

Belief propagation for minimum weight many-to-one matchings in the random complete graph*

Mustafa Khandwawala[†]

Abstract

In a complete bipartite graph with vertex sets of cardinalities n and n' , assign random weights from exponential distribution with mean 1, independently to each edge. We show that, as $n \rightarrow \infty$, with $n' = \lceil n/\alpha \rceil$ for any fixed $\alpha > 1$, the minimum weight of many-to-one matchings converges to a constant (depending on α). Many-to-one matching arises as an optimization step in an algorithm for genome sequencing and as a measure of distance between finite sets. We prove that a belief propagation (BP) algorithm converges asymptotically to the optimal solution. We use the objective method of Aldous to prove our results. We build on previous works on minimum weight matching and minimum weight edge cover problems to extend the objective method and to further the applicability of belief propagation to random combinatorial optimization problems.

Keywords: belief propagation ; local weak convergence ; many-to-one matching ; objective method ; random graph.

AMS MSC 2010: Primary 60C05, Secondary 68Q87.

Submitted to EJP on May 3, 2014, final version accepted on December 6, 2014.

Supersedes arXiv:1405.1292v1.

1 Introduction

We address here an optimization problem over bipartite graphs related to the matching problem – called the many-to-one matching problem, which has applications in an algorithm for genome sequencing and as a measure of distance between finite sets. Given two sets A and B , with $|A| \geq |B|$, consider a bipartite graph G with vertex set $V = A \cup B$, and edge set $E \subset A \times B$. A many-to-one matching in G is a subgraph \mathcal{M} such that each vertex in A – called the *one side* – has degree 1 in \mathcal{M} , and each vertex in B – called the *many side* – has degree 1 or more in \mathcal{M} . A many-to-one matching can be viewed as an onto function from A to B . Each edge $e \in E$ has a weight $\xi_e \in \mathbf{R}_+ = [0, \infty)$. The cost of the matching \mathcal{M} is the sum of the weights of the edges in \mathcal{M} . We focus here on the minimum cost many-to-one matchings on complete bipartite graphs where the edge-weights are independent and identically distributed (i.i.d.) random variables.

*Support: Department of Science and Technology, Government of India and by a TCS fellowship grant.

[†]Indian Institute of Science, Bangalore, India. Currently at INRIA, Paris, France.

E-mail: mustafa.khandwawala@inria.fr

For several optimization problems over graphs with random edge-weights, where the objective is to minimize the sum weight of a subset of edges under some constraints, the expected optimal cost converges to a constant as $n \rightarrow \infty$. Such results are known for the minimum spanning tree [12], matching [1, 4], edge cover [15, 16], and traveling salesman problem (TSP) [24] over complete graphs or complete bipartite graphs with i.i.d. edge-weights. Furthermore, simple, iterative *belief propagation* algorithms can be used to find asymptotically optimal solutions for matching [22] and edge cover [16].

We prove a similar result for many-to-one matching. Let us be a little more specific. Fix a real number $\alpha > 1$. In the complete bipartite graph $K_{n,n'}$ with vertex sets A and B of cardinality n and $n' = \lceil n/\alpha \rceil$ respectively, we take the edge-weights to be i.i.d. random variables having the exponential distribution with mean 1. Denote by M_n^α the minimum cost of many-to-one matching on $K_{n,n'}$. We show that the expected value of M_n^α converges to a constant (which depends on α and can be computed) as $n \rightarrow \infty$. Further, we show that a belief propagation (BP) algorithm finds the asymptotically optimal many-to-one matching in $O(n^2)$ steps.

We proceed to show these results via the objective method following [4] for matching and [16] for edge cover. Before we give the background and overview of the methods, we describe two practical applications of the many-to-one matching problem.

1.1 Applications of many-to-one matching

1.1.1 Restriction scaffold problem

Ben-Dor et al. [8] introduced the *restriction scaffold problem* for assisting the finishing step in a shotgun genome sequencing project. Shotgun sequencing refers to the process of generating several short clones of a large target genome and then sequencing these individual clones. Sequencing a short genome is much easier than sequencing the entire genome, and the cloning process is automated and fast, but the locations of clones are random. This process, which is like a random sampling, leaves us with several sequenced segments of the target. It is possible to construct the target sequence by assembling these short sequences by matching their overlaps, provided that we have sufficient redundancy in the samples to cover the entire sequence with enough overlap.

Usually, asking for enough redundancy – that this assembly based on overlaps reconstructs the whole sequence – will result in excessive cloning and sequencing tasks. As a tradeoff, we can give up on this high redundancy, and accept as the output of this assembly task large contiguous sequences (Ben-Dor et al. call these *contigs*), which together cover a large part of the target sequence. If the gaps – the positions in the target sequence that are not covered by these contigs – are small, those areas of the genome can be targeted manually to finish the sequencing project. However, the relative position and orientation of the contigs, and the sizes of gaps between successive contigs are unknown. We need some information to figure out how these contigs align along the target genome. See the top part of Figure 1.

Ben-Dor et al. propose the use of restriction enzymes to generate this information. Restriction enzymes cut the DNA at specific recognition sequences (differing according to the enzyme) called the restriction site of the enzyme. These restriction sites form a scaffold over which the contigs can be assembled. Again, see Figure 1. They measure the approximate sizes of the fragments obtained by digesting the target DNA with a restriction enzyme. On the other hand, by assuming a particular ordering and orientation of the contigs and the gaps between successive contigs the sizes of the fragments can be computed by reading out the recognition sequence of the enzyme along the contigs (their sequences are known). Comparing these computed fragment sizes with the measured fragment sizes for several different restriction enzymes, gives them a way to obtain a

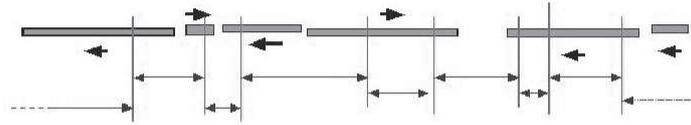


Figure 1: “Fraction of a configuration and one enzyme data. Top: the contigs; bold arrow denote the orientation of each contig. Vertical lines: restrictions sites. Bottom: restriction fragments between neighboring sites.”
Image and the quoted caption from [8].

new estimate of the ordering, orientation and gaps. This is the basis of their iterative algorithm for solving the restriction scaffold problem.

The step of the iterative algorithm involving the comparison of the computed fragment sizes with the measured fragment sizes is formulated as a many-to-one matching problem – the computed fragment sizes on the “one” side, and the measured fragment sizes on the “many” side. The measurement step process does not indicate the number of fragments corresponding to a particular measurement. The many-to-one nature of the association captures the notion that a measurement may correspond to one or more physical fragments, and each physical fragment must correspond to one measurement. The weight of an edge captures how much the two sizes differ. Specifically, the weight of an edge joining a computed fragment size $s(c)$ to a measured fragment size $s(m)$ is $|\log(s(c)) - \log(s(m))|$. They solve the many-to-one matching problem using this linear property: the vertices are represented as points on a line, a vertex corresponding to a size s is at position $\log(s)$ on the line, and the cost of matching two points is the distance between the points.

We do not assume this linear structure in our formulation of the many-to-one problem on the complete bipartite graph. However the i.i.d. weight assumption is the mean-field approximation of a geometric structure in the following sense. Take a set A of n points and a set B of $n' = \lceil n/\alpha \rceil$ points independently and uniformly at random in the unit sphere in \mathbf{R}^d . When n and n' are large, the probability that the distance between a typical point of A and a typical point of B is less than r is approximately r^d for small r . We can represent the distances between points of the two sets by forming a complete bipartite graph and taking the weight of an edge as the distance between the corresponding points. For a mean-field model we ignore the geometric structure of \mathbf{R}^d (and triangular inequality) and only take into account the interpoint distances, by taking the edge-weights to be i.i.d. copies of a nonnegative random variable ξ satisfying $\Pr(\xi < r) \approx r^d$ as $r \rightarrow 0$.

1.1.2 A measure of distance between sets of points

A concept of distance between finite sets of points is useful in many areas like machine learning, computational geometry, and comparison of theories. Such a distance is derived from a given distance function (or a metric) between the points. For example, in a clustering process over a set of examples, suppose we are given a function d such that $d(e_1, e_2)$ corresponds to the distance between two examples e_1 and e_2 . Clustering methods such as TIC [10] rely on a function d_1 that specifies the distance $d_1(C_1, C_2)$ for two clusters (which are sets of examples) C_1 and C_2 .

Eiter and Mannila [11] use the term *surjection distance* for a measure of similarity, based on optimal many-to-one matching, between two finite sets of points in a metric space. The surjection distance is the minimum cost of a many-to-one matching, with the

points of the larger set forming the vertices of the one side, the points of the smaller set forming the vertices of the many side of a bipartite graph, and the edge-weights taken as the distance between the corresponding points. By reduction to the minimum weight matching problem they find an $O(n^3)$ algorithm for many-to-one matching. Belief propagation (in the random setting we consider here) yields an asymptotically optimal solution in $O(n^2)$ steps.

1.2 Background and overview of the methods

Several of the results for random combinatorial optimization problems have originated as conjectures supported by heuristic statistical mechanical methods such as the replica method and the cavity method [17, 18, 19, 20, 21]. Establishing the validity of these methods is a challenge even for mathematically simple models.

Aldous [1, 4] provided a rigorous proof of the $\zeta(2)$ limit conjecture for the random matching (or assignment) problem. The method he used is called the *objective method*. A survey of this method and its application to a few other problems was given in [3]. The objective method provides a rigorous counterpart to the *cavity method*, but has had success only in a few settings.

The crux of the objective method is to obtain a suitable distributional limit for a sequence of finite random weighted graphs of growing vertex set size. The space of weighted graphs is endowed with a metric that captures the proximity based on the local neighborhood around a vertex. The weak convergence of probability measures on this space is called *local weak convergence* to emphasize that the convergence applies to events that are essentially characterized by the local structure around a typical vertex, and not to events that crucially depend on the global structure of the graph. This form of convergence has appeared in varying forms in [14, 9, 3]. Once we have this convergence, we can relate the optimization problem on finite graphs to an optimization problem on the limit object. For example, Aldous [4] showed that the limit object for the sequence of random weighted complete bipartite graphs $K_{n,n}$ is the *Poisson weighted infinite tree (PWIT)*. Aldous then writes the optimization problem in terms of local conditions at the vertices of the tree, and using the symmetry in the structure of the PWIT, these relations result in a distributional identity called the *recursive distributional equation* [5]. A solution to this equation is then used to describe asymptotically optimal matchings on $K_{n,n}$.

One can exploit the recursive distributional equation to construct an iterative decentralized message passing algorithm, a version of which is called belief propagation (BP) in the computer science literature. BP algorithms are known to converge to the correct solution on graphs without cycles. For graphs with cycles, provable guarantees are known for BP only for certain problems. For example, Bayati, Shah and Sharma [7] proved that the BP algorithm for maximum weight matching on bipartite graphs converges to the correct value as long as the maximum weight matching is unique, and Sanghavi, Malioutov and Willsky [23] proved that BP for matching is as powerful as LP relaxation. Salez and Shah [22] studied belief propagation for the random assignment problem, and showed that a BP algorithm on $K_{n,n}$ converges to an update rule on the limit object PWIT. The iterates on the PWIT converge in distribution to the minimum cost assignment. The iterates are near the optimal solution in $O(n^2)$ steps whereas the worst case optimal algorithm on bipartite graphs is $O(n^3)$ (expected time $O(n^2 \log n)$ for i.i.d. edge capacities); see Salez and Shah [22] and references therein.

In message passing algorithms like belief propagation, the calculations are performed at each vertex of a graph from the messages received from its neighbors. This local nature makes feasible its analysis via the objective method. The author and Sundaresan, in [16], extended the objective method by combining local weak convergence and belief

propagation to prove and characterize the limiting expected minimum cost for the edge cover problem. In this paper, we implement this general program for the many-to-one matching problem. The proof relies on two properties: (1) the neighborhoods around most vertices are tree-like, and (2) the effect of the boundary on the belief propagation messages at the root of the tree vanishes as the boundary is taken to infinity. The second property is formalized as “endogeny” (see [5]) and is handled in Section 7.

1.3 Contributions of this paper

This paper extends the line of work on combinatorial optimization problems over graphs with i.i.d. edge-weights. Although the methods are similar to those in [4, 22, 16], there are key differences from earlier work in addressing the many-to-one matching problem. We now highlight them. For example, we must deal with two types of vertices corresponding to two different constraints on the degrees. The different viewpoints from a vertex in A and a vertex in B makes the limit object for the sequence of graphs $K_{n,n'}$ different from the Poisson weighted infinite tree (PWIT) obtained as the limit of bipartite graphs $K_{n,n}$ with vertex sets of the same cardinality. The difference is in the distribution of the weights of the edges at odd and even levels from the root. As a consequence, the recursive distributional equation (RDE) – guiding the construction of an optimal many-to-one matching on the limit object – is expressed in terms of a two-step recursion.

We prove that this recursive distributional equation has a unique solution, and that the many-to-one matching on the limit object generated by the solution of this equation is optimal among all *involution invariant* many-to-one matchings. To establish correctness of BP, we establish that the fixed-point solution of the RDE has a full domain of attraction, and that the stationary process with the fixed-point marginal distribution on the limit object satisfies endogeny.

2 Main results

From now on we will write $K_{n,n/\alpha}$ to mean $K_{n,n'}$, where $n' = \lceil n/\alpha \rceil$. Recall that A and B denote the *one* side and the *many* side in the bipartite graph $K_{n,n/\alpha}$. The weights of the edges are i.i.d. random variables with the exponential distribution, and for simplicity in the subsequent analysis, we take the mean of the edge-weights to be n . This only changes the cost of the optimal solution by a factor of n , and so we have a rescaling in the left hand side of (2.1). Also note that though we deal here explicitly with the exponential distribution, the results go through unchanged – as is standard in the mean-field optimization literature – for any distribution that has a distribution function F satisfying $F(0) = 0$ and $F'(0) = 1$ (only the right-derivative).

Our first result establishes the limit of the scaled expected minimum cost of the random many-to-one matching problem.

Theorem 2.1. *For $\alpha > 1$, the minimum cost M_n^α of many-to-one matching on $K_{n,n/\alpha}$, satisfies*

$$\begin{aligned} \lim_{n \rightarrow \infty} n^{-1} \mathbb{E} M_n^\alpha &= -\text{Li}_2(-\gamma^{-1}) - \frac{1}{2} \log^2(1 + \gamma^{-1}) + w_o \log(1 + \gamma^{-1}) + w_o \\ &=: c_*^\alpha, \end{aligned} \tag{2.1}$$

where Li_2 is the *dilogarithm*: $\text{Li}_2(z) = -\int_0^z \frac{\log(1-t)}{t} dt = \sum_{k=1}^\infty \frac{z^k}{k^2}$; w_o is the *positive solution* to $\alpha = w_o + e^{-w_o}$; γ equals $w_o e^{w_o}$.

Our second result shows that a belief propagation algorithm gives a many-to-one matching that is asymptotically optimal as $n \rightarrow \infty$. The BP algorithm is based on the corresponding BP algorithms developed in [22, 7] for matching and in [16] for edge

cover. The proof uses the technique of showing that the update rule of BP converges to an update rule on a limit infinite tree.

We now define a BP algorithm on an arbitrary bipartite graph $G = (A \cup B, E)$ with edge-costs. First, some notation. We write $w \sim v$ to mean w is a neighbor of v , i.e., $\{v, w\} \in E$. For an edge $e = \{v, w\} \in E$, we write its cost as $\xi_G(e)$ or $\xi_G(v, w)$ (with the ordering irrelevant). For each vertex $v \in V$, we associate a nonempty subset $\pi_G(v)$ of its neighbors such that $\pi_G(v)$ has cardinality 1 if $v \in A$. Write $\mathcal{M}(\pi_G) = \{\{v, w\} \mid v \in A \cup B, w \in \pi_G(v)\}$.

The BP algorithm is an iterative message passing algorithm. In each iteration $k \geq 0$, every vertex $v \in V$ sends a message $X_G^k(w, v)$ to each neighbor $w \sim v$ according to the following rules:

Initialization:

$$X_G^0(w, v) = 0. \tag{2.2}$$

Update rule:

$$X_G^{k+1}(w, v) = \begin{cases} \min_{u \sim v, u \neq w} \{ \xi_G(v, u) - X_G^k(v, u) \} & \text{if } v \in A, \\ \min_{u \sim v, u \neq w} \{ (\xi_G(v, u) - X_G^k(v, u))^+ \} & \text{if } v \in B. \end{cases} \tag{2.3}$$

Decision rule:

$$\pi_G^k(v) = \begin{cases} \arg \min_{u \sim v} \{ \xi_G(v, u) - X_G^k(v, u) \} & \text{if } v \in A, \\ \arg \min_{u \sim v} \{ (\xi_G(v, u) - X_G^k(v, u))^+ \} & \text{if } v \in B. \end{cases} \tag{2.4}$$

Note that the subset $\mathcal{M}(\pi_G^k(v))$ is not necessarily a many-to-one matching, but it can be modified to be a many-to-one matching \mathcal{M}_n^k by some patching. We also remark that while $\xi_G(v, u) = \xi_G(u, v)$, the messages $X_G^k(w, v)$ and $X_G^k(v, w)$ can differ.

We analyze the belief propagation algorithm for $G = K_{n, n/\alpha}$ and i.i.d. exponential random edge-costs, and prove that after sufficiently large number of iterates, the expected cost of the many-to-one matching given by the BP algorithm is close to the limit value in Theorem 2.1.

Theorem 2.2. *On $K_{n, n/\alpha}$, the output $\mathcal{M}(\pi_{K_{n, n/\alpha}}^k)$ of the BP algorithm can be patched to get a many-to-one matching \mathcal{M}_n^k that satisfies*

$$\lim_{k \rightarrow \infty} \lim_{n \rightarrow \infty} n^{-1} \mathbb{E} \left[\sum_{e \in \mathcal{M}_n^k} \xi_{K_{n, n/\alpha}}(e) \right] = c_*^\alpha. \tag{2.5}$$

3 Local weak convergence

First we recollect the terminology for defining convergence of graphs, borrowed from [3].

3.1 Rooted geometric networks

A graph $G = (V, E)$ along with a length function $l : E \rightarrow (0, \infty]$ and label $l : V \rightarrow \mathcal{L}$ is called a *network*. \mathcal{L} is a finite set of labels. The *distance* between two vertices in the network is the infimum of the sum of lengths of the edges of a path connecting the two vertices, the infimum being taken over all such paths. We call the network a *geometric network* if for each vertex $v \in V$ and positive real ρ , the number of vertices within a distance ρ of v is finite. We denote the space of geometric networks by \mathcal{G} .

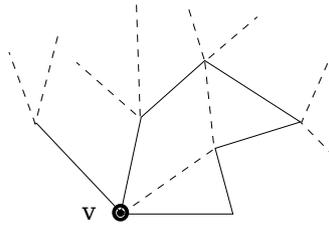


Figure 2: Neighborhood $\mathcal{N}_\rho(G)$ of graph G . The solid edges form the neighborhood, and form paths of length at most ρ from the root v . Dashed edges are the other edges of G .

A geometric network with a distinguished vertex v is called a *rooted geometric network* with root v . We denote the space of all connected rooted geometric networks by \mathcal{G}_* . It is important to note that in \mathcal{G}_* we do not distinguish between rooted isomorphisms of the same network – isomorphisms that preserve the graph structure along with the root, vertex labels, and edge-lengths. We will use the notation (G, ϕ) to denote an element of \mathcal{G}_* which is the isomorphism class of rooted networks with underlying network G and root ϕ .

3.2 Local weak convergence

We call a positive real number ρ a *continuity point* of G if no vertex of G is exactly at a distance ρ from the root of G . Let $\mathcal{N}_\rho(G)$ denote the neighborhood of the root of G up to distance ρ . $\mathcal{N}_\rho(G)$ contains all vertices of G which are within a distance ρ from the root of G (Figure 2). We take $\mathcal{N}_\rho(G)$ to be an element of \mathcal{G}_* by inheriting the same length function l as G , same label function as G , and the same root as that of G .

We say that a sequence of rooted geometric networks $G_n, n \geq 1$ *converges locally* to an element G_∞ in \mathcal{G}_* if for each continuity point ρ of G_∞ , there is an n_ρ such that for all $n \geq n_\rho$, there exists a graph isomorphism $\gamma_{n,\rho}$ from $\mathcal{N}_\rho(G_\infty)$ to $\mathcal{N}_\rho(G_n)$ that maps the root of the former to the root of the latter, preserves the labels of the vertices, and for each edge e of $\mathcal{N}_\rho(G_\infty)$, the length of $\gamma_{n,\rho}(e)$ converges to the length of e as $n \rightarrow \infty$.

The space \mathcal{G}_* can be suitably metrized to make it a separable and complete metric space. See [2] for details on the metric. One can then consider probability measures on this space and endow that space with the topology of weak convergence of measures. This notion of convergence is called *local weak convergence*.

In our setting of complete graphs $K_{n,n/\alpha} = (A \cup B, E)$ with random i.i.d. edge-costs $\{\xi_e, e \in E\}$, we regard the edge-costs to be the lengths of the edges, and declare a vertex φ of $K_{n,n/\alpha}$ chosen uniformly at random as the root of $K_{n,n/\alpha}$. Assign the label o to the vertices in A , and the label m to the vertices in B ; the label of a vertex will indicate its permitted degrees in a many-to-one matching: o for one and m for many. This makes $(K_{n,n/\alpha}, \varphi)$, the isomorphism class of $K_{n,n/\alpha}$ and the root φ , a random element of \mathcal{G}_* . We will show (Theorem 3.1 below) that the sequence of random rooted geometric networks $K_{n,n/\alpha}$ converges in the local weak sense to an element of \mathcal{G}_* , which we define below.

3.3 A mixture of two inhomogeneous Poisson weighted infinite trees

The graph that we define here is a variant of the Poisson weighted infinite tree that appears in [3]. We use the notation from [22] to define the variant. This object is a random element of \mathcal{G}_* whose distribution is a mixture of the distributions of two random trees: \mathcal{T}_α^o and \mathcal{T}_α^m . Let \mathcal{V} denote the set of all finite strings over the alphabet $\mathbf{N} = \{1, 2, 3, \dots\}$. Let ϕ denote the empty string. Let $|v|$ denote the length of the string v . For a string v with $|v| \geq 1$, let \dot{v} denote the parent of v , i.e., the string with the last

letter deleted. Let $\mathcal{E} = \{\{v, vi\}, v \in \mathcal{V}, i \in \mathbb{N}\}$. Both \mathcal{T}_α^o and \mathcal{T}_α^m have the same underlying graph $(\mathcal{V}, \mathcal{E})$ with root ϕ . Each vertex v in \mathcal{V} has a label from the set $\mathcal{L} = \{o, m\}$: v has label o in \mathcal{T}_α^o (respectively \mathcal{T}_α^m) if $|v|$ is even (respectively odd), and label m if $|v|$ is odd (respectively even). In particular, the root ϕ has label o in \mathcal{T}_α^o , and label m in \mathcal{T}_α^m . Construct a family of independent Poisson processes, with intensity 1, on \mathbf{R}_+ :

$$\{\xi^v = (\xi_1^v, \xi_2^v, \dots), v \in \mathcal{V}\}.$$

Take the length of an edge $\{v, vi\}$ in \mathcal{E} to be ξ_i^v if v has label m , and $\alpha\xi_i^v$ if v has label o . If we denote the probability measure of \mathcal{T}_α^o and \mathcal{T}_α^m by ν_o and ν_m respectively, then the probability measure of \mathcal{T}_α is the mixture

$$\nu = \frac{\alpha}{1 + \alpha}\nu_o + \frac{1}{1 + \alpha}\nu_m. \tag{3.1}$$

We now prove that this distribution is the local weak limit of the sequence of distributions of $(K_{n,n/\alpha}, \varphi)$, where φ is chosen uniformly at random from the vertices of $K_{n,n/\alpha}$.

Theorem 3.1. *The sequence $(K_{n,n/\alpha}, \varphi)$ having edge-weights that are independent random variables with exponential mean n distribution converges to \mathcal{T}_α (with root ϕ) as $n \rightarrow \infty$ in the local weak sense.*

Proof. In $K_{n,n/\alpha}$, let A denote the larger set with cardinality n and let B denote the set with smaller cardinality $n' = \lceil n/\alpha \rceil$. Assign the label o to vertices in A (the one side), and label m to the vertices in B (the many side). Condition the root φ to have label o : this occurs with probability $n/(n + \lceil n/\alpha \rceil)$. Fix a positive integer N , and for $n > N$ identify vertices of $K_{n,n/\alpha}$ with strings in \mathcal{V} restricted to the alphabet $\{1, 2, \dots, N\}$. The empty string ϕ identifies the root φ . The strings $1, 2, \dots, N$ denote the first N neighbors of the root φ arranged in increasing order of the weights of edges incident on the root. If $v \in \mathcal{V} \setminus \phi$ denotes some vertex in $K_{n,n/\alpha}$ then $v1, v2, \dots, vN$ denote the first N neighbors of v excluding \dot{v} arranged in increasing order of the weights of edges incident on v . This generates a subtree of width N of $(\mathcal{V}, \mathcal{E})$. Assign weights and vertex labels to the edges and vertices of the subtree according to those in $K_{n,n/\alpha}$. Call this weighted tree $T_{N,n}$. Note that several strings in \mathcal{V} may correspond to a single vertex of $K_{n,n/\alpha}$. Fix another positive integer H . Call the restriction of $T_{N,n}$ to vertices v with $|v| \leq H$ the weighted tree $T_{N,H,n}$.

Let (the scripted) \mathcal{T}_N and $\mathcal{T}_{N,H}$ denote the corresponding restrictions of \mathcal{T}_α^o . Clearly, $T_{N,H,n}$ and $\mathcal{T}_{N,H}$ are isomorphic finite graphs with the same vertex labels (the roots of both are labeled o). We can view the edge-weights as elements of $\mathbf{R}^{\mathcal{E}(\mathcal{T}_{N,H})}$, where $\mathcal{E}(\mathcal{T}_{N,H})$ denotes the edges in $\mathcal{T}_{N,H}$. Write $\nu_{N,H,n}$ and $\nu_{N,H}$ for the measure of $T_{N,H,n}$ and $\mathcal{T}_{N,H}$ respectively.

We bound the likelihood ratio $\frac{d\nu_{N,H,n}}{d\nu_{N,H}}$ for the restriction that each vertex of $T_{N,H,n}$ corresponds to a unique vertex in $K_{n,n/\alpha}$. For any vertex v , the vertex $v1$ is taken to be the w that minimizes $\xi(v, w)$ among all $w \neq \dot{v}$, and $w \sim v$. Suppose that we have associated vertices upto vi for some $1 \leq i < N$ to vertices of $K_{n,n/\alpha}$. If v has label o the next vertex $v(i + 1)$ is the vertex w for which the difference $\xi(v, w) - \xi(v, vi)$, $w \notin \{\dot{v}, v1, \dots, vi\}$, is the minimum. The density of the difference $\xi(v, v(i + 1)) - \xi(v, vi)$ for $i \geq 1$ (respectively, $\xi(v, v1)$ for $i = 0$) at $x > 0$ is at least

$$\left(1 - \frac{|\mathcal{T}_{N,H}|}{\lceil n/\alpha \rceil - N}\right) \left[\frac{\lceil n/\alpha \rceil - 1 - i}{n} \exp\left(-\frac{\lceil n/\alpha \rceil - 1 - i}{n}x\right)\right], \tag{3.2}$$

where $|\mathcal{T}_{N,H}|$ is the number of vertices in $\mathcal{T}_{N,H}$. To see this, observe that the term with square brackets is, using the memoryless property of the exponential distribution,

the density of the minimum among the $\lceil n/\alpha \rceil - (i + 1)$ exponential random variables corresponding to nodes on the many side (excluding $\{v, v_1, \dots, v_i\}$). The term outside the square brackets is a lower bound on the probability that the minimum is not at one of the vertices of $K_{n,n/\alpha}$ already picked. This establishes (3.2). The density of the corresponding difference in $\mathcal{T}_{N,H}$ is

$$\frac{1}{\alpha} \exp\left(-\frac{x}{\alpha}\right).$$

The ratio of the conditional densities is at least

$$\left(1 - \frac{|\mathcal{T}_{N,H}|}{n/\alpha - N}\right) \left(1 - \frac{N}{n/\alpha}\right).$$

Doing this for all vertices of $\mathcal{T}_{N,H}$, and using the fact that under the restriction all edges are distinct, we get

$$\frac{d\nu_{N,H,n}}{d\nu_{N,H}} \geq \left(\left(1 - \frac{|\mathcal{T}_{N,H}|}{n/\alpha - N}\right) \left(1 - \frac{N}{n/\alpha}\right)\right)^{|\mathcal{T}_{N,H}|}.$$

The lower bound converges to 1 as $n \rightarrow \infty$. This implies that

$$\nu_{N,H,n} \xrightarrow{w} \nu_{N,H}, \tag{3.3}$$

as $n \rightarrow \infty$.

Now to get the local weak convergence result, fix a $\rho > 0$ and $\epsilon > 0$, and take N and H large enough such that uniformly for all $n > |\mathcal{T}_{N,H}|$ the probability that the ρ -neighborhood of the root of $K_{n,n/\alpha}$ is a subset of $T_{N,H,n}$ is greater than $1 - \epsilon$; under this event the ρ -neighborhood of the root of $K_{n,n/\alpha}$ converges weakly to the ρ -neighborhood of the root of \mathcal{T}_α^o by (3.3). This shows that conditional on the root having label o , \mathcal{T}_α^o is the local weak limit of $K_{n,n/\alpha}$. Similarly we can show that conditional on the root having label m , the local weak limit is \mathcal{T}_α^m . The probability of the label of the root to be o and m converges respectively to $\alpha/(1 + \alpha)$ and $1/(1 + \alpha)$. This completes the proof. \square

The above theorem is a generalization of Aldous’s result [1, Lemma 10] for the local weak limit of complete bipartite graphs in which both the vertex subsets have the same cardinality. A similar result was earlier established by Hajek [14, Sec. IV] for a class of sparse Erdős-Rényi random graphs.

The above theorem says that if we look at an arbitrary large but fixed neighborhood of the root of $K_{n,n/\alpha}$, then for large n , it looks like the corresponding neighborhood of the root of \mathcal{T}_α . This suggests that if boundary conditions can be ignored we may be able to relate optimal many-to-one matchings on $K_{n,n/\alpha}$ with an appropriate many-to-one matching on \mathcal{T}_α (to be precise, an optimal *involution invariant* many-to-one matching on \mathcal{T}_α). Furthermore, the local neighborhood of the root of $K_{n,n/\alpha}$ is a tree for large enough n (with high probability). So we may expect belief propagation on $K_{n,n/\alpha}$ to converge. Both these observations form the basis of the proof technique here, just like in the case of edge cover in [16]. In the matching case these ideas were developed by Aldous in [1, 4] to obtain the limit, and later by Salez and Shah [22] to show the applicability of belief propagation.

First, let us see that the expected cost of an optimal many-to-one matching on $K_{n,n/\alpha}$ can be written in terms of the weights of the edges from the root:

$$\mathbb{E} \left[\sum_{e \in M_n^*} \xi_{K_{n,n/\alpha}}(e) \right] = n \mathbb{E} \left[\sum_{\{\varphi, v\} \in M_n^*} \xi(\varphi, v) \mid \text{label}(\varphi) = o \right].$$

This is because the root, conditioned to have label o , is uniform among the n vertices of the subset A . We seek to relate the expectation on the right with the expectation corresponding to optimal many-to-one matchings on \mathcal{T}_α .

4 Recursive distributional equation

For a labeled tree \mathcal{T} with root ϕ , a many-to-one matching is a collection of edges such that vertices with label o have degree 1 and the vertices with label m have degree at least 1. The cost is the sum of the weights of the edges in the collection. This is well defined for a finite \mathcal{T} . Write $C(\mathcal{T})$ for the minimum cost of many-to-one matching on \mathcal{T} , and write $C(\mathcal{T} \setminus \phi)$ for the minimum cost of many-to-one matching on $\mathcal{T} \setminus \phi$, which is the forest obtained after removing the root ϕ from \mathcal{T} . Assume that the size of the tree is such that such many-to-one matchings exist. We derive a recursive relation involving these quantities.

If v is a child of the root ϕ of \mathcal{T} let \mathcal{T}_v denote the subtree containing v and all its descendants, with v as the root. Irrespective of the label of ϕ , we have

$$C(\mathcal{T} \setminus \phi) = \sum_{v \sim \phi} C(\mathcal{T}_v). \tag{4.1}$$

If ϕ has label o then among all matchings which match ϕ to a vertex u the minimum cost is

$$\xi(\phi, u) + \min \{C(\mathcal{T}_u), C(\mathcal{T}_u \setminus u)\} + \sum_{v \sim \phi, v \neq u} C(\mathcal{T}_v).$$

$C(\mathcal{T})$ is the minimum of the above over all children of the root:

$$C(\mathcal{T}) = \min_{u \sim \phi} \left\{ \xi(\phi, u) + \min \{C(\mathcal{T}_u), C(\mathcal{T}_u \setminus u)\} + \sum_{v \sim \phi, v \neq u} C(\mathcal{T}_v) \right\}. \tag{4.2}$$

Subtracting (4.1) from (4.2), we get

$$C(\mathcal{T}) - C(\mathcal{T} \setminus \phi) = \min_{u \sim \phi} \left\{ \xi(\phi, u) - (C(\mathcal{T}_u) - C(\mathcal{T}_u \setminus u))^+ \right\}. \tag{4.3}$$

If ϕ has label m then among all matchings which match ϕ to vertices in a nonempty set I the minimum cost is

$$\sum_{u \in I} (\xi(\phi, u) + C(\mathcal{T}_u \setminus u)) + \sum_{v \sim \phi, v \notin I} C(\mathcal{T}_v).$$

$C(\mathcal{T})$ is the minimum of the above over all nonempty I :

$$C(\mathcal{T}) = \min_{I \text{ nonempty}} \left\{ \sum_{u \in I} (\xi(\phi, u) + C(\mathcal{T}_u \setminus u)) + \sum_{v \sim \phi, v \notin I} C(\mathcal{T}_v) \right\}. \tag{4.4}$$

Subtracting (4.1) from (4.4), we get

$$C(\mathcal{T}) - C(\mathcal{T} \setminus \phi) = \min_{I \text{ nonempty}} \left\{ \sum_{u \in I} [\xi(\phi, u) - (C(\mathcal{T}_u) - C(\mathcal{T}_u \setminus u))] \right\},$$

the positive part of which is

$$(C(\mathcal{T}) - C(\mathcal{T} \setminus \phi))^+ = \min_{u \sim \phi} \left\{ [\xi(\phi, u) - (C(\mathcal{T}_u) - C(\mathcal{T}_u \setminus u))]^+ \right\}. \tag{4.5}$$

Now consider what happens when we take \mathcal{T} to be the infinite tree \mathcal{T}_α . The quantities $C(\mathcal{T}_\alpha)$ and $C(\mathcal{T}_\alpha \setminus \phi)$ become infinite, and the relations (4.3) and (4.5) are meaningless. Observe that conditional on the label of the root ϕ the subtrees $\mathcal{T}_{\alpha,i}, i \sim \phi$ are i.i.d., and the conditional distribution is same as the distribution of \mathcal{T}_α^m (\mathcal{T}_α^o) given that ϕ has label

$o(m)$. With this distribution structure of the subtrees, (4.3) and (4.5) motivate us to write a two-step *recursive distributional equation* (RDE): an equation in two distributions μ^o on \mathbf{R} and μ^m on \mathbf{R}_+ that satisfy

$$\begin{aligned} X^o &\stackrel{D}{=} \min_i \{\alpha \xi_i - X_i^m\}, \\ X^m &\stackrel{D}{=} \min_i \left\{ (\xi_i - X_i^o)^+ \right\}, \end{aligned} \tag{4.6}$$

such that

- (a) $\{\xi_i, i \geq 1\}$ is a Poisson process of rate 1 on $[0, \infty)$,
- (b) X^o has distribution μ^o ,
- (c) $X_i^m, i \geq 1$ are i.i.d. random variables with distribution μ^m , independent of the Poisson process $\{\xi_i\}$
- (d) X^m has distribution μ^m ,
- (e) $X_i^o, i \geq 1$ are i.i.d. random variables with distribution μ^o , independent of the Poisson process $\{\xi_i\}$.

We will call (μ^o, μ^m) the solution to the RDE (4.6).

A recursive distributional equation is the primary tool that leads to the optimal solution on the infinite tree \mathcal{T}_α . The limit value of the expected optimal cost is obtained by calculations from the solution of the RDE. Recursive distributional equations come up in the related problems of random matching [4] and edge cover. Aldous and Bandyopadhyay [5] present a survey of this type of RDEs and introduce several RDE concepts that we will use later.

If $\mathcal{P}(S)$ denotes the space of probability measures on a space S , a recursive distributional equation (RDE) is defined in [5] as a fixed-point equation on $\mathcal{P}(S)$ of the form

$$X \stackrel{D}{=} g(\xi; (X_j, 1 \leq j < N)), \tag{4.7}$$

where $X_j, j \geq 1$, are i.i.d. S -valued random variables having the same distribution as X , and are independent of the pair (ξ, N) , ξ is a random variable on some space, and N is a random variable on $\mathbf{N} \cup \{+\infty\}$. g is a given S -valued function. A solution to the RDE is a common distribution of $X, X_j, j \geq 1$ satisfying (4.7).

The two-step equation (4.6) can easily be expressed in the form (4.7) by composing the two steps of (4.6) into a single step. One can then solve for a fixed point distribution, say μ^o , which then automatically yields μ^m via one of the steps of (4.6).

We can use the relation (4.7) to construct a tree indexed stochastic process, say $X_i, i \in \mathcal{V}$, which is called a recursive tree process (RTP) [5]. Associate to each vertex $i \in \mathcal{V}$, an independent copy (ξ_i, N_i) of the pair (ξ, N) , and require X_i to satisfy

$$X_i \stackrel{D}{=} g(\xi_i; (X_{i,j}, 1 \leq j < N_i)),$$

with X_i independent of $\{(\xi_{i'}, N_{i'}) \mid |i'| < |i|\}$. If $\mu \in \mathcal{P}(S)$ is a solution to the RDE (4.7), there exists a stationary RTP, i.e., each X_i is distributed as μ . Such a process is called an invariant RTP with marginal distribution μ .

4.1 Solution to the RDE

Write F and G for the survival functions of the distributions μ^o and μ^m respectively, i.e., $F(t) = P\{X^o > t\}$ and $G(t) = P\{X^m > t\}$. F and G satisfy the following:

$$F(t) = \exp\left(-\int_{-t}^{\infty} \frac{G(z)}{\alpha} dz\right), \quad t \in \mathbf{R}; \tag{4.8}$$

$$G(t) = \begin{cases} \exp\left(-\int_{-t}^{\infty} F(z) dz\right) & \text{if } t \geq 0, \\ 1 & \text{if } t < 0. \end{cases} \tag{4.9}$$

For $t \geq 0$, we can write

$$\begin{aligned} F(t) &= \exp\left(-\frac{1}{\alpha}\left(\int_{-t}^0 dz - \int_0^\infty G(z) dz\right)\right) \\ &= e^{-w_m/\alpha} e^{-t/\alpha}, \end{aligned} \tag{4.10}$$

where $w_m = EX^m$. Let $w_o = E(X^o)^+$. By (4.10), we have $w_o = \alpha e^{-w_m/\alpha}$. We can now rewrite (4.10) using w_o as

$$F(t) = \frac{w_o}{\alpha} e^{-t/\alpha}, \quad t \geq 0. \tag{4.11}$$

From (4.8) and (4.9), F and G are differentiable for all $t \in \mathbf{R} \setminus \{0\}$. Differentiating (4.8) and (4.9), we get

$$F'(-t) = -\frac{1}{\alpha} F(-t)G(t), \quad t \neq 0, \tag{4.12}$$

$$G'(t) = -G(t)F(-t), \quad t > 0. \tag{4.13}$$

These two equations imply that

$$\alpha F(-t) + G(t) = \text{constant}, \quad t > 0.$$

The validity of this equation can be extended to $t = 0$ because F is continuous for all t and G is right-continuous with left limits at $t = 0$. Taking $t \rightarrow \infty$ in the above relation shows that the constant is equal to α , and so

$$F(-t) = 1 - \frac{G(t)}{\alpha}, \quad t \geq 0. \tag{4.14}$$

Substituting this in (4.13) results in the differential equation

$$G'(t) = -G(t) \left(1 - \frac{G(t)}{\alpha}\right), \quad t > 0. \tag{4.15}$$

To get the boundary conditions take $t = 0$ in (4.9), (4.11), and (4.14) to get

$$\begin{aligned} G(0) &= e^{-w_o}, \\ F(0) &= w_o/\alpha, \\ F(0) &= 1 - G(0)/\alpha. \end{aligned}$$

These imply $G(0) = \alpha - w_o$ and $\alpha = w_o + e^{-w_o}$. Since $\alpha > 1$, the latter equation determines a unique $w_o > 0$.

Integration of the differential equation (4.15) gives

$$\begin{aligned} G(t) &= \alpha \frac{G(0)e^{-t}}{\alpha - G(0) + G(0)e^{-t}} \\ &= \frac{\alpha}{1 + w_o e^{w_o} e^t} \end{aligned}$$

for $t \geq 0$. Using (4.14), we get F for $(-\infty, 0]$ as well. We thus obtain the following expressions for the distributions F and G :

$$F(t) = \begin{cases} \frac{w_o}{\alpha} e^{-t/\alpha} & \text{if } t \geq 0, \\ 1 - \frac{1}{1 + \gamma e^{-t}} & \text{if } t < 0; \end{cases} \tag{4.16}$$

$$G(t) = \begin{cases} \frac{\alpha}{1 + \gamma e^t} & \text{if } t \geq 0, \\ 1 & \text{if } t < 0, \end{cases} \tag{4.17}$$

where $\gamma = w_o e^{w_o}$.

For future use, we note that the density of μ^o is:

$$f(t) = \begin{cases} \frac{w_o}{\alpha^2} e^{-t/\alpha} & \text{if } t > 0, \\ \frac{\gamma e^{-t}}{(1+\gamma e^{-t})^2} & \text{if } t \leq 0. \end{cases} \tag{4.18}$$

5 Optimal involution invariant many-to-one matching on the infinite tree

The RDE (4.6) guides us in the construction of a many-to-one matching on \mathcal{T}_α ; the following Lemma forms the first step.

Lemma 5.1. *There exists a process*

$$\left(\mathcal{T}_\alpha, (\xi_e, e \in E(\mathcal{T}_\alpha)), (X(\vec{e}), \vec{e} \in \vec{E}(\mathcal{T}_\alpha)) \right)$$

where \mathcal{T}_α is an infinite tree with edge-lengths $\{\xi_e, e \in E(\mathcal{T}_\alpha)\}$, generated from the distribution (3.1), and $\{X(\vec{e}), \vec{e} \in \vec{E}(\mathcal{T}_\alpha)\}$ is a stochastic process satisfying the following properties.

(a) For each directed edge $(u, v) \in \vec{E}(\mathcal{T}_\alpha)$,

$$X(u, v) = \begin{cases} \min \left\{ \xi(v, w) - X(v, w) : (v, w) \in \vec{E}(\mathcal{T}_\alpha), w \neq u \right\} & \text{if label}(v) = o, \\ \min \left\{ (\xi(v, w) - X(v, w))^+ : (v, w) \in \vec{E}(\mathcal{T}_\alpha), w \neq u \right\} & \text{if label}(v) = m. \end{cases} \tag{5.1}$$

(b) If $(u, v) \in \vec{E}(\mathcal{T}_\alpha)$ is directed away from the root of \mathcal{T}_α , then $X(u, v)$ has distribution F conditioned on v having label o , and distribution G conditioned on v having label m ; F and G are as in (4.16) and (4.17).

(c) If $(u, v) \in \vec{E}(\mathcal{T}_\alpha)$ the random variables $X(u, v)$ and $X(v, u)$ are independent.

(d) For a fixed $z > 0$, conditional on the event that there exists an edge of length z at the root, say $\{\phi, v_z\}$, and conditional on the root having label m , the random variables $X(\phi, v_z)$ and $X(v_z, \phi)$ are independent random variables having distributions F and G respectively; if the root is conditioned to have label o , the distributions are reversed.

Proof. This Lemma is the analogue of Lemma 5.8 of [3] and Lemma 1 of [16]. It has essentially the same proof, with appropriate changes to handle vertex labels. We omit the details. \square

We use the process $\{X(\vec{e})\}$ to construct a many-to-one matching \mathcal{M}_{opt} on \mathcal{T}_α . For each vertex v of \mathcal{T}_α , define a set

$$\mathcal{M}_{\text{opt}}(v) = \begin{cases} \arg \min_{y \sim v} \{ \xi(v, y) - X(v, y) \} & \text{if label}(v) = o, \\ \arg \min_{y \sim v} \{ (\xi(v, y) - X(v, y))^+ \} & \text{if label}(v) = m. \end{cases} \tag{5.2}$$

By continuity of the edge-weight distribution, when v has label o , almost surely $\mathcal{M}_{\text{opt}}(v)$ will be a singleton, whereas when v has label m , $\mathcal{M}_{\text{opt}}(v)$ may have more than one element – e.g., all vertices y for which the difference $\xi(v, y) - X(v, y)$ is negative.

Define

$$\mathcal{M}_{\text{opt}} = \bigcup_v \{ \{v, w\} : w \in \mathcal{M}_{\text{opt}}(v) \}.$$

The following lemma proves that the collection \mathcal{M}_{opt} yields a consistent many-to-one matching – if v with label o picks w with label m , then w is the only m -labeled vertex that picks v .

Lemma 5.2. *For any two vertices v, w of \mathcal{T}_α , we have*

$$w \in \mathcal{M}_{\text{opt}}(v) \iff \xi(v, w) < X(v, w) + X(w, v).$$

As a consequence,

$$w \in \mathcal{M}_{\text{opt}}(v) \iff v \in \mathcal{M}_{\text{opt}}(w).$$

Proof. Let $w \in \mathcal{M}_{\text{opt}}(v)$. Suppose v has label o . Then

$$w = \arg \min \{ \xi(v, y) - X(v, y), y \sim v \}.$$

This implies

$$\begin{aligned} \xi(v, w) - X(v, w) &= \min \{ \xi(v, y) - X(v, y), y \sim v \} \\ &< \min \{ \xi(v, y) - X(v, y), y \sim v, y \neq w \} = X(w, v). \end{aligned}$$

Conversely, if $\xi(v, w) < X(v, w) + X(w, v)$ then using the definition of $X(w, v)$, it follows that $w = \arg \min \{ \xi(v, y) - X(v, y), y \sim v \}$, and so $w \in \mathcal{M}_{\text{opt}}(v)$.

Now suppose v has label m . Then there are two cases. (1) $\xi(v, w) - X(v, w) < 0$. In this case $\xi(v, w) < X(v, w) + X(w, v)$ since $X(w, v) \geq 0$. Alternatively, (2) $\xi(v, w) - X(v, w) > 0$, in which case, by (5.2), w is the unique vertex where the minimum is attained. Hence

$$0 < \xi(v, w) - X(v, w) < \xi(v, y) - X(v, y) \quad \text{for all } y \sim v, y \neq w,$$

and so

$$\xi(v, w) - X(v, w) < \min \{ \xi(v, y) - X(v, y), y \sim v, y \neq w \} = X(w, v),$$

which then yields $\xi(v, w) < X(v, w) + X(w, v)$. Again, it is easy to see that if $\xi(v, w) < X(v, w) + X(w, v)$ then $w \in \arg \min \{ (\xi(v, y) - X(v, y))^+, y \sim v \}$, which implies $w \in \mathcal{M}_{\text{opt}}(v)$.

Thus, we have established the first statement of the lemma, which is

$$w \in \mathcal{M}_{\text{opt}}(v) \iff \xi(v, w) < X(v, w) + X(w, v).$$

The condition on the right-hand side above is symmetric in v, w , and hence the second statement of the lemma is proved. \square

5.1 Evaluating the cost

The next result is merely to demonstrate that the limiting expected minimum cost can be computed using numerical methods.

Define

$$c_*^\alpha = -\text{Li}_2(-\gamma^{-1}) - \frac{1}{2} \log^2(1 + \gamma^{-1}) + w_o \log(1 + \gamma^{-1}) + w_o,$$

where Li_2 is the dilogarithm: $\text{Li}_2(z) = -\int_0^z \frac{\log(1-t)}{t} dt = \sum_{k=1}^\infty \frac{z^k}{k^2} = z \int_0^\infty t/(e^t - z) dt$, $\gamma = w_o e^{w_o}$, and w_o is the positive solution to $w_o + e^{-w_o} = \alpha$.

Theorem 5.3.

$$\mathbb{E} \left[\sum_{v \in \mathcal{M}_{\text{opt}}(\phi)} \xi(\phi, v) \mid \text{label}(\phi) = o \right] = c_*^\alpha.$$

Proof. Condition on the event that the root ϕ has label o . Now fix a $z > 0$, and condition on the event that there is a neighbor v_z of ϕ with $\xi(\phi, v_z) = z$. Call this event E_z . If we condition a Poisson process to have a point at some location, then the conditional process on removing this point is again a Poisson process with the same intensity. This shows that under E_z , $X(\phi, v_z)$ and $X(v_z, \phi)$ have the distributions G and F respectively. They are also independent. Using these observations, the conditional expected cost can be written as

$$\begin{aligned}
 & \mathbb{E} \left[\sum_{v \in \mathcal{M}_{\text{opt}}(\phi)} \xi(\phi, v) \mid \text{label}(\phi) = o \right] \\
 &= \int_{z=0}^{\infty} z \mathbb{P} \{z < X(\phi, v_z) + X(v_z, \phi)\} \frac{1}{\alpha} dz \\
 &= \int_{z=0}^{\infty} \int_{x=-\infty}^{\infty} z G(z-x) f(x) \frac{1}{\alpha} dx dz \\
 &= \int_{z=0}^{\infty} \int_{x=-\infty}^0 z \frac{1}{1 + \gamma e^{z-x}} \frac{\gamma e^{-x}}{(1 + \gamma e^{-x})^2} dx dz \\
 &\quad + \int_{z=0}^{\infty} \int_{x=0}^z z \frac{1}{1 + \gamma e^{z-x}} \frac{w_o}{\alpha^2} e^{-x/\alpha} dx dz + \int_{z=0}^{\infty} \int_{x=z}^{\infty} z f(x) \frac{1}{\alpha} dx dz \\
 &\hspace{15em} \text{(using the expression for the density of } \mu^o \text{ (4.18))} \\
 &= \int_{x=-\infty}^0 -\frac{\gamma e^{-x}}{(1 + \gamma e^{-x})^2} \text{Li}_2(-\gamma^{-1} e^x) dx + \frac{1}{\alpha^2} \int_{x=0}^{\infty} \int_{t=0}^{\infty} \frac{t+x}{1 + \gamma e^t} w_o e^{-x/\alpha} dt dx \\
 &\quad + \frac{1}{\alpha} \int_{z=0}^{\infty} z F(z) dz \quad \text{(taking } t = z - x) \\
 &= \int_{s=\gamma}^{\infty} -\frac{1}{(1+s)^2} \text{Li}_2(-s^{-1}) ds \\
 &\quad + \frac{1}{\alpha^2} \int_{x=0}^{\infty} w_o e^{-x/\alpha} (-\text{Li}_2(-\gamma^{-1}) + x \log(1 + \gamma^{-1})) dx \\
 &\quad + \int_{z=0}^{\infty} z \frac{w_o}{\alpha^2} e^{-z/\alpha} dz \quad \text{(taking } s = \gamma e^{-x}) \\
 &= \frac{\text{Li}_2(-\gamma^{-1})}{1 + \gamma} + \frac{1}{2} \log^2(1 + \gamma^{-1}) - \frac{w_o}{\alpha} \text{Li}_2(-\gamma^{-1}) \\
 &\quad + w_o \log(1 + \gamma^{-1}) + w_o \\
 &= -\text{Li}_2(-\gamma^{-1}) - \frac{1}{2} \log^2(1 + \gamma^{-1}) + w_o \log(1 + \gamma^{-1}) + w_o \\
 &\quad \left(\text{using } \frac{\alpha}{1 + \gamma} = \alpha - w_o \right). \quad \square
 \end{aligned}$$

Note that as $\alpha \rightarrow 1$, the expected cost converges to $\pi^2/6$, which is the limiting expected minimum cost of a standard matching on $K_{n,n}$.

5.2 Optimality in the class of involution invariant many-to-one matchings

The local weak limit of a sequence of uniformly rooted random network satisfies a property called involution invariance [3, 2]. Informally, involution invariance and its equivalent property called unimodularity specify that a distribution on the space of rooted networks \mathcal{G}_* does not depend on the choice of the root.

Refer to Section 5 of [16] for a discussion on involution invariance as it relates to our setting; a very general study of involution invariance is central to [2].

Let \mathcal{G}_{**} be the space of isomorphism classes of connected geometric networks with a distinguished directed edge. An element of \mathcal{G}_{**} (G, o, x) will denote the isomorphism class corresponding to a network G and distinguished vertex pair (o, x) , where $\{o, x\}$ is an edge in G . A probability measure μ on \mathcal{G}_{**} is called involution invariant if it satisfies the following for all Borel $f : \mathcal{G}_{**} \rightarrow [0, \infty]$

$$\int \sum_{x \in V(G)} f(G, o, x) d\mu(G, o) = \int \sum_{x \in V(G)} f(G, x, o) d\mu(G, o).$$

We can represent a many-to-one matching \mathcal{M} on a graph G by a membership map on the edge-set of G : $e \mapsto \mathbb{1}_{\{e \in \mathcal{M}\}}$. We call a random many-to-one matching on a random graph G involution invariant if the distribution of G with the above map on its edges is involution invariant. Following the argument in the last paragraph of Section 5 of [16] we can restrict our attention to involution invariant many-to-one matchings on \mathcal{T}_α . We now show that our candidate many-to-one matching \mathcal{M}_{opt} has the minimum expected cost among such many-to-one matchings.

Theorem 5.4. *Let \mathcal{M} be an involution invariant many-to-one matching on \mathcal{T}_α . Write $\mathcal{M}(\phi)$ for the set of vertices of \mathcal{T}_α adjacent to the root ϕ in \mathcal{M} (This set will be a singleton if label of ϕ is o). Then*

$$\mathbb{E} \left[\sum_{v \in \mathcal{M}(\phi)} \xi(\phi, v) \right] \geq \mathbb{E} \left[\sum_{v \in \mathcal{M}_{\text{opt}}(\phi)} \xi(\phi, v) \right].$$

Before we start to prove Theorem 5.4, we examine a consequence of Theorem 5.4.

Corollary 5.5.

$$\mathbb{E} \left[\sum_{v \in \mathcal{M}(\phi)} \xi(\phi, v) \mid \text{label}(\phi) = o \right] \geq \mathbb{E} \left[\sum_{v \in \mathcal{M}_{\text{opt}}(\phi)} \xi(\phi, v) \mid \text{label}(\phi) = o \right].$$

Proof. Involution invariance of \mathcal{M} implies that

$$\mathbb{E} \left[\sum_{v \in \mathcal{M}(\phi)} \xi(\phi, v) \mathbb{1}_{\{\text{label}(\phi)=o\}} \right] = \mathbb{E} \left[\sum_{v \in \mathcal{M}(\phi)} \xi(\phi, v) \mathbb{1}_{\{\text{label}(\phi)=m\}} \right].$$

Using that $\mathbb{P} \{ \text{label}(\phi) = o \} = \frac{\alpha}{1+\alpha}$ and $\mathbb{P} \{ \text{label}(\phi) = m \} = \frac{1}{1+\alpha}$, we get

$$\mathbb{E} \left[\sum_{v \in \mathcal{M}(\phi)} \xi(\phi, v) \right] = \frac{2\alpha}{1+\alpha} \mathbb{E} \left[\sum_{v \in \mathcal{M}(\phi)} \xi(\phi, v) \mid \text{label}(\phi) = o \right].$$

The same holds for \mathcal{M}_{opt} too, and so the corollary follows from Theorem 5.4. □

We will need some definitions for the proof of Theorem 5.4. For each directed edge (v, w) of \mathcal{T}_α , if v has label o and w has label m , define a random variable

$$Y(v, w) = \min \left\{ \sum_{y \in A} (\xi(w, y) - X(w, y)) \mid \begin{array}{l} A \subset N_w \setminus \{v\}, \\ A \text{ nonempty} \end{array} \right\}, \tag{5.3}$$

where N_w is the set of neighbors of w . It is easy to see that this random variable can be written as

$$Y(v, w) = \begin{cases} \min_{y \sim w, y \neq v} \{ \xi(w, y) - X(w, y) \} & \text{if } \xi(w, y) - X(w, y) \geq 0 \text{ for all } y \sim w, y \neq v, \\ \sum_{y \sim w, y \neq v} (\xi(w, y) - X(w, y)) \mathbb{1}_{\{\xi(w, y) - X(w, y) < 0\}} & \text{otherwise.} \end{cases}$$

Note that $(Y(v, w))^+ = X(v, w)$. If v has label o and w has label m , define

$$Y(w, v) = X(w, v).$$

Suppose that $E \left[\sum_{v \in \mathcal{M}(\phi)} \xi(\phi, v) \right] < \infty$. Then $\mathcal{M}(\phi)$ is a finite set with probability 1 because $\{\xi(\phi, v), v \sim \phi\}$ are points of a Poisson process of rate 1 or $1/\alpha$. For such a many-to-one matching \mathcal{M} , define

$$A(\mathcal{M}) = \sum_{v \in \mathcal{M}(\phi)} X(\phi, v) + \max_{v \notin \mathcal{M}(\phi), v \sim \phi} Y(v, \phi). \tag{5.4}$$

The following two lemmas will be used to prove Theorem 5.4.

Lemma 5.6. *Let \mathcal{M} be a many-to-one matching on \mathcal{T}_α such that*

$$E \left[\sum_{v \in \mathcal{M}(\phi)} \xi(\phi, v) \right] < \infty.$$

Then almost surely,

$$\sum_{v \in \mathcal{M}(\phi)} \xi(\phi, v) \geq A(\mathcal{M}).$$

Furthermore,

$$\sum_{v \in \mathcal{M}_{\text{opt}}(\phi)} \xi(\phi, v) = A(\mathcal{M}_{\text{opt}}).$$

Lemma 5.7. *Let \mathcal{M} be a many-to-one matching on \mathcal{T}_α such that*

$$E \left[\sum_{v \in \mathcal{M}(\phi)} \xi(\phi, v) \right] < \infty.$$

If \mathcal{M} is involution invariant, we have $E[A(\mathcal{M})] \geq E[A(\mathcal{M}_{\text{opt}})]$.

Proof of Theorem 5.4. If $E \left[\sum_{v \in \mathcal{M}(\phi)} \xi(\phi, v) \right] = \infty$ the statement of the theorem is trivially true. Assume that it is finite. We are now in a position to apply Lemmas 5.6 and 5.7 as follows to get the result:

$$\begin{aligned} E \left[\sum_{v \in \mathcal{M}(\phi)} \xi(\phi, v) \right] &\geq E[A(\mathcal{M})] && \text{(Lemma 5.6)} \\ &\geq E[A(\mathcal{M}_{\text{opt}})] && \text{(Lemma 5.7)} \\ &= E \left[\sum_{v \in \mathcal{M}_{\text{opt}}(\phi)} \xi(\phi, v) \right]. && \text{(Lemma 5.6)} \quad \square \end{aligned}$$

Let us now complete the proofs of Lemmas 5.6 and 5.7.

Proof of Lemma 5.6. Suppose ϕ has label o . Then $\mathcal{M}(\phi)$ is a singleton, say $\{u\}$. For $v \sim \phi$, we have by (5.2) and (5.1),

$$Y(v, \phi) = X(v, \phi) \leq \xi(\phi, y) - X(\phi, y), y \sim \phi, y \neq v.$$

If $v \neq u$, replacing y with u on the right-hand side, we get

$$Y(v, \phi) \leq \xi(\phi, u) - X(\phi, u).$$

Therefore,

$$\max_{v \notin \mathcal{M}(\phi), v \sim \phi} Y(v, \phi) \leq \xi(\phi, u) - X(\phi, u).$$

This gives

$$\sum_{v \in \mathcal{M}(\phi)} \xi(\phi, v) \geq A(\mathcal{M}).$$

If $\mathcal{M}_{\text{opt}}(\phi) = \{u^*\}$ then

$$u^* = \arg \min_{y \sim \phi} \{\xi(\phi, y) - X(\phi, y)\}$$

implies that, for any $v \neq u^*, v \sim \phi$,

$$Y(v, \phi) = X(v, \phi) = \xi(\phi, u^*) - X(\phi, u^*).$$

Therefore,

$$\max_{v \notin \mathcal{M}_{\text{opt}}(\phi), v \sim \phi} Y(v, \phi) = \xi(\phi, u^*) - X(\phi, u^*),$$

and so

$$\sum_{v \in \mathcal{M}_{\text{opt}}(\phi)} \xi(\phi, v) = A(\mathcal{M}_{\text{opt}}).$$

Now suppose ϕ has label m . Then by (5.3),

$$Y(v, \phi) \leq \sum_{y \in A} (\xi(\phi, y) - X(\phi, y))$$

for all $A \subset N_\phi \setminus \{v\}$, A nonempty. In particular, for any $v \notin \mathcal{M}(\phi)$, we can choose $A = \mathcal{M}(\phi)$ to obtain

$$Y(v, \phi) \leq \sum_{y \in \mathcal{M}(\phi)} (\xi(\phi, y) - X(\phi, y)).$$

This implies

$$\max_{v \notin \mathcal{M}(\phi), v \sim \phi} Y(v, \phi) \leq \sum_{y \in \mathcal{M}(\phi)} (\xi(\phi, y) - X(\phi, y)). \quad (5.5)$$

Thanks to the finite expectation assumption in the lemma, $\mathcal{M}(\phi)$ is a finite set almost surely, and so $\sum_{y \in \mathcal{M}(\phi)} X(\phi, y)$ is finite almost surely. Rearrangement of (5.5) then yields

$$\sum_{v \in \mathcal{M}(\phi)} \xi(\phi, v) \geq A(\mathcal{M}).$$

We can write $\mathcal{M}_{\text{opt}}(\phi)$ as

$$\mathcal{M}_{\text{opt}}(\phi) = \arg \min_A \left\{ \sum_{y \in A} (\xi(\phi, y) - X(\phi, y)) : A \subset N_\phi, A \text{ nonempty} \right\}. \quad (5.6)$$

From (5.3) and (5.6), for any $v \notin \mathcal{M}_{\text{opt}}(\phi)$, we have

$$Y(v, \phi) = \sum_{y \in \mathcal{M}_{\text{opt}}(\phi)} (\xi(\phi, y) - X(\phi, y)),$$

and hence

$$\max_{v \notin \mathcal{M}_{\text{opt}}(\phi), v \sim \phi} Y(v, \phi) = \sum_{y \in \mathcal{M}_{\text{opt}}(\phi)} (\xi(\phi, y) - X(\phi, y)).$$

It follows by rearrangement that

$$\sum_{v \in \mathcal{M}_{\text{opt}}(\phi)} \xi(\phi, v) = A(\mathcal{M}_{\text{opt}}). \quad \square$$

Proof of Lemma 5.7. Define

$$\tilde{A}(\mathcal{M}) = \sum_{v \in \mathcal{M}(\phi)} X(v, \phi) + \max_{v \notin \mathcal{M}(\phi), v \sim \phi} Y(v, \phi). \quad (5.7)$$

$\tilde{A}(\mathcal{M})$ is similar to $A(\mathcal{M})$, but with the arguments of X alone reversed. We will prove Lemma 5.7 by showing the following two results:

(a) For an involution invariant many-to-one matching \mathcal{M}

$$\mathbb{E} [\tilde{A}(\mathcal{M})] = \mathbb{E} [A(\mathcal{M})]. \quad (5.8)$$

(b) Almost surely,

$$\tilde{A}(\mathcal{M}) \geq \tilde{A}(\mathcal{M}_{\text{opt}}). \quad (5.9)$$

We first prove (5.8). First, by involution invariance of \mathcal{M} , we have

$$\mathbb{E} \left[\sum_{v \in \mathcal{M}(\phi)} X(\phi, v) \right] = \mathbb{E} \left[\sum_{v \in \mathcal{M}(\phi)} X(v, \phi) \right]. \quad (5.10)$$

Indeed, the left-hand side equals

$$\int_{\mathcal{G}_*} \sum_{v \sim \phi} X(\phi, v) \, d\mu_{\mathcal{M}}([G, \phi]),$$

where $\mu_{\mathcal{M}}$ is the probability measure on \mathcal{G}_* corresponding to \mathcal{M} . By involution invariance, this equals

$$\int_{\mathcal{G}_*} \sum_{v \sim \phi} X(v, \phi) \, d\mu_{\mathcal{M}}([G, \phi]),$$

which is equal to the right-hand side of (5.10). Thanks to the finite expectation assumption of the lemma, we saw in the proof of Lemma 5.6 that the term $\max_{v \notin \mathcal{M}(\phi), v \sim \phi} Y(v, \phi)$ is finite almost surely. Now observe that $A(\mathcal{M})$ (respectively $\tilde{A}(\mathcal{M})$) is obtained by adding the almost surely finite random variable $\max_{v \notin \mathcal{M}(\phi), v \sim \phi} Y(v, \phi)$ to the random variable which is the argument of the expectation in the left-hand side of (5.10) (respectively the right-hand side of (5.10)). Taking expectation and using the equality in (5.10), we get (5.8).

Now we will prove (5.9). Suppose that ϕ has label o . Both $\mathcal{M}(\phi)$ and $\mathcal{M}_{\text{opt}}(\phi)$ are singleton sets: suppose $\mathcal{M}(\phi) = \{u\}$ and $\mathcal{M}_{\text{opt}}(\phi) = \{u^*\}$. If $u = u^*$ then $\tilde{A}(\mathcal{M}) = \tilde{A}(\mathcal{M}_{\text{opt}})$. Now suppose $u \neq u^*$. Then

$$Y(v, \phi) = X(v, \phi) = \xi(\phi, u^*) - X(\phi, u^*)$$

for all $v \neq u^*, v \sim \phi$. For all such v ,

$$Y(v, \phi) = X(u, \phi) \leq X(u^*, \phi) = Y(u^*, \phi).$$

Consequently,

$$\begin{aligned} \max_{v \neq u^*, v \sim \phi} Y(v, \phi) &= X(u, \phi), \\ \max_{v \neq u, v \sim \phi} Y(v, \phi) &= X(u^*, \phi). \end{aligned}$$

This implies that

$$\begin{aligned}\tilde{A}(\mathcal{M}_{\text{opt}}) &= X(u^*, \phi) + X(u, \phi), \\ \tilde{A}(\mathcal{M}) &= X(u, \phi) + X(u^*, \phi).\end{aligned}$$

Thus $\tilde{A}(\mathcal{M}) = \tilde{A}(\mathcal{M}_{\text{opt}})$.

Now suppose that ϕ has label m . First condition on the event

$$L_1 = \{|\mathcal{M}_{\text{opt}}(\phi)| > 1\}.$$

Observe that, under L_1 , $\xi(\phi, y) - X(\phi, y) < 0$, $y \sim \phi$ if and only if $y \in \mathcal{M}_{\text{opt}}(\phi)$, and there are at least two such y . Then, by (5.1),

$$X(v, \phi) = 0 \text{ for all } v \sim \phi. \tag{5.11}$$

Also from (5.3) and (5.6), for any $v \sim \phi$, we have

$$Y(v, \phi) \geq \sum_{y \in \mathcal{M}_{\text{opt}}(\phi)} (\xi(\phi, y) - X(\phi, y)) = Y(w, \phi)$$

if $w \notin \mathcal{M}_{\text{opt}}(\phi)$. This implies

$$Y(v, \phi) \geq \max_{w \notin \mathcal{M}_{\text{opt}}(\phi), w \sim \phi} Y(w, \phi) \text{ for all } v \sim \phi.$$

In particular,

$$\max_{v \notin \mathcal{M}(\phi), v \sim \phi} Y(v, \phi) \geq \max_{w \notin \mathcal{M}_{\text{opt}}(\phi), w \sim \phi} Y(w, \phi). \tag{5.12}$$

Combining (5.11) and (5.12) gives

$$\sum_{v \in \mathcal{M}(\phi)} X(v, \phi) + \max_{v \notin \mathcal{M}(\phi), v \sim \phi} Y(v, \phi) \geq \sum_{v \in \mathcal{M}_{\text{opt}}(\phi)} X(v, \phi) + \max_{v \notin \mathcal{M}_{\text{opt}}(\phi), v \sim \phi} Y(v, \phi).$$

Thus $\tilde{A}(\mathcal{M}) \geq \tilde{A}(\mathcal{M}_{\text{opt}})$ under L_1 .

Now consider the event $L_2 = \{|\mathcal{M}_{\text{opt}}(\phi)| = 1\}$. Let

$$\begin{aligned}X_\phi^{(1)} &= \min_{v \sim \phi} (\xi(\phi, v) - X(\phi, v)), \\ X_\phi^{(2)} &= \min_{v \sim \phi}^{(2)} (\xi(\phi, v) - X(\phi, v)),\end{aligned}$$

where $\min^{(2)}$ stands for the second minimum.

Let $\mathcal{M}_{\text{opt}}(\phi) = \{u^*\}$. Then $X(u^*, \phi) = X_\phi^{(2)}$ (the second minimum is nonnegative), and for $v \in \mathcal{M}(\phi) \setminus \mathcal{M}_{\text{opt}}(\phi)$, $X(v, \phi) = (X_\phi^{(1)})^+$. So we get

$$\sum_{v \in \mathcal{M}(\phi)} X(v, \phi) - \sum_{v \in \mathcal{M}_{\text{opt}}(\phi)} X(v, \phi) = \sum_{v \in \mathcal{M}(\phi) \setminus \mathcal{M}_{\text{opt}}(\phi)} \left(X_\phi^{(1)} \right)^+ - X_\phi^{(2)} \mathbb{1}_{\{u^* \notin \mathcal{M}(\phi)\}}. \tag{5.13}$$

From (5.3), if $v \neq u^*$, then $Y(v, \phi) = X_\phi^{(1)}$. Also $Y(u^*, \phi) = X_\phi^{(2)}$. Since $X_\phi^{(2)} \geq X_\phi^{(1)}$, we get

$$\max_{v \notin \mathcal{M}(\phi), v \sim \phi} Y(v, \phi) = X_\phi^{(2)} \mathbb{1}_{\{u^* \notin \mathcal{M}(\phi)\}} + X_\phi^{(1)} \mathbb{1}_{\{u^* \in \mathcal{M}(\phi)\}},$$

and

$$\max_{v \notin \mathcal{M}_{\text{opt}}(\phi), v \sim \phi} Y(v, \phi) = X_\phi^{(1)}.$$

Therefore,

$$\max_{v \notin \mathcal{M}(\phi), v \sim \phi} Y(v, \phi) - \max_{v \notin \mathcal{M}_{\text{opt}}(\phi), v \sim \phi} Y(v, \phi) = \left(X_\phi^{(2)} - X_\phi^{(1)} \right) \mathbb{1}_{\{u^* \notin \mathcal{M}(\phi)\}}. \quad (5.14)$$

Adding (5.13) and (5.14), and canceling $X_\phi^{(2)} \mathbb{1}_{\{u^* \notin \mathcal{M}(\phi)\}}$, we get

$$\begin{aligned} & \sum_{v \in \mathcal{M}(\phi)} X(v, \phi) + \max_{v \notin \mathcal{M}(\phi), v \sim \phi} Y(v, \phi) \\ & \quad - \sum_{v \in \mathcal{M}_{\text{opt}}(\phi)} X(v, \phi) - \max_{v \notin \mathcal{M}_{\text{opt}}(\phi), v \sim \phi} Y(v, \phi) \\ & = \sum_{v \in \mathcal{M}(\phi) \setminus \mathcal{M}_{\text{opt}}(\phi)} \left(X_\phi^{(1)} \right)^+ - X_\phi^{(1)} \mathbb{1}_{\{u^* \notin \mathcal{M}(\phi)\}} \\ & \geq 0, \end{aligned}$$

where the last inequality follows because there exists a $v \in \mathcal{M}(\phi) \setminus \mathcal{M}_{\text{opt}}(\phi)$ by virtue of our assumption that $\mathcal{M}(\phi) \neq \mathcal{M}_{\text{opt}}(\phi)$. Thus $\tilde{A}(\mathcal{M}) \geq \tilde{A}(\mathcal{M}_{\text{opt}})$ under L_2 as well. \square

6 Completing the lower bound

Theorem 6.1. *Let M_n^* be the optimal many-to-one matching on $K_{n,n/\alpha}$. Then*

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \left[\sum_{e \in M_n^*} \xi_{K_{n,n/\alpha}}(e) \right] \geq c_\alpha^*.$$

Proof. First observe that

$$\mathbb{E} \left[\sum_{e \in M_n^*} \xi_{K_{n,n/\alpha}}(e) \right] = n \mathbb{E} \left[\sum_{\{\varphi, v\} \in M_n^*} \xi(\varphi, v) \mid \text{label}(\varphi) = o \right],$$

because the root, conditioned to have label o , is uniform among the n vertices of the subset A .

The theorem can now be proved from Corollary 5.5 by following the exact steps of the proof of Theorem 9 in [16]. \square

7 Endogeny

As we now move to establish the upper bound, our aim will be to show that BP run for large enough iterations on $K_{n,n/\alpha}$ gives us a many-to-one matching that has cost close to the optimal for large n . We will show this by relating BP on $K_{n,n/\alpha}$ with BP on \mathcal{T}_α through local weak convergence. It turns out that the X process of Lemma 5.1 generated from the RDE (4.6) arises as the limit of the message process on \mathcal{T}_α (Theorem 9.3). This result depends on an important property of the X process called *endogeny*, which was defined in [5].

Definition 7.1. *An invariant RTP with marginal distribution μ is said to be endogenous if the root variable X_ϕ is almost surely measurable with respect to the σ -algebra $\sigma(\{(\xi_i, N_i) \mid i \in \mathcal{V}\})$.*

Conceptually this means that the random variable at the root X_ϕ is a function of the “data” $\{(\xi_i, N_i) \mid i \in \mathcal{V}\}$, and there is no external randomness. This measurability also implies that if we restrict attention to a large neighborhood of the root, the root variable X_ϕ is almost deterministic given this neighborhood: the effect of the boundary vanishes.

This sort of long range independence allows the use of belief propagation in conjunction with local weak convergence. When we run BP updates on a graph the number of iterations correspond to the size of the neighborhood around a vertex involved in the computation of the messages at that vertex. Endogeny gives us a way to say that this information is sufficient to make decisions at each vertex and jointly yield an almost optimal solution.

For a general RDE (4.7), write $T : \mathcal{P}(S) \rightarrow \mathcal{P}(S)$ for the map induced by the function g on the space of probability measures on S . We now define a bivariate map $T^{(2)} : \mathcal{P}(S \times S) \rightarrow \mathcal{P}(S \times S)$, which maps a distribution $\mu^{(2)} \in \mathcal{P}(S \times S)$ to the joint distribution of

$$\left(g \left(\xi; \left(X_j^{(1)}, 1 \leq j < N \right) \right), g \left(\xi; \left(X_j^{(2)}, 1 \leq j < N \right) \right) \right),$$

where $(X_j^{(1)}, X_j^{(2)})_{j \geq 1}$ are independent with joint distribution $\mu^{(2)}$ on $S \times S$, and the family of random variables $(X_j^{(1)}, X_j^{(2)})_{j \geq 1}$ are independent of the pair (ξ, N) .

It is easy to see that if μ is a fixed point of the RDE then the associated *diagonal measure* $\mu^{\nearrow} := \text{Law}(X, X)$ where $X \sim \mu$ is a fixed point of the operator $T^{(2)}$.

Theorem 11c of [5], reproduced below, provides a way of proving endogeny.

Theorem 7.2 (Theorem 11c of [5]). *For a Polish space S , an invariant recursive tree process with marginal distribution μ is endogenous if and only if*

$$T^{(2)n}(\mu \otimes \mu) \xrightarrow{D} \mu^{\nearrow},$$

where $\mu \otimes \mu$ is the product measure.

We will use the above theorem to prove the endogeny of the RTP for our problem. The proof of Theorem 7.3 follows closely the method of [6], where the author proves the endogeny of the logistic RDE (which corresponds to the minimum weight matching problem). Similar to [6], we will use a transformation to get a fixed point equation for decreasing functions on \mathbf{R} . In our case the operator is a composition of two similar operators, in contrast with a single operator in [6].

Theorem 7.3. *The invariant recursive tree process of Lemma 5.1 arising from the solution to the RDE (4.6) is endogenous.*

Proof. We first write an RDE for X^o . Let $g_1 \left((\xi_i)_{i \geq 1}; (x_i^o)_{i \geq 1} \right) = \min_i (\xi_i - x_i^o)^+$. Also define $g_2 \left((\xi_i)_{i \geq 1}; (x_i^m)_{i \geq 1} \right) = \min_i (\xi_i - x_i^m)$. Then define the map

$$g_o \left(\left((\xi_i)_{i \geq 1}, (\xi_{ij})_{i,j \geq 1} \right); (x_{ij}^o)_{i,j \geq 1} \right) = g_2 \left((\xi_i)_{i \geq 1}; \left(g_1 \left((\xi_{ij})_{j \geq 1}; (x_{ij}^o)_{j \geq 1} \right) \right)_{i \geq 1} \right).$$

We now have the RDE

$$X^o \stackrel{D}{=} g_o \left(\left((\xi_i)_{i \geq 1}, (\xi_{ij})_{i,j \geq 1} \right); (X_{ij}^o)_{i,j \geq 1} \right),$$

where

- (a) $(\xi_i)_{i \geq 1}$ is a Poisson process of rate $1/\alpha$ on \mathbf{R}_+ ,
- (b) for each $i \geq 1$, $(\xi_{ij})_{j \geq 1}$ is Poisson process of rate 1 on \mathbf{R}_+ ,
- (c) $X_{ij}^o, i, j \geq 1$ are independent random variables having the same distribution as X^o , and
- (d) the Poisson processes in (a) and (b) and the random variables $X_{i,j}$ in (c) are independent of each other.

Call the induced mapping from $\mathcal{P}(\mathbf{R}_+)$ to $\mathcal{P}(\mathbf{R}_+)$ as T . The fixed point of this map is the distribution μ , with survival function F as in (4.16). We shall argue that

$$T^{(2)n}(\mu \otimes \mu) \xrightarrow{D} \mu^{\nearrow}. \tag{7.1}$$

Thus the invariant RTP associated with g_o is endogenous. A similar argument shows that one can define an RDE g_m and show that the corresponding invariant RTP associated with g_m is endogenous. Recognizing that the layers of the two RTPs interlace, we easily see that the invariant RTP of Lemma 1 is endogenous. Let us now show (7.1).

Consider the bivariate distributional equations

$$\begin{aligned} \begin{pmatrix} X_{(1)}^m \\ X_{(2)}^m \end{pmatrix} &\stackrel{D}{=} \begin{pmatrix} \min_i (\xi_i - X_{i(1)}^o)^+ \\ \min_i (\xi_i - X_{i(2)}^o)^+ \end{pmatrix}, \\ \begin{pmatrix} X_{(1)}^o \\ X_{(2)}^o \end{pmatrix} &\stackrel{D}{=} \begin{pmatrix} \min_i (\alpha \xi_i - X_{i(1)}^m) \\ \min_i (\alpha \xi_i - X_{i(2)}^m) \end{pmatrix}, \end{aligned}$$

where $\{\xi_i, i \geq 1\}$ is a Poisson process of rate 1 on \mathbf{R}_+ , $\{(X_{i(1)}^o, X_{i(2)}^o), i \geq 1\}$ are i.i.d. random variables with joint survival function $F_0^{(2)}$ (that is, for every $x, y \in \mathbf{R}$, $F_0^{(2)}(x, y) = P\{X_{i(1)}^o > x, X_{i(2)}^o > y\}$) independent of the Poisson process, $\{(X_{i(1)}^m, X_{i(2)}^m), i \geq 1\}$ are i.i.d. random variables with joint survival function $G_0^{(2)}$ independent of the Poisson process. Let $\mathcal{P}^{(2)}$ denote the space of joint survival functions of $\mathbf{R} \times \mathbf{R}$ valued random variables. Write $\Gamma^{(2)} : \mathcal{P}^{(2)} \rightarrow \mathcal{P}^{(2)}$ and $\varphi^{(2)} : \mathcal{P}^{(2)} \rightarrow \mathcal{P}^{(2)}$ for the maps defined by

$$\begin{aligned} \Gamma^{(2)}(F_0^{(2)})(x, y) &= P\{X_{(1)}^m > x, X_{(2)}^m > y\}, \\ \varphi^{(2)}(G_0^{(2)})(x, y) &= P\{X_{(1)}^o > x, X_{(2)}^o > y\}. \end{aligned}$$

Let $F_0^{(2)}(x, y) = F(x)F(y)$ for all $x, y \in \mathbf{R}$, and define two sequences of joint distributions $\{G_k^{(2)}, k \geq 0\}$ and $\{F_k^{(2)}, k \geq 0\}$:

$$\begin{aligned} G_k^{(2)} &= \Gamma^{(2)}(F_k^{(2)}), \\ F_{k+1}^{(2)} &= \varphi^{(2)}(G_k^{(2)}). \end{aligned}$$

It is clear that the marginal distributions corresponding to $F_k^{(2)}$ and $G_k^{(2)}$ are F and G respectively for each k . By Theorem 7.2, the invariant RTP associated with g_o is endogenous if $\{F_k^{(2)}, k \geq 1\}$ converges to the joint distribution of the degenerate random variable that has both components equal. This is equivalent to the condition $F_k^{(2)}(x, x) \rightarrow F(x)$ as $k \rightarrow \infty$ for all $x \in \mathbf{R}$.

From the bivariate RDE, if the joint distribution of $(X_{i(1)}^o, X_{i(2)}^o)$ is $F_k^{(2)}$ then for

$x, y \geq 0$,

$$\begin{aligned} G_k^{(2)}(x, y) &= \mathbb{P} \left\{ X_{(1)}^m > x, X_{(2)}^m > y \right\} \\ &= \mathbb{P} \left\{ \text{No point of } \left(\xi_i; X_{i(1)}^o, X_{i(2)}^o \right) \text{ in } \{(z; u, v) : z - u \leq x \text{ or } z - v \leq y\} \right\} \\ &= \exp \left(- \int_{z=0}^{\infty} \mathbb{P} \{ z - X \leq x \text{ or } z - Y \leq y \} dz \right) \\ &\quad \left((X, Y) \text{ has distribution } F_k^{(2)} \right) \\ &= \exp \left(- \int_{z=0}^{\infty} (F(z - x) + F(z - y) - F_k^{(2)}(z - x, z - y)) dz \right) \\ &= G(x) \exp \left(- \int_{z=0}^{\infty} (F(z - y) - F_k^{(2)}(z - x, z - y)) dz \right). \end{aligned}$$

Write $g_k(x) = G_k^{(2)}(x, x)$ and $f_k(x) = F_k^{(2)}(x, x), x \in \mathbf{R}$. Then

$$g_k(x) = \begin{cases} G(x) \exp(-\int_{-x}^{\infty} (F(z) - f_k(z)) dz), & x \geq 0, \\ 1, & x < 0. \end{cases}$$

Similar analysis with the equation for $(X_{(1)}^o, X_{(2)}^o)$ gives

$$f_{k+1}(x) = F(x) \exp \left(- \int_{-x}^{\infty} \frac{1}{\alpha} (G(z) - g_k(z)) dz \right), x \in \mathbf{R}. \tag{7.2}$$

Define for $k \geq 0$,

$$\beta_k(x) = \int_{-x}^{\infty} (F(z) - f_k(z)) dz, \tag{7.3}$$

$$\gamma_k(x) = \int_{-x}^{\infty} \frac{1}{\alpha} (G(z) - g_k(z)). \tag{7.4}$$

Then

$$g_k(x) = \begin{cases} G(x) \exp(-\beta_k(x)), & x \geq 0 \\ 1, & x < 0, \end{cases} \tag{7.5}$$

$$f_{k+1}(x) = F(x) \exp(-\gamma_k(x)), k \geq 0. \tag{7.6}$$

We now prove certain relations for the functions β_k and γ_k .

Lemma 7.4. For all $k \geq 0$:

$$\begin{aligned} \beta_k(x) &\geq 0, \\ \gamma_k(x) &\geq 0, \\ \beta_{k+1}(x) &\leq \beta_k(x), \text{ and} \\ \gamma_{k+1}(x) &\leq \gamma_k(x) \text{ for all } x. \end{aligned} \tag{7.7}$$

Proof. We will prove the Lemma by induction.

$f_0(z) = F^2(z) \leq F(z)$ for all z implies that $\beta_0(x) \geq 0$ for all x . This implies $g_0(z) \leq G(z)$ for all z , and hence $\gamma_0(x) \geq 0$ for all x . By (7.2), we have

$$f_1(x) \geq F(x) \exp(-\int_{-x}^{\infty} \frac{1}{\alpha} G(z) dz) = f_0(x) \text{ for all } x,$$

$$\beta_1(x) \leq \beta_0(x) \text{ for all } x \text{ by (7.3),}$$

$$g_1(x) \geq g_0(x) \text{ for all } x \text{ by (7.5),}$$

$$\gamma_1(x) \leq \gamma_0(x) \text{ for all } x \text{ by (7.4).}$$

We have showed that the equations (7.7) hold for $k = 0$, now assume that they hold for some $k \geq 0$.

$$\begin{aligned} \gamma_k(x) \geq 0 &\Rightarrow f_{k+1}(x) \leq F(x) \text{ for all } x && \text{by (7.6)} \\ &\Rightarrow \beta_{k+1}(x) \geq 0 \text{ for all } x && \text{by (7.3)} \\ &\Rightarrow g_{k+1}(x) \leq G(x) \text{ for all } x && \text{by (7.5)} \\ &\Rightarrow \gamma_{k+1}(x) \geq 0 \text{ for all } x && \text{by (7.4)}. \end{aligned}$$

$$\begin{aligned} \gamma_{k+1}(x) \leq \gamma_k(x) &\Rightarrow f_{k+2}(x) \geq f_{k+1}(x) \text{ for all } x && \text{by (7.6)} \\ &\Rightarrow \beta_{k+2}(x) \leq \beta_{k+1}(x) \text{ for all } x && \text{by (7.3)} \\ &\Rightarrow g_{k+2}(x) \geq g_{k+1}(x) \text{ for all } x && \text{by (7.5)} \\ &\Rightarrow \gamma_{k+2}(x) \leq \gamma_{k+1}(x) \text{ for all } x && \text{by (7.4)}. \end{aligned}$$

By induction, the properties (7.7) hold for all $k \geq 0$. □

Because for each x , the sequences $\{\beta_k(x), k \geq 0\}$ and $\{\gamma_k(x), k \geq 0\}$ are decreasing, the sequences of functions $\{\beta_k, k \geq 0\}$ and $\{\gamma_k, k \geq 0\}$ converge pointwise to functions, say β and γ , respectively, as $k \rightarrow \infty$. The requirement for endogeneity that $f_k(x) \rightarrow F(x)$ for all x as $k \rightarrow \infty$ reduces to showing that $\gamma(x) = 0$ for all x . We will show that both β and γ are identically 0.

Using (7.3) and (7.6), we have

$$\beta_{k+1}(x) = \int_{-x}^{\infty} F(z) \left(1 - e^{-\gamma_k(z)}\right) dz.$$

Since the term inside the integrand is bounded above by $F(z)$, and $\int_{-x}^{\infty} F(z) dz$ exists for all x , by the dominated convergence theorem, we conclude that

$$\beta(x) = \int_{-x}^{\infty} F(z) \left(1 - e^{-\gamma(z)}\right) dz. \tag{7.8}$$

Similarly from (7.4) and (7.5), we have

$$\gamma(x) = \begin{cases} \int_{-x}^{\infty} \frac{G(z)}{\alpha} \left(1 - e^{-\beta(z)}\right) dz, & x < 0, \\ \int_0^{\infty} \frac{G(z)}{\alpha} \left(1 - e^{-\beta(z)}\right) dz, & x \geq 0. \end{cases} \tag{7.9}$$

Differentiating (7.8) and (7.9), we get

$$\begin{aligned} \beta'(x) &= F(-x) \left(1 - e^{-\gamma(-x)}\right), \quad x \in \mathbf{R}, \\ \gamma'(x) &= \frac{G(-x)}{\alpha} \left(1 - e^{-\beta(-x)}\right), \quad x < 0. \end{aligned}$$

Recall that for $x \geq 0$ (see (4.12), (4.13))

$$\begin{aligned} F'(-x) &= -\frac{1}{\alpha} F(-x)G(x), \\ G'(x) &= -G(x)F(-x). \end{aligned}$$

Define a function $\eta : \mathbf{R}_+ \rightarrow \mathbf{R}_+$ as

$$\eta(x) = \alpha F(-x)e^{\gamma(-x)} + G(x)e^{-\beta(x)}, \quad x \geq 0. \tag{7.10}$$

This function is continuous and differentiable on $(0, \infty)$. Differentiating, we get (for $x > 0$)

$$\begin{aligned} \eta'(x) &= -\alpha F(-x)\gamma'(-x)e^{\gamma(-x)} - \alpha F'(-x)e^{\gamma(-x)} \\ &\quad - G(x)\beta'(x)e^{-\beta(x)} + G'(x)e^{-\beta(x)} \\ &= -F(-x)G(x)\left(1 - e^{-\beta(x)}\right)e^{\gamma(-x)} + F(-x)G(x)e^{\gamma(-x)} \\ &\quad - G(x)F(-x)\left(1 - e^{-\gamma(-x)}\right)e^{-\beta(x)} - G(x)F(-x)e^{-\beta(x)} \\ &= F(-x)G(x)e^{-\beta(x)}\left((e^{\gamma(-x)} + e^{-\gamma(-x)} - 2)\right). \end{aligned}$$

For any $y \in \mathbf{R}$, $e^y + e^{-y} \geq 2$, hence

$$\eta'(x) \geq 0$$

for all $x > 0$.

Now observe that $\lim_{x \rightarrow \infty} \eta(x) = \alpha$. By (7.9), $\gamma(x) = \gamma(0)$ for $x \geq 0$. Using this in (7.8) gives

$$\beta(0) = \int_0^\infty F(z)\left(1 - e^{-\gamma(0)}\right) dz = \left(1 - e^{-\gamma(0)}\right) w_o.$$

Substituting values of $F(0)$ and $G(0)$ in (7.10), and using the above relation between $\beta(0)$ and $\gamma(0)$, we get

$$\begin{aligned} \eta(0) &= w_o e^{\gamma(0)} + e^{-w_o} e^{-\beta(0)} \\ &= w_o e^{\gamma(0)} + e^{-w_o} e^{-(1 - e^{-\gamma(0)})w_o} \\ &= w_o e^{\gamma(0)} + e^{-2w_o} e^{w_o e^{-\gamma(0)}}. \end{aligned}$$

Using $\gamma(0) \geq 0$, it can be seen that $\eta(0) \geq w_o + e^{-w_o} = \alpha$, with equality if and only if $\gamma(0) = 0$. We see that η is a continuous increasing function on $[0, \infty)$ with $\eta(0) \geq \alpha$ and $\lim_{x \rightarrow \infty} \eta(x) = \alpha$. η must then be a constant function with $\eta(0) = \alpha$, and so $\gamma(0) = 0$. For any $x \in \mathbf{R}$, $0 \leq \gamma(x) \leq \gamma(0)$ implies that γ must be identically 0. This proves that $f_k(x)$ converges to $F(x)$ for all $x \in \mathbf{R}$, and so the invariant RTP is endogenous. \square

8 Domain of attraction of the RDE map

Consider the maps

$$\varphi G_0(t) = \exp\left(-\int_{-t}^\infty \frac{G_0(z)}{\alpha} dz\right), \quad t \in \mathbf{R}, \tag{8.1}$$

$$\Gamma F_0(t) = \begin{cases} \exp\left(-\int_{-t}^\infty F_0(z) dz\right), & t \geq 0, \\ 1, & t < 0. \end{cases} \tag{8.2}$$

Write $T = \Gamma\varphi$. We have $F = \varphi G$, $G = \Gamma F$, and $G = TG$ is the unique fixed point of T .

We are interested in the domain of attraction of G . Specifically, we investigate the convergence of $T^k G_0$ for an arbitrary survival function G_0 . We will see that the domain of attraction is all those G_0 satisfying $\int_0^\infty G_0(z) dz < \infty$. We proceed via a sequence of Lemmas.

Lemma 8.1. *Suppose G_0 is such that $\int_0^\infty G_0(t) dt < \infty$. Then,*

- (a) $TG_0(t) = 1$ for $t < 0$,
- (b) $TG_0(0) < 1$,
- (c) $TG_0(t)$ is strictly decreasing for $t \geq 0$,

- (d) $TG_0(\cdot)$ is continuous on $(0, \infty)$ (right-continuous at 0), and
- (e) $\int_0^\infty TG_0(t) dt < \infty$.

Proof. (a), (b), (c), and (d) follow directly from the definitions (8.1) and (8.2).

To get (e), observe that $F_0 = \varphi G_0$ satisfies, for $z \leq 0$, $F_0(z) \geq F_0(0) > 0$ by the assumption that $\int_0^\infty G_0(t) dt < \infty$. Thus for $t \geq 0$,

$$\int_{-t}^\infty F_0(z) dz \geq tF_0(0).$$

This implies, for $t \geq 0$,

$$TG_0(t) = \exp\left(-\int_{-t}^\infty F_0(z) dz\right) \leq e^{-tF_0(0)},$$

which is integrable. □

Since, after one step of the iteration, TG_0 has properties (a), (b), (c), (d), and (e), we might assume, without loss of generality, that G_0 itself has these properties.

Define the transformation

$$\widehat{G}_0(x) = x - \left[\log\left(\frac{\alpha - G_0(x)}{\gamma G_0(x)}\right) \right]^+, \quad x \geq 0. \tag{8.3}$$

$\widehat{G}_0(x)$ is the largest shift at x such that $G(x - \text{shift}) \leq G_0(x)$. It has the following properties.

Lemma 8.2.

- (a) $\widehat{G}_0(x) \leq x$ with equality if and only if $G_0(x) \geq G(0)$.
- (b) $G_0(x) \geq G(x - \widehat{G}_0(x))$ with equality if and only if $G_0(x) \leq G(0)$.

Proof. (a) Clearly $\widehat{G}_0(x) \leq x$, because the second term that is subtracted in (8.3) is ≥ 0 . Equality holds if and only if the logarithm in (8.3) is ≤ 0

$$\begin{aligned} \iff \frac{\gamma G_0(x)}{\alpha - G_0(x)} &\geq 1 \\ \iff G_0(x) &\geq G(0) = \frac{\alpha}{1 + \gamma}. \end{aligned}$$

(b) When $G_0(x) \geq G(0)$, we saw in (a) that $\widehat{G}_0(x) = x$. Hence in this case $G_0(x) \geq G(0) = G(x - \widehat{G}_0(x))$. We deal with the other case and the second part of (a) simultaneously by showing

$$G_0(x) \leq G(0) \iff G_0(x) = G(x - \widehat{G}_0(x)).$$

This follows from the implications:

$$\begin{aligned} G_0(x) \leq G(0) &\iff \widehat{G}_0(x) = x - \log\left(\frac{\alpha - G_0(x)}{\gamma G_0(x)}\right) \\ &\iff x - \widehat{G}_0(x) = \log\left(\frac{\alpha - G_0(x)}{\gamma G_0(x)}\right) \\ &\iff G(x - \widehat{G}_0(x)) = \frac{\alpha}{1 + \gamma\left(\frac{\alpha - G_0(x)}{\gamma G_0(x)}\right)} = G_0(x). \end{aligned} \quad \square$$

Lemma 8.3. Let G_0 satisfy the properties in Lemma 8.1. Let $x \geq 0$.

- (a) For $b \leq 0$, $\widehat{G}_0(x) \geq b$ if and only if $G_0(x) \geq G(x - b)$.

(b) For $B \geq 0$, $\widehat{G}_0(x) \leq B$ if and only if $G_0(x) \leq G(x - B)$.

The above statements also hold when all the inequalities are replaced by strict inequalities.

Proof. Suppose that $\widehat{G}_0(x) \geq b$. Since G is decreasing, we have $G(x - b) \leq G(x - \widehat{G}_0(x)) \leq G_0(x)$ (the last inequality by Lemma 8.2(b)).

Conversely, when $G_0(x) \geq G(x - b)$, either $G_0(x) = G(x - \widehat{G}_0(x))$, which implies $\widehat{G}_0(x) \geq b$, or by Lemma 8.2(b) $G_0(x) > G(0)$, which by Lemma 8.2(a) gives $\widehat{G}_0(x) = x \geq b$, because $x \geq 0$ and $b \leq 0$.

Now let $\widehat{G}_0(x) \leq B, B \geq 0$. Then for $x < B, G_0(x) \leq 1 = G(x - B)$. For $x > B \geq \widehat{G}_0(x)$, by Lemma 8.2(b), $G_0(x) < G(0)$, and so by Lemma 8.2(a), $G_0(x) = G(x - \widehat{G}_0(x)) \leq G(x - B)$. The inequality extends to $x = B$ by right-continuity of G_0 and G at B and 0 respectively.

Conversely, when $G_0(x) \leq G(x - B)$, Lemma 8.2(b) gives $G(x - \widehat{G}_0(x)) \leq G_0(x) \leq G(x - B)$, and so $\widehat{G}_0(x) \leq B$.

Also by strict monotonicity of $G_0(x)$ for $x \geq 0$, all the inequalities can be made strict. □

Lemma 8.4. Suppose G_0 satisfies the properties in Lemma 8.1, and suppose further that there exist real numbers $b \leq 0$ and $B \geq 0$ such that $b \leq \widehat{G}_0(x) \leq B$ for all $x \geq 0$. Then $b \leq \widehat{TG}_0(x) \leq B$ for all $x \geq 0$.

Proof. By Lemma 8.3, we have $G(z - b) \leq G_0(z) \leq G(z - B)$ for $z \geq 0$. But this implies $G(z - b) \leq G_0(z) \leq G(z - B)$ for all $z \in \mathbf{R}$ as well, because for $z < 0, G_0(z) = G(z - B) = 1$. Using these in (8.1), for all $t \in \mathbf{R}$, we have

$$\varphi G_0(t) \leq \exp\left(-\int_{-t}^{\infty} \frac{G(z - b)}{\alpha} dz\right) = F(t + b), \tag{8.4}$$

$$\varphi G_0(t) \geq \exp\left(-\int_{-t}^{\infty} \frac{G(z - B)}{\alpha} dz\right) = F(t + B). \tag{8.5}$$

Then by (8.2), for $x \geq 0$, we have

$$\begin{aligned} TG_0(x) = \Gamma\varphi G_0(x) &\geq \exp\left(-\int_{-x}^{\infty} F(t + b) dt\right) = G(x - b), \\ TG_0(x) = \Gamma\varphi G_0(x) &\leq \exp\left(-\int_{-x}^{\infty} F(t + B) dt\right) \leq G(x - B), \end{aligned}$$

where the last inequality follows from (4.9) to handle the case $x < B$.

Lemma 8.3 now completes the proof. □

The above Lemma shows that if \widehat{G}_0 is bounded then $\inf \widehat{T^k G_0}$ is increasing in k and $\left(\sup \widehat{T^k G_0}\right)^+$ is decreasing in k , and so they converge to some b^* and B^* respectively ($b^* \leq 0 \leq B^*$):

$$\inf \widehat{T^k G_0} \uparrow b^* \text{ and } \left(\sup \widehat{T^k G_0}\right)^+ \downarrow B^* \text{ as } k \rightarrow \infty. \tag{8.6}$$

We will show that $b^* = B^* = 0$. (This will prove that both $\inf \widehat{T^k G_0}$ and $\sup \widehat{T^k G_0}$ converge to 0.)

First we show that the terms $\inf \widehat{T^k G_0}$ and $\left(\sup \widehat{T^k G_0}\right)^+$ are strictly monotone in k unless they are 0.

Lemma 8.5. Suppose G_0 satisfies the properties in Lemma 8.1. If $-\infty < \inf \widehat{G}_0 < 0$ then $\inf \widehat{TG}_0 > \inf \widehat{G}_0$. If $0 < \sup \widehat{G}_0 < \infty$ then $\sup \widehat{TG}_0 < \sup \widehat{G}_0$.

Proof. Let $B = \sup \widehat{G}_0 > 0$. Fix $a \in (0, B)$. By Lemma 8.3, $G_0(x) \leq G(x - B)$. Using this in (8.1), we have, for $x \geq 0$

$$\begin{aligned} \varphi G_0(x) &\geq \exp\left(-\frac{1}{\alpha}\left(\int_{-x}^0 G(z - B) dz + \int_0^a G_0(z) dz + \int_a^\infty G(z - B) dz\right)\right) \\ &= \exp\left(-\int_{-x}^\infty \frac{G(z - B)}{\alpha} dz\right) \exp\left(\int_0^a \frac{G(z - B) - G_0(z)}{\alpha} dz\right) \\ &= F(x + B) \exp\left(\int_0^a \frac{1 - G_0(z)}{\alpha} dz\right) \\ &= \kappa F(x + B), \end{aligned} \tag{8.7}$$

where $\kappa > 1$ (because $G_0(z) < 1$ for $z \geq 0$).

Let $x \geq B$. Using (8.5) to bound $\varphi G_0(z)$ for $z \in [-x, 0)$ and (8.7) to bound $\varphi G_0(z)$ for $z \geq 0$ in (8.2), we get

$$\begin{aligned} TG_0(x) &= \Gamma \varphi G_0(x) \\ &\leq \exp\left(-\int_{-x}^0 F(z + B) dz - \int_0^\infty \kappa F(z + B) dz\right) \\ &= \exp\left(-\int_{-x}^\infty F(z + B) dz\right) \exp\left(-(\kappa - 1) \int_0^\infty F(z + B) dz\right) \\ &= \kappa' G(x - B), \end{aligned} \tag{8.8}$$

where $\kappa' < 1$. Thus $TG_0(x) < G(x - B)$, and so by Lemma 8.3(b) for strict inequalities, $\widehat{TG}_0(x) < B$.

For $x \in [0, B)$, by Lemma 8.2(a), we have $\widehat{TG}_0(x) \leq x < B$. By continuity of \widehat{TG}_0 , the supremum of $\widehat{TG}_0(x)$ over any compact subset of $[0, \infty)$ is strictly less than B . We now need to verify the strict inequality as $x \rightarrow \infty$. To show this, we have the following:

$$\begin{aligned} \lim_{x \rightarrow \infty} \widehat{TG}_0(x) &= \lim_{x \rightarrow \infty} \left(x - \left(\log\left(\frac{\alpha - TG_0(x)}{\gamma TG_0(x)}\right)\right)^+\right) \quad (\text{by definition (8.3)}) \\ &= \lim_{x \rightarrow \infty} \left(x - \log\left(\frac{\alpha - TG_0(x)}{\gamma TG_0(x)}\right)\right) \quad (\text{by } \lim_{x \rightarrow \infty} TG_0(x) = 0 < G(0)) \\ &\leq \lim_{x \rightarrow \infty} \left(x - \log\left(\frac{\alpha - \kappa' G(x - B)}{\gamma \kappa' G(x - B)}\right)\right) \quad (\text{by (8.8)}) \\ &= \lim_{x \rightarrow \infty} \left(x - \log\left(\frac{\alpha - \kappa' \frac{\alpha}{1 + \gamma e^{x-B}}}{\gamma \kappa' \frac{\alpha}{1 + \gamma e^{x-B}}}\right)\right) \\ &= \lim_{x \rightarrow \infty} \left(x - \log\left(\frac{1 + \gamma e^{x-B}}{\gamma \kappa'} - \frac{1}{\gamma}\right)\right) \\ &= B + \log \kappa' < B. \end{aligned}$$

We use the same technique for the infimum. Suppose $b = \inf \widehat{G}_0 < 0$. Fix $c \in (b, 0)$. Using $G_0(x) \geq G(x - b)$ for $x \geq 0$ (Lemma 8.3(a)) in (8.1), we have for $x \geq 0$

$$\begin{aligned} \varphi G_0(x) &\leq \exp\left(-\frac{1}{\alpha}\left(\int_{-x}^c G(z - b) dz + \int_c^0 G_0(z) dz + \int_0^\infty G(z - b) dz\right)\right) \\ &= \exp\left(-\int_{-x}^\infty \frac{G(z - b)}{\alpha} dz\right) \exp\left(\int_c^0 \frac{G(z - b) - G_0(z)}{\alpha} dz\right) \\ &= F(x + b) \exp\left(-\int_c^0 \frac{1 - G(z - b)}{\alpha} dz\right) \\ &= \delta F(x + B), \end{aligned} \tag{8.9}$$

where $\delta < 1$. Using (8.4) to bound $\varphi G_0(z)$ for $z < 0$ and (8.9) to bound $\varphi G_0(z)$ for $z \geq 0$ in (8.2), we get for $x \geq 0$

$$\begin{aligned} TG_0(x) &= \Gamma \varphi G_0(x) \geq \exp\left(-\int_{-x}^0 F(z+b) dz - \int_0^\infty \delta F(z+b) dz\right) \\ &= \exp\left(-\int_{-x}^\infty F(z+b) dz\right) \exp\left((1-\delta) \int_0^\infty F(z+b) dz\right) \\ &= \delta' G(x-b), \end{aligned}$$

where $\delta' > 1$. This implies the strict inequality $\widehat{TG_0}(x) > b$ for all $x \geq 0$ (Lemma 8.3(a)). We now verify the inequality for $x \rightarrow \infty$; as in the case of supremum,

$$\begin{aligned} \lim_{x \rightarrow \infty} \widehat{TG_0}(x) &= \lim_{x \rightarrow \infty} x - \left(\log\left(\frac{\alpha - TG_0(x)}{\gamma TG_0(x)}\right)\right)^+ \\ &\geq \lim_{x \rightarrow \infty} x - \log\left(\frac{\alpha - \delta' G(x-b)}{\gamma \delta' G(x-b)}\right) \\ &= b + \log \delta' > b. \end{aligned} \quad \square$$

Lemma 8.6. *Suppose G_0 satisfies the properties in Lemma 8.1. Then $\widehat{T^2G_0}$ is bounded.*

Proof. Using $\int_0^\infty G_0(z) dz < \infty$ and $G_0(z) = 1$ for $z < 0$ in (8.1), we have

$$\varphi G_0(x) = O(e^{-x/\alpha}) \text{ as } x \rightarrow +\infty.$$

Then $\int_0^\infty \varphi G_0(z) dz < \infty$, and by (8.2)

$$TG_0(x) = \Theta\left(e^{-\int_{-x}^0 \varphi G_0(z) dz}\right) \text{ as } x \rightarrow +\infty. \tag{8.10}$$

Since $\varphi G_0(z)$ is decreasing (and not identically 0), we have

$$TG_0(x) = O(e^{-ux}) \text{ as } x \rightarrow +\infty,$$

$u = \varphi G_0(0) > 0$. Using this in (8.1) gives

$$\varphi TG_0(x) = 1 - O(e^{ux}) \text{ as } x \rightarrow -\infty.$$

Now replacing φG_0 by φTG_0 in (8.10), we get

$$T^2G_0(x) = O(e^{-x}) \text{ as } x \rightarrow +\infty.$$

Furthermore, using $\varphi G_0(z) \leq 1$ for $z < 0$ in (8.10), we also get

$$T^2G_0(x) = \Omega(e^{-x}) \text{ as } x \rightarrow +\infty,$$

and so

$$T^2G_0(x) = \Theta(e^{-x}) \text{ as } x \rightarrow +\infty. \tag{8.11}$$

Let us now choose K sufficiently large such that

$$G(x+K) = \frac{\alpha}{1 + \gamma e^{x+K}} \leq T^2G_0(x) \leq \frac{\alpha}{1 + \gamma e^{x-K}} = G(x-K)$$

for all sufficiently large x . This, by Lemma 8.3, proves that $\widehat{T^2G_0}(x)$ is bounded. □

For future reference, let us also note that, by (8.1),

$$\varphi T^2G_0(x) = 1 - \Omega(e^x) \text{ as } x \rightarrow -\infty. \tag{8.12}$$

Lemma 8.7. *Suppose \widehat{G}_0 is bounded and G_0 satisfies the properties in Lemma 8.1. Then $\widehat{T^k G_0}$ is differentiable on $(0, \infty)$ except possibly at one point, and the derivatives are uniformly integrable for $k \geq 3$:*

$$\sup_{k \geq 3} \int_{x > B} |\widehat{T^k G_0}'(x)| dx \xrightarrow{B \rightarrow \infty} 0.$$

Proof. Let $x_0^k = \inf \{x \geq 0 : T^k G_0(x) \leq G(0)\}$. Then $\widehat{T^k G_0}(x) = x$ for $x \in [0, x_0^k]$, and $\widehat{T^k G_0}'(x) = 1$, $x \in (0, x_0^k)$. The boundedness of $\widehat{G}_0(x)$ by B , say, implies by Lemma 8.4 that $\widehat{T^k G_0}(x)$ are uniformly upper bounded by B for all x , and by Lemma 8.2(a) $x_0^k \leq B$. For $x > x_0^k$, by definition of x_0^k , we have $T^k G_0(x) \leq G(0)$, in which case

$$\widehat{T^k G_0}(x) = x - \log \left(\frac{\alpha - G_0(x)}{\gamma G_0(x)} \right).$$

Consequently,

$$\widehat{T^k G_0}'(x) = 1 + \frac{\alpha (T^k G_0)'(x)}{T^k G_0(x) (\alpha - T^k G_0(x))}.$$

By (8.2),

$$\widehat{T^k G_0}'(x) = 1 - \frac{\alpha (\varphi T^{k-1} G_0)(-x)}{\alpha - T^k G_0(x)}.$$

As (8.11) and (8.12) hold respectively for $T^k G_0$ and $\varphi T^k G_0$ for all $k \geq 2$, the above equation gives

$$\widehat{T^k G_0}'(x) = O(e^{-x}).$$

This gives the uniform integrability of the derivatives. □

Theorem 8.8. *For any survival function G_0 satisfying $\int_0^\infty G_0(z) dz < \infty$,*

$$T^k G_0 \rightarrow G \text{ pointwise as } k \rightarrow \infty.$$

The expectations also converge:

$$\lim_{k \rightarrow \infty} \int_0^\infty T^k G_0(x) dx = \int_0^\infty G(x) dx.$$

Proof. By Lemmas 8.1, 8.4, 8.6, and 8.7, for $k \geq 5$, $\widehat{T^k G_0}$ are uniformly bounded and have derivatives that are uniformly integrable.

We prove that the limits in (8.6) satisfy $b^* = B^* = 0$. The functions $T^k G_0, k \geq 1$ are bounded and 1-Lipshitz on $(0, \infty)$. By Arzela-Ascoli theorem this sequence is relatively compact with respect to compact convergence. There exists a subsequence that converges:

$$T^{n_k} G_0 \xrightarrow{k \rightarrow \infty} G_\infty. \tag{8.13}$$

The uniform continuity of the function $y \mapsto \log(\frac{\alpha - y}{\gamma y})$ on every compact subset of $(0, 1)$ implies that the transforms (which are uniformly bounded) converge (compact convergence):

$$\widehat{T^{n_k} G_0} \xrightarrow{k \rightarrow \infty} \widehat{G}_\infty.$$

The uniform integrability of the derivatives in Lemma 8.7 says that for any $\epsilon > 0$, there exists B sufficiently large such that $\widehat{T^{n_k} G_0}(x) \in (\widehat{T^{n_k} G_0}(B) - \epsilon, \widehat{T^{n_k} G_0}(B) + \epsilon)$ for all

$x > B$ and for all $k \geq 3$. This extends the uniform convergence of $\widehat{T^{n_k}G_0} \rightarrow \widehat{G_\infty}$ in the compact set $[0, B]$ to $[0, \infty)$. The uniform convergence implies

$$\begin{aligned} \inf \widehat{G_\infty} &= \lim_{k \rightarrow \infty} \inf \widehat{T^{n_k}G_0} = b^*, \\ \sup \left(\widehat{G_\infty}\right)^+ &= \lim_{k \rightarrow \infty} \left(\sup \widehat{T^{n_k}G_0}\right)^+ = B^*. \end{aligned}$$

We can apply the pointwise convergence (8.13) in the maps (8.1) and (8.2), thanks to the dominated convergence theorem, to conclude

$$T^{n_k+1}G_0 \xrightarrow{k \rightarrow \infty} TG_\infty.$$

The same arguments now conclude that

$$\begin{aligned} \inf \widehat{TG_\infty} &= b^*, \\ \sup \left(\widehat{TG_\infty}\right)^+ &= B^*. \end{aligned}$$

Lemma 8.5 allows this only if $b^* = B^* = 0$, and so $T^kG_0(x) \xrightarrow{k \rightarrow \infty} G(x)$ for all x .

Furthermore, there exists $B > 0$ such that for all $k \geq 2$, $\sup \widehat{T^kG_0} \leq B$. This implies that $T^kG_0(x) \leq G(x - B)$ for all $x \geq 0$ (Lemma 8.3), which by dominated convergence theorem gives convergence of the expectations:

$$\lim_{k \rightarrow \infty} \int_0^\infty T^kG_0(x) dx = \int_0^\infty G(x) dx. \quad \square$$

9 Belief propagation

We now show that the belief propagation (BP) algorithm on $K_{n,n/\alpha}$ converges to a many-to-one matching solution that is asymptotically optimal as $n \rightarrow \infty$.

9.1 Convergence of BP on the infinite tree

In this section we will prove that the messages on \mathcal{T}_α converge, and relate the resulting matching with the matching \mathcal{M}_{opt} of Section 5.

The message process can essentially be written as

$$X_{\mathcal{T}_\alpha}^{k+1}(\dot{v}, v) = \begin{cases} \min_{i \geq 1} \{ \xi_{\mathcal{T}_\alpha}(v, vi) - X_{\mathcal{T}_\alpha}^k(v, vi) \}, & \text{if label}(v) = o \\ \min_{i \geq 1} \{ (\xi_{\mathcal{T}_\alpha}(v, vi) - X_{\mathcal{T}_\alpha}^k(v, vi))^+ \}, & \text{if label}(v) = m, \end{cases} \quad (9.1)$$

where the initial messages $X_{\mathcal{T}_\alpha}^0(\dot{v}, v)$ are i.i.d. random variables (zero in the case of our algorithm; see (2.2)).

By the structure of \mathcal{T}_α it is clear that – the initial distribution being fixed – the distribution of the messages $X_{\mathcal{T}_\alpha}^k(\dot{v}, v)$, $v \in \mathcal{V}$ depends only on k and the label of the corresponding vertex v . Also, it can be seen from the analysis of RDE (4.6) in Section 4 that if we denote the survival function of the distribution of messages at some step k' by F_0 for vertices with label o and G_0 for vertices with label m , then after one update the respective survival functions will be given by the maps (8.1) and (8.2): φG_0 and ΓF_0 . Considering the messages $X_{\mathcal{T}_\alpha}^k(\dot{v}, v)$ only for vertices v having label m , the common distribution G_0 at step k' evolves to TG_0 at step $k' + 2$ ($T = \Gamma\varphi$); the distribution ΓF_0 at step $k' + 1$ changes to $T\Gamma F_0$ at $k' + 3$. The distribution sequence generated by the BP iterations (for $k \geq k'$) is obtained by interleaving the two sequences generated by applying the map T iteratively to F_0 and ΓF_0 . By Section 8, the distributions along both

the odd and even subsequences converge to G , and consequently from the nature of the two step RDE, we conclude that the common distribution of $X_{\mathcal{T}_\alpha}^k(\dot{v}, v)$, v having label m converges to G , and the common distribution of $X_{\mathcal{T}_\alpha}^k(\dot{v}, v)$, v having label o converges to F as $k \rightarrow \infty$.

By the endogeny (Section 7) of the RDE – considering the evolution of messages at vertices of the same label – the following result follows in the same way as Lemma 6 in [16].

Lemma 9.1. *If the initial messages $X_{\mathcal{T}_\alpha}^0(\dot{v}, v)$, $\text{label}(v) = m$, are i.i.d. random variables with survival function G then the message process (9.1) converges in L^2 to the process X as $k \rightarrow \infty$.*

We now prove that if the initial values are i.i.d. random variables with some arbitrary distribution (not necessarily one with survival function G), then the message process (9.1) does indeed converge to the unique stationary configuration. Of course, the initial condition of particular interest to us is the all zero initial condition (2.2), but we will prove a more general result.

The following lemma will allow us to interchange limit and minimization while working with the updates on \mathcal{T}_α .

Lemma 9.2. *Let $\{X_{\mathcal{T}_\alpha}^0(\dot{v}, v) : \text{label}(v) = o\}$ be initialized to i.i.d. random variables with a distribution on \mathbf{R} having survival function F_0 , and let $\{X_{\mathcal{T}_\alpha}^0(\dot{v}, v) : \text{label}(v) = m\}$ be initialized to i.i.d. random variables with a distribution on \mathbf{R}_+ having survival function G_0 . Then for v having label o , the map*

$$\pi_{\mathcal{T}_\alpha}^k(v) = \arg \min_{u \sim v} \{ \xi_{\mathcal{T}_\alpha}(v, u) - X_{\mathcal{T}_\alpha}^k(v, u) \}$$

is a.s. well defined and singleton for all $k \geq 1$, and

$$\sup_{k \geq 1} \mathbb{P} \left\{ \arg \min_{i \geq 1} \{ \xi_{\mathcal{T}_\alpha}(v, vi) - X_{\mathcal{T}_\alpha}^k(v, vi) \} \geq i_0 \right\} \rightarrow 0 \quad \text{as } i_0 \rightarrow \infty.$$

Correspondingly, for v having label m , the map

$$\pi_{\mathcal{T}_\alpha}^k(v) = \arg \min_{u \sim v} \left\{ \left(\xi_{\mathcal{T}_\alpha}(v, u) - X_{\mathcal{T}_\alpha}^k(v, u) \right)^+ \right\}$$

is a.s. well defined and finite for all $k \geq 1$, and

$$\sup_{k \geq 1} \mathbb{P} \left\{ \max_{i \geq 1} \arg \min_{i \geq 1} \left\{ \left(\xi_{\mathcal{T}_\alpha}(v, vi) - X_{\mathcal{T}_\alpha}^k(v, vi) \right)^+ \right\} \geq i_0 \right\} \rightarrow 0 \quad \text{as } i_0 \rightarrow \infty.$$

The proof of this Lemma follows from the proofs of Lemma 7 of [16] and Lemma 5.4 of [22].

We are now in a position to prove the required convergence.

Theorem 9.3. *The recursive tree process defined by (9.1) with i.i.d. initial messages (according to labels) converges to the unique stationary configuration in the following sense. For every $v \in \mathcal{V}$*

$$X_{\mathcal{T}_\alpha}^k(v, vi) \xrightarrow{L^2} X(v, vi) \quad \text{as } k \rightarrow \infty.$$

Also, the decisions at the root converge, i.e., $\mathbb{P} \{ \pi_{\mathcal{T}_\alpha}^k(\phi) \neq \mathcal{M}_{\text{opt}}(\phi) \} \rightarrow 0$ as $k \rightarrow \infty$.

This can be proved by applying the proof of Theorem 13 of [16] and Theorem 5.2 of [22] separately to the vertices having label o and m .

9.2 Convergence of the update rule on $K_{n,n/\alpha}$ to the update rule on \mathcal{T}_α

We use from [22] the modified definition of local convergence applied to geometric networks with edge labels, i.e., networks in which each directed edge (v, w) has a label $\lambda(v, w)$ taking values in some Polish space. For local convergence of a sequence of such labeled networks G_1, G_2, \dots to a labeled geometric network G_∞ , we add the additional requirement that the rooted graph isomorphisms $\gamma_{n,\rho}$ satisfy

$$\lim_{n \rightarrow \infty} \lambda_{G_n}(\gamma_{n,\rho}(v, w)) = \lambda_{G_\infty}(v, w)$$

for each directed edge (v, w) in $\mathcal{N}_\rho(G_\infty)$.

Now we view the configuration of BP on a graph G at the k^{th} iteration as a labeled geometric network with the label on edge (v, w) given by the pair

$$\left(X_G^k(v, w), \mathbb{1}_{\{v \in \pi_G^k(w)\}} \right).$$

With this definition, our convergence result can be written as the following theorem.

Theorem 9.4. *For every fixed $k \geq 0$, the k^{th} step configuration of BP on $K_{n,n/\alpha}$ converges in the local weak sense to the k^{th} step configuration of BP on \mathcal{T}_α :*

$$\left(K_{n,n/\alpha}, X_{K_{n,n/\alpha}}^k(v, w), \mathbb{1}_{\{v \in \pi_{K_{n,n/\alpha}}^k(w)\}} \right) \xrightarrow{\text{l.w.}} \left(\mathcal{T}_\alpha, X_{\mathcal{T}_\alpha}^k(v, w), \mathbb{1}_{\{v \in \pi_{\mathcal{T}_\alpha}^k(w)\}} \right). \quad (9.2)$$

Proof. The proof of this theorem proceeds along the lines of the proof of Theorem 4.1 of [22].

Consider an almost sure realization of the convergence $K_{n,n/\alpha} \rightarrow \mathcal{T}_\alpha$. For such convergence to hold, the labels of the roots must match eventually, i.e., almost surely label of the root of $K_{n,n/\alpha}$ equals the label of the root of \mathcal{T}_α for sufficiently large n .

Recall from Section 3 the enumeration of the vertices of \mathcal{T}_α from the set \mathcal{V} . We now recursively enumerate the vertices of $K_{n,n/\alpha}$ with multiple members of \mathcal{V} . Denote the root by ϕ . If $v \in \mathcal{V}$ denotes a vertex of $K_{n,n/\alpha}$, then $(v_1, v_2, \dots, v(h_v - 1))$ denote the neighbors of v in $K_{n,n/\alpha}$ ordered by increasing lengths of the corresponding edge with v ; h_v is the number of neighbors of v ($\lceil n/\alpha \rceil$ if v has label o and n if v has label m). Then the convergence in (9.2) is shown if we argue that

$$\forall \{v, w\} \in \mathcal{E} \quad X_{K_{n,n/\alpha}}^k(v, w) \xrightarrow{p} X_{\mathcal{T}_\alpha}^k(v, w) \quad \text{and} \quad \forall v \in \mathcal{V} \quad \pi_{K_{n,n/\alpha}}^k(v) \xrightarrow{p} \pi_{\mathcal{T}_\alpha}^k(v) \quad \text{as } n \rightarrow \infty.$$

The above is trivially true for $k = 0$. We write the update and decision rules, when v has label o , as

$$X_{K_{n,n/\alpha}}^{k+1}(w, v) = \min_{u \in \{v_1, \dots, v(\lceil n/\alpha \rceil - 1), \dot{v}\} \setminus \{w\}} \left\{ \xi_{K_{n,n/\alpha}}(v, u) - X_{K_{n,n/\alpha}}^k(v, u) \right\} \quad \text{and}$$

$$\pi_{K_{n,n/\alpha}}^k(v) = \arg \min_{u \in \{v_1, \dots, v(\lceil n/\alpha \rceil - 1), \dot{v}\}} \left\{ \xi_{K_{n,n/\alpha}}(v, u) - X_{K_{n,n/\alpha}}^k(v, u) \right\},$$

and when v has label m , as

$$X_{K_{n,n/\alpha}}^{k+1}(w, v) = \min_{u \in \{v_1, \dots, v(n-1), \dot{v}\} \setminus \{w\}} \left\{ \left(\xi_{K_{n,n/\alpha}}(v, u) - X_{K_{n,n/\alpha}}^k(v, u) \right)^+ \right\} \quad \text{and}$$

$$\pi_{K_{n,n/\alpha}}^k(v) = \arg \min_{u \in \{v_1, \dots, v(n-1), \dot{v}\}} \left\{ \left(\xi_{K_{n,n/\alpha}}(v, u) - X_{K_{n,n/\alpha}}^k(v, u) \right)^+ \right\},$$

We want to use the convergence of each term on the right-hand side inductively to conclude the convergence of the term on the left. The minimum taken over an unbounded number of terms as $n \rightarrow \infty$ creates a problem. However the following lemma allows us to restrict attention to a uniformly bounded number of terms for each n with probability as high as desired, and hence obtain convergence in probability for each $k \geq 0$. \square

Lemma 9.5. *For all $k \geq 0$, when $v \in \mathcal{V}$ has label o*

$$\lim_{i_0 \rightarrow \infty} \limsup_{n \rightarrow \infty} \mathbb{P} \left\{ \arg \min_{1 \leq i \leq \lceil n/\alpha \rceil - 1} \left\{ \xi_{K_{n,n/\alpha}}(v, vi) - X_{K_{n,n/\alpha}}^k(v, vi) \right\} \geq i_0 \right\} = 0,$$

and when v has label m

$$\lim_{i_0 \rightarrow \infty} \limsup_{n \rightarrow \infty} \mathbb{P} \left\{ \max_{1 \leq i \leq n-1} \arg \min \left\{ \left(\xi_{K_{n,n/\alpha}}(v, vi) - X_{K_{n,n/\alpha}}^k(v, vi) \right)^+ \right\} \geq i_0 \right\} = 0.$$

Proof. The proof is the same as the proof of Lemma 4.1 of [22]. Also see Lemma 8 in [16]. \square

9.3 Completing the upper bound - Proof of Theorem 2.2

By Theorem 9.3, $\pi_{\mathcal{T}_\alpha}^k(\phi) \xrightarrow{p} \mathcal{M}_{\text{opt}}(\phi)$ as $k \rightarrow \infty$. It follows that

$$\sum_{v \in \pi_{\mathcal{T}_\alpha}^k(\phi)} \xi_{\mathcal{T}_\alpha}(\phi, v) \xrightarrow{p} \sum_{v \in \mathcal{M}_{\text{opt}}(\phi)} \xi_{\mathcal{T}_\alpha}(\phi, v) \quad \text{as } k \rightarrow \infty. \tag{9.3}$$

We now prove convergence in expectation. Observe that

$$v \in \pi_{\mathcal{T}_\alpha}^k(\phi) \Rightarrow \xi_{\mathcal{T}_\alpha}(\phi, v) - X_{\mathcal{T}_\alpha}^k(\phi, v) \leq (\xi_{\mathcal{T}_\alpha}(\phi, 1) - X_{\mathcal{T}_\alpha}^k(\phi, 1))^+ \leq \xi_{\mathcal{T}_\alpha}(\phi, 1).$$

By (9.1), $X_{\mathcal{T}_\alpha}^k(\phi, v) \leq \xi_{\mathcal{T}_\alpha}(v, v1)$. Thus,

$$v \in \pi_{\mathcal{T}_\alpha}^k(\phi) \Rightarrow \xi_{\mathcal{T}_\alpha}(\phi, v) \leq \xi_{\mathcal{T}_\alpha}(\phi, 1) + \xi_{\mathcal{T}_\alpha}(v, v1). \tag{9.4}$$

This implies

$$\sum_{v \in \pi_{\mathcal{T}_\alpha}^k(\phi)} \xi_{\mathcal{T}_\alpha}(\phi, v) \leq \xi_{\mathcal{T}_\alpha}(\phi, 1) + \sum_{i \geq 2} \xi_{\mathcal{T}_\alpha}(\phi, i) \mathbb{1}_{\{\xi_{\mathcal{T}_\alpha}(\phi, i) \leq \xi_{\mathcal{T}_\alpha}(\phi, 1) + \xi_{\mathcal{T}_\alpha}(i, i1)\}}.$$

The right-hand side in the above equation is an integrable random variable. For future use let us compute its explicit expectation under the condition that ϕ has label o . First decompose the sum term as

$$\begin{aligned} & \sum_{i \geq 2} \xi_{\mathcal{T}_\alpha}(\phi, 1) \mathbb{1}_{\{\xi_{\mathcal{T}_\alpha}(\phi, i) \leq \xi_{\mathcal{T}_\alpha}(\phi, 1) + \xi_{\mathcal{T}_\alpha}(i, i1)\}} \\ & \quad + \sum_{i \geq 2} (\xi_{\mathcal{T}_\alpha}(\phi, i) - \xi_{\mathcal{T}_\alpha}(\phi, 1)) \mathbb{1}_{\{\xi_{\mathcal{T}_\alpha}(\phi, i) \leq \xi_{\mathcal{T}_\alpha}(\phi, 1) + \xi_{\mathcal{T}_\alpha}(i, i1)\}}. \end{aligned}$$

By exchanging the order of expectation and sum (we can freely do so as the terms are nonnegative), we can replace the random variables $\xi_{\mathcal{T}_\alpha}(i, i1)$ with a single exponential random variable ξ' with mean 1. Using independent increments of the Poisson process we can write the expectation as

$$\mathbb{E} \left[\sum_{i \geq 1} \alpha \mathbb{1}_{\{\xi_i \leq \xi'\}} \right] + \mathbb{E} \left[\sum_{i \geq 1} \xi_i \mathbb{1}_{\{\xi_i \leq \xi'\}} \right],$$

where $(\xi_i)_{i \geq 1}$ is a Poisson process of rate $1/\alpha$, independent of ξ' . It is easy to see that this expectation equals $1 + 1/\alpha$. Thus we have,

$$\mathbb{E} \left[\xi_{\mathcal{T}_\alpha}(\phi, 1) + \sum_{i \geq 2} \xi_{\mathcal{T}_\alpha}(\phi, i) \mathbb{1}_{\{\xi_{\mathcal{T}_\alpha}(\phi, i) \leq \xi_{\mathcal{T}_\alpha}(\phi, 1) + \xi_{\mathcal{T}_\alpha}(i, i1)\}} \mid \text{label}(\phi) = o \right] = \alpha + 1 + 1/\alpha. \tag{9.5}$$

The case when ϕ is conditioned to have label m can be analyzed similarly, but we do not need the explicit value of the expectation.

Equation (9.3) and the dominated convergence theorem give

$$\lim_{k \rightarrow \infty} \mathbb{E} \left[\sum_{v \in \pi_{\mathcal{T}_\alpha}^k(\phi)} \xi_{\mathcal{T}_\alpha}(\phi, v) \right] = \mathbb{E} \left[\sum_{v \in \mathcal{M}_{\text{opt}}(\phi)} \xi_{\mathcal{T}_\alpha}(\phi, v) \right].$$

Involution invariance gives (see Corollary 5.5)

$$\begin{aligned} \lim_{k \rightarrow \infty} \mathbb{E} \left[\sum_{v \in \pi_{\mathcal{T}_\alpha}^k(\phi)} \xi_{\mathcal{T}_\alpha}(\phi, v) \mid \text{label}(\phi) = o \right] &= \mathbb{E} \left[\sum_{v \in \mathcal{M}_{\text{opt}}(\phi)} \xi_{\mathcal{T}_\alpha}(\phi, v) \mid \text{label}(\phi) = o \right] \\ &= c_\alpha^*, \end{aligned} \tag{9.6}$$

where the last equality follows from Theorem 5.3.

By Theorem 9.4 and Lemma 9.5, using the definition of local weak convergence, we have

$$\sum_{v \in \pi_{K_{n,n/\alpha}}^k(\phi)} \xi_{K_{n,n/\alpha}}(\phi, v) \xrightarrow{P} \sum_{v \in \pi_{\mathcal{T}_\alpha}^k(\phi)} \xi_{\mathcal{T}_\alpha}(\phi, v) \quad \text{as } n \rightarrow \infty. \tag{9.7}$$

We now apply the arguments that lead to (9.4) to $\pi_{K_{n,n/\alpha}}^k(\phi)$, and obtain

$$v \in \pi_{K_{n,n/\alpha}}^k(\phi) \Rightarrow \xi_{K_{n,n/\alpha}}(\phi, v) \leq \xi_{K_{n,n/\alpha}}(\phi, 1) + \xi_{K_{n,n/\alpha}}(v, v1).$$

For any two vertices u, v of $K_{n,n/\alpha}$, define $S_n(u, v) = \min_{w \neq u, v} \xi_{K_{n,n/\alpha}}(u, w)$. Then for a vertex v of $K_{n,n/\alpha}$, $\xi_{K_{n,n/\alpha}}(\phi, 1) \leq S_n(\phi, v)$ and $\xi_{K_{n,n/\alpha}}(v, v1) \leq S_n(v, \phi)$. This gives

$$v \in \pi_{K_{n,n/\alpha}}^k(\phi) \Rightarrow \xi_{K_{n,n/\alpha}}(\phi, v) \leq S_n(\phi, v) + S_n(v, \phi).$$

Consequently,

$$\sum_{v \in \pi_{K_{n,n/\alpha}}^k(\phi)} \xi_{K_{n,n/\alpha}}(\phi, v) \leq \sum_v \xi_{K_{n,n/\alpha}}(\phi, v) \mathbb{1}_{\{\xi_{K_{n,n/\alpha}}(\phi, v) \leq S_n(\phi, v) + S_n(v, \phi)\}}. \tag{9.8}$$

Condition that ϕ has label o . Observe that $\xi_{K_{n,n/\alpha}}(\phi, v)$, $S_n(\phi, v)$, and $S_n(v, \phi)$ are independent exponential random variables with means $n, n/(\lceil n/\alpha \rceil - 1)$, and $n/(n - 1)$ respectively. So we can write

$$\begin{aligned} &\mathbb{E} \left[\xi_{K_{n,n/\alpha}}(\phi, v) \mathbb{1}_{\{\xi_{K_{n,n/\alpha}}(\phi, v) \leq S_n(\phi, v) + S_n(v, \phi)\}} \mid \text{label}(\phi) = o \right] \\ &= \int_0^\infty \int_0^x \frac{t}{n} e^{-t/n} dt \left(\frac{(\lceil n/\alpha \rceil - 1)(n - 1)}{n(n - \lceil n/\alpha \rceil)} \right) \left(e^{-\frac{\lceil n/\alpha \rceil - 1}{n}x} - e^{-\frac{n-1}{n}x} \right) dx \\ &= \left(\frac{n}{\lceil n/\alpha \rceil} + 1 \right) \left(\frac{1}{\lceil n/\alpha \rceil} + \frac{1}{n} - \frac{1}{n\lceil n/\alpha \rceil} \right) - \frac{1}{\lceil n/\alpha \rceil}. \end{aligned}$$

Summing over all neighbors of ϕ , we get

$$\begin{aligned} \mathbb{E} \left[\sum_v \xi_{K_{n,n/\alpha}}(\phi, v) \mathbb{1}_{\{\xi_{K_{n,n/\alpha}}(\phi, v) \leq S_n(\phi, v) + S_n(v, \phi)\}} \right] \\ = \left(\frac{n}{\lceil n/\alpha \rceil} + 1 \right) \left(1 + \frac{\lceil n/\alpha \rceil}{n} - \frac{1}{n} \right) - 1, \end{aligned} \quad (9.9)$$

which converges to $1 + \alpha + 1/\alpha$ as $n \rightarrow \infty$.

Using local weak convergence, we can see that

$$\begin{aligned} \sum_v \xi_{K_{n,n/\alpha}}(\phi, v) \mathbb{1}_{\{\xi_{K_{n,n/\alpha}}(\phi, v) \leq S_n(\phi, v) + S_n(v, \phi)\}} \\ \xrightarrow{P} \xi_{\mathcal{T}_\alpha}(\phi, 1) + \sum_{i \geq 2} \xi_{\mathcal{T}_\alpha}(\phi, i) \mathbb{1}_{\{\xi_{\mathcal{T}_\alpha}(\phi, i) \leq \xi_{\mathcal{T}_\alpha}(\phi, 1) + \xi_{\mathcal{T}_\alpha}(i, i1)\}}. \end{aligned}$$

By (9.5), the expectation of the limit above equals $1 + \alpha + 1/\alpha$. Using this with (9.7), (9.8), and (9.9), the generalized dominated convergence theorem yields

$$\begin{aligned} \lim_{n \rightarrow \infty} \mathbb{E} \left[\sum_{v \in \pi_{K_{n,n/\alpha}}^k(\phi)} \xi_{K_{n,n/\alpha}}(\phi, v) \mid \text{label}(\phi) = o \right] \\ = \mathbb{E} \left[\sum_{v \in \pi_{\mathcal{T}_\alpha}^k(\phi)} \xi_{\mathcal{T}_\alpha}(\phi, v) \mid \text{label}(\phi) = o \right]. \end{aligned} \quad (9.10)$$

Combining (9.10) and (9.6) gives

$$\lim_{k \rightarrow \infty} \lim_{n \rightarrow \infty} \mathbb{E} \left[\sum_{v \in \pi_{K_{n,n/\alpha}}^k(\phi)} \xi_{K_{n,n/\alpha}}(\phi, v) \mid \text{label}(\phi) = o \right] = c_\alpha^*. \quad (9.11)$$

Use of

$$\mathbb{E} \left[\sum_{e \in \mathcal{M}(\pi_{K_{n,n/\alpha}}^k)} \xi_{K_{n,n/\alpha}}(e) \right] = n \mathbb{E} \left[\sum_{v \in \pi_{K_{n,n/\alpha}}^k(\phi)} \xi_{K_{n,n/\alpha}}(\phi, v) \mid \text{label}(\phi) = o \right]$$

with (9.11) completes the proof of Theorem 2.2. □

9.4 Completing the proof of Theorem 2.1

As remarked earlier, $\mathcal{M}(\pi_{K_{n,n/\alpha}}^k)$, unlike its limit \mathcal{M}_{opt} , is not a many-to-one matching. The matching condition is violated for a vertex $v \in A$ if there are more than one vertices w such that $v \in \pi_{K_{n,n/\alpha}}^k(w)$. We now show how to construct a valid many-to-one matching from $\mathcal{M}(\pi_{K_{n,n/\alpha}}^k)$ with little increase in cost.

Theorem 9.6. *Fix $\epsilon > 0$. There exist many-to-one matchings \mathcal{M}_n on $K_{n,n/\alpha}$ such that for all sufficiently large k*

$$\limsup_{n \rightarrow \infty} n^{-1} \mathbb{E} \left[\sum_{e \in \mathcal{M}_n} \xi_{K_{n,n/\alpha}}(e) \right] \leq \lim_{n \rightarrow \infty} n^{-1} \mathbb{E} \left[\sum_{e \in \mathcal{M}(\pi_{K_{n,n/\alpha}}^k)} \xi_{K_{n,n/\alpha}}(e) \right] + \epsilon.$$

Proof. Define a partial matching $\tilde{\mathcal{M}}(\pi_{K_{n,n/\alpha}}^k) = \{\{v, w\} \mid w \in \pi_{K_{n,n/\alpha}}^k(v), v \in A\}$. Because it is a subset of $\mathcal{M}(\pi_{K_{n,n/\alpha}}^k)$, the cost of $\tilde{\mathcal{M}}(\pi_{K_{n,n/\alpha}}^k)$ denoted cost_n^k is bounded above by the cost of $\mathcal{M}(\pi_{K_{n,n/\alpha}}^k)$. Every vertex in A has exactly one edge in $\tilde{\mathcal{M}}(\pi_{K_{n,n/\alpha}}^k)$, but some vertices in B may be unmatched. Call this set of unmatched vertices U_n^k . By Theorem 9.4, for a vertex $w \in B$,

$$\mathbb{P}\{w \in U_n^k\} \rightarrow \mathbb{P}\left\{\phi \notin \bigcup_{v \sim \phi} \pi_{\mathcal{T}_\alpha}^k(v) \mid \text{label}(\phi) = m\right\} =: r$$

as $n \rightarrow \infty$. We have $r \rightarrow 0$ as $k \rightarrow \infty$ (Theorem 9.3). For any $\delta > 0$, we then have $\mathbb{E}|U_n^k| \leq (r + \delta)n/\alpha$ for all sufficiently large n .

Section 3.3 of [25] describes a method (which is based on [13]) for constructing a perfect matching in the random complete graph K_n from a *diluted* matching that has a small fraction of vertices unmatched, with a small increase in cost. Our case is much simpler because more than one edge is allowed for a vertex in B . We start with the first step of the proof in [25], and deviate afterwards.

In $K_{n,n/\alpha}$ substitute the single edge between every pair of vertices with a countably infinite number of edges and take their weights according to a Poisson process on \mathbb{R}_+ with rate $1/n$. The weight of the smallest edge between every pair of vertices is still distributed as the original (exponential with mean n), and so the optimal many-to-one matching in this new network is same as that in $K_{n,n/\alpha}$. Independently color every edge red with probability $1 - p$ and green with probability p (we will later take p to 0). Obtain a partial matching $\tilde{\mathcal{M}}(\pi_{K_{n,n/\alpha}}^k)$ on this network using only red edges. The restriction to red edges just affects the scaling of the cost: the weight of the smallest red edge between a vertex pair has exponential distribution with mean $n/(1 - p)$, and so the cost of this partial matching is $\text{cost}_n^k/(1 - p)$.

For every matched vertex w in $B \setminus U_n^k$ leave out one vertex in $\pi_{K_{n,n/\alpha}}^k(w)$ (choice made arbitrarily if size more than 1) and collect the remaining vertices in a set S_n^k . There are exactly $n - (\lceil n/\alpha \rceil - |U_n^k|)$ vertices in this set. We can complete the many-to-one matching by greedily matching every vertex in U_n^k with a distinct vertex in S_n^k and dropping the original matched edge of that vertex. We will do this by using only green edges. The conditional expected cost of these green edges, given the coloring and the weight of the red edges, is at most $\frac{|U_n^k|}{|S_n^k| - |U_n^k| + 1} \frac{n}{p} \leq \frac{|U_n^k|}{(1 - 1/\alpha)p}$. We get this because the choice of the sets U_n^k and S_n^k is governed by the red edges, whose weights are independent of those of the green edges.

The expected cost of the constructed many-to-one matching is at most

$$\mathbb{E}\left[\frac{\text{cost}_n^k}{(1 - p)} + \frac{|U_n^k|}{(1 - 1/\alpha)p}\right] \leq \frac{\mathbb{E}[\text{cost}_n^k]}{(1 - p)} + \frac{(r + \delta)n}{(\alpha - 1)p}$$

for large n .

Since $r \rightarrow 0$ as $k \rightarrow \infty$ and $\delta > 0, p > 0$ are arbitrary, we can bound the expected cost by $\mathbb{E}[\text{cost}_n^k] + \epsilon n$ for large k and large n . \square

By Theorem 2.2 for any $\epsilon > 0$, we can find k large such that

$$\lim_{n \rightarrow \infty} n^{-1} \mathbb{E}\left[\sum_{e \in \mathcal{M}(\pi_{K_{n,n/\alpha}}^k)} \xi_{K_{n,n/\alpha}}(e)\right] \leq c_\alpha^* + \epsilon.$$

This combined with Theorem 9.6 gives

$$\limsup_{n \rightarrow \infty} n^{-1} \mathbb{E} M_n \leq c_\alpha^* + 2\epsilon.$$

Since ϵ is arbitrary, we get the upper bound

$$\limsup_{n \rightarrow \infty} n^{-1} \mathbb{E} M_n \leq c_\alpha^*.$$

With the lower bound from Theorem 6.1, we have proved Theorem 2.1. \square

Observe that for any $\epsilon > 0$, there exist K_ϵ and N_ϵ such that for all $k \geq K_\epsilon$ and $n \geq N_\epsilon$, we have

$$n^{-1} \mathbb{E} \left[\sum_{e \in \mathcal{M} \left(\binom{k}{K_{n,n/\alpha}} \right)} \xi_{K_{n,n/\alpha}}(e) \right] \leq c_\alpha^* + \epsilon.$$

Thus, for large n , the BP algorithm gives a solution with cost within ϵ of the optimal value in K_ϵ iterations. In an iteration, the algorithm requires $O(n)$ computations at every vertex. This gives an $O(K_\epsilon n^2)$ running time for the BP algorithm to compute an ϵ -approximate solution.

References

- [1] David Aldous, *Asymptotics in the random assignment problem*, Probab. Theory Related Fields **93** (1992), no. 4, 507–534. MR-1183889
- [2] David Aldous and Russell Lyons, *Processes on unimodular random networks*, Electron. J. Probab. **12** (2007), no. 54, 1454–1508. MR-2354165
- [3] David Aldous and J. Michael Steele, *The objective method: probabilistic combinatorial optimization and local weak convergence*, Encyclopaedia Math. Sci., vol. 110, pp. 1–72, Springer, Berlin, 2004. MR-2023650
- [4] David J. Aldous, *The $\zeta(2)$ limit in the random assignment problem*, Random Structures Algorithms **18** (2001), no. 4, 381–418. MR-1839499
- [5] David J. Aldous and Antar Bandyopadhyay, *A survey of max-type recursive distributional equations*, Ann. Appl. Probab. **15** (2005), no. 2, 1047–1110. MR-2134098
- [6] Antar Bandyopadhyay, *Endogeny for the logistic recursive distributional equation*, Z. Anal. Anwend. **30** (2011), no. 2, 237–251. MR-2793003
- [7] Mohsen Bayati, Devavrat Shah, and Mayank Sharma, *Max-product for maximum weight matching: convergence, correctness, and LP duality*, IEEE Trans. Inform. Theory **54** (2008), no. 3, 1241–1251. MR-2445062
- [8] Amir Ben-Dor, Richard M Karp, Benno Schwikowski, and Ron Shamir, *The restriction scaffold problem*, Journal of Computational Biology **10** (2003), no. 3-4, 385–398.
- [9] Itai Benjamini and Oded Schramm, *Recurrence of distributional limits of finite planar graphs*, Electron. J. Probab. **6** (2001), no. 23, 13 pp. (electronic). MR-1873300
- [10] Hendrik Blockeel, Luc De Raedt, and Jan Ramong, *Top-down induction of clustering trees*, Proceedings of the 15th International Conference on Machine Learning, Morgan Kaufmann, 1998, pp. 55–63.
- [11] Thomas Eiter and Heikki Mannila, *Distance measures for point sets and their computation*, Acta Inform. **34** (1997), no. 2, 109–133. MR-1465033
- [12] A. M. Frieze, *On the value of a random minimum spanning tree problem*, Discrete Appl. Math. **10** (1985), no. 1, 47–56. MR-770868
- [13] Alan Frieze, *On random symmetric travelling salesman problems*, Math. Oper. Res. **29** (2004), no. 4, 878–890. MR-2104159
- [14] Bruce Hajek, *Performance of global load balancing by local adjustment*, IEEE Trans. Inform. Theory **36** (1990), no. 6, 1398–1414. MR-1080823

- [15] Martin Hessler and Johan Wästlund, *Edge cover and polymatroid flow problems*, Electron. J. Probab. **15** (2010), no. 72, 2200–2219. MR-2748403
- [16] Mustafa Khandwawala and Rajesh Sundaresan, *Belief propagation for optimal edge cover in the random complete graph*, Ann. Appl. Probab. **24** (2014), no. 6, 2414–2454. MR-3262507
- [17] Marc Mézard and Giorgio Parisi, *Replicas and optimization*, Journal de Physique Lettres **46** (1985), no. 17, 771–778.
- [18] ———, *Mean-field equations for the matching and the travelling salesman problems*, EPL (Europhysics Letters) **2** (1986), no. 12, 913.
- [19] ———, *On the solution of the random link matching problem*, J.Physique **48** (1987), 1451–1459.
- [20] Marc Mézard, Giorgio Parisi, and Miguel Angel Virasoro, *Spin glass theory and beyond*, World Scientific Lecture Notes in Physics, vol. 9, World Scientific Publishing Co. Inc., Teaneck, NJ, 1987. MR-1026102
- [21] Giorgio Parisi, *A conjecture on random bipartite matching*, (1998), arXiv:cond-mat/9801176 [cond-mat.dis-nn] <http://arxiv.org/abs/cond-mat/9801176>.
- [22] Justin Salez and Devavrat Shah, *Belief propagation: an asymptotically optimal algorithm for the random assignment problem*, Math. Oper. Res. **34** (2009), no. 2, 468–480. MR-2554069
- [23] Sujay Sanghavi, Dmitry Malioutov, and Alan Willsky, *Belief propagation and LP relaxation for weighted matching in general graphs*, IEEE Trans. Inform. Theory **57** (2011), no. 4, 2203–2212. MR-2760242
- [24] Johan Wästlund, *The mean field traveling salesman and related problems*, Acta Math. **204** (2010), no. 1, 91–150. MR-2600434
- [25] ———, *Replica symmetry of the minimum matching*, Ann. of Math. (2) **175** (2012), no. 3, 1061–1091. MR-2912702

Acknowledgments. This work resulted from discussions with Rajesh Sundaresan. I also thank him for suggestions that helped improve the presentation.