



HAL
open science

The densest subgraph problem in sparse random graphs

Venkat Anantharam, Justin Salez

► **To cite this version:**

Venkat Anantharam, Justin Salez. The densest subgraph problem in sparse random graphs. The Annals of Applied Probability, 2016, 26 (1), pp.305-327. 10.1214/14-AAP1091 . hal-00919079

HAL Id: hal-00919079

<https://hal.science/hal-00919079>

Submitted on 16 Dec 2013

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.

The densest subgraph problem in sparse random graphs

Venkat Anantharam* and Justin Salez†

December 16, 2013

Abstract

We determine the asymptotic behavior of the maximum subgraph density of large random graphs with a prescribed degree sequence. The result applies in particular to the Erdős-Rényi model, where it settles a conjecture of Hajek (1990). Our proof consists in extending the notion of balanced loads from finite graphs to their local weak limits, using unimodularity. This is a new illustration of the objective method described by Aldous and Steele (2004).

Keywords: maximum subgraph density; load balancing; local weak convergence; objective method; unimodularity; pairing model.

2010 MSC: 60C05, 05C80, 90B15.

1 Introduction

Balanced allocations Let $G = (V, E)$ be a simple undirected locally finite graph. Write \vec{E} for the set of oriented edges, i.e. ordered pairs of adjacent vertices. An *allocation* on G is a map $\theta: \vec{E} \rightarrow [0, 1]$ satisfying $\theta(i, j) + \theta(j, i) = 1$ for every $\{i, j\} \in E$. The *load* induced by θ at a vertex $o \in V$ is

$$\partial\theta(o) := \sum_{i \sim o} \theta(i, o),$$

where \sim means adjacency in G . θ is *balanced* if for every $(i, j) \in \vec{E}$,

$$\partial\theta(i) < \partial\theta(j) \implies \theta(i, j) = 0. \tag{1}$$

*EECS Department, University of California, Berkeley

†Laboratoire de Probabilités et Modèles Aléatoires, Université Paris Diderot

Intuitively, one may think of each edge as carrying a unit amount of load, which has to be distributed over its end-points in such a way that the total load is as much balanced as possible across the graph. In that respect, (1) is a *local optimality* criterion: modifying the allocation along an edge cannot further reduce the load imbalance between its end-points. When G is finite, this condition happens to guarantee *global optimality* in a very strong sense. Specifically, the following conditions are equivalent (see [16]).

- (i) θ is balanced.
- (ii) θ minimizes $\sum_{o \in V} f(\partial\theta(o))$, for some strictly convex $f: [0, \infty) \rightarrow \mathbb{R}$.
- (iii) θ minimizes $\sum_{o \in V} f(\partial\theta(o))$, for every convex $f: [0, \infty) \rightarrow \mathbb{R}$.

In particular, balanced allocations exist on G and they all induce the same loads $\partial\theta: V \rightarrow [0, \infty)$.

The densest subgraph problem. The load $\partial\theta(o)$ induced at a vertex $o \in V$ by some (hence every) balanced allocation θ on G has a remarkable graph-theoretical interpretation: it measures the *local density* of G at o . In particular, it was shown in [16] that the vertices receiving the highest load in G solve the classical densest subgraph problem: the value $\max \partial\theta$ coincides with the *maximum subgraph density* of a subgraph in G ,

$$\varrho(G) := \max_{\emptyset \subsetneq H \subseteq V} \frac{|E(H)|}{|H|},$$

and the set $H = \operatorname{argmax} \partial\theta$ is the largest set achieving this maximum. This surprising connection with a well-known and important graph parameter justifies a deeper study of balanced loads in large graphs. For this purpose, it is convenient to encode the various loads of G into a probability measure on \mathbb{R} , called the *empirical load distribution* of G :

$$\mathcal{L}_G = \frac{1}{|V|} \sum_{o \in V} \delta_{\partial\theta(o)}.$$

Conjecture in the Erdős-Rényi case. Motivated by the above connection, Hajek [16] studied the asymptotic behavior of \mathcal{L}_G on the popular Erdős-Rényi model, where the graph $G = G_n$ is chosen uniformly at random among all graphs with $m = \lfloor \alpha n \rfloor$ edges on $V = \{1, \dots, n\}$. In the regime where the density parameter $\alpha \geq 0$ is kept fixed while $n \rightarrow \infty$, he conjectured that the

empirical load distribution \mathcal{L}_{G_n} should concentrate around a deterministic probability measure $\mathcal{L} \in \mathcal{P}(\mathbb{R})$, in the sense that

$$\mathcal{L}_{G_n} \xrightarrow[n \rightarrow \infty]{\mathcal{P}(\mathbb{R})} \mathcal{L}.$$

Coming back to the densest subgraph problem and despite the non-continuity of the essential supremum w.r.t. to weak convergence, Hajek conjectured that

$$\varrho(G_n) \xrightarrow[n \rightarrow \infty]{\mathbb{P}} \varrho := \sup\{t \in \mathbb{R} : \mathcal{L}([0, t]) < 1\}.$$

Finally, using a non-rigorous analogy with the case of finite trees, Hajek proposed a description of \mathcal{L} and ϱ in terms of the solutions to a distributional fixed-point equation which will be given later. In this paper, we establish this triple conjecture together with its analogue for various other sparse random graphs, using the unifying framework of *local weak convergence*.

2 The framework of local weak convergence

This section gives a brief account of the framework of local weak convergence. For more details, we refer to the seminal paper [6] and to the surveys [3, 2].

Rooted graphs. A *rooted graph* (G, o) is a graph $G = (V, E)$ together with a distinguished vertex $o \in V$, called the *root*. We let \mathcal{G}_\star denote the set of all locally finite connected rooted graphs considered up to *rooted isomorphism*, i.e. $(G, o) \equiv (G', o')$ if there exists a bijection $\gamma: V \rightarrow V'$ that preserves roots ($\gamma(o) = o'$) and adjacency ($\{i, j\} \in E \iff \{\gamma(i), \gamma(j)\} \in E'$). We write $[G, o]_h$ for the (finite) rooted subgraph induced by the vertices lying at graph-distance at most $h \in \mathbb{N}$ from o . The distance

$$\text{DIST}((G, o), (G', o')) := \frac{1}{1+r} \text{ where } r = \sup\{h \in \mathbb{N} : [G, o]_h \equiv [G', o']_h\},$$

turns \mathcal{G}_\star into a complete separable metric space, see [2].

Local weak limits. Let $\mathcal{P}(\mathcal{G}_\star)$ denote the set of Borel probability measures on \mathcal{G}_\star , equipped with the usual topology of weak convergence (see e.g. [7]). Given a finite deterministic graph $G = (V, E)$, we construct a random element of \mathcal{G}_\star by choosing uniformly at random a vertex $o \in V$ to be the root, and restricting G to the connected component of o . The resulting law is denoted by $\mathcal{U}(G)$. If $\{G_n\}_{n \geq 1}$ is a sequence of finite graphs such that $\{\mathcal{U}(G_n)\}_{n \geq 1}$ admits a weak limit $\mu \in \mathcal{P}(\mathcal{G}_\star)$, we call μ the *local weak limit* of $\{G_n\}_{n \geq 1}$.

Edge-rooted graphs. Let \mathcal{G}_{**} denote the space of locally finite connected graphs with a distinguished oriented edge, taken up to the natural isomorphism relation and equipped with the natural distance, which turns it into a complete separable metric space. With any function $f: \mathcal{G}_{**} \rightarrow \mathbb{R}$ is naturally associated a function $\partial f: \mathcal{G}_* \rightarrow \mathbb{R}$, defined by

$$\partial f(G, o) = \sum_{i \sim o} f(G, i, o).$$

Dually, with any measure $\mu \in \mathcal{P}(\mathcal{G}_*)$ is naturally associated a non-negative measure $\vec{\mu}$ on \mathcal{G}_{**} , defined as follows: for any Borel function $f: \mathcal{G}_{**} \rightarrow [0, \infty)$,

$$\int_{\mathcal{G}_{**}} f d\vec{\mu} = \int_{\mathcal{G}_*} (\partial f) d\mu.$$

Note that the total mass $\vec{\mu}(\mathcal{G}_{**})$ of the measure $\vec{\mu}$ is precisely

$$\text{deg}(\mu) := \int_{\mathcal{G}_*} \text{deg}(G, o) d\mu(G, o).$$

Unimodularity. It was shown in [2] that any $\mu \in \mathcal{P}(\mathcal{G}_*)$ arising as the local weak limit of some sequence of finite graphs satisfies

$$\int_{\mathcal{G}_{**}} f d\vec{\mu} = \int_{\mathcal{G}_{**}} f^* d\vec{\mu},$$

for any Borel $f: \mathcal{G}_{**} \rightarrow [0, \infty)$. Here, $f^*: \mathcal{G}_{**} \rightarrow \mathbb{R}$ denotes the *reversal* of f :

$$f^*(G, i, o) = f(G, o, i).$$

A measure $\mu \in \mathcal{P}(\mathcal{G}_*)$ satisfying this invariance is called *unimodular*, and the set of all unimodular probability measures on \mathcal{G}_* is denoted by \mathcal{U} .

Marks on oriented edges. It will sometimes be convenient to work with *networks*, i.e. graphs equipped with a map from \vec{E} to some fixed complete separable metric space Ξ . The above definitions extend naturally, see [2].

Unimodular Galton-Watson trees. Let $\pi = \{\pi_n\}_{n \geq 0}$ be a probability distribution on \mathbb{N} with non-zero finite mean, and let $\hat{\pi} = \{\hat{\pi}_n\}_{n \geq 0}$ denote its size-biased version:

$$\hat{\pi}_n = \frac{(n+1)\pi_{n+1}}{\sum_k k\pi_k} \quad (n \in \mathbb{N}) \quad (2)$$

A *unimodular Galton-Watson tree* with degree distribution π is a random rooted tree \mathbb{T} obtained by a Galton-Watson branching process where the root has offspring distribution π and all its descendants have offspring distribution $\hat{\pi}$. The law of \mathbb{T} is unimodular, and is denoted by $\text{UGWT}(\pi)$. Such trees play a distinguished role as they are the local weak limits of many natural sequences of random graphs, including those produced by the pairing model.

The pairing model. Given a sequence $\mathbf{d} = \{d(i)\}_{1 \leq i \leq n}$ of non-negative integers whose sum is even, the pairing model [8, 19] generates a random graph $\mathbb{G}[\mathbf{d}]$ on $V = \{1, \dots, n\}$ as follows: $d(i)$ *half-edges* are attached to each $i \in V$, and the $2m = d(1) + \dots + d(n)$ half-edges are paired uniformly at random to form m edges. Loops and multiple edges are removed (a few variants exist, see [18], but they are equivalent for our purpose). Now, consider a degree sequence $\mathbf{d}_n = \{d_n(i)\}_{1 \leq i \leq n}$ for each $n \geq 1$ and assume that

$$\forall k \in \mathbb{N}, \quad \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\{d_n(i)=k\}} \xrightarrow[n \rightarrow \infty]{} \pi_k, \quad (3)$$

for some probability distribution $\pi = \{\pi_k\}_{k \in \mathbb{N}}$ on \mathbb{N} with finite, non-zero mean. Under the additional assumption that

$$\sup_{n \geq 1} \left\{ \frac{1}{n} \sum_{i=1}^n d_n^2(i) \right\} < \infty,$$

the local weak limit of $\{\mathbb{G}[\mathbf{d}_n]\}_{n \geq 1}$ is $\mu := \text{UGWT}(\pi)$ almost-surely, see [9].

3 Main results

Balanced loads on unimodular random graphs. Our first main result is that the notion of balanced allocations can be extended from finite graphs to their local weak limits, in such a way that the induced loads behave continuously with respect to local weak convergence. Define a *Borel allocation* as a measurable function $\Theta: \mathcal{G}_{**} \rightarrow [0, 1]$ such that $\Theta + \Theta^* = 1$, and call it *balanced* on $\mu \in \mathcal{U}$ if for $\vec{\mu}$ -a-e $(G, i, o) \in \mathcal{G}_{**}$,

$$\partial\Theta(G, i) < \partial\Theta(G, o) \implies \Theta(G, i, o) = 0.$$

This natural definition is the right analogue of (1) when finite graphs are replaced by unimodular measures, as demonstrated by the following result.

Theorem 1. *Let $\mu \in \mathcal{U}$ be such that $\deg(\mu) < \infty$. Then,*

1. Existence. *There is a Borel allocation Θ_0 that is balanced on μ .*
2. Optimality. *For any Borel allocation Θ , the following are equivalent:*
 - (i) Θ is balanced on μ .
 - (ii) Θ minimizes $\int f \circ \partial\Theta d\mu$ for some strictly convex $f: [0, \infty) \rightarrow \mathbb{R}$.
 - (iii) Θ minimizes $\int f \circ \partial\Theta d\mu$ for every convex $f: [0, \infty) \rightarrow \mathbb{R}$.
 - (iv) $\partial\Theta = \partial\Theta_0$, μ -a-e.
3. Continuity. *For any sequence $\{G_n\}_{n \geq 1}$ with local weak limit μ ,*

$$\mathcal{L}_{G_n} \xrightarrow[n \rightarrow \infty]{\mathcal{P}(\mathbb{R})} \mathcal{L},$$

where \mathcal{L} denotes the law of the random variable $\partial\Theta_0 \in L^1(\mathcal{G}_*, \mu)$.

4. Variational characterization. *The mean-excess function of the random variable $\partial\Theta_0$, namely $\Phi_\mu: t \mapsto \int_{\mathcal{G}_*} (\partial\Theta_0 - t)^+ d\mu$, is given by*

$$\Phi_\mu(t) = \max_{\substack{f: \mathcal{G}_* \rightarrow [0, 1] \\ \text{Borel}}} \left\{ \frac{1}{2} \int_{\mathcal{G}_{**}} \widehat{f} d\vec{\mu} - t \int_{\mathcal{G}_*} f d\mu \right\}, \quad (t \in \mathbb{R})$$

where $\widehat{f}(G, i, o) := f(G, o) \wedge f(G, i)$.

The special case of unimodular Galton-Watson trees. Our second main result is an explicit resolution of the above variational problem in the important special case where $\mu = \text{UGWT}(\pi)$, for an arbitrary degree distribution $\pi = \{\pi_n\}_{n \geq 0}$ on \mathbb{N} with finite, non-zero mean. Throughout the paper, we let $[x]_0^1$ denote the closest point to $x \in \mathbb{R}$ in the interval $[0, 1]$, i.e.

$$[x]_0^1 := \begin{cases} 0 & \text{if } x \leq 0 \\ x & \text{if } x \in [0, 1] \\ 1 & \text{if } x \geq 1. \end{cases}$$

Given $t \in \mathbb{R}$ and $Q \in \mathcal{P}([0, 1])$, we let $F_{\pi, t}(Q) \in \mathcal{P}([0, 1])$ denote the law of

$$[1 - t + \xi_1 + \cdots + \xi_{\widehat{D}}]_0^1,$$

where \widehat{D} follows the size-biased distribution $\widehat{\pi}$ defined at (2), and where $\{\xi_n\}_{n \geq 1}$ are IID with law Q , independent of \widehat{D} . As it turns out, the solutions to the distributional fixed-point equation $Q = F_{\pi, t}(Q)$ determine $\Phi_\mu(t)$.

Theorem 2. *When $\mu = \text{UGWT}(\pi)$, we have for every $t \in \mathbb{R}$:*

$$\Phi_\mu(t) = \max_{Q=F_{\pi,t}(Q)} \left\{ \frac{\mathbb{E}[D]}{2} \mathbb{P}(\xi_1 + \xi_2 > 1) - t \mathbb{P}(\xi_1 + \dots + \xi_D > t) \right\},$$

where $D \sim \pi$ and where $\{\xi_n\}_{n \geq 1}$ are IID with law Q , independent of D . The maximum is over all choices of $Q \in \mathcal{P}([0, 1])$ subject to $Q = F_{\pi,t}(Q)$.

Back to the densest subgraph problem. By analogy with the case of finite graphs, we define the maximum subgraph density of a measure $\mu \in \mathcal{U}$ with $\deg(\mu) < \infty$ as the essential supremum of the random variable $\partial\Theta_0$ constructed in Theorem 1. In other words,

$$\varrho(\mu) := \sup\{t \in \mathbb{R} : \Phi_\mu(t) > 0\}.$$

In light of Theorem 1, it is natural to seek a continuity principle of the form

$$\left(G_n \xrightarrow[n \rightarrow \infty]{\text{LWC}} \mu \right) \implies \left(\varrho(G_n) \xrightarrow[n \rightarrow \infty]{} \varrho(\mu) \right). \quad (4)$$

However, a moment of thought shows that the graph parameter $\varrho(G)$ is too sensitive to be captured by local weak convergence. Indeed, if $|V(G_n)| \rightarrow \infty$, then adding a large but fixed clique to G_n will arbitrarily boost the value of $\varrho(G_n)$ without affecting the local weak limit of $\{G_n\}_{n \geq 1}$. Nevertheless, our third main result states that (4) holds for graphs produced by the pairing model, under a mild exponential moment assumption.

Theorem 3. *Consider a degree sequence $\mathbf{d}_n = \{d_n(i)\}_{1 \leq i \leq n}$ for each $n \geq 1$. Assume that (3) holds for some $\pi = \{\pi_k\}_{k \geq 1}$ with $\pi_0 + \pi_1 < 1$, and that*

$$\sup_{n \geq 1} \left\{ \frac{1}{n} \sum_{i=1}^n e^{\theta d_n(i)} \right\} < \infty, \quad (5)$$

for some $\theta > 0$. Then, $\varrho(\mathbb{G}[\mathbf{d}_n]) \xrightarrow[n \rightarrow \infty]{\mathbb{P}} \varrho(\mu)$ with $\mu = \text{UGWT}(\pi)$.

Note that this result applies in particular to the Erdős-Rényi random graph \mathbb{G}_n with n vertices and $m = \lfloor \alpha n \rfloor$ edges. Indeed, the conditional law of \mathbb{G}_n given its (random) degree sequence \mathbf{d}_n is precisely $\mathbb{G}[\mathbf{d}_n]$, and $\{\mathbf{d}_n\}_{n \geq 1}$ satisfies a-s the conditions (3) and (5) with $\pi = \text{Poisson}(2\alpha)$. Therefore, Theorems 1, 2 and 3 settle and generalize the conjectures of Hajek [17]. Since a graph G is k -orientable ($k \in \mathbb{N}$) if and only if $\varrho(G) < k$, Theorem 3 also extends recent results on the k -orientability of the Erdős-Rényi random graph [12, 11]. See also [15, 13, 22, 21] for various generalizations.

4 Proof outline and related work.

The objective method. This work is a new illustration of the general principles exposed in the *objective method* by Aldous and Steele [3]. The latter provides a powerful framework for the unified study of sparse random graphs and has already led to several remarkable results. Two prototypical examples are the celebrated $\zeta(2)$ limit in the random assignment problem due to Aldous [4], and the asymptotic enumeration of spanning trees in large graphs by Lyons [24]. Since then, the method has been successfully applied to various other combinatorial enumeration/optimization problems on graphs, including (but not limited to) [28, 14, 26, 10, 22, 25, 21, 20].

Lack of correlation decay. In the problem considered here, there is a major obstacle to a straightforward application of the objective method: the balanced load at a vertex is *not* determined by the local environment around that vertex. For example, a ball of radius h in a d -regular graph with girth h is indistinguishable from that of the root in a d -regular tree with height h . However, the induced load is $\frac{d}{2}$ in the first case and $1 - \frac{1}{(d-1)^{h-1}d}$ in the second case. This long-range dependence gives rise to non-uniqueness issues when one tries to properly extend the notion of balanced loads from finite to infinite graphs. We refer the reader to [17] for a detailed study of this phenomenon, therein called *load percolation*, as well as several related questions.

Relaxation. To overcome the lack of correlation decay, we introduce a suitable relaxation of the balancing condition (1), which we call ε -balancing. Remarkably enough, any positive value of the perturbative parameter ε suffices to annihilate the long-range dependences described above. This allows us to define a unique ε -balanced Borel allocation $\Theta_\varepsilon: \mathcal{G}_{**} \rightarrow [0, 1]$ and to establish the continuity of the induced load $\partial\Theta_\varepsilon: \mathcal{G}_* \rightarrow [0, \infty)$ with respect to local convergence (Section 5). In the limit where ε tends to 0, we further prove that Θ_ε converges in a certain sense, and that the limiting Borel allocation Θ_0 is balanced (Section 6). This quickly leads to a proof of Theorem 1 (Section 7). In spirit, the role of the perturbative parameter $\varepsilon > 0$ is comparable to that of the temperature in [10], although no Gibbs-Boltzmann measure is involved in the present work.

Recursion on trees. As many other graph-theoretical problems, load balancing has a simple recursive structure when considered on trees. Indeed, once the value of the allocation along a given edge $\{i, j\}$ has been fixed, the problem naturally decomposes into two independent sub-problems, cor-

responding to the two disjoint subtrees formed by removing $\{i, j\}$. Note, however, that the loads of i and j must be shifted by a suitable amount to take into account the contribution of the removed edge. The precise effect of this shift on the loads induced at i and j defines what we call the *response functions* of the two subtrees (Section 8). It is those response functions that satisfy a recursion (Section 9). Recursions on trees automatically give rise to distributional fixed-point equations when specialized to Galton-Watson trees. Such equations are a common ingredient in the objective method, see [5]. In our case this leads to the proof of Theorem 2 (Section 10).

Dense subgraphs in the pairing model. Finally, the proof of Theorem 3 (Section 11) relies on a property of random graphs with a prescribed degree sequence that might be of independent interest: under an exponential moment assumption, we use the first-moment method to prove that dense subgraphs are extensively large with high probability. See Proposition 12 for the precise statement, and [23, Lemma 6] for a result in the same direction.

5 ε -balancing

In all this section, $G = (V, E)$ is a locally finite graph and $\varepsilon > 0$ is a fixed parameter. An allocation θ on G is called ε -balanced if for every $(i, j) \in \vec{E}$,

$$\theta(i, j) = \left[\frac{1}{2} + \frac{\partial\theta(i) - \partial\theta(j)}{2\varepsilon} \right]_0^1. \quad (6)$$

This can be viewed as a relaxed version of (1). Its interest lies in the fact that it fixes the non-uniqueness issue on infinite graphs.

Proposition 1 (Existence, uniqueness and monotony). *If G has bounded degrees, then there is a unique ε -balanced allocation θ on G . If moreover $E' \subseteq E$, then the ε -balanced allocation θ' on $G' = (V, E')$ satisfies $\partial\theta' \leq \partial\theta$.*

Proof. Existence is a consequence of Schauder's fixed-point Theorem, see e.g. [1, Theorem 8.2]. Now, consider $E' \subseteq E$ and let θ, θ' be ε -balanced allocations on G, G' respectively. Fix $o \in V$ and set

$$I := \{i \in V : \{i, o\} \in E', \theta'(i, o) > \theta(i, o)\}.$$

Clearly,

$$\partial\theta'(o) - \partial\theta(o) \leq \sum_{i \in I} (\theta'(i, o) - \theta(i, o)).$$

On the other-hand, since the map $x \mapsto \left[\frac{1}{2} + \frac{x}{2\varepsilon}\right]_0^1$ is non-decreasing and Lipschitz with constant $\frac{1}{2\varepsilon}$, our assumption on θ, θ' implies that for every $i \in I$,

$$\theta'(i, o) - \theta(i, o) \leq \frac{1}{2\varepsilon} (\partial\theta'(i) - \partial\theta(i) - \partial\theta'(o) + \partial\theta(o)).$$

Injecting this into the above inequality and rearranging, we obtain

$$\begin{aligned} \partial\theta'(o) - \partial\theta(o) &\leq \frac{1}{|I| + 2\varepsilon} \sum_{i \in I} (\partial\theta'(i) - \partial\theta(i)) \\ &\leq \frac{\Delta}{\Delta + 2\varepsilon} \max_{i \in I} (\partial\theta'(i) - \partial\theta(i)), \end{aligned} \quad (7)$$

where Δ denotes the maximum degree in G . Now, observe that $\partial\theta, \partial\theta'$ are $[0, \Delta]$ -valued, so that $M := \sup_V (\partial\theta' - \partial\theta)$ is finite. Property (7) forces $M \leq 0$, which proves the monotony $E' \subseteq E \implies \partial\theta' \leq \partial\theta$. In particular, $E' = E$ implies $\partial\theta' = \partial\theta$, which in turns forces $\theta' = \theta$, thanks to (6). \square

We now remove the bounded-degree assumption as follows. Fix $\Delta \in \mathbb{N}$, and consider the truncated graph $G^\Delta = (V, E^\Delta)$ obtained from G by isolating all nodes having degree more than Δ , i.e.

$$E^\Delta = \{\{i, j\} \in E : \deg(G, i) \vee \deg(G, j) \leq \Delta\}.$$

By construction, G^Δ has degree at most Δ , and we let $\Theta_\varepsilon^\Delta(G, i, j)$ denote the mass sent along $(i, j) \in \vec{E}$ in the unique ε -balanced allocation on G^Δ , with the understanding that $\Theta_\varepsilon^\Delta(G, i, j) = 0$ if $\{i, j\} \notin E^\Delta$. By uniqueness, this quantity depends only on the isomorphism class of the edge-rooted graph (G, i, j) , so that we have a well-defined map $\Theta_\varepsilon^\Delta: \mathcal{G}_{**} \rightarrow [0, 1]$. By an immediate induction on $r \in \mathbb{N}$, the local contraction (7) yields

$$[G, o]_r \equiv [G', o']_r \implies |\partial\Theta_\varepsilon^\Delta(G, o) - \partial\Theta_\varepsilon^\Delta(G', o')| \leq \Delta \left(1 + \frac{2\varepsilon}{\Delta}\right)^{-r}.$$

Since the map $x \mapsto \left[\frac{1}{2} + \frac{x}{2\varepsilon}\right]_0^1$ is Lipschitz with constant $\frac{1}{2\varepsilon}$, it follows that

$$[G, i, j]_r \equiv [G', i', j']_r \implies |\Theta_\varepsilon^\Delta(G, i, j) - \Theta_\varepsilon^\Delta(G', i', j')| \leq \frac{\Delta}{2\varepsilon} \left(1 + \frac{2\varepsilon}{\Delta}\right)^{-r}.$$

Thus, the map $\Theta_\varepsilon^\Delta$ is equicontinuous. Now, the sequence of sets $\{E_\Delta\}_{\Delta \geq 1}$ increases to E , so the monotony in Proposition 1 guarantees that $\{\partial\Theta_\varepsilon^\Delta\}_{\Delta \geq 1}$ converges pointwise on \mathcal{G}_* . Moreover, any given $\{i, j\} \in E$ belongs to E^Δ for large enough Δ , and the definition of ε -balancing yields

$$\Theta_\varepsilon^\Delta(G, i, j) = \left[\frac{1}{2} + \frac{\partial\Theta_\varepsilon^\Delta(G, i) - \partial\Theta_\varepsilon^\Delta(G, j)}{2\varepsilon}\right]_0^1.$$

Consequently, the pointwise limit $\Theta_\varepsilon := \lim_{\Delta \rightarrow \infty} \Theta_\varepsilon^\Delta$ exists in $[0, 1]^{\mathcal{G}_{**}}$. It clearly satisfies $\Theta_\varepsilon + \Theta_\varepsilon^* = 1$ and it is Borel as the pointwise limit of continuous maps. Thus, it is a Borel allocation. Moreover, letting $\Delta \rightarrow \infty$ above yields

$$\Theta_\varepsilon(G, i, j) = \left[\frac{1}{2} + \frac{\partial\Theta_\varepsilon(G, i) - \partial\Theta_\varepsilon(G, j)}{2\varepsilon} \right]_0^1. \quad (8)$$

6 The $\varepsilon \rightarrow 0$ limit

In this section, we fix $\mu \in \mathcal{U}$ with $\deg(\mu) < \infty$. We write $\|f\|_p$ for the norm in both $L^p(\mu)$ and $L^p(\bar{\mu})$, as which one is meant should be clear from the context. Note that by unimodularity, we have for any Borel allocation Θ ,

$$\|\Theta\|_1 = \int_{\mathcal{G}_{**}} \Theta d\bar{\mu} = \int_{\mathcal{G}_{**}} \frac{\Theta + \Theta^*}{2} d\bar{\mu} = \frac{\deg(\mu)}{2}. \quad (9)$$

Proposition 2. *The limit $\Theta_0 := \lim_{\varepsilon \rightarrow 0} \Theta_\varepsilon$ exists in $L^2(\bar{\mu})$ and is a balanced Borel allocation on μ .*

Proof. We will establish the following Cauchy property: for $0 < \varepsilon \leq \varepsilon'$,

$$\|\Theta_{\varepsilon'} - \Theta_\varepsilon\|_2^2 \leq \|\Theta_\varepsilon\|_2^2 - \|\Theta_{\varepsilon'}\|_2^2. \quad (10)$$

This guarantees the existence of $\Theta_0 = \lim_{\varepsilon \rightarrow 0} \Theta_\varepsilon$ in $L^2(\bar{\mu})$. The rest of the claim follows, since Borel allocations are closed in $L^2(\bar{\mu})$ and letting $\varepsilon \rightarrow 0$ in (8) shows that Θ_0 is balanced on μ . Let us first prove (10) under the extra assumption that $\mu(\{(G, o) : \deg(G, o) \leq \Delta\}) = 1$ for some $\Delta \in \mathbb{N}$. This ensures that $f \in L^2(\bar{\mu})$, where

$$f(G, i, o) := \partial\Theta_\varepsilon(G, o) + \varepsilon\Theta_\varepsilon(G, i, o).$$

A straightforward manipulation of (8) shows that

$$f(G, i, o) > f(G, o, i) \implies \Theta_\varepsilon(G, i, o) = 0.$$

This implies $\Theta_\varepsilon f + \Theta_\varepsilon^* f^* = f \wedge f^*$. On the other hand, $f \wedge f^* \leq \Theta_{\varepsilon'} f + \Theta_{\varepsilon'}^* f^*$ since $\Theta_{\varepsilon'} + \Theta_{\varepsilon'}^* = 1$. Thus, $\Theta_\varepsilon f + \Theta_\varepsilon^* f^* \leq \Theta_{\varepsilon'} f + \Theta_{\varepsilon'}^* f^*$. Integrating against $\bar{\mu}$ and invoking unimodularity, we get $\langle \Theta_\varepsilon - \Theta_{\varepsilon'}, f \rangle_{L^2(\bar{\mu})} \leq 0$ or more explicitly,

$$\langle \partial\Theta_\varepsilon - \partial\Theta_{\varepsilon'}, \partial\Theta_\varepsilon \rangle_{L^2(\mu)} + \varepsilon \langle \Theta_\varepsilon - \Theta_{\varepsilon'}, \Theta_\varepsilon \rangle_{L^2(\bar{\mu})} \leq 0.$$

But we have not yet used $\varepsilon \leq \varepsilon'$, so we may exchange $\varepsilon, \varepsilon'$ to get

$$\langle \partial\Theta_{\varepsilon'} - \partial\Theta_\varepsilon, \partial\Theta_{\varepsilon'} \rangle_{L^2(\mu)} + \varepsilon' \langle \Theta_{\varepsilon'} - \Theta_\varepsilon, \Theta_{\varepsilon'} \rangle_{L^2(\bar{\mu})} \leq 0.$$

Adding-up those inequalities and rearranging, we finally arrive at

$$(\varepsilon' - \varepsilon)\langle \Theta_\varepsilon - \Theta_{\varepsilon'}, \Theta_{\varepsilon'} \rangle_{L^2(\bar{\mu})} \geq \|\partial\Theta_\varepsilon - \partial\Theta_{\varepsilon'}\|_2^2 + \varepsilon\|\Theta_\varepsilon - \Theta_{\varepsilon'}\|_2^2.$$

In particular, $\langle \Theta_\varepsilon, \Theta_{\varepsilon'} \rangle_{L^2(\bar{\mu})} \geq \|\Theta_{\varepsilon'}\|_2^2$ and (10) follows since

$$\|\Theta_{\varepsilon'} - \Theta_\varepsilon\|_2^2 = \|\Theta_{\varepsilon'}\|_2^2 + \|\Theta_\varepsilon\|_2^2 - 2\langle \Theta_\varepsilon, \Theta_{\varepsilon'} \rangle_{L^2(\bar{\mu})}.$$

Finally, if our extra assumption is dropped, we may use (10) with $\Theta_\varepsilon, \Theta_{\varepsilon'}$ replaced by $\Theta_\varepsilon^\Delta, \Theta_{\varepsilon'}^\Delta$ and let then $\Delta \rightarrow \infty$. By construction, $\Theta_\varepsilon^\Delta \rightarrow \Theta_\varepsilon$ and $\Theta_{\varepsilon'}^\Delta \rightarrow \Theta_{\varepsilon'}$ pointwise, and (10) follows by dominated convergence. \square

Proposition 3. *Let $\{G_n\}_{n \geq 1}$ be finite graphs with local weak limit μ . Then,*

$$\mathcal{L}_{G_n} \xrightarrow[n \rightarrow \infty]{\mathcal{P}(\mathbb{R})} \mathcal{L},$$

where $\mathcal{L} = \mathcal{L}_\mu$ is the law of the random variable $\partial\Theta_0 \in L^1(\mu)$.

Proof. For each $n \geq 1$ we let \widehat{G}_n denote the network obtained by encoding a balanced allocation θ_n as $[0, 1]$ -valued marks on the oriented edges of G_n . The sequence $\{\mathcal{U}(\widehat{G}_n)\}_{n \geq 1}$ is tight, because $\{\mathcal{U}(G_n)\}_{n \geq 1}$ converges weakly and the marks are $[0, 1]$ -valued. Consider any subsequential weak limit (\mathbb{G}, o, θ) . By construction, (\mathbb{G}, o) has law μ and θ is a-s a balanced allocation on \mathbb{G} . Our goal is to establish that a-s, $\partial\theta(o) = \partial\Theta_0(\mathbb{G}, o)$. Set $\theta'(i, j) := \Theta_0(\mathbb{G}, i, j)$. Note that the random rooted network $(\mathbb{G}, o, \theta, \theta')$ is unimodular, since (\mathbb{G}, o, θ) is a local weak limit of finite networks and Θ_0 is Borel. Now,

$$\begin{aligned} \mathbb{E} \left[(\partial\theta(o) - \partial\theta'(o))^+ \right] &= \mathbb{E} \left[\sum_{i \sim o} (\theta(i, o) - \theta'(i, o)) \mathbf{1}_{\partial\theta(o) > \partial\theta'(o)} \right] \\ &= \mathbb{E} \left[\sum_{i \sim o} (\theta(o, i) - \theta'(o, i)) \mathbf{1}_{\partial\theta(o) > \partial\theta'(o)} \right] \\ &= \mathbb{E} \left[\sum_{i \sim o} (\theta'(i, o) - \theta(i, o)) \mathbf{1}_{\partial\theta(o) > \partial\theta'(o)} \right], \end{aligned}$$

where the second equality follows from unimodularity and the third one from the identities $\theta(o, i) = 1 - \theta(i, o)$ and $\theta'(o, i) = 1 - \theta'(i, o)$. Combining the first and last lines, we see that $\mathbb{E} \left[(\partial\theta(o) - \partial\theta'(o))^+ \right]$ equals

$$\frac{1}{2} \mathbb{E} \left[\sum_{i \sim o} (\theta(i, o) - \theta'(i, o)) (\mathbf{1}_{\partial\theta(o) > \partial\theta'(o)} - \mathbf{1}_{\partial\theta(i) > \partial\theta'(i)}) \right].$$

The fact that θ, θ' are balanced on the edge $\{i, o\}$ easily implies that $\theta(i, o) - \theta'(i, o)$ and $\mathbf{1}_{\partial\theta(o) > \partial\theta'(o)} - \mathbf{1}_{\partial\theta(i) > \partial\theta'(i)}$ can neither be simultaneously positive, nor simultaneously negative. Therefore, $\mathbb{E} \left[(\partial\theta(o) - \partial\theta'(o))^+ \right] \leq 0$. Exchanging the roles of θ, θ' yields $\partial\theta(o) = \partial\theta'(o)$ a-s, as desired. \square

7 Proof of Theorem 1

Proposition 4. *Let Θ be a Borel allocation. Then for all $t \in \mathbb{R}$,*

$$\int_{\mathcal{G}_*} (\partial\Theta - t)^+ d\mu \geq \sup_{\substack{f: \mathcal{G}_* \rightarrow [0,1] \\ \text{Borel}}} \left\{ \frac{1}{2} \int_{\mathcal{G}_{**}} \widehat{f} d\bar{\mu} - t \int_{\mathcal{G}_*} f d\mu \right\},$$

with equality for all $t \in \mathbb{R}$ if and only if Θ is balanced on μ .

Proof. Fix a Borel $f: \mathcal{G}_* \rightarrow [0, 1]$. Clearly, $(\partial\Theta - t)^+ \geq (\partial\Theta - t)f$ and hence

$$\int_{\mathcal{G}_*} (\partial\Theta - t)^+ d\mu \geq \int_{\mathcal{G}_*} f \partial\Theta d\mu - t \int_{\mathcal{G}_*} f d\mu. \quad (11)$$

Using the unimodularity of μ and the identity $\Theta + \Theta^* = 1$, we have

$$\begin{aligned} \int_{\mathcal{G}_*} f \partial\Theta d\mu &= \frac{1}{2} \int_{\mathcal{G}_{**}} (f(G, o)\Theta(G, i, o) + f(G, i)\Theta(G, o, i)) d\bar{\mu}(G, i, o) \\ &\geq \frac{1}{2} \int_{\mathcal{G}_{**}} (f(G, o) \wedge f(G, i)) d\bar{\mu}(G, i, o). \end{aligned} \quad (12)$$

Combining (11) and (12) proves the inequality. Let us now examine the equality case. First, equality holds in (11) if and only if for μ -a-e $(G, o) \in \mathcal{G}_*$,

$$\begin{aligned} \partial\Theta(G, o) > t &\implies f(G, o) = 1 \\ \partial\Theta(G, o) < t &\implies f(G, o) = 0. \end{aligned}$$

Second, equality holds in (12) if and only if for $\bar{\mu}$ -a-e $(G, i, o) \in \mathcal{G}_{**}$,

$$f(G, i) < f(G, o) \implies \Theta(G, i, o) = 0.$$

If Θ is balanced on μ , then the choice $f = \mathbf{1}_{\{\partial\Theta > t\}}$ clearly satisfies all those requirements, so that equality holds for each $t \in \mathbb{R}$ in the Proposition. This proves the *if* part and shows that the supremum in Proposition 4 is attained, since at least one balanced allocation exists by Proposition 2. Now, for the *only if* part, suppose that equality is achieved in Proposition 4. Then the above requirements imply that for $\bar{\mu}$ -a-e $(G, i, o) \in \mathcal{G}_{**}$,

$$\partial\Theta(G, o) < t < \partial\Theta(G, i) \implies \Theta(G, o, i) = 1.$$

Since this must be true for all $t \in \mathbb{Q}$, it follows that Θ is balanced on μ . \square

Proof of Theorem 1. The existence, the continuity and the variational characterization were established in Proposition 2, 3 and 4, respectively. Now, let Θ, Θ' be Borel allocations, and assume that Θ is balanced. Applying Proposition 4 to Θ and Θ' shows that for all $t \in \mathbb{R}$,

$$\int_{\mathcal{G}_*} (\partial\Theta - t)^+ d\mu \leq \int_{\mathcal{G}_*} (\partial\Theta' - t)^+ d\mu.$$

On the other-hand, (9) guarantees that $\partial\Theta, \partial\Theta'$ have the same mean. Those two conditions together are well-known to be equivalent to the convex ordering $\partial\Theta \preceq_{cx} \partial\Theta'$ (see e.g. [27]), meaning that for any convex $f: [0, \infty) \rightarrow \mathbb{R}$,

$$\int_{\mathcal{G}_*} (f \circ \partial\Theta) d\mu \leq \int_{\mathcal{G}_*} (f \circ \partial\Theta') d\mu.$$

We have just proved (i) \implies (iii), and (iii) \implies (ii) is obvious. In particular, Θ_0 satisfies (ii) and (iii). The *only if* part of Proposition 4 shows that (iii) \implies (i). The implication (iv) \implies (iii) is obvious given that Θ_0 satisfies (iii). Thus, it only remains to prove (ii) \implies (iv). Assume that Θ minimizes $\int (f \circ \partial\Theta) d\mu$ for some strictly convex function $f: [0, \infty) \rightarrow \mathbb{R}$, and let m denote the value of this minimum. Since Θ_0 satisfies (ii), we also have $\int (f \circ \partial\Theta_0) d\mu = m$. But then $\Theta' := (\Theta_0 + \Theta)/2$ is an allocation and by convexity,

$$\int_{\mathcal{G}_*} (f \circ \partial\Theta') d\mu \leq \int_{\mathcal{G}_*} \frac{(f \circ \partial\Theta) + (f \circ \partial\Theta_0)}{2} d\mu = m.$$

This inequality contradicts the definition of m , unless it is an equality. Since f is strictly convex, this forces $\partial\Theta = \partial\Theta_0$ μ -a-e. \square

8 Response functions

As many other graph-theoretical problems, load balancing has a simple recursive structure when specialized to trees. However, the exact formulation of this recursion requires the possibility to *condition* the allocation to take a certain value at a given edge, and we first need to give a proper meaning to this operation. Let $G = (V, E)$ be a locally finite graph and $b: V \rightarrow \mathbb{R}$ a function called the *baseload*. An allocation θ is *balanced with respect to b* if

$$b(i) + \partial\theta(i) < b(j) + \partial\theta(j) \implies \theta(i, j) = 0,$$

for all $(i, j) \in \vec{E}$. This is precisely the definition of balancing, except that the load *felt* by each vertex $i \in V$ is shifted by a certain amount $b(i)$. Similarly,

θ is ε -balanced with respect to b if for all $(i, j) \in \vec{E}$,

$$\theta(i, j) = \left[\frac{1}{2} + \frac{b(i) + \partial\theta(i) - b(j) - \partial\theta(j)}{2\varepsilon} \right]_0^1. \quad (13)$$

The arguments used in Proposition 1 are easily extended to this situation.

Proposition 5 (Existence, uniqueness and monotony). *If G has bounded degree and if b is bounded, then there is a unique ε -balanced allocation with baseload b . Moreover, if $b' \leq b$ is bounded and if $E' \subseteq E$, then the ε -balanced allocation θ' on $G' = (V, E')$ with baseload b' satisfies $b' + \partial\theta' \leq b + \partial\theta$ on V .*

As in Section 5, we then define an ε -balanced allocation in the general case by considering the truncated graph G^Δ with baseload the truncation of b to $[-\Delta, \Delta]$, and let then $\Delta \rightarrow \infty$. Monotony guarantees the existence of a limiting ε -balanced allocation. We shall need the following property.

Proposition 6 (Non-expansion). *Let θ, θ' be the ε -balanced allocations with baseloads $b, b' : V \rightarrow \mathbb{R}$. Set $f = \partial\theta + b$ and $f' = \partial\theta' + b'$. Then,*

$$\|f' - f\|_{\ell^1(V)} \leq \|b' - b\|_{\ell^1(V)}.$$

Proof. By considering $b'' = b \wedge b'$ and using the triangle inequality, we may assume that $b \leq b'$. Note that this implies $f \leq f'$, thanks to Proposition 5. When G is finite, the claim trivially follows from conservation of mass:

$$\sum_{o \in V} (f'(o) - f(o)) = \sum_{o \in V} (b'(o) - b(o)).$$

This then extends to the case where G has bounded degrees with b, b' bounded as follows: choose finite subsets $V_1 \subseteq V_2 \subseteq \dots$ such that $\cup_{n \geq 1} V_n = V$. For each $n \geq 1$, let θ_n, θ'_n denote the ε -balanced allocations on the subgraph induced by V_n , with baseloads the restrictions of b, b' to V_n . Then $\theta_n \rightarrow \theta$ and $\theta'_n \rightarrow \theta'$ pointwise, by compactness and uniqueness. Now, any finite $K \subseteq V$ is contained in V_n for large enough n , and since V_n is finite we know that $f_n := \partial\theta_n + b$ and $f'_n := \partial\theta'_n + b'$ satisfy

$$\sum_{i \in K} |f'_n(i) - f_n(i)| \leq \sum_{i \in V_n} |b'(i) - b(i)|.$$

Letting $n \rightarrow \infty$ yields the desired result, since K is arbitrary. Finally, for the general case, we may apply the result to the truncated graph G^Δ with baseloads the truncation of b, b' to $[-\Delta, \Delta]$, and let then $\Delta \rightarrow \infty$. \square

Although the uniqueness established in Proposition 5 does not extend to the $\varepsilon = 0$ case, the following weaker result will be useful in the next Section.

Proposition 7 (Weak uniqueness). *Assume that θ, θ' are balanced with respect to b and that $\|\partial\theta - \partial\theta'\|_{\ell^1(V)} < \infty$. Then, $\partial\theta = \partial\theta'$.*

Proof. Fix $\delta > 0$. Then the level set $S := \{j \in V : \partial\theta'(j) - \partial\theta(j) > \delta\}$ must be finite. Therefore, it satisfies the conservation of mass :

$$\sum_{j \in S} \partial\theta'(j) - \partial\theta(j) = \sum_{(i,j) \in E(V-S,S)} \theta'(i,j) - \theta(i,j). \quad (14)$$

Now, if $(i,j) \in E(V-S,S)$ then clearly, $\partial\theta'(i) - \partial\theta(i) < \partial\theta'(j) - \partial\theta(j)$. Consequently, at least one of the following inequalities must hold :

$$b(j) - b(i) < \partial\theta(i) - \partial\theta(j) \quad \text{or} \quad b(j) - b(i) > \partial\theta'(i) - \partial\theta'(j).$$

The first one implies $\theta(i,j) = 1$ and the second $\theta'(i,j) = 0$, since θ, θ' are balanced with respect to b . In either case, we have $\theta'(i,j) \leq \theta(i,j)$. Thus, the right-hand side of (14) is non-positive, hence so must the left-hand side be. This contradicts the definition of S unless $S = \emptyset$, i.e. $\partial\theta' \leq \partial\theta + \delta$. Since δ is arbitrary, we conclude that $\partial\theta' \leq \partial\theta$. Equality follows by symmetry. \square

Given $o \in V$ and $x \in \mathbb{R}$, we set $f_{(G,o)}^\varepsilon(x) = x + \partial\theta(o)$ where θ is the ε -balanced allocation with baseload x at o and 0 elsewhere. We call $f_{(G,o)}^\varepsilon : \mathbb{R} \rightarrow \mathbb{R}$ the *response function* of the rooted graph (G, o) . Propositions 5 and 6 guarantee that $f_{(G,o)}^\varepsilon$ is non-decreasing and non-expansive, i.e.

$$x \leq y \implies 0 \leq f_{(G,o)}^\varepsilon(y) - f_{(G,o)}^\varepsilon(x) \leq y - x. \quad (15)$$

Note for future use that the definition of $f_{(G,o)}^\varepsilon(x)$ also implies

$$0 \leq f_{(G,o)}^\varepsilon(x) - x \leq \deg(G, o). \quad (16)$$

When G is a tree, response functions turn out to satisfy a simple recursion.

9 A recursion on trees

We are now ready to formulate the promised recursion. Fix a tree $T = (V, E)$. Deleting $\{i, j\} \in E$ creates two disjoint subtrees, which we view as rooted at i and j and denote $T_{i \rightarrow j}$ and $T_{j \rightarrow i}$, respectively.

Proposition 8. For any $o \in V$, the response function $\mathfrak{f}_{(T,o)}^\varepsilon$ is invertible and

$$\{\mathfrak{f}_{(T,o)}^\varepsilon\}^{-1} = \text{Id} - \sum_{i \sim o} \left[1 - \{\mathfrak{f}_{T_{i \rightarrow o}}^\varepsilon + \varepsilon(2\text{Id} - 1)\}^{-1} \right]_0^1. \quad (17)$$

Proof. $\mathfrak{f}_{T_{i \rightarrow o}}^\varepsilon + \varepsilon(2\text{Id} - 1)$ increases continuously from \mathbb{R} onto \mathbb{R} , so its inverse $\{\mathfrak{f}_{T_{i \rightarrow o}}^\varepsilon + \varepsilon(2\text{Id} - 1)\}^{-1}$ exists and increases continuously from \mathbb{R} onto \mathbb{R} . Consequently, the function $g: \mathbb{R} \rightarrow \mathbb{R}$ appearing in the right-hand side of (17) is itself continuously increasing from \mathbb{R} onto \mathbb{R} , hence invertible. Given $x \in \mathbb{R}$, it now remains to prove that $t := \mathfrak{f}_{(T,o)}^\varepsilon(x)$ satisfies $g(t) = x$. By definition,

$$t = x + \partial\theta(o), \quad (18)$$

where θ denotes the ε -balanced allocation on T with baseload x at o and 0 elsewhere. Now fix $i \sim o$. The restriction of θ to $T_{i \rightarrow o}$ is clearly an ε -balanced allocation on $T_{i \rightarrow o}$ with baseload $\theta(o, i)$ at i and 0 elsewhere. This is precisely the allocation appearing in the definition of $\mathfrak{f}_{T_{i \rightarrow o}}^\varepsilon(\theta(o, i))$, hence

$$\mathfrak{f}_{T_{i \rightarrow o}}^\varepsilon(\theta(o, i)) = \partial\theta(i).$$

Thus, the fact that θ is ε -balanced along (o, i) may now be rewritten as

$$\theta(o, i) = \left[\frac{1}{2} + \frac{t - \mathfrak{f}_{T_{i \rightarrow o}}^\varepsilon(\theta(o, i))}{2\varepsilon} \right]_0^1. \quad (19)$$

But by definition, $x_i := \{\mathfrak{f}_{T_{i \rightarrow o}}^\varepsilon + \varepsilon(2\text{Id} - 1)\}^{-1}(t)$ is the unique solution to

$$x_i = \frac{1}{2} + \frac{t - \mathfrak{f}_{T_{i \rightarrow o}}^\varepsilon(x_i)}{2\varepsilon}. \quad (20)$$

Comparing (19) and (20), we see that $\theta(o, i) = [x_i]_0^1$, i.e. $\theta(i, o) = [1 - x_i]_0^1$. Re-injecting this into (18), we arrive exactly at the desired $x = g(t)$. \square

In the remainder of this section, we fix a vanishing sequence $\{\varepsilon_n\}_{n \geq 1}$ and study the pointwise limit $\mathfrak{f} = \lim_{n \rightarrow \infty} \mathfrak{f}_{(T,o)}^{\varepsilon_n}$, when it exists. Note that \mathfrak{f} needs not be invertible. However, (15) and (16) guarantee that \mathfrak{f} is non-decreasing with $\mathfrak{f}(\pm\infty) = \pm\infty$, so that it admits a well-defined right-continuous inverse

$$\mathfrak{f}^{-1}(t) := \sup\{x \in \mathbb{R} : \mathfrak{f}(x) \leq t\} \quad (t \in \mathbb{R}).$$

Proposition 9. Assume that $\ell_o := \lim_{n \rightarrow \infty} \partial\Theta_{\varepsilon_n}(T, o)$ exists for each $o \in V$. Then $\mathfrak{f}_{T_{i \rightarrow j}} := \lim_{n \rightarrow \infty} \mathfrak{f}_{T_{i \rightarrow j}}^{\varepsilon_n}$ exists pointwise for each $(i, j) \in \vec{E}$, and

$$\mathfrak{f}_{T_{i \rightarrow j}}^{-1}(t) = t - \sum_{k \sim i, k \neq j} \left[1 - \mathfrak{f}_{T_{k \rightarrow i}}^{-1}(t) \right]_0^1, \quad (21)$$

for every $t \in \mathbb{R}$. Moreover, for every $o \in V$,

$$\ell_o > t \iff \sum_{i \sim o} [1 - \mathfrak{f}_{T_{i \rightarrow o}}^{-1}(t)]_0^1 > t. \quad (22)$$

Proof. Fix $(i, j) \in \vec{E}$, $x \in \mathbb{R}$ and let us show that $\mathfrak{f}_{T_{i \rightarrow j}}(x) := \lim_{n \rightarrow \infty} \mathfrak{f}_{T_{i \rightarrow j}}^{\varepsilon_n}(x)$ exists. By definition, $\mathfrak{f}_{T_{i \rightarrow j}}^\varepsilon(x) = x + \partial\theta_\varepsilon(i)$, where θ_ε is the ε -balanced allocation on $T_{i \rightarrow j}$ with baseload x at i and 0 elsewhere. Since the set of allocations on $T_{i \rightarrow j}$ is compact, it is enough to consider two subsequential limits θ, θ' of $\{\theta_{\varepsilon_n}\}_{n \geq 1}$ and prove that $\partial\theta = \partial\theta'$. Passing to the limit in (13), we know that θ, θ' are balanced with respect to the above baseload. Writing $V_{i \rightarrow j}$ for the vertex set of $T_{i \rightarrow j}$, Lemma 7 reduces our task to proving

$$\|\partial\theta - \partial\theta'\|_{\ell^1(V_{i \rightarrow j})} < \infty. \quad (23)$$

Let θ_ε^* be the restriction of Θ_ε to $T_{i \rightarrow j}$. Thus, θ_ε^* is an allocation on $T_{i \rightarrow j}$ and it is ε -balanced with baseload $\theta_\varepsilon^*(j, i)$ at i and 0 elsewhere. Consequently, Proposition 6 guarantees that for any finite $K \subseteq V_{i \rightarrow j} \setminus \{i\}$,

$$\|\partial\theta_\varepsilon - \partial\theta_\varepsilon^*\|_{\ell^1(K)} \leq |x| + 1.$$

Applying this to $\varepsilon, \varepsilon' > 0$ and using the triangle inequality, we obtain

$$\|\partial\theta_\varepsilon - \partial\theta_{\varepsilon'}\|_{\ell^1(K)} \leq 2|x| + 2 + \|\partial\theta_\varepsilon^* - \partial\theta_{\varepsilon'}^*\|_{\ell^1(K)}.$$

Since $\{\partial\theta_{\varepsilon_n}^*\}_{n \geq 1}$ converges by assumption, we may pass to the limit to obtain $\|\partial\theta - \partial\theta'\|_{\ell^1(K)} \leq 2|x| + 2$. But K is arbitrary, so (23) follows. This shows that $\mathfrak{f}_{T_{i \rightarrow j}} := \lim_{n \rightarrow \infty} \mathfrak{f}_{T_{i \rightarrow j}}^{\varepsilon_n}$ exists pointwise. We now recall two classical facts about non-decreasing functions $\mathfrak{f}: \mathbb{R} \rightarrow \mathbb{R}$ with $\mathfrak{f}(\pm\infty) = \pm\infty$. First, \mathfrak{f}^{-1} is non-decreasing, so that its discontinuity set $\mathcal{D}(\mathfrak{f}^{-1})$ is countable. Second, the pointwise convergence $\mathfrak{f}_n \rightarrow \mathfrak{f}$ implies $\mathfrak{f}_n^{-1}(t) \rightarrow \mathfrak{f}^{-1}(t)$ for every $t \in \mathbb{R} \setminus \mathcal{D}(\mathfrak{f}^{-1})$. Consequently, letting $\varepsilon \rightarrow 0$ in (17) proves (21) for $t \notin \mathcal{D} := \mathcal{D}(\mathfrak{f}_{T_{i \rightarrow j}}^{-1}) \cup \bigcup_{k \sim i} \mathcal{D}(\mathfrak{f}_{T_{k \rightarrow i}}^{-1})$. The equality then extends to \mathbb{R} since \mathcal{D} is countable and both sides of (21) are right-continuous in t . Replacing $T_{i \rightarrow j}$ with (T, o) in the above argument shows that $\mathfrak{f}_{(T, o)} := \lim_{n \rightarrow \infty} \mathfrak{f}_{(T, o)}^{\varepsilon_n}$ exists and satisfies

$$\mathfrak{f}_{(T, o)}^{-1}(t) = t - \sum_{i \sim o} [1 - \mathfrak{f}_{T_{i \rightarrow o}}^{-1}(t)]_0^1 \quad (t \in \mathbb{R}).$$

Finally, recall that $\mathfrak{f}_{(T, o)}^{\varepsilon_n}(0) = \partial\Theta_{\varepsilon_n}(T, o)$ for all $n \geq 1$, so that $\mathfrak{f}_{(T, o)}(0) = \ell_o$. But $\mathfrak{f}_{(T, o)}(0) > t \iff \mathfrak{f}_{(T, o)}^{-1}(t) < 0$ by definition of $\mathfrak{f}_{(T, o)}^{-1}$, so (22) follows. \square

10 Proof of Theorem 2

In all this section, we consider networks rather than graphs, where each oriented edge (i, j) is equipped with a mark $\xi(i, j) \in [0, 1]$. Given $t \in \mathbb{R}$, we shall be interested in marks that satisfy the local recursion

$$\xi(i, j) = \left[1 - t + \sum_{k \sim i, k \neq j} \xi(k, i) \right]_0^1, \quad (i, j) \in \vec{E}. \quad (24)$$

We start with a simple Lemma.

Lemma 1. *Under (24), $\partial\xi(i) \wedge \partial\xi(j) > t \iff \xi(i, j) + \xi(j, i) > 1$.*

Proof. We prove the equivalence case by case. Note that by assumption,

$$\xi(i, j) = [1 - t + \partial\xi(i) - \xi(j, i)]_0^1 \quad (25)$$

$$\xi(j, i) = [1 - t + \partial\xi(j) - \xi(i, j)]_0^1. \quad (26)$$

- If $0 < \xi(i, j), \xi(j, i) < 1$, then the equivalence trivially holds since we may safely remove the truncation $[\cdot]_0^1$ from (25)-(26) to obtain

$$\partial\xi(i) - t = \xi(i, j) + \xi(j, i) - 1 = \partial\xi(j) - t.$$

- If $\xi(j, i) = 0$, then we have $1 - t + \partial\xi(j) - \xi(i, j) \leq 0$ thanks to (26), and hence $\partial\xi(j) \leq t$. Thus, both sides of the equivalence are false.
- If $\xi(i, j) = 1, \xi(j, i) > 0$, then using $\xi(i, j) = 1$ in (25) gives $\partial\xi(i) - t \geq \xi(j, i)$ and since $\xi(j, i) > 0$ we obtain $\partial\xi(i) > t$. On the other hand, using $\xi(j, i) > 0$ in (26) gives $\partial\xi(j) > t + \xi(i, j) - 1$ and since $\xi(i, j) = 1$ we obtain $\partial\xi(j) > t$. Thus, both sides of the equivalence are true.

The other possible cases follow by exchanging $\xi(i, j)$ and $\xi(j, i)$. \square

We are ready for the proof of Theorem 2, which we divide into two parts.

Proposition 10. *Let $Q \in \mathcal{P}([0, 1])$ be any solution to $Q = F_{\pi, t}(Q)$. Then,*

$$\Phi_\mu(t) \geq \frac{\mathbb{E}[D]}{2} \mathbb{P}(\xi_1 + \xi_2 > 1) - t \mathbb{P}(\xi_1 + \dots + \xi_D > t),$$

where $D \sim \pi$ and where $\{\xi_n\}_{n \geq 1}$ are i.i.d. with law Q , independent of D .

Proof. Kolmogorov's extension Theorem allows us to convert the consistency equation $Q = F_{\pi,t}(Q)$ into a random rooted tree $\mathbb{T} \sim \text{UGWT}(\pi)$ equipped with marks satisfying (24) a-s, such that conditionally on the structure of $[\mathbb{T}, o]_h$, the marks from generation h to $h-1$ are IID with law Q . This random rooted network is easily checked to be unimodular. Thus, we may apply Proposition 4 with $f = \mathbf{1}_{\partial\xi > t}$. By Lemma 1, we have $\widehat{f} = \mathbf{1}_{\xi + \xi^* > 1}$ and hence

$$\Phi_\mu(t) \geq \frac{1}{2}\vec{\mu}(\xi + \xi^* > 1) - t\mu(\partial\xi > t).$$

This is precisely the desired result, since we have by construction

$$\mu(\partial\xi > t) = \mathbb{P}(\xi_1 + \dots + \xi_D > t), \quad \vec{\mu}(\xi + \xi^* > 1) = \mathbb{E}[D]\mathbb{P}(\xi_1 + \xi_2 > 1)$$

where $D \sim \pi$ and ξ_1, ξ_2, \dots are IID with law Q , independent of D . \square

Proposition 11. *There is a solution $Q \in \mathcal{P}([0, 1])$ to $Q = F_{\pi,t}(Q)$ such that*

$$\Phi_\mu(t) = \frac{\mathbb{E}[D]}{2}\mathbb{P}(\xi_1 + \xi_2 > 1) - t\mathbb{P}(\xi_1 + \dots + \xi_D > t).$$

where $D \sim \pi$ and where $\{\xi_n\}_{n \geq 1}$ are IID with law Q , independent of D .

Proof. Let $\mathbb{T} \sim \text{UGWT}(\pi)$. Thanks to Proposition 2, we have

$$\partial\Theta_\varepsilon(\mathbb{T}, o) \xrightarrow[\varepsilon \rightarrow 0]{L^2} \partial\Theta_0(\mathbb{T}, o).$$

In particular, there is a deterministic vanishing sequence $\varepsilon_1, \varepsilon_2, \dots$ along which the convergence holds almost surely. This almost-sure convergence automatically extends from the root to all vertices, since *everything shows up at the root* of a unimodular random network [2, Lemma 2.3]. Therefore, \mathbb{T} satisfies almost-surely the assumption of Proposition 9. Consequently, the marks $\xi(i, j) := [1 - \mathfrak{f}_{\mathbb{T}_{i \rightarrow j}}^{-1}(t)]_0^1$ satisfy (24) almost-surely, and

$$\partial\Theta_0(\mathbb{T}, o) > t \iff \partial\xi(o) > t.$$

This ensures that $f = \mathbf{1}_{\partial\xi > t}$ satisfies the requirement for equality in Proposition 4, and we may then use Lemma 1 to rewrite the conclusion as

$$\Phi_\mu(t) = \frac{1}{2}\vec{\mu}(\xi + \xi^* > 1) - t\mu(\partial\xi > t).$$

Now, $D = \deg(\mathbb{T}, o)$ has law π , and conditionally on D , the subtrees $\{\mathbb{T}_{i \rightarrow o}\}_{i \sim o}$ are IID copies of a homogenous Galton-Watson tree $\widehat{\mathbb{T}}$ with offspring distribution $\widehat{\pi}$. Since $\xi(i, o)$ depends only on the subtree $\mathbb{T}_{i \rightarrow o}$, we obtain

$$\mu(\partial\xi > t) = \mathbb{P}(\xi_1 + \dots + \xi_D > t), \quad \vec{\mu}(\xi + \xi^* > 1) = \mathbb{E}[D]\mathbb{P}(\xi_1 + \xi_2 > 1)$$

where ξ_1, ξ_2, \dots are IID copies of $[1 - \mathfrak{f}_{\widehat{\mathbb{T}}}^{-1}(t)]_0^1$, independent of D . In turn, removing the root of $\widehat{\mathbb{T}}$ splits it into a $\widehat{\pi}$ -distributed number of IID copies of $\widehat{\mathbb{T}}$, so that the law Q of $[1 - \mathfrak{f}_{\widehat{\mathbb{T}}}^{-1}(t)]_0^1$ satisfies $Q = F_{\pi,t}(Q)$. \square

11 Proof of Theorem 3

Fix a degree sequence $\mathbf{d} = \{d(i)\}_{1 \leq i \leq n}$ and set $2m = \sum_{i=1}^n d(i)$.

Lemma 2. *The number of edges of $\mathbb{G}[\mathbf{d}]$ with both end-points in $S \subseteq \{1, \dots, n\}$ is stochastically dominated by a Binomial with mean $\frac{1}{m} \left(\sum_{i \in S} d_i\right)^2$.*

Proof. We assume that $s := \sum_{i \in S} d_i < m$, otherwise the claim is trivial. It is classical that $\mathbb{G}[\mathbf{d}]$ can be generated sequentially: at each step $1 \leq t \leq m$, a half-edge is selected and paired with a uniformly chosen other half-edge. The selection rule is arbitrary, and we choose to give priority to half-edges whose end-point lies in S . Let X_t be the number of edges with both end-points in S after t steps. Then $\{X_t\}_{0 \leq t \leq m}$ is a Markov chain with $X_0 = 0$ and transitions

$$X_{t+1} := \begin{cases} X_t + 1 & \text{with conditional probability } \frac{(s - X_t - t - 1)^+}{2m - 2t - 1} \\ X_t & \text{otherwise.} \end{cases}$$

For every $0 \leq t < m$, the fact that $X_t \geq 0$ ensures that

$$\frac{(s - X_t - t - 1)^+}{2m - 2t - 1} \leq \frac{s - t - 1}{2m - 2t - 1} \mathbf{1}_{(t < s)} \leq \frac{s}{2m} \mathbf{1}_{(t < s)},$$

where the second inequality uses the condition $s < m$. This shows that X_m is in fact stochastically dominated by a Binomial $(s, \frac{s}{2m})$, which is enough. \square

Lemma 3. *Let $X_{k,r}$ be the number of induced subgraphs with k vertices and at least r edges in $\mathbb{G}[\mathbf{d}]$. Then, for any $\theta > 0$,*

$$\mathbb{E}[X_{k,r}] \leq \left(\frac{2r}{\theta^2 m}\right)^r \left(\frac{e}{k} \sum_{i=1}^n e^{\theta d_i}\right)^k.$$

Proof. First observe that if $Z \sim \text{Bin}(n, p)$ then by a simple union-bound,

$$\mathbb{P}(Z \geq r) \leq \binom{n}{r} p^r \leq \frac{n^r p^r}{r!} = \frac{\mathbb{E}[Z]^r}{r!}.$$

Thus, by Lemma 2, the number Z_S of edges with both end-points in S satisfies

$$\mathbb{P}(Z_S \geq r) \leq \frac{1}{r! m^r} \left(\sum_{i \in S} d_i\right)^{2r} \leq \left(\frac{2r}{\theta^2 m}\right)^r \prod_{i \in S} e^{\theta d_i},$$

where we have used the crude bounds $x^{2r} \leq (2r)! e^x$ and $(2r)!/r! \leq (2r)^r$. The result follows by summing over all S with $|S| = k$ and observing that

$$\sum_{|S|=k} \prod_{i \in S} e^{\theta d_i} \leq \frac{1}{k!} \left(\sum_{i=1}^n e^{\theta d_i}\right)^k \leq \left(\frac{k}{e} \sum_{i=1}^n e^{\theta d_i}\right)^k.$$

The second inequality follows from the classical lower-bound $k! \geq \left(\frac{k}{e}\right)^k$. \square

We now fix $\{\mathbf{d}_n\}_{n \geq 1}$ as in Theorem 3. Let $Z_{\delta,t}^{(n)}$ be the number of subsets $\emptyset \subsetneq S \subseteq \{1, \dots, n\}$ such that $|S| \leq \delta n$ and $|E(S)| \geq t|S|$ in $\mathbb{G}_n := \mathbb{G}[\mathbf{d}_n]$.

Proposition 12. *For each $t > 1$, there is $\delta > 0$ and $\kappa < \infty$ such that*

$$\mathbb{E} \left[Z_{\delta,t}^{(n)} \right] \leq \kappa \left(\frac{\ln n}{n} \right)^{t-1},$$

uniformly in $n \geq 1$. In particular, $Z_{\delta,t}^{(n)} = 0$ w.h.p. as $n \rightarrow \infty$.

Proof. The assumptions of Theorem 3 guarantee that for some $\theta > 0$,

$$\alpha := \inf_{n \geq 1} \left\{ \frac{1}{n} \sum_{i=1}^n d_n(i) \right\} > 0 \quad \text{and} \quad \lambda := \sup_{n \geq 1} \left\{ \frac{1}{n} \sum_{i \in V} e^{\theta d_n(i)} \right\} < \infty.$$

Now, fix $t > 1$ and choose $\delta > 0$ small enough so that $f(\delta) < 1$, where

$$f(\delta) := \left(1 \vee \frac{2(1+t)}{\alpha \theta^2} \right)^{t+1} e \lambda \delta^{t-1}.$$

Using Lemma 3 and the trivial inequality $kt \leq [kt] \leq k(t+1)$, we have

$$\mathbb{E} \left[X_{k,[kt]}^{(n)} \right] \leq \left(\frac{2[kt]}{\theta^2 k \alpha} \right)^{[kt]} (e \lambda)^k \left(\frac{k}{n} \right)^{[kt]-k} \leq f^k \left(\frac{k}{n} \right).$$

Since f is increasing, we see that for any $1 \leq m \leq \delta n$,

$$\begin{aligned} \mathbb{E} \left[Z_{\delta,t}^{(n)} \right] &= \sum_{k=1}^{\lfloor \delta n \rfloor} \mathbb{E} \left[X_{k,[kt]}^{(n)} \right] \leq \sum_{k=1}^{m-1} f^k \left(\frac{m}{n} \right) + \sum_{k=m}^{\lfloor \delta n \rfloor} f^k(\delta) \\ &\leq \frac{f\left(\frac{m}{n}\right)}{1 - f\left(\frac{m}{n}\right)} + \frac{f(\delta)^m}{1 - f(\delta)}. \end{aligned}$$

Choose $m \sim c \ln n$ with c fixed. As $n \rightarrow \infty$, the first term is of order $\left(\frac{\ln n}{n}\right)^{t-1}$ while the second is of order $f(\delta)^{c \ln n} \ll \left(\frac{\ln n}{n}\right)^{t-1}$, if c is large enough. \square

Proof of Theorem 3. The assumptions on $\{\mathbf{d}_n\}_{n \geq 1}$ are more than sufficient to guarantee that a-s, the local weak limit of $\{\mathbb{G}_n\}_{n \geq 1}$ is $\mu := \text{UGWT}(\pi)$ (see e.g [9]). Thus, the weak convergence $\mathcal{L}_{\mathbb{G}_n} \rightarrow \mathcal{L}$ holds a-s, where \mathcal{L} is the law of $\partial\Theta_0$ under μ . Now, if $t < \varrho(\mu)$ then $\mathcal{L}((t, \infty)) > 0$, so the Portmanteau Theorem ensures that $\liminf_n \mathcal{L}_{\mathbb{G}_n}((t, \infty)) > 0$ a-s. Consequently,

$$\mathbb{P}(\varrho(\mathbb{G}_n) \leq t) = \mathbb{P}(\mathcal{L}_{\mathbb{G}_n}((t, \infty)) = 0) \xrightarrow[n \rightarrow \infty]{} 0.$$

On the other-hand, if $t > \varrho(\mu)$ then $\mathcal{L}([t, \infty)) = 0$, so the Portmanteau Theorem gives $\mathcal{L}_{\mathbb{G}_n}([t, \infty)) \rightarrow 0$ a-s. Thus, with δ as in Proposition 12,

$$\mathbb{P}(\varrho(\mathbb{G}_n) > t) \leq \mathbb{P}(\mathcal{L}_{\mathbb{G}_n}([t, \infty)) > \delta) + \mathbb{P}\left(Z_{\delta,t}^{(n)} > 0\right) \xrightarrow{n \rightarrow \infty} 0.$$

Note that the requirement $t > 1$ is fulfilled, since $\varrho(\mu) \geq 1$. Indeed, every node in a tree of size n has load $1 - \frac{1}{n}$, and the assumption $\pi_0 + \pi_1 < 1$ guarantees that the size of the random tree $\mathbb{T} \sim \text{UGWT}(\pi)$ is unbounded. \square

Acknowledgements

This research was supported by the ARO MURI grant W911NF-08-1-0233, *Tools for the Analysis and Design of Complex Multi-Scale Networks*, the NSF grant CNS-0910702, and the NSF Science & Technology Center grant CCF-0939370, *Science of Information*. The work was initiated at the follow-up meeting of the Newton Institute programme №86 on *Stochastic Processes in Communication Sciences*. The authors thank M. Lelarge and R. Sundaresan for fruitful discussions.

References

- [1] Ravi P. Agarwal, Maria Meehan, and Donal O'Regan. *Fixed point theory and applications*, volume 141 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 2001.
- [2] D. Aldous and R. Lyons. Processes on unimodular random networks. *Electronic Journal of Probability*, 12:1454–1508, 2007.
- [3] David Aldous and J. Michael Steele. The objective method: probabilistic combinatorial optimization and local weak convergence. In *Probability on discrete structures*, volume 110 of *Encyclopaedia Math. Sci.*, pages 1–72. Springer, Berlin, 2004.
- [4] David J. Aldous. The $\zeta(2)$ limit in the random assignment problem. *Random Structures Algorithms*, 18(4):381–418, 2001.
- [5] David J. Aldous and Antar Bandyopadhyay. A survey of max-type recursive distributional equations. *Ann. Appl. Probab.*, 15(2):1047–1110, 2005.

- [6] Itai Benjamini and Oded Schramm. Recurrence of distributional limits of finite planar graphs. *Electron. J. Probab.*, 6:no. 23, 13 pp. (electronic), 2001.
- [7] Patrick Billingsley. *Convergence of probability measures*. Wiley Series in Probability and Statistics: Probability and Statistics. John Wiley & Sons Inc., New York, second edition, 1999. A Wiley-Interscience Publication.
- [8] Béla Bollobás. A probabilistic proof of an asymptotic formula for the number of labelled regular graphs. *European J. Combin.*, 1(4):311–316, 1980.
- [9] Charles Bordenave. Lecture notes on random graphs and probabilistic combinatorial optimization. <http://www.math.univ-toulouse.fr/~bordenave/coursRG.pdf>, 2012.
- [10] Charles Bordenave, Marc Lelarge, and Justin Salez. Matchings on infinite graphs. *Probab. Theory Related Fields*, 157(1-2):183–208, 2013.
- [11] Julie Anne Cain, Peter Sanders, and Nick Wormald. The random graph threshold for k -orientability and a fast algorithm for optimal multiple-choice allocation. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '07, pages 469–476, Philadelphia, PA, USA, 2007. Society for Industrial and Applied Mathematics.
- [12] Daniel Fernholz and Vijaya Ramachandran. The k -orientability thresholds for gn, p . In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '07, pages 459–468, Philadelphia, PA, USA, 2007. Society for Industrial and Applied Mathematics.
- [13] Nikolaos Fountoulakis, Megha Khosla, and Konstantinos Panagiotou. The multiple-orientability thresholds for random hypergraphs. *arXiv:1309.6772*, September 2013.
- [14] David Gamarnik, Tomasz Nowicki, and Grzegorz Swirszcz. Maximum weight independent sets and matchings in sparse random graphs. Exact results using the local weak convergence method. *Random Structures Algorithms*, 28(1):76–106, 2006.
- [15] Pu Gao and Nicholas C. Wormald. Load balancing and orientability thresholds for random hypergraphs [extended abstract]. In *STOC'10—Proceedings of the 2010 ACM International Symposium on Theory of Computing*, pages 97–103. ACM, New York, 2010.

- [16] Bruce Hajek. Performance of global load balancing by local adjustment. *IEEE Trans. Inform. Theory*, 36(6):1398–1414, 1990.
- [17] Bruce Hajek. Balanced loads in infinite networks. *Ann. Appl. Probab.*, 6(1):48–75, 1996.
- [18] Remco van der Hofstad Hofstad. *Random Graphs and Complex Networks*. 2013. Course notes available at <http://www.win.tue.nl/~rhofstad/NotesRGCN.html>.
- [19] Svante Janson. The probability that a random multigraph is simple. *Combin. Probab. Comput.*, 18(1-2):205–225, 2009.
- [20] Mustafa Khandwawala and Rajesh Sundaresan. Belief propagation for optimal edge-cover in the random complete graph. *arXiv:1212.6027*, 2012.
- [21] Mathieu Leconte, Marc Lelarge, and Laurent Massoulié. Convergence of multivariate belief propagation, with applications to cuckoo hashing and load balancing. *arXiv:1207.1659*, July 2012.
- [22] Marc Lelarge. A new approach to the orientation of random hypergraphs. *CoRR*, abs/1201.5335, 2012.
- [23] Tomasz Łuczak. Sparse random graphs with a given degree sequence. In *Random graphs, Vol. 2 (Poznań, 1989)*, Wiley-Intersci. Publ., pages 165–182. Wiley, New York, 1992.
- [24] Russell Lyons. Asymptotic enumeration of spanning trees. *Combin. Probab. Comput.*, 14(4):491–522, 2005.
- [25] Justin Salez. Weighted enumeration of spanning subgraphs in locally tree-like graphs. *Random Structures Algorithms*, 43(3):377–397, 2013.
- [26] Justin Salez and Devavrat Shah. Belief propagation: an asymptotically optimal algorithm for the random assignment problem. *Math. Oper. Res.*, 34(2):468–480, 2009.
- [27] Moshe Shaked and J. George Shanthikumar. *Stochastic orders and their applications*. Probability and Mathematical Statistics. Academic Press Inc., Boston, MA, 1994.
- [28] J. Michael Steele. Minimal spanning trees for graphs with random edge lengths. In *Mathematics and computer science, II (Versailles, 2002)*, Trends Math., pages 223–245. Birkhäuser, Basel, 2002.