MÉZARD AND ANDREA MONTANARI: *Information, Physics, and Computation*. Oxford Univ. Press, 2009. [doi:10.1093/acprof:oso/9780198570837.001.0001] 3, 5, 6

[37] ELCHANAN MOSSEL, JOE NEEMAN, AND ALLAN SLY: Stochastic block models and reconstruction, 2012. [arXiv:1202.1499] 4

[38] ELCHANAN MOSSEL, JOE NEEMAN, AND ALLAN SLY: Consistency thresholds for the planted bisection model. In *Proc. 47th STOC*, pp. 69–75. ACM Press, 2015. [doi:10.1145/2746539.2746603, arXiv:1407.1591] 2, 4

[39] ELCHANAN MOSSEL, JOE NEEMAN, AND ALLAN SLY: Belief propagation, robust reconstruction and optimal recovery of block models. *Ann. Appl. Probab.*, 26(4):2211–2256, 2016. Preliminary version in COLT'14. [doi:10.1214/15-AAP1145, arXiv:1309.1380] 4

[40] ELCHANAN MOSSEL, JOE NEEMAN, AND ALLAN SLY: A proof of the block model threshold conjecture. *Combinatorica*, pp. 1–44, 2017. [doi:10.1007/s00493-016-3238-8, arXiv:1311.4115] 4

[41] RALPH NEININGER AND LUDGER RÜSCHENDORF: A general limit theorem for recursive algorithms and combinatorial structures. *Ann. Appl. Probab.*, 14(1):378–418, 2004. [doi:10.1214/aoap/1075828056] 5

[42] HARALD RÄCKE: Optimal hierarchical decompositions for congestion minimization in networks. In *Proc. 40th STOC*, pp. 255–264. ACM Press, 2008. [doi:10.1145/1374376.1374415] 2, 4

[43] MICHEL TALAGRAND: The Parisi formula. *Ann. Math.*, 163(1):221–263, 2006. [doi:10.4007/annals.2006.163.221] 4

[44] VAN H. VU: Spectral norm of random matrices. *Combinatorica*, 27(6):721–736, 2007. Preliminary version in STOC'05. [doi:10.1007/s00493-007-2190-z] 2

AUTHORS

Amin Coja-Oghlan
Professor
Goethe University Frankfurt, Germany
acoghlan@math.uni-frankfurt.de
http://www.uni-frankfurt.de/52107833

Oliver Cooley
Assistant professor
Graz University of Technology, Austria
cooley@math.tugraz.at

Mihyun Kang
Professor
Graz University of Technology, Austria
kang@math.tugraz.at
http://www.math.tugraz.at/~kang/

Kathrin Skubch
Ph. D. student
Goethe University Frankfurt, Germany
skubch@math.uni-frankfurt.de
http://www.uni-frankfurt.de/53778787

ABOUT THE AUTHORS

AMIN COJA-OGHLAN studied Mathematics and Computer Science in Hamburg and Berlin. He graduated from the University of Hamburg with a Ph. D. in Mathematics in 2001 under the supervision of Johannes Michaliček and obtained a Habilitation from Humboldt University Berlin in 2005. After visiting Carnegie Mellon University in 2007, he held faculty positions at the University of Edinburgh and the University of Warwick before joining Goethe University Frankfurt in 2012.

OLIVER COOLEY received a Bachelor and Master of Mathematics from the University of Cambridge in 2004 and 2005. He graduated from the University of Birmingham with a Ph. D. in Mathematics in 2010 under the supervision of Daniela Kühn and Deryk Osthus. His research interests include random graphs and hypergraphs, extremal graph theory and extremal hypergraph theory.

MIHYUN KANG graduated with a Ph. D. in Mathematics from KAIST (Korea Advanced Institute of Science and Technology) in 2001 under the supervision of Geon Ho Choe. Her research fields include random graphs, random hypergraphs and random graphs on surfaces.

KATHRIN SKUBCH is a Ph. D. student of Amin Coja-Oghlan at Goethe University Frankfurt. Her research fields include probabilistic combinatorics, random graphs and hypergraphs and phase transitions in discrete structures.