A Proof of the Conjecture due to Parisi for the Finite Random Assignment Problem

Chandra Nair
Dept. of EE
Stanford University
mchandra@stanford.edu

Balaji Prabhakar
Depts. of EE and CS
Stanford University
balaji@stanford.edu

Mayank Sharma
Dept. of EE
Stanford University
msharma@stanford.edu

Abstract

Suppose that there are n jobs and n machines and it costs c_{ij} to execute job i on machine j. The assignment problem concerns the determination of a one-to-one assignment of jobs onto machines so as to minimize the cost of executing all the jobs. When the c_{ij} are independent and identically distributed exponentials of mean 1, Parisi has made the beautiful conjecture that the average minimum cost equals $\sum_{i=1}^{n} \frac{1}{i^2}$. Coppersmith and Sorkin (1999) have generalized Parisi's conjecture to the average value of the smallest k-assignment when there are n jobs and m machines. In this paper, we develop combinatorial and probabilistic arguments and resolve a conjecture of Sharma and Prabhakar (2002). This implies a proof of Parisi's conjecture. Our arguments also resolve a conjecture due to Nair (2002), and this yields a proof of the Coppersmith-Sorkin conjecture.

1 Introduction

Suppose there are n jobs and n machines and it costs c_{ij} to execute job i on machine j. An assignment (or a matching) is a one-to-one mapping of jobs onto machines. Representing an assignment as a permutation $\pi : \{1, ..., n\} \rightarrow \{1, ..., n\}$, the cost of the assignment π equals $\sum_{i=1}^{n} c_{i,\pi(i)}$. The assignment problem consists of finding the assignment with the lowest cost. Let

$$A_n = \min_{\pi} \sum_{i=1}^{n} c_{i,\pi(i)}$$

represent the cost of the minimizing assignment. In the *random* assignment problem the c_{ij} are i.i.d. random variables drawn from some distribution. A quantity of interest in the random assignment problem is the expected minimum cost, $\mathbb{E}(A_n)$.

When the c_{ij} are i.i.d. exp(1) variables, Parisi [Pa 98] makes the following conjecture:

$$\mathbb{E}(A_n) = \sum_{i=1}^{n} \frac{1}{i^2}.$$

Coppersmith and Sorkin [CS 99] propose a larger class of conjectures which state that the expected cost of the minimum k-assignment in an $m \times n$ matrix of i.i.d. exp(1) entries is:

$$F(k, m, n) = \sum_{i,j \geq 0, i+j<k} \frac{1}{(m-i)(n-j)}.$$

By definition, $F(n, n, n) = \mathbb{E}(A_n)$ and their expression coincides with Parisi's conjecture.

In this paper, we prove Parisi's conjecture by two different but related strategies. Both involve establishing the exponentiality of the increments of the weights of matchings. The first builds on the work of Sharma and Prabhakar [SP 02] and establishes Parisi's conjecture by showing that certain increments of weights of matchings are exponentially distributed with a given rate and are independent. The second method builds on Nair [Na 02] and establishes the Coppersmith-Sorkin conjectures. It does this by showing that certain other increments are exponentials with given rates; the increments are not required to be (and, in fact, aren't) independent.

The two methods mentioned above use a common set of combinatorial and probabilistic arguments. For ease of exposition, we choose to present the proof of the conjectures in [SP 02] first. We then show how the arguments also resolve the conjectures in [Na 02].

Before surveying the literature on this problem, it is important to mention that simultaneously and independently of our work Linusson and Wästlund [LW 03] have also announced a proof of the Parisi and Coppersmith-Sorkin conjectures based on a quite different approach.

1.1 Background and related work

There has been a lot of work on determining bounds for the expected minimum cost and calculating its asymptotic value. Assuming, for now, that $\lim_n \mathbb{E}(A_n)$ exists, let us denote it by A^*. We survey some of the work; more details can be found in [St 97, CS 99]. Early work uses feasible solutions to the dual linear programming (LP) formulation of the assignment problem for obtaining the following lower bounds for A^*: $(1 + 1/e)$ by Lazarus [La 93], 1.441 by Goemans and Kodialam [GK 93], and 1.51 by Olin [Ol 92]. The first upper bound of 3 was given by Walkup [Wa 79], who thus demonstrated that $\limsup_n E(A_n)$ is finite. Walkup's argument was later made constructive by Karp *et al* [KKV 94]. Karp [Ka 84, Ka 87] made a subtle use of LP duality to obtain a better upper bound of 2. Coppersmith and Sorkin [CS 99] have further improved the bound to 1.94.

Meanwhile, it had been observed through simulations that for large n, $E(A_n) \approx 1.642$ [BKMP 86]. Mézard and Parisi [MP 87] used the *replica method* [MPV 87] of statistical physics to argue that $A^* = \frac{\pi^2}{6}$. (Thus, Parisi's conjecture for the finite random assignment problem is an elegant restriction, for i.i.d. exp(1) costs, of the asymptotic result to the first n terms in the expansion: $\frac{\pi^2}{6} = \sum_{i=1}^{\infty} \frac{1}{i^2}$.) More interestingly, their method allowed them to determine the density of the edge-weight distribution of the limiting optimal matching. These sharp (but non-rigorous) asymptotic results, and others of a similar flavor that they obtained in several combinatorial optimization problems, sparked interest in the replica method and in the random assignment problem.

Aldous [Al 92] proved that A^* exists by identifying the limit as the value of a minimum-cost matching problem on a certain random weighted infinite tree. In the same work he also established that the distribution of c_{ij} affects A^* only through the value of its density function at 0 (provided it exists and is strictly positive). Thus, as far as the value of A^* is concerned, the distributions $U[0, 1]$ and exp(1) are equivalent. More recently, Aldous [Al 01] has established that $A^* = \pi^2/6$, and obtained the same limiting optimal edge-weight distribution as [MP 87]. He also obtains a number of other interesting results such as the asymptotic essential uniqueness (AEU) property—which roughly states that almost-optimal matchings have almost all their edges equal to those of the optimal matching.

Generalizations of Parisi's conjecture have also been made in other ways. Linusson and Wästlund [LW 00] conjecture an expression for the expected cost of the minimum k-assignment in an $m \times n$ matrix consisting of zeroes at some specified positions and exp(1) entries at all other places. Indeed, it is by establishing this conjecture in their recent work [LW 03] that they obtain proofs of the Parisi and Coppersmith-Sorkin conjectures. Buck, Chan and Robbins [BCR 02] generalize the Coppersmith-Sorkin conjecture for matrices $M \equiv [m_{ij}]$ with $m_{ij} \sim \exp(r_i c_j)$ for $r_i, c_j > 0$.

Alm and Sorkin [AS 02] verify the Coppersmith-Sorkin conjecture when $k \leq 4$, $k = m = 5$ and $k = m = n = 6$; and Coppersmith and Sorkin [CS 02] study the expected incremental cost of going from the smallest $(m-1)$-assignment in an $(m-1) \times n$ matrix to the smallest m-assignment in an

$m \times n$ matrix.

2 Preliminaries

We begin by recalling the notation and the essentials of the development in [SP 02]. Let $C = [c_{ij}]$ be an $n \times n$ cost matrix with i.i.d. exp(1) entries. Delete the top row of C to obtain the rectangular matrix L of dimensions $(n-1) \times n$. For each $i = 1, \ldots, n$, let S_i be the cost of the minimum-cost permutation in the sub-matrix obtained by deleting the i^{th} column of L. These quantities are illustrated below, taking L to be the following 2×3 matrix.

$$L: \begin{array}{|c|c|c|} \hline 3 & 6 & 11 \\ \hline 9 & 2 & 20 \\ \hline \end{array}$$

$$\begin{array}{|c|c|} \hline 6 & \mathbf{11} \\ \hline \mathbf{2} & 20 \\ \hline \end{array} \Rightarrow S_1 = 13; \qquad \begin{array}{|c|c|} \hline 3 & \mathbf{11} \\ \hline \mathbf{9} & 20 \\ \hline \end{array} \Rightarrow S_2 = 20; \qquad \begin{array}{|c|c|} \hline \mathbf{3} & 6 \\ \hline 9 & \mathbf{2} \\ \hline \end{array} \Rightarrow S_3 = 5.$$

Denote the ordered sequence of S_i by $T_i,\ i = 1, \ldots, n$. That is, let σ be the random permutation of $\{1, \ldots, n\}$ such that $S_{\sigma(1)} \leq \ldots \leq S_{\sigma(n)}$. Then $T_i = S_{\sigma(i)}$. In the above example, $T_1 = 5$, $T_2 = 13$ and $T_3 = 20$.

Conventions
(1) Throughout this paper, if the minimum weight permutation is not unique, we shall consistently break the tie in favor of one of these permutations.
(2) To avoid an explosion of the notation, we shall some times use the same symbol for both the *name* of a matching and for its *weight*. For example, the T_1 defined above might refer to the smallest matching as well as to its weight.

Conjecture 1 *[SP 02] For $j = 1, \ldots, n-1$, $T_{j+1} - T_j \sim \exp(j(n-j))$ and these increments are independent of each other.*[1]

Remarks: It was shown in [SP 02] that Conjecture 1 implies Parisi's conjecture: $E(A_n) = \sum_{i=1}^{n} \frac{1}{i^2}$. Further, it was established that $T_2 - T_1 \sim \exp(n-1)$ and that $T_3 - T_2 \perp\!\!\!\perp T_2 - T_1$. As noted in [SP 02], a natural generalization of their argument for showing $T_2 - T_1 \sim \exp(n-1)$ fails for higher increments. However, we use a different generalization in this paper, involving a subtle randomization step, to prove Conjecture 1.

2.1 A Sketch of the Proof

We first introduce some notation that will be used in the proof. For an $(n-j) \times n$ matrix, L_{n-j}, of i.i.d. exp(1) random variables define T_1^{n-j} to be the weight of the smallest matching of size $n-j$.

Let K_1 be the matrix formed by the $n-j$ columns used by T_1^{n-j}. Consider the set of all matchings in L_{n-j} of size $n-j$ which use exactly $n-j-1$ columns of K_1. Define T_2^{n-j} to be the weight of the smallest matching in this set.

Now let K_{12} be the matrix of size $(n-j) \times (n-j-1)$ consisting of the columns used by *both* T_1^{n-j} and T_2^{n-j}. Consider the set of all matchings in L_{n-j} of size $n-j$ which use exactly $n-j-2$ columns of K_{12}. Let T_3^{n-j} be the weight of the smallest matching in this set.

[1] The symbol '$\sim$' stands for 'is distributed as' and the symbol ' $\perp\!\!\!\perp$ ' stands for 'is independent of'.

In this way we can recursively define $K_{12\ldots i}$, $i = 3, \ldots, n-j$ to be the matrix of size $(n-j) \times (n-j-i+1)$ consisting of the columns present in all the matchings $\{T_1^{n-j}, T_2^{n-j}, \ldots, T_i^{n-j}\}$, and obtain T_4^{n-j}, ..., T_{n-j+1}^{n-j} as above. We shall refer to the matchings $\{T_1^{n-j}, ..., T_{n-j+1}^{n-j}\}$ as the *T-matchings* of the matrix L_{n-j}.

Note: When $j = 1$ then $L_{n-1} = L$ and the *T-matchings* of the matrix L are exactly equal to the sequence $\{T_i\}_1^n$ defined above, i.e $T_i = T_i^{n-1}, \ i = 1, \ldots, n$.

Our proof of Conjecture 1 is inductive, and follows the steps below.

1. At the j^{th} step, we shall start with a matrix L_{n-j} of size $(n-j) \times n$ containing i.i.d. exp(1) random variables such that $T_{k+1} - T_k = T_{k-j+2}^{n-j} - T_{k-j+1}^{n-j}, k \geq j$.

   We shall prove that $T_2^{n-j} - T_1^{n-j} \sim \exp(j(n-j))$ and hence conclude that $T_{j+1} - T_j \sim \exp(j(n-j))$.

2. Next, we obtain L_{n-j-1}, a matrix of size $(n-j-1) \times n$, from L_{n-j} such that $T_{k+1} - T_k = T_{k-j+1}^{n-j-1} - T_{k-j}^{n-j-1}, k \geq j+1$.

3. We shall show that L_{n-j-1} contains i.i.d. exp(1) random variables, thus completing the induction from L_{n-j} to L_{n-j-1}.

   We will also show that $T_{j+1} - T_j$ is independent of L_{n-j-1} and observe that the higher increments $T_{k+1} - T_k$, for $k > j$, are functions of L_{n-j-1}. This will allow us to conclude that $T_{j+1} - T_j$ is independent of $T_{k+1} - T_k$ for $k > j$.

Remark: The randomization procedure alluded to earlier is used in obtaining L_{n-j-1} from L_{n-j} and ensures that L_{n-j-1} has i.i.d. exp(1) entries.

We state some combinatorial properties regarding matchings in Section 3 that will be useful for the rest of the paper. The proofs of these lemmas are provided in the Appendix. Section 4 establishes the above induction, thus completing the proof of Conjecture 1. For completeness, in Section 5, we recall the argument in [SP 02] and show that Conjecture 1 implies Parisi's conjecture. We extend the method used in Section 4 to prove the Coppersmith-Sorkin conjecture in Section 6. We conclude in Section 7.

3 Some combinatorial properties of matchings

Lemma 1 *Consider an $n \times (n+l)$ matrix M and let K_1 be the set of columns used by T_1^n. Then for any column $\mathbf{c} \in K_1$, the smallest matching of size n that doesn't use $\mathbf{c}$, denoted by $\mathcal{M}_{\mathbf{c}}$, contains exactly one element outside the columns in K_1.*

Lemma 2 *Consider an $n \times (n+l)$ matrix M in which T_1^n is the smallest matching of size n. Suppose there exists a collection of n columns in M, denoted by K, with the property that the smallest matching $\mathcal{M}$ of size n in K is lighter than $\mathcal{M}'$: the smallest among all matchings of size n that have exactly one element outside K. Then $\mathcal{M} = T_1^n$.*

Lemma 3 *Let $\mathcal{M}$ be the smallest matching of size n in an $n \times (n+l)$ matrix M. Let M_{ext} be any $(n+r) \times (n+l)$ matrix $(r \leq l)$ whose first n rows equal M (i.e. M_{ext} is an extension of M). Let $\mathcal{M}'$ denote the smallest matching of size $n+r$ in M_{ext}. Then the set of columns in which the elements of $\mathcal{M}$ lie is a subset of the columns in which the elements of $\mathcal{M}'$ lie.*

The following is a useful alternate description of $T_i^{n-j}, \ \ i \geq 2$. Compute the weight of the smallest matching in L_{n-j} after removing its i^{th} column and denote the weight of that matching by S_i^{n-j}.

Note that for all $i \notin K_1$, $S_i^{n-j} = T_1^{n-j}$. Thus we shall only consider S_i^{n-j} when $i \in K_1$ and arrange them in increasing order to get $\{R_i^{n-j}\}_1^{n-j}$. Then we get the following lemma.

Lemma 4 $T_{i+1}^{n-j} = R_i^{n-j}, \quad i = 1, ..., n-j.$

Corollary 1 *The matchings* T_i^{n-j}, $i = 2, ..., n-j+1$ *contain exactly one element outside* K_1.

Corollary 2 *Given a matrix* L_{n-j}, *arranging the* S_i^{n-j} *for* $i \in K_{12\ldots k}$ *in increasing order gives the sequence* $T_{k+1}^{n-j}, T_{k+2}^{n-j}, \ldots, T_{n-j+1}^{n-j}$.

Lemma 5 *[SP 03] Consider all the elements of an* $n \times (n+1)$ *matrix* M *that participate in at least one of the matchings* T_i *(or, equivalently, in one of the* S_j*). The total number of such elements is* $2n$ *and exactly* 2 *such elements are present in every row of* M.

Proof See Sharma and Prabhakar [SP 03]. ∎

Lemma 6 *Consider an* $n \times (n+l)$ *matrix* M *and mark all the elements that participate in any one of the matchings* $T_1^n, T_2^n, \ldots, T_{n+1}^n$ *defined in Section 1. Then exactly two elements will be marked in each row.*

4 Proof of Conjecture 1

We shall execute the three steps mentioned earlier. Thus, we shall begin with an $(n-j) \times n$ matrix, L_{n-j}, which has i.i.d. exp(1) entries. Recall that when $j = 1$, $L_{n-1} = L$ and $T_{k+1} - T_k = T_{k+1}^{n-1} - T_k^{n-1}$, $k = 1, \ldots, n-1$, since by definition $T_k = T_k^{n-1}\ \forall\, k$. This starts off the induction.

Now consider the general step j. Suppose we have that L_{n-j} contains i.i.d. exp(1) entries and $T_{k+1} - T_k = T_{k-j+2}^{n-j} - T_{k-j+1}^{n-j}$ for $k \geq j$, where the T_i^{n-j} are as defined in Section 1.

Step 1: Distribution of $T_2^{n-j} - T_1^{n-j} = T_{j+1} - T_j$

Theorem 1 *[Na 02]* $T_2^{n-j} - T_1^{n-j} \sim \exp(j(n-j))$.

Proof Let K be an $(n-j) \times (n-j)$ matrix consisting of any $n-j$ columns of L_{n-j}. Let T_1^K be the smallest matching of size $n-j$ in K. For any element c not in K, let Δ_c be the weight of the smallest matching of size $n-j-1$ from K that does not have an element in the same row as c. Thus, $c + \Delta_c$ will be the weight of the smallest matching of size $n-j$ that contains c and precisely $n-j-1$ elements from K. As an illustration, the crosses in the figure below constitute Δ_c, and the elements marked by the dots constitute T_1^K.

•			×				
	•					c	
	×	•					
×			•				

Let $\bar{K}$ be the $(n-j) \times j$ matrix formed by the j columns of L_{n-j} not present in K. Let $c_1, c_2, ..., c_{j(n-j)}$ be an enumeration of the elements of $\bar{K}$. We claim that $d_i \triangleq T_1^K - \Delta_{c_i} \geq 0$. This follows from two points: (i) T_1^K has $n-j-1$ elements that are not in the row of c_i, and so c_i has the option of choosing them as Δ_{c_i}, and (ii) all the entries of L_{n-j} are positive. Further, note that $\sigma(d_i) \subset \sigma(K),\ i = 1, ..., j(n-j)$, where $\sigma(X)$ denotes the sigma-algebra generated by X.

We proceed by stating the following fact, whose proof follows from a simple application of the memoryless property of exponentials.

Fact 1 *Suppose $X_i,\ i = 1, ..., m$, are independent* $\exp(1)$ *random variables. Let $Y_i \geq 0,\ i = 1, ..., m$ be random variables such that $\sigma(Y_1, ..., Y_m) \subset \mathcal{F}$ for some sigma algebra $\mathcal{F}$. If $X_i \perp\!\!\!\perp \mathcal{F}$ for each i, then on the event $\{X_i \geq Y_i,\ i = 1, ..., m\}$, $X_i - Y_i$ are independent* $\exp(1)$ *random variables and independent of $\mathcal{F}$.*

Consider the event $E = \{K$ contains the smallest matching of size $n-j$ in $L_{n-j}\}$. We will show that E is same as the event $F = \{c_i \geq d_i,\ i = 1, ..., j(n-j)\}$. Since $T_1^K = T_1^{n-j}$ by definition of E, the minimality of T_1^{n-j} implies that $c_i + \Delta_{c_i} \geq T_1^K$. Hence $E \subset F$. Conversely, given F, Lemma 2 implies that $T_1^K = T_1^{n-j}$. Therefore $F \subset E$.

Given this, Fact 1 implies $c_i - d_i$ are independent $\exp(1)$ random variables on the event E. Conditioned on E, observe that

$$T_2^{n-j} - T_1^{n-j} = \min_i \{c_i - d_i\}. \tag{1}$$

Being the minimum of $j(n-j)$ independent $\exp(1)$ random variables, $T_2^{n-j} - T_1^{n-j} \sim \exp(j(n-j))$ on E.

Now K was any $(n-j) \times (n-j)$ sub-matrix of L_{n-j}. Repeating the argument for all possible choices of K proves the theorem. ■

Remark: The theorem in [Na 02] considers a slightly more general setting of matchings of size k in an $m \times n$ matrix. The argument used in the above theorem is an extension of the argument in [SP 02] for showing that $T_2 - T_1 \sim \exp(n-1)$.

Corollary 3 *Let K_1 be the columns used by the smallest matching in L_{n-j}. For each $c_i \notin K_1$, $c_i - (T_2^{n-j} - T_1^{n-j}) \geq 0$.*

Proof This follows from equation (1) and the fact that $d_i \geq 0$, where d_i are defined in the proof of Theorem 1. ■

Step 2: Obtaining L_{n-j-1} from L_{n-j}

The matrix L_{n-j-1} is obtained from L_{n-j} by applying the series of operations Φ, Λ and Π, as depicted below

$$L_{n-j} \xrightarrow{\Phi} L^*_{n-j} \xrightarrow{\Lambda} \tilde{L}_{n-j} \xrightarrow{\Pi} L_{n-j-1}.$$

It would be natural to represent the *T-matchings* corresponding to each of the intermediate matrices L^*_{n-j} and $\tilde{L}_{n-j}$ using appropriate superscripts. For example, the *T-matchings* of L^*_{n-j} can be denoted $\{T_i^{*(n-j)}\}$. However, this creates a needless clutter of symbols. We will instead denote L^*_{n-j} and $\{T_i^{*(n-j)}\}$ simply as L^* and $\{T_i^*\}$, respectively. Table 1 summarizes the notations we shall use.

We now specify the operations Φ, Λ and Π.

Φ : In the matrix L_{n-j}, subtract the value $T_2^{n-j} - T_1^{n-j}$ from each entry not in the sub-matrix K_1. (Recall from Section 1 that K_1 denotes the $(n-j) \times (n-j)$ sub-matrix of L_{n-j} whose columns contain the entries used by T_1^{n-j}.) Let the resultant matrix of size $(n-j) \times n$ be L^*.

$\mathbf{\Lambda}$: Consider the matrix L^*. Generate a random variable X with $\mathbb{P}(X = 1) = \mathbb{P}(X = 2) = \frac{1}{2}$. Define Y so that $Y = 1$ when $X = 2$ and $Y = 2$ when $X = 1$. Denote by e_1 the unique entry of the matching T^*_X outside the matrix K_{12}. Denote by e_2 the unique entry of the matching T^*_Y outside the matrix K_{12}. Now remove the row of L^* in which e_1 is present and append it to the bottom of the matrix.[2] Call the resultant matrix $\tilde{L}$. Note that $\tilde{L}$ is a row permutation of L^*.

$\mathbf{\Pi}$: Remove the last row from $\tilde{L}$ (i.e., the row containing e_1 and appended to the bottom in operation Λ) to obtain the matrix L_{n-j-1}.

Lemma 7 *The entries of L^* are all non-negative.*

Proof Entries in the columns of K_1 are left unchanged by Φ; hence they are non-negative. Corollary 3 establishes the non-negativity of entries in the other columns. ■

Lemma 8 *The following statements hold:*
(i) $T^*_2 = T^*_1 = T^{n-j}_1$.
(ii) For $i \geq 2$, $T^*_{i+1} - T^*_i = T^{n-j}_{i+1} - T^{n-j}_i$.
(iii) For $i \geq 1$, $\tilde{T}_i = T^*_i$.

Proof Since T^{n-j}_1 is entirely contained in the sub-matrix K_1, its weight remains the same in L^*. Let $\mathcal{S}$ be the set of all matchings of size $n - j$ in L^* that contain exactly one element outside the columns of K_1. From the definition of L^*, it is clear that every matching in $\mathcal{S}$ is lighter by exactly $T^{n-j}_2 - T^{n-j}_1$ compared to its weight in L_{n-j}. From the definition of T^{n-j}_2 we know that every matching in $\mathcal{S}$ had a weight larger than (or equal to) T^{n-j}_2 in the matrix L_{n-j}. Therefore, every matching in $\mathcal{S}$ has a weight larger than (or equal to) $T^{n-j}_2 - (T^{n-j}_2 - T^{n-j}_1)$ in L^*. So every matching that has exactly one element outside the columns of K_1 in L^* has a weight larger than (or equal to) T^{n-j}_1. Therefore, from Lemma 2 it follows that T^{n-j}_1 is the smallest matching in L^*. Thus, we have $T^*_1 = T^{n-j}_1$.

From Corollary 1 we know that T^*_i has exactly one element outside the columns of K_1 for $i \geq 2$. Since every matching in $\mathcal{S}$ is lighter by $T^{n-j}_2 - T^{n-j}_1$ from its weight in L_{n-j}, it follows that $T^*_i = T^{n-j}_i - (T^{n-j}_2 - T^{n-j}_1)$ for $i \geq 2$. Substituting $i = 2$, we obtain $T^*_2 = T^{n-j}_1$. This proves part (i). And considering the differences $T^*_{i+1} - T^*_i$ establishes part (ii).

Since the values of *T-matchings* are invariant under row and column permutations, part (iii) follows from the fact that $\tilde{L}$ is a row permutation of L^*. ■

To complete Step 2 of the induction we need to establish that L_{n-j-1} has the following properties.

Lemma 9 $T_{k+1} - T_k = T^{n-j-1}_{k-j+1} - T^{n-j-1}_{k-j}$, $j + 1 \leq k \leq n - 1$.

[2] The random variable X is used to break the tie between the two matchings T^*_1 and T^*_2, both of which have the same weight (see Lemma 8). This randomized tie-breaking is essential for ensuring that L_{n-j-1} has i.i.d. exp(1) entries; indeed, if we were to choose the entry in T^*_1 (or, for that matter, in T^*_2) with probability 1, then the corresponding L_{n-j-1} will not have i.i.d. exp(1) entries.

Matrix	L_{n-j}	$L^* = L^*_{n-j}$	$\tilde{L} = \tilde{L}_{n-j}$	L_{n-j-1}
Dimensions	$(n-j) \times n$	$(n-j) \times n$	$(n-j) \times n$	$(n-j-1) \times n$
T-matchings	$\{T^{n-j}_i\}$	$\{T^*_i\}$	$\{\tilde{T}_i\}$	$\{T^{n-j-1}_i\}$

Table 1: Operations to transform L_{n-j} to L_{n-j-1}.

Proof The proof of the lemma consists of establishing the following: For $k = j+1, \ldots, n-1$

$$\begin{aligned} T_{k+1} - T_k &\overset{(a)}{=} T^{n-j}_{k-j+2} - T^{n-j}_{k-j+1} \\ &\overset{(b)}{=} \tilde{T}_{k-j+2} - \tilde{T}_{k-j+1} \\ &\overset{(c)}{=} T^{n-j-1}_{k-j+1} - T^{n-j-1}_{k-j}, \end{aligned}$$

where (a) follows from the induction hypothesis on $T^{n-j}_{i+1} - T^{n-j}_{i}$, and (b) follows from Lemma 8. We shall establish (c) by showing that

$$\tilde{T}_i = T^{n-j-1}_{i-1} + v, \;\; i = 2, \ldots, n-j+1 \tag{2}$$

for some appropriately defined constant v.

Two cases arise: For e_1 and e_2 as defined in the operation Λ, Case 1 is when e_1 and e_2 are present in the last row of $\tilde{L}$, and Case 2 is when e_1 is present in the last row of $\tilde{L}$ and e_2 is present in a different row.

Case 1: We claim that the values of e_1 and e_2 are equal, say to v. This is because e_1 and e_2 choose the same matching of size $n-j-1$ from the matrix K_{12} (call this matching $\mathcal{M}$) to form the matchings $\tilde{T}_1$ and $\tilde{T}_2$. But $\tilde{T}_1 = \tilde{T}_2 = T^{n-j}_1$, from Lemma 8. Therefore $e_1 = e_2$.

Consider all matchings of size $n-j-1$ in L_{n-j-1} that have exactly one entry outside the columns defined by K_{12}. Clearly, one (or possibly both) of the entries e_1 and e_2 could have chosen these matchings to form candidates for $\tilde{T}_1$. The fact that the weight of these candidates is larger than $\tilde{T}_1$ indicates that the weight of these size $n-j-1$ matchings are larger than the weight of $\mathcal{M}$. Thus, from Lemma 2, we have that $\mathcal{M}$ equals T^{n-j-1}_1: the smallest matching of size $n-j-1$ in L_{n-j-1}. Therefore, $T^{n-j}_1 = \tilde{T}_1 = \tilde{T}_2 = T^{n-j-1}_1 + v$.

Now consider $\tilde{S}_i$, the smallest matching in $\tilde{L}$ obtained by deleting the i^{th} column in K_{12}. Since this is $\tilde{T}_k$ for some $k \geq 3$, $\tilde{S}_i$ must use one of the entries e_1 or e_2, according to Lemma 6. Therefore, $\tilde{S}_i \geq S^{n-j-1}_i + v$, where S^{n-j-1}_i is the weight of the best matching in L_{n-j-1} that doesn't use the i^{th} column in K_{12}.

However, from Corollary 1 applied to L_{n-j-1}, we have that S^{n-j-1}_i has exactly one element outside K_{12} since K_{12} defines the columns in which the smallest matching of L_{n-j-1} is present. Therefore S^{n-j-1}_i can pick at least one of the two entries e_1 or e_2 to form a candidate for $\tilde{S}^{n-j}_i$ that has weight $S^{n-j-1}_i + v$. This implies $\tilde{S}_i \leq S^{n-j-1}_i + v$.

This shows $\tilde{S}_i = S^{n-j-1}_i + v$.

But, arranging $\tilde{S}_i$, for all $i \in K_{12}$, in increasing order gives us $\tilde{T}_3, \ldots, \tilde{T}_{n-j+1}$. And arranging S^{n-j-1}_i in increasing order gives us $T^{n-j-1}_2, \ldots, T^{n-j-1}_{n-j}$. These observations follow from Corollary 2. This verifies equation (2) and completes Case 1.

Case 2: In this case, only e_1 is present in the last row of $\tilde{L}$. Let v be its value. Let $\mathcal{M}$ denote the matching of size $n-j-1$ in K_{12} that e_1 goes with to form $\tilde{T}_1$. Consider the column, say $\mathbf{c_2}$, in which e_2 is present, and further consider the entry in this column in the last row of $\tilde{L}$ (i.e in the same row as e_1). Denote this entry by e'_2 and let v_1 be its value.

Since e'_2 when combined with the matching $\mathcal{M}$ in K_{12} has a weight larger than $\tilde{T}_1$, it must be that $v_1 = v + x$ for some $x > 0$.

Now, given any $\epsilon > 0$ let $\tilde{L}^\epsilon$ be a matrix identical to $\tilde{L}$ in every entry except e'_2. The value of e'_2 is changed to $v + \epsilon$ from $v_1 = v + x$. Let $\{\tilde{T}^\epsilon_i\}$ denote the *T-matchings* of $\tilde{L}^\epsilon$.

Claim 1 $\tilde{T}^\epsilon_i = \tilde{T}_i$ *for every* i.

Proof It is clearly sufficient to show that e_2' is not used by any of the matchings $\tilde{T}_i^\epsilon$ or $\tilde{T}_i$. Since $\tilde{T}_2$ uses an entry in the last row that is different from e_1 and e_2' (it doesn't use e_1 since $\tilde{T}_2$ cannot use the column in which e_1 is present; it cannot use e_2' since it uses e_2 in the same column), it follows from Lemma 6 that e_2' cannot be present in any of the $\tilde{T}_i$.

To see that e_2' cannot be present in any of the $\tilde{T}_i^\epsilon$, consider the following. Compute $\tilde{S}_{\mathbf{c_1}}^\epsilon$, the smallest matching in $\tilde{L}^\epsilon$ that avoids column $\mathbf{c_1}$, the column in which e_1 is present. We shall now show that $\tilde{S}_{\mathbf{c_1}}^\epsilon$ uses e_2 and has weight $\tilde{T}_1$ ($= \tilde{T}_2$, by Lemma 8).

Consider the set, F, of all matchings in $\tilde{L}^\epsilon$ that do not use $\mathbf{c_1}$ or e_2'. Since all the entries that give rise to matchings in F have the same weight in $\tilde{L}$ and $\tilde{L}^\epsilon$, it follows that the best matching in F uses e_2 and has weight $\tilde{T}_2 = \tilde{T}_1$. Now consider the set F' of all matchings in $\tilde{L}^\epsilon$ that do not use $\mathbf{c_1}$ but *necessarily* use e_2'. Write a matching in F' as $e_2' + \mathcal{M}'$, where $\mathcal{M}'$ is a member of the set, G, of all matchings of size $n-j-1$ that do not use $\mathbf{c_1}$, $\mathbf{c_2}$ and the row containing e_2' (the last row). We know that $\tilde{T}_1 = e_1 + \mathcal{M}$. From the minimality of $\tilde{T}_1$, it follows that $\mathcal{M}$ is the lightest matching in the set, G', of all matchings of size $n-j-1$ that do not use $\mathbf{c_1}$ and the row containing e_1 (which also contains e_2'). Since $G \subset G'$, this implies that $\mathcal{M}$ is the lightest matching in G, and $e_2' + \mathcal{M}$ is the lightest matching in F'. (Note that, $\mathcal{M}$ uses the columns K_{12} and hence is in G.) Therefore, the weight of $e_2' + \mathcal{M}$ equals $\epsilon + \tilde{T}_1 > \tilde{T}_1 = \tilde{T}_2$; thus, it is heavier than the lightest matching in F. Combining these, we have $\tilde{S}_{\mathbf{c_1}}^\epsilon$ uses e_2 and has weight $\tilde{T}_1$.

Now, $\tilde{S}_{\mathbf{c_1}}^\epsilon$ must use an entry, say e', in the last row that is not e_1 (it avoids the column $\mathbf{c_1}$) or e_2' (it uses e_2). Therefore by Lemma 6, it follows that every matching $\tilde{T}_i^\epsilon$ uses one of the entries e_1 or e'. This implies e_2' cannot be present in any of the $\tilde{T}_i^\epsilon$, and the proof of the claim is complete. ■

Define a matrix $\tilde{L}^{-\epsilon}$ which is identical to the matrix $\tilde{L}^\epsilon$ except for the entries e_1 and e_2'. We change the values of e_1 and e_2' to $v - \epsilon$. Let the *T-matchings* of $\tilde{L}^{-\epsilon}$ be denoted by $\{\tilde{T}_i^{-\epsilon}\}$.

Compute $\tilde{S}_{\mathbf{c_1}}^{-\epsilon}$, the smallest matching in $\tilde{L}^\epsilon$ that avoids column $\mathbf{c_1}$. An argument similar to the one before gives us the following: the smallest matching in F uses e_2 and has weight $\tilde{T}_1$; the smallest matching in F' uses $e_2' + \mathcal{M}$ and has weight $\tilde{T}_1 - \epsilon$. Combining these, we deduce that $\tilde{S}_{\mathbf{c_1}}^{-\epsilon}$ uses e_2' and has weight $\tilde{T}_1 - \epsilon$.

Now compute $\tilde{S}_{\mathbf{c_2}}^{-\epsilon}$, the smallest matching in $\tilde{L}^{-\epsilon}$ which avoids the column $\mathbf{c_2}$. It is easy to see that the $\tilde{S}_{\mathbf{c_2}}^{-\epsilon}$ uses e_1 and $\mathcal{M}$ and has weight $\tilde{T}_1 - \epsilon$.

Hence from Lemma 6, every $\tilde{T}_i^{-\epsilon}$ must use one of e_1 or e_2'. An argument similar to Case 1 now yields the following:

$$\tilde{T}_1^{-\epsilon} = \tilde{T}_2^{-\epsilon} = T_1^{n-j-1} + v - \epsilon, \text{ and for all } i \geq 2,\ \tilde{T}_i^{-\epsilon} = T_i^{n-j-1} + v - \epsilon. \tag{3}$$

(Recall the argument: First show that the equation holds for the $\tilde{S}_i^{-\epsilon}$ and then arrange these in increasing order to get the $\tilde{T}_i^{-\epsilon}$. To complete, note that the first $n-j-1$ rows of $\tilde{L}^{-\epsilon}$ have entries identical in value to those of $\tilde{L}$.)

As ϵ tends to zero, we have that entry by entry the matrix $\tilde{L}^{-\epsilon}$ tends to $\tilde{L}^\epsilon$. Hence the weights $\{\tilde{T}_i^{-\epsilon}\}$ and $\{\tilde{T}_i^\epsilon\}$ approach one another. Letting ϵ tending to zero in Equation (3) and using Claim 1 we obtain the following:

$$T_i^{n-j-1} + v = \lim_{\epsilon \to 0} \tilde{T}_i^{-\epsilon} = \lim_{\epsilon \to 0} \tilde{T}_i^\epsilon = \tilde{T}_i.$$

This verifies Equation (2) in Case 2 and completes the proof of the Lemma 9. ■

Corollary 4 *Let $\mathcal{M}$ be the smallest matching of size $n-j-1$ in $\tilde{L}$, contained in the columns of K_{12}, that e_1 goes with to form $\tilde{T}_1 = T_1^{n-j}$. Then $\mathcal{M} = T_1^{n-j-1}$, the smallest matching of size $n-j-1$ in L_{n-j-1}.*

Proof Letting v be the weight of e_1, we note that $\tilde{T}_1$ is formed by e_1 and $\mathcal{M}$. From Equation (2), $\tilde{T}_1$ has a weight equal to $v + T_1^{n-j-1}$. Hence the weight of $\mathcal{M}$ equals T_1^{n-j-1}. ■

Step 3: L_{n-j-1} has i.i.d. $\exp(1)$ entries

We compute the joint distribution of the entries of L_{n-j-1} and verify that they are i.i.d. exp(1) variables. To do this, we identify the set, $\mathcal{D}$, of all $(n-j)\times n$ matrices, L_{n-j}, that have a non-zero probability of mapping to a particular realization of L_{n-j-1} under the operations Φ, Λ and Π. By induction, the entries of L_{n-j} are i.i.d exp(1) random variables. We integrate L_{n-j} over $\mathcal{D}$ to obtain the joint distribution of the entries of L_{n-j-1}.

Accordingly, fix a realization of $L_{n-j-1} \in \mathbb{R}_+^{(n-j-1)\times n}$ and represent it as below

$$L_{n-j-1} = \begin{bmatrix} l_{1,1} & l_{1,2} & \cdot\ \cdot\ \cdot\ \cdot & l_{1,n-2} & l_{1,n-1} & l_{1,n} \\ \cdot & \cdot & \cdot\ \cdot\ \cdot\ \cdot\ \cdot & & \cdot & \cdot \\ \cdot & \cdot & \cdot\ \cdot\ \cdot\ \cdot\ \cdot & & \cdot & \cdot \\ l_{n-j-1,1} & l_{n-j-1,2} & \cdot\ \cdot\ \cdot\ \cdot & l_{n-j-1,n-2} & l_{n-j-1,n-1} & l_{n-j-1,n} \end{bmatrix}. \tag{4}$$

Let $\mathcal{D}_\Pi = \Pi^{-1}(L_{n-j-1})$ be the pre-image of L_{n-j-1} under Π, and let $\mathcal{D}_\Lambda = \Lambda^{-1}\Pi^{-1}(L_{n-j-1})$. Now Λ is a random map, whose action depends on the value taken by X. In turn, this is related to whether e_1 and e_2 are on the same row or not. Therefore we may write $\mathcal{D}_\Lambda$ as the disjoint union of the sets $\mathcal{D}_\Lambda^s$ and $\mathcal{D}_\Lambda^d$, which respectively correspond to e_1 and e_2 belonging to the same and to different rows. Finally, $\mathcal{D} = \Phi^{-1}\Lambda^{-1}\Pi^{-1}(L_{n-j-1})$, the set we're trying to determine.

Consider an $M' \in \mathbb{R}_+^{(n-j)\times n}$, such as represented below

$$M' = \begin{bmatrix} & & & L_{n-j-1} & & & \\ r_1 & r_2 & \cdot\ \cdot\ \cdot & r_{n-j-1} & x_1 & \cdot\ \cdot & x_{j+1} \end{bmatrix}.$$

In the above representation $r_1, ..., r_{n-j-1}$ denote the elements in the columns of K_1^{n-j-1}, the sub-matrix of L_{n-j-1} which contains T_1^{n-j-1}.

We shall show that $\mathcal{D}_\Pi$ is precisely the set of all such M' for which the vector $(r_1, ..., r_{n-j-1}, x_1, ..., x_{j+1})$ satisfies the conditions stated in Lemma 10.

Consider an element d outside the columns of K_1^{n-j-1} in L_{n-j-1}. Let Δ_d be the cost of the smallest matching, say $\mathcal{M}_d$, of size $n-j-1$ with the following property: The entries of $\mathcal{M}_d$ are in $K_1^{n-j-1} \cup (r_1, .., r_{n-j-1})$ and no entry is present in the same row as d. Clearly $d \cup \mathcal{M}_d$ is a matching of size $n-j$ in the matrix M'. Let d_o be that entry outside the columns of K_1^{n-j-1} in L_{n-j-1} which minimizes $d + \Delta_d$. Let $J = d_o + \Delta_{d_o}$, and let j_o denote the column in which d_o occurs.

Given any $\vec{r} = (r_1, ..., r_{n-j-1}) \in \mathbb{R}_+^{n-j-1}$, the following lemma stipulates conditions that $(x_1, ..., x_{j+1})$ must satisfy so that M' is in $\mathcal{D}_\Pi$.

Lemma 10 *For any $\vec{r} \in \mathbb{R}_+^{n-j-1}$, let $\mathcal{S}_\Pi(\vec{r})$ be the collection of all M' such that one of the following two conditions hold:*

- *(i) There exist i and k such that $x_i = x_k$, $x_i + T_1^{n-j-1} < J$ and $x_m > x_i$ for all $m \neq i, k$.*
- *(ii) There exists $i \neq j_o$ such that $x_m > x_i$ for all $m \neq i$ and $x_i + T_1^{n-j-1} = J$.*

Then $\mathcal{D}_\Pi = \mathcal{S}_\Pi \stackrel{\triangle}{=} \bigcup_{\vec{r}\in\mathbb{R}_+^{n-j-1}} \mathcal{S}_\Pi(\vec{r})$.

Proof (α) $\mathcal{D}_\Pi \subset \mathcal{S}_\Pi$: Let $M \in \mathcal{D}_\Pi$ be any matrix such that $\Pi(M) = L_{n-j-1}$. Therefore the sub-matrix consisting of the first $n-j-1$ rows must be L_{n-j-1}. Let $r_1, .., r_{n-j-1}$ be the entries in the last row of M in the columns of K_1^{n-j-1}. We shall show that $M \in \mathcal{S}_\Pi(\vec{r})$.

By the definition of Λ we know that the entry e_1 occurs in the last row. From Corollary 4 we know that e_1 chooses the matching T_1^{n-j-1} to form the matching $\tilde{T}_X$. Hence e_1 must be one of the x_i's. By definition, e_2 lies outside the columns of K_{12}. From Corollary 4 we know that the columns of K_{12} are the same as the columns of K_1^{n-j-1}. We also know, from Lemma 8 and the fact that X and Y only take the values 1 or 2, that $\tilde{T}_X = \tilde{T}_Y = \tilde{T}_1 = \tilde{T}_2$.

Two cases occur.

(a): e_1 and e_2 are in the last row. In this case, we know from the proof of Lemma 9 that $e_1 = e_2$. (This is Case 1 in the proof of Lemma 9 and there we showed that $e_1 = e_2 = v$.) Since both e_1 and e_2 are in the last row and outside columns of K_1^{n-j-1}, we know that $e_1 = x_i$ and $e_2 = x_k$ for some $i \neq k$. Therefore, $x_i = x_k$. We know that $e_1 + T_1^{n-j-1} = \tilde{T}_X$, hence from the minimality of $\tilde{T}_X$ we have $x_i + T_1^{n-j-1} < J$. Also, $x_i + T_1^{n-j-1} < x_m + T_1^{n-j-1}$ for $m \neq i, k$ for the same reason. This implies M satisfies the first condition of Lemma 10. Therefore, under (a) it follows that $M \in \mathcal{S}_\Pi(\vec{r})$.

(b): e_1 is in the last row and e_2 is not. In this case, from the minimality of $d_o + \Delta_{d_o}$, it follows that $e_2 + \Delta_{e_2} \geq d_o + \Delta_{d_o} = J_1$. But we also know that e_2 together with a $n-j-1$ size matching from the columns of K_{12} (which are the same as the columns of K_1^{n-j-1}) forms $\tilde{T}_Y$. Therefore, $e_2 + \Delta_{e_2} = \tilde{T}_Y = \tilde{T}_1$. The minimality of $\tilde{T}_1$ now implies $e_2 + \Delta_{e_2} \leq d_o + \Delta_{d_o} = J$. Therefore we conclude that e_2 must in fact be d_o. (Combinatorially, $e_2 + \Delta_{e_2}$ and $d_o + \Delta_{d_o}$ could just as well be two different matchings with the same weight. But we ignore this zero probability event.)

Thus, $\tilde{T}_1 = e_2 + \Delta_{e_2} = J$. We know that e_1 is one of the x's and e_1 and e_2 occur in different columns, hence $e_1 = x_i$ for some $i \neq j_o$. Thus, $x_i + T_1^{n-j-1} = \tilde{T}_1$ and from the minimality of $\tilde{T}_1$, we also have that $x_i + T_1^{n-j-1} > x_m + T_1^{n-j-1}$ for $m \neq i$. Thus, M satisfies the second condition of Lemma 10 and hence $M \in \mathcal{S}_\Pi(\vec{r})$.

(β) $\mathcal{S}_\Pi \subset \mathcal{D}_\Pi$: Let $M \in \mathcal{S}_\Pi(\vec{r})$ for some $\vec{r}$. Clearly $\Pi(M) = L_{n-j-1}$ since Π operates by removing the last row of the matrix. Therefore, we just need to establish the existence of an L_{n-j} that has a non-zero probability of mapping to the matrix M. Let $\{T_i^M\}$ denote the *T-matchings* of M.

The matrix M satisfies one of Conditions (i) or (ii) of Lemma 10. Accordingly, this gives rise to two cases.

(a): M satisfies Condition (i). We claim that $\Lambda(\Phi(M)) = M$. From Lemma 3 we have that T_1^M must use all the columns of T_1^{n-j-1}. This implies that exactly one entry of T_1^M lies outside the columns of K_1^{n-j-1}. But, Condition (i) has it that $x_i + T_1^{n-j-1} \leq \min\{x_m + T_1^{n-j-1}, J\} = \min\{x_m + T_1^{n-j-1}, d + \Delta_d\}$. Since the last minimization is over all possible choices of this lone entry, d, that T_1^M chooses outside the columns of K_1^{n-j-1}, it follows that $T_1^M = x_i + T_1^{n-j-1}$. Condition (i) also has it that $x_k = x_i$. Hence $T_1^M = T_2^M = x_k + T_1^{n-j-1}$.

This implies $T_2^M - T_1^M = 0$ and hence $\Phi(M) = M$. Since x_i and x_k are the entries of T_1^M and T_2^M outside the columns of K_{12}, this implies e_1 and e_2 are x_i and x_k in some order. Observe that Λ removes the row in which e_1 is present and appends it to the bottom of the matrix. However, in this case, as e_1 is already present in the last row, we obtain $\Lambda(M) = M$. Thus $\Lambda\Phi(M) = M$ and this shows that M is in the range of $\Lambda\Phi(\cdot)$.

(b): M satisfies Condition (ii). Again starting with $L_{n-j} = M$, we claim that $\Lambda\Phi(M) = M$ with probability $\frac{1}{2}$.

A similar argument as in Case (a) yields $x_i + T_1^{n-j-1} = T_1^M = T_2^M = J = d_o + \Delta_{d_o}$. Hence $T_2^M - T_1^M = 0$ and we get that $\Phi(M) = M$. Note that e_1 and e_2 are decided by the outcome of X. Hence $\mathbb{P}(e_1 = x_i, e_2 = d_o) = \frac{1}{2} = \mathbb{P}(e_2 = x_i, e_1 = d_o)$. When $e_1 = x_i$, by the definition of Λ we get that $\Lambda(M) = M$. When $e_1 = d_o$ the row that moves to the bottom is the one in which d_o occurs; hence $\Lambda(M) \neq M$ in this case. Therefore, with probability $\frac{1}{2}$ we will obtain M as the result of the operation $\Lambda\Phi(M)$. This implies M is in the range of $\Lambda\Phi(\cdot)$.

Taken together, both cases imply that if $M \in \mathcal{S}_\Pi(\vec{r})$ then $M \in \mathcal{D}_\Pi$ for any $\vec{r} \in \mathbb{R}_+^{n-j-1}$. Therefore, $\mathcal{S}_\Pi \subset \mathcal{D}_\Pi$. ■

Let $\mathcal{D}^s_\Pi$ and $\mathcal{D}^d_\Pi$ be the subsets of matrices in $\mathcal{D}_\Pi$ that satisfy Conditions (i) and (ii) of Lemma 10, respectively. Clearly $\mathcal{D}_\Pi$ is the disjoint union of $\mathcal{D}^s_\Pi$ and $\mathcal{D}^d_\Pi$. (The superscripts s and d are mnemonic for whether e_1 and e_2 occurred in the same row or in different rows.)

Now that we have identified $\mathcal{D}_\Pi$ explicitly in Lemma 10, $\mathcal{D}_\Lambda$ can be identified following manner: Pick a matrix $M_1 \in \mathcal{D}_\Pi$ and form $n-j$ matrices by removing the last row of M_1 and placing it back as the ith row, for $i = 1, ..., n-j$. Call this collection of $n-j$ matrices $\mathcal{S}_\Lambda(M_1)$. Define

$$\mathcal{S}_\Lambda = \bigcup_{M_1 \in \mathcal{D}_\Pi} \mathcal{S}_\Lambda(M_1).$$

Lemma 11 $\mathcal{S}_\Lambda = \mathcal{D}_\Lambda$.

Proof Let $M \in \mathcal{D}_\Lambda$, then we have $\Lambda(M) = M_1$ for some $M_1 \in \mathcal{D}_\Pi$. The map Λ operates on M by removing a row and appending it to the bottom. Hence if $\Lambda(M) = M_1$ then $M \in \mathcal{S}_\Lambda(M_1) \subset \mathcal{S}_\Lambda$. Therefore $\mathcal{D}_\Lambda \subset \mathcal{S}_\Lambda$.

Conversely, consider an $M \in \mathcal{S}_\Lambda$. We will show that $M \in \mathcal{D}_\Lambda$ by establishing two things: M belongs to the range of Φ (hence we may apply Λ to it), and $\Pi\Lambda(M) = L_{n-j-1}$ with a non-zero probability.

Consider an $M \in \mathcal{S}_\Lambda(M_1)$ for some $M_1 \in \mathcal{D}_\Pi$. Now, from Lemma 8 we know that $T_1^{M_1} = T_2^{M_1}$ for any $M_1 \in \mathcal{D}_\Pi$. Since row permutations do not change the values of *T-matchings*, it follows that $T_1^M = T_2^M$. This implies $\Phi(M) = M$, or M is in the range of Φ.

To complete the proof, we again consider any $M \in \mathcal{S}_\Lambda$ and a corresponding $M_1 \in \mathcal{D}_\Pi$ such that $M \in \mathcal{S}_\Lambda(M_1)$. As before, there are two possibilities: (a) $M_1 \in \mathcal{D}^s_\Pi$ or (b) $M_1 \in \mathcal{D}^d_\Pi$.

Under (a), e_1 and e_2 are in the same row of M_1. Since M and M_1 are row permutations of one another, e_1 and e_2 are also in the same row of M. Therefore, Λ will map M back into M_1 with probability 1, by simply moving this common row to the bottom.

By contrast, under (b), e_1 and e_2 are on different rows of M_1 and hence of M. Thus, this time around Λ will map M back into M_1 with probability of $\frac{1}{2}$. This is because Λ has a choice of two rows (depending on the outcome of X) that could be moved down with equal probability. But only one of these choices will yield M_1.

This shows $\mathcal{S}_\Lambda \subset \mathcal{D}_\Lambda$ and completes the proof of Lemma 11. ∎

Partition $\mathcal{D}_\Lambda$ into $\mathcal{D}^s_\Lambda = \mathcal{S}_\Lambda(\mathcal{D}^s_\Pi)$ and $\mathcal{D}^d_\Lambda = \mathcal{S}_\Lambda(\mathcal{D}^d_\Pi)$; these partitions correspond to the cases where e_1 and e_2 are in the same row and in different rows. Also observe that when $M \in \mathcal{D}^s_\Lambda$ we have $\Pi\Lambda(M) = L_{n-j-1}$ with probability one and when $M \in \mathcal{D}^d_\Lambda$ we have $\Pi\Lambda(M) = L_{n-j-1}$ with probability $\frac{1}{2}$.

We are finally ready to characterize $\mathcal{D}$, the set of L_{n-j}'s that map to a particular realization of L_{n-j-1} with non-zero probability.

Consider any $M \in \mathcal{D}_\Lambda$. Let θ be any positive number. Consider the column, say $\mathbf{c}_1$ in M which contains x_i. (Recall, from Lemma 10, that x_i is the smallest of the x_m's in the last row deleted by Π.) Let K be the union of the columns of K_1^{n-j-1} and $\mathbf{c}_1$. Add θ to every entry in M outside the columns K. Denote the resultant matrix by $S_1(\theta, M)$. Let

$$\mathcal{S}_1 = \bigcup_{\theta>0, M \in \mathcal{D}_\Lambda} S_1(\theta, M).$$

Now consider the column, say $\mathbf{c}_2$ in M where the entry x_k or d_o (depending on whether $\Lambda(M)$ satisfies condition (i) or condition (ii) of Lemma 10) is present. Let K' be the union of columns of K_1^{n-j-1} and $\mathbf{c}_2$. Now add θ to every entry in M outside the columns K'. Call this matrix $S_2(\theta, M)$. Let

$$\mathcal{S}_2 = \bigcup_{\theta>0, M \in \mathcal{D}_\Lambda} S_2(\theta, M),$$

and note that $\mathcal{S}_1$ and $\mathcal{S}_2$ are disjoint since $\mathbf{c}_1 \neq \mathbf{c}_2$.
Remark: Note that θ is added to precisely $j(n-j)$ entries in M in each of the two cases above.

Lemma 12 *$\mathcal{D} = \mathcal{S}_1 \cup \mathcal{S}_2$.*

Proof Consider $M \in \mathcal{D}$. From Lemmas 10 and 11 we know that $\Phi(M) \in \mathcal{D}_\Lambda$. Let $T_2^M - T_1^M = \theta$. Let K_1 denote the columns that contain T_1^M. Subtracting θ from the entries of M outside the columns of K_1 leaves us with $M_1 \in \mathcal{D}_\Lambda$. From the proof of Lemma 8 we know that under Φ the locations of the entries of *T-matchings* do not change; only the weights of $T_i^M,\ i \geq 2$ are reduced by $T_2^M - T_1^M = \theta$. Therefore, in the matrix M, exactly one of x_i or y (where y is x_k or d_o depending on which condition of Lemma 10 $\Lambda(M_1)$ satisfies) was part of T_1^M. If it is x_i, then we see that $M = S_1(\theta, M_1)$. If it is y, then $M = S_2(\theta, M_1)$. This implies $M \in \mathcal{S}_1 \cup \mathcal{S}_2$.

For the converse, consider an $M = S_1(\theta, M_1)$ for some $M_1 \in \mathcal{D}_\Lambda$ and $\theta > 0$. Since $x_i + T_1^{n-j-1}$ was the smallest matching in M_1, and M dominates M_1 entry-by-entry, $x_i + T_1^{n-j-1}$ continues to be the smallest matching in M. Consider every matching that contains exactly one element outside the columns defined by the matching $x_i \cup T_1^{n-j-1}$. By construction, the weight of these matchings in M exceeds the weight of the same matchings in M_1 by an amount precisely equal to θ. Hence, by Corollary 1, we have that $T_i^M - T_i^{M_1} = \theta$ for $i \geq 2$. This further implies $T_2^M - T_1^M = T_2^{M_1} - T_1^{M_1} + \theta$.

But we know that for any $M_1 \in \mathcal{D}_\Lambda$, $T_2^{M_1} = T_1^{M_1}$. Therefore $T_2^M - T_1^M = \theta$, and $\Phi(M)$ is the matrix that results from subtracting θ from each entry outside the columns containing the matching $x_i \cup T_1^{n-j-1}$. But, by the definition of $S_1(\theta, M_1)$, $\Phi(M)$ is none other than the matrix M_1. Therefore $M \in \mathcal{D}$, and $\mathcal{S}_1 \subset \mathcal{D}$.

Now consider an $M = S_2(\theta, M_1)$. In this case $x_k + T_1^{n-j-1}$ (or $d_o + \Delta_{d_o}$) continues to remain the smallest matching in M. An argument identical to the one above establishes that $\Phi(M) = M_1$. Hence, $M \in \mathcal{D}$ and $\mathcal{S}_2 \subset \mathcal{D}$ and the proof of the lemma is complete. ∎

Remark: Note that the variable θ used in the characterization of $\mathcal{D}$ precisely equals the value of $T_2^M - T_1^M$, as shown in the proof of Lemma 12.

Continuing, we can partition $\mathcal{D}$ into the two sets $\mathcal{D}^s$ and $\mathcal{D}^d$ as below:

$$\mathcal{D}^s = S_1(\mathbb{R}_+, \mathcal{D}_\Lambda^s) \cup S_2(\mathbb{R}_+, \mathcal{D}_\Lambda^s) \quad \text{and} \quad \mathcal{D}^d = S_1(\mathbb{R}_+, \mathcal{D}_\Lambda^d) \cup S_2(\mathbb{R}_+, \mathcal{D}_\Lambda^d). \tag{5}$$

Observe that whenever $M \in \mathcal{D}^s$, we have $\Phi(M) \in \mathcal{D}_\Lambda^s$ and hence $\Pi\Lambda\Phi(M) = L_{n-j-1}$ with probability 1. For $M \in \mathcal{D}^d$, $\Phi(M) \in \mathcal{D}_\Lambda^d$. Hence $\Pi\Lambda\Phi(M) = L_{n-j-1}$ with probability $\frac{1}{2}$.

Now that we have characterized $\mathcal{D}$, we "integrate out the marginals" $(r_1, ..., r_{n-j-1}, x_1, ..., x_{j+1})$ and θ by setting

$$\vec{v} = (L_{n-j-1}, \vec{r}, \theta) \quad \text{and} \quad \vec{w} = (\vec{v}, \vec{x}),$$

where $L_{n-j-1} \in \mathbb{R}_+^{(n-j-1)\times n}$ is as defined at equation (4). We will evaluate $f_v(\vec{v}) = \int_{\mathcal{R}_1} f_w(\vec{v}, \vec{x}) d\vec{x} + \int_{\mathcal{R}_2} f_w(\vec{v}, \vec{x}) d\vec{x}$, to obtain the marginal density of $\vec{v}$. The regions $\mathcal{R}_1$ and $\mathcal{R}_2$ are defined by the set of all $\vec{x}$s that satisfy (i) and (ii) of Lemma 10, respectively.

On $\mathcal{R}_1$, we have that $x_i = x_k < J - T_1^{n-j-1}$ for J as in Lemma 10. We shall set $U = J - T_1^{n-j-1}$, and $u_m = x_m - x_i$ for $m \neq i, k$. Finally, define

$$s_v = l_{1,1} + ... + l_{n-j-1,n} + r_1 + ... + r_{n-j-1} + j(n-j)\theta.$$

Thus, s_v denotes the sum of all of the entries of L_{n-j} except those in $\vec{x}$. As noted in the remark preceding Lemma 12, the value θ was added to precisely $j(n-j)$ entries. We have

$$\begin{aligned}
\int_{\mathcal{R}_1} f_w(\vec{v}, \vec{x}) d\vec{x} &\overset{(a)}{=} 2(n-j)\binom{j+1}{2} \int_0^U \int_0^\infty ... \int_0^\infty e^{-(s_v + (j+1)x_i + u_1 + ... + u_{j-1})}\, du_1 \ldots du_{j-1}\, dx_i \\
&= j(n-j)e^{-s_v}\left(1 - e^{-(j+1)U}\right).
\end{aligned} \tag{6}$$

The factor $\binom{j+1}{2}$ at equality (a) comes from the possible choices for i and k from 1,...,$j+1$, the factor $(n-j)$ comes from the row choices available to e_1 as in Lemma 11, and the factor 2 corresponds to the partition, $\mathcal{S}_1$ or $\mathcal{S}_2$, that L_{n-j} belongs to.

Similarly, on $\mathcal{R}_2$, we have that $x_i = J - T_1^{n-j-1} \stackrel{\triangle}{=} U$ and we shall set $u_m = x_m - x_i$ for $m \neq i$ to obtain

$$\begin{aligned}\int_{\mathcal{R}_2} f_w(\vec{v},\vec{x})d\vec{x} &\stackrel{(b)}{=} \frac{1}{2}\left[2(n-j)j\int_0^\infty ... \int_0^\infty e^{-(s_v+(j+1)U+u_1+...+u_j)}\, du_1 \, ... \, du_j\right] \\ &= j(n-j)e^{-s_v}e^{-(j+1)U}. \end{aligned} \tag{7}$$

In equality (b) above, the factors j, $(n-j)$ and 2 come, respectively, from the choice of positions available to x_i,[3] the row choices available to e_1 and the partition, $\mathcal{S}_1$ or $\mathcal{S}_2$, that L_{n-j} belongs to. The factor $\frac{1}{2}$ comes from the fact that on $\mathcal{R}_2$, e_1 and e_2 occur on different rows. Therefore, L_{n-j} is in $\mathcal{D}^d$ and will map to the desired L_{n-j-1} with probability $\frac{1}{2}$.

Putting (6) and (7) together, we obtain

$$f_v(\vec{v}) = j(n-j)e^{-s_v} = e^{-(l_{1,1}+...+l_{n-j-1,n})} \times j(n-j)e^{-j(n-j)\theta} \times e^{-(r_1+...+r_{n-j-1})}.$$

We summarize the above in the following lemma.

Lemma 13 *The following hold:*
(i) L_{n-j-1} consists of i.i.d. $\exp(1)$ *variables.*
(ii) $\theta = T_2^{n-j} - T_1^{n-j}$ is an $\exp(j(n-j))$ *random variable.*
(iii) $\vec{r}$ consists of i.i.d. $\exp(1)$ *variables.*
(iv) L_{n-j-1}, $(T_2^{n-j} - T_1^{n-j})$, and $\vec{r}$ are independent.

Remark: It is worth noting that the above lemma provides an alternate proof of Theorem 1. That is, it shows that $T_2^{n-j} - T_1^{n-j} = T_{j+1} - T_j$ is an $\exp(j(n-j))$ random variable.

From Lemma 9 we know that the increments $\{T_{k+1} - T_k, k > j\}$ are a function of the entries of L_{n-j-1}. Given this and the independence of L_{n-j-1} and $T_{j+1} - T_j$ from the above lemma, we get the following corollary.

Corollary 5 *$T_{j+1} - T_j$ is independent of $T_{k+1} - T_k$ for $k > j$.*

In conjunction with Theorem 1, Corollary 5 completes the proof of Conjecture 1. The next section recalls the argument (from [SP 02]) that establishing Conjecture 1 leads to the resolution of Parisi's conjecture.

5 Resolution of Parisi's conjecture

Let C be a matrix of i.i.d exp(1) entries of dimension $n \times n$. Denote by L the sub-matrix obtained by removing the first row of C. Recall that for $i = 1, \ldots, n$, S_i is the cost of the minimum-cost permutation in the sub-matrix of L obtained by deleting it's i^{th} column. As before, denote the ordered sequence of the S_i's by $\{T_i\}_1^n$.

[3]Note that there are only j choices available to x_i since it has to occur in a column other than the one in which d_o occurs.

Now consider $\mathbb{E}(A_n)$, abbreviated to E_n.

$$\begin{aligned} E_n &= \int_0^\infty P(A_n > x)\, dx \\ &= \int_0^\infty P(A_n(\pi) > x, \forall \pi)\, dx \\ &= \int_0^\infty P(c_{1j} > x - S_j, \forall j)\, dx \\ &= \mathbb{E}_L\left(\int_0^\infty P(c_{1j} > x - s_j, \forall j)\, dx \mid L\right) \end{aligned}$$

where $\mathbb{E}_L(.|L)$ denotes the conditional expectation with respect to the matrix L. Note that conditioned on L, the S_i's are constant and are therefore denoted by s_i. Next, consider

$$\begin{aligned} I &= \int_0^\infty P(c_{1j} > x - s_j, \forall j)\, dx \\ &= \int_0^\infty \prod_{j=1}^n P(c_{1j} > x - s_j)\, dx \qquad \text{(independence of } c_{1j}\text{)} \\ &= \int_0^{t_1} dx + \int_{t_1}^{t_2} e^{-(x-t_1)}\, dx + \ldots + \int_{t_{n-1}}^{t_n} e^{-((n-1)x - t_1 - \cdots - t_{n-1})}\, dx + \int_{t_n}^\infty e^{-(nx - t_1 - \cdots - t_n)}\, dx \\ &\quad \text{(where the } t_i \text{ are obtained by ordering the } s_i\text{)} \\ &= t_1 + \left(1 - e^{-(t_2 - t_1)}\right) + \frac{1}{2}\left(e^{-(t_2-t_1)} - e^{-(2t_3 - t_1 - t_2)}\right) + \ldots + \\ &\quad \frac{1}{n-1}\left(e^{-((n-2)t_{n-1} - t_1 - \cdots - t_{n-2})} - e^{-((n-1)t_n - t_1 - \cdots - t_{n-1})}\right) + \frac{1}{n} e^{-((n-1)t_n - t_1 - \cdots - t_{n-1})} \\ &= t_1 + 1 - \frac{1}{2} e^{-(t_2 - t_1)} - \frac{1}{6} e^{-(2t_3 - t_1 - t_2)} - \cdots - \frac{1}{n(n-1)} e^{-((n-1)t_n - t_1 - \cdots - t_{n-1})}. \end{aligned}$$

Therefore,

$$\begin{aligned} E_n &= \mathbb{E}(T_1) + 1 - \sum_{i=2}^n \frac{1}{i(i-1)} \mathbb{E}\left(e^{-((i-1)T_i - T_1 - \cdots - T_{i-1})}\right) \\ &= \mathbb{E}(T_1) + 1 - \sum_{i=2}^n \frac{1}{i(i-1)} \mathbb{E}\left(e^{\sum_{j=1}^{i-1} -j(T_{j+1} - T_j)}\right). \end{aligned} \tag{8}$$

We proceed by using the following from Section 4:

Theorem 2 *For* $j = 1, \ldots, n-1$, $\quad T_{j+1} - T_j \sim \exp(j(n-j))$ *and these increments are independent of each other.*

Using this theorem we obtain

$$\mathbb{E}\left(e^{\sum_{j=1}^{i-1} -j(T_{j+1} - T_j)}\right) = \prod_{j=1}^{i-1} \mathbb{E}\left(e^{-j(T_{j+1} - T_j)}\right) = \prod_{j=1}^{i-1} \frac{n-j}{n-j+1} = \frac{n-i+1}{n}.$$

Substituting this in (8) gives

$$E_n = \mathbb{E}(T_1) + \frac{1}{n^2} + \frac{1}{n} \sum_{i=1}^{n-1} \frac{1}{i}. \tag{9}$$

We are left with having to evaluate $\mathbb{E}(T_1)$. It turns out that Theorem 2 proves useful in doing this as well.

First, for $j = 2, \ldots, n$,

$$\begin{aligned}
\mathbb{E}(T_j) &= \mathbb{E}(T_1) + \sum_{k=2}^{j} \mathbb{E}(T_k - T_{k-1}) \\
&= \mathbb{E}(T_1) + \sum_{k=2}^{j} \frac{1}{(k-1)(n-k+1)} \quad \text{(by Theorem 2).}
\end{aligned}$$

Moreover, the random variables $S_1, \ldots, S_n$ are all distributed as A_{n-1}. Therefore, $\mathbb{E}(S_i) = \mathbb{E}_{n-1}, \; \forall\, i$. Further, since S_1 is one of the T_i's chosen uniformly at random, it is straightforward to note that

$$\begin{aligned}
\mathbb{E}(S_1) &= \frac{1}{n} \sum_{j=1}^{n} \mathbb{E}(T_j) \\
&= \frac{1}{n} \sum_{j=1}^{n} \left(\mathbb{E}(T_1) + \sum_{k=2}^{j} \frac{1}{(k-1)(n-k+1)} \right) \\
&= \mathbb{E}(T_1) + \frac{1}{n} \sum_{j=1}^{n-1} \frac{1}{j}. \qquad (10)
\end{aligned}$$

Assuming, by induction, that $E_{n-1} = \sum_{j=1}^{n-1} \frac{1}{j^2} = \mathbb{E}(S_1)$, we obtain from (10) that

$$\mathbb{E}(T_1) = \sum_{j=1}^{n-1} \left(\frac{1}{j^2} - \frac{1}{nj} \right). \qquad (11)$$

Substituting $\mathbb{E}(T_1)$ obtained above in (9) we get

$$E_n = \sum_{i=1}^{n-1} \left(\frac{1}{i^2} - \frac{1}{ni} \right) + \frac{1}{n^2} + \frac{1}{n} \sum_{i=1}^{n-1} \frac{1}{i} = \sum_{i=1}^{n} \frac{1}{i^2}. \qquad (12)$$

This establishes Parisi's conjecture.

6 The Coppersmith-Sorkin Conjecture

As mentioned in the introduction, Coppersmith and Sorkin [CS 99] have conjectured that the expected cost of the minimum k-assignment in an $m \times n$ matrix, P, of i.i.d. exp(1) entries is:

$$F(k, m, n) = \sum_{i,j \geq 0, i+j<k} \frac{1}{(m-i)(n-j)}. \qquad (13)$$

Nair [Na 02] has proposed a larger set of conjectures that identifies each term in equation (13) as the expected value of an exponentially distributed random variable corresponding to an increment of appropriately sized matchings in the $m \times n$ matrix. We prove this larger set of conjectures using the machinery developed in Section 4 and therefore establish the Coppersmith-Sorkin conjecture.

Consider the matrix P and w.l.o.g. assume that $m \leq n$. For $1 \leq p \leq m$, let V_1^p be the weight of the smallest matching of size p and denote the columns it occupies by K_1^p. Next, for $1 \leq p \leq m-1$, define

$V_2^p, \ldots, V_{p+1}^p$ to be the ordered sequence (increasingly by weight) of the smallest matching obtained by deleting one column of K_1^p at a time. We shall refer to the set $V_1^m \bigcup \left(\bigcup_{p=1}^{m-1} \{V_1^p, \ldots, V_{p+1}^p\} \right)$ as the *V-matchings* of the matrix P. Nair [Na 02] has made the following distributional conjectures regarding the increments of the *V-matchings*.

Conjecture 2 *[Na 02] The following hold:*

$$
\begin{array}{rcl}
V_1^1 & \sim & \exp(mn) \\
------ & ------ & ------ \\
V_2^1 - V_1^1 & \sim & \exp(m(n-1)) \\
V_1^2 - V_2^1 & \sim & \exp((m-1)n) \\
------ & ------ & ------ \\
 & . & \\
 & . & \\
 & . & \\
------ & ------ & ------ \\
V_2^i - V_1^i & \sim & \exp(m(n-i)) \\
V_3^i - V_2^i & \sim & \exp((m-1)(n-i+1)) \\
. & . & . \\
V_{p+1}^i - V_p^i & \sim & \exp((m-p+1)(n-i+p-1)) \\
. & . & . \\
V_{i+1}^i - V_i^i & \sim & \exp((m-i+1)(n-1)) \\
V_1^{i+1} - V_{i+1}^i & \sim & \exp((m-i)n) \\
------ & ------ & ------ \\
 & . & \\
 & . & \\
------ & ------ & ------ \\
V_2^{m-1} - V_1^{m-1} & \sim & \exp(m(n-m+1)) \\
. & . & . \\
V_m^{m-1} - V_{m-1}^{m-1} & \sim & \exp(2(n-1)) \\
V_1^m - V_m^{m-1} & \sim & \exp(n)
\end{array}
$$

We have grouped the increments according to size. That is, the i^{th} group consists of the differences in the weights: $V_1^i, V_2^i, \ldots, V_{i+1}^i, V_1^{i+1}$, and V_1^i is the matching of size i, V_2^i is the second smallest of size i, etc, until V_1^{i+1}—the smallest matching of size $i+1$. Note that according to Conjecture 2 the telescopic sum $V_1^{i+1} - V_1^i$ has expected value $F(i+1, m, n) - F(i, m, n)$. Note also that $V_1^1 \sim \exp(mn)$, being the smallest of $m \times n$ independent $\exp(1)$ variables.

The rest of the section is devoted to establishing Conjecture 2 for the i^{th} group.

Proof of Conjecture 2

We will establish the conjectures for the i^{th} group inductively, as in Section 4. Consider a matrix, P_{m-j+1}, of size $(m-j+1) \times n$ and let $\{V_p^{q,m-j+1}\}$ denote[4] its *V-matchings*. The induction consists

[4]We regret the cumbersome notation; but we must keep track of three indices: one for the number of rows in the matrix (of size $m-j+1 \times n$), one for the size of the matching, q, and one for the rank of the matching, p, among matchings of size q.

of the following two steps:

Inductive Hypothesis:

- Assume the increments satisfy the following combinatorial identities
$$\begin{array}{rcl} V_2^{i-j+1,m-j+1} - V_1^{i-j+1,m-j+1} & = & V_{j+1}^i - V_j^i \\ V_3^{i-j+1,m-j+1} - V_2^{i-j+1,m-j+1} & = & V_{j+2}^i - V_{j+1}^i \\ \cdots & \cdots & \cdots \\ V_{i-j+2}^{i-j+1,m-j+1} - V_{i-j+1}^{i-j+1,m-j+1} & = & V_{i+1}^i - V_i^i \\ V_1^{i-j+2,m-j+1} - V_{i-j+2}^{i-j+1,m-j+1} & = & V_1^{i+1} - V_{i+1}^i. \end{array} \tag{14}$$
- The entries of P_{m-j+1} are i.i.d. exp(1) random variables.

Induction Step:

Step 1: From P_{m-j+1}, form a matrix P_{m-j} of size $(m-j) \times n$ with the property that
$$\begin{array}{rcl} V_2^{i-j,m-j} - V_1^{i-j,m-j} & = & V_{j+2}^i - V_{j+1}^i \\ V_3^{i-j,m-j} - V_2^{i-j,m-j} & = & V_{j+3}^i - V_{j+2}^i \\ \cdots & \cdots & \cdots \\ V_{i-j+1}^{i-j,m-j} - V_{i-j}^{i-j,m-j} & = & V_{i+1}^i - V_i^i \\ V_1^{i-j+1,m-j} - V_{i-j+1}^{i-j,m-j} & = & V_1^{i+1} - V_{i+1}^i. \end{array}$$

Step 2: Establish that the entries of P_{m-j} are i.i.d. exp(1) random variables.

This completes the induction step since P_{m-j} satisfies the induction hypothesis for the next iteration.

In Step 2 we also show that $V_2^{i-j+1,m-j+1} - V_1^{i-j+1,m-j+1} \sim \exp((m-j+1)(n-i+j-1))$ and hence conclude from equation (14) that $V_{j+1}^i - V_j^i \sim \exp((m-j+1)(n-i+j-1))$.

The induction starts at $j = 1$ and terminates at $j = i-1$. Observe that the matrix P satisfies the inductive hypothesis for $j = 1$ trivially.

Proof of the Induction:

Step 1: Form the matrix $\bar{P}_{m-j+1}$ of size $(m-j+1) \times (m-i+n)$ by adding $m-i$ columns of zeroes to the left of P_{m-j+1} as below
$$\bar{P}_{m-j+1} = [\mathbf{0}\,|P_{m-j+1}\,].$$

Let $\bar{T}_1^{m-j+1}, ..., \bar{T}_{m-j+2}^{m-j+1}$ denote the weight of the *T-matchings* of the matrix $\bar{P}_{m-j+1}$. Then we make the following claim.

Claim 2
$$\begin{array}{rcl} \bar{T}_1^{m-j+1} & = & V_1^{i-j+1,m-j+1} \\ \bar{T}_2^{m-j+1} & = & V_2^{i-j+1,m-j+1} \\ & \cdots & \\ \bar{T}_{i-j+2}^{m-j+1} & = & V_{i-j+2}^{i-j+1,m-j+1} \end{array}$$

and
$$\bar{T}_{i-j+3}^{m-j+1} = \bar{T}_{i-j+4}^{m-j+1} = \cdots = \bar{T}_{m-j+2}^{m-j+1} = V_1^{i-j+2,m-j+1}.$$

Proof Note that any matching of size $m-j+1$ in $\bar{P}_{m-j+1}$ can have at most $m-i$ zeroes. Therefore, it is clear that the smallest matching of size $m-j+1$ in $\bar{P}_{m-j+1}$ is formed by picking $m-i$ zeroes along with the smallest matching of size $i-j+1$ in P_{m-j+1}. Thus, $\bar{T}_1^{m-j+1} = V_1^{i-j+1,m-j+1}$.

If we remove any of the columns containing zeroes we get the smallest matching of size $m-j+1$ in the rest of the matrix $\bar{P}_{m-j+1}$ by combining $m-i-1$ zeroes with the smallest matching of size $i-j+2$ in P_{m-j+1}. Hence $m-i$ of the $\bar{T}$'s, corresponding to each column of zeroes, have weight equal to $V_1^{i-j+2,m-j+1}$.

If we remove any column containing $V_1^{i-j+1,m-j+1}$, then the smallest matching of size $m-j+1$ in $\bar{P}_{m-j+1}$ is obtained by $m-i$ zeroes and the smallest matching of size $i-j+1$ in P_{m-j+1} that avoids this column. Hence they have weights $V_p^{i-j+1,m-j+1}$ for $p \in \{2,3,...,i-j+2\}$.

We claim that $V_1^{i-j+2,m-j+1}$ is larger than $V_p^{i-j+1,m-j+1}$ for p in $\{1,2,3,...,i-j+2\}$. Clearly $V_1^{i-j+2,m-j+1} > V_1^{i-j+1,m-j+1}$. Further, for $p \geq 2$, we have a matching of size $i-j+1$ in $V_1^{i-j+2,m-j+1}$ that avoids the same column that $V_p^{i-j+1,m-j+1}$ avoids in $V_1^{i-j+1,m-j+1}$. But $V_p^{i-j+1,m-j+1}$ is the smallest matching of size $i-j+1$ that avoids this column. Hence we conclude that $V_1^{i-j+2,m-j+1} > V_p^{i-j+1,m-j+1}$.

Hence arranging the weights (in increasing order) of the smallest matchings of size $m-j+1$ in $\bar{P}_{m-j+1}$, obtained by removing one column of $\bar{T}_1^{m-j+1}$ at a time, establishes the claim. ∎

From the above it is clear that the matchings $\bar{T}_1^{m-j+1}$ and $\bar{T}_2^{m-j+1}$ are formed by $m-i$ zeroes and the matchings $V_1^{i-j+1,m-j+1}$ and $V_2^{i-j+1,m-j+1}$ respectively. Hence, as in Section 4, we have two elements, one each of $\bar{T}_1^{m-j+1}$ and $\bar{T}_2^{m-j+1}$ that lie outside the common columns of $\bar{T}_1^{m-j+1}$ and $\bar{T}_2^{m-j+1}$. Let $\bar{K}_{12}$ denote these common columns. (Note that $\bar{K}_{12}$ necessarily includes the $m-i$ columns of zeros).

We now proceed to perform the procedure outlined in Section 4 for obtaining P_{m-j} from P_{m-j+1} by working through the matrix $\bar{P}_{m-j+1}$ which is an extension of P_{m-j+1}.

Accordingly, form the matrix $\bar{P}^*$ by removing the value $\bar{T}_2^{m-j+1} - \bar{T}_1^{m-j+1}$ from all the entries in $\bar{P}_{m-j+1}$ that lie outside $\bar{K}_{12}$. Generate a random variable X, as before, with $\mathbb{P}(X=1) = \mathbb{P}(X=2) = \frac{1}{2}$. Let $Y=1$ when $X=2$ and let $Y=2$ when $X=1$. Denote by e_1 the unique entry of the matching $\bar{T}_X^*$ outside the matrix $\bar{K}_{12}$. Denote by e_2 the unique entry of the matching $\bar{T}_Y^*$ outside the matrix $\bar{K}_{12}$. Remove the row containing e_1 and call this matrix $\bar{P}_{m-j}$. Now remove the $m-i$ columns of zeroes to obtain the matrix P_{m-j} of size $(m-j)\times n$.

Let $\bar{T}_1^{m-j}, ..., \bar{T}_{m-j+1}^{m-j}$ denote the weight of the *T-matchings* of the matrix $\bar{P}_{m-j}$ and $V_p^{q,m-j}$ denote the *V-matchings* of the matrix P_{m-j}. We make the following claim.

Claim 3

$$\begin{aligned} \bar{T}_1^{m-j} &= V_1^{i-j,m-j} \\ \bar{T}_2^{m-j} &= V_2^{i-j,m-j} \\ &\cdots \\ \bar{T}_{i-j+1}^{m-j} &= V_{i-j+1}^{i-j,m-j} \end{aligned}$$

and

$$\bar{T}_{i-j+2}^{m-j} = \bar{T}_{i-j+3}^{m-j} = \cdots = \bar{T}_{m-j+1}^{m-j} = V_1^{i-j+1,m-j}.$$

Proof The proof is identical to that of Claim 2. ∎

Now from Lemma 9 in Section 4 we know that

$$\bar{T}_{p+2}^{m-j+1} - \bar{T}_{p+1}^{m-j+1} = \bar{T}_{p+1}^{m-j} - \bar{T}_p^{m-j} \text{ for } 1 \leq p \leq m-j. \tag{15}$$

Finally, combining Equation (15), Claim 2, Claim 3 and the inductive hypothesis on P_{m-j+1} we obtain:

$$\begin{array}{rcl} V_2^{i-j,m-j} - V_1^{i-j,m-j} & = & V_{j+2}^i - V_{j+1}^i \\ V_3^{i-j,m-j} - V_2^{i-j,m-j} & = & V_{j+3}^i - V_{j+2}^i \\ \cdots & \cdots & \cdots \\ V_{i-j+1}^{i-j,m-j} - V_{i-j}^{i-j,m-j} & = & V_{i+1}^i - V_i^i \\ V_1^{i-j+1,m-j} - V_{i-j+1}^{i-j,m-j} & = & V_1^{i+1} - V_{i+1}^i. \end{array}$$

This completes Step 1 of the induction.

Step 2: Again we reduce the problem to the one in Section 4 by working with the matrices $\bar{P}_{m-j+1}$ and $\bar{P}_{m-j}$ instead of the matrices P_{m-j+1} and P_{m-j}. (Note that the necessary and sufficient conditions for a P_{m-j+1} to be in the pre-image of a particular realization of P_{m-j} is exactly same as the necessary and sufficient conditions for a $\bar{P}_{m-j+1}$ to be in the pre-image of a particular realization of $\bar{P}_{m-j}$.)

Let $\mathcal{R}_1$ denote all matrices $\bar{P}_{m-j+1}$, of size $m-j+1$, that map to a particular realization of $\bar{P}_{m-j}$ with e_1 and e_2 in the same row. Let $\mathcal{R}_2$ denote all matrices $\bar{P}_{m-j+1}$ that map to a particular realization of $\bar{P}_{m-j}$ with e_1 and e_2 in different rows. Observe that in $\mathcal{R}_2$, $\bar{P}_{m-j+1}$ will map to the particular realization of $\bar{P}_{m-j}$ with probability $\frac{1}{2}$ as in Section 4. We borrow the notation from Section 4 for the rest of the proof.

(Before proceeding, it helps to make some remarks relating the quantities in this section to their counterparts in Section 4. The matrix L_{n-j} had dimensions $(n-j) \times n$; its counterpart $\bar{P}_{m-j+1}$ has dimensions $(m-j+1) \times (m-i+n)$. The number of columns in L_{n-j} outside the columns of T_1^{n-j} equalled j; now the number of columns of $\bar{P}_{m-j+1}$ outside the columns of $\bar{T}_1^{m-j+1}$ equals $n-i+j-1$. This implies that the value $\theta = \bar{T}_2^{m-j+1} - \bar{T}_1^{m-j+1}$ will be subtracted from precisely $(m-j+1)(n-i+j-1)$ elements of $\bar{P}_{m-j+1}$. Note also that the vector $\vec{r}$, of length $m-j$, has exactly $m-i$ zeroes and $i-j$ non-zero elements.)

Let $P_{m-j} = [p_{k,l}]$ denote a particular realization of P_{m-j}. We proceed by setting, as in Section 4,

$$\vec{v} = (P_{m-j}, \vec{r}, \theta) \quad \text{and} \quad \vec{w} = (\vec{v}, \vec{x}).$$

We will evaluate $f_v(\vec{v}) = \int_{\mathcal{R}_1} f_w(\vec{v}, \vec{x}) d\vec{x} + \int_{\mathcal{R}_2} f_w(\vec{v}, \vec{x}) d\vec{x}$, to obtain the marginal density of $\vec{v}$.

On $\mathcal{R}_1$, we have that $x_a = x_b < U$ for U as in Section 4. (The counterparts of x_a and x_b in Section 4 were x_i and x_k, and these were defined according to Lemma 10.) We shall set $u_l = x_l - x_a$ for $l \neq a, b$. Finally, define

$$s_v = p_{1,1} + \ldots + p_{m-j,n} + r_1 + \ldots + r_{i-j} + (m-j+1)(n-i+j-1)\theta.$$

Thus, s_v denotes the sum of all of the entries of P_{m-j+1} except those in $\vec{x}$. We have

$$\begin{aligned} & \int_{\mathcal{R}_1} f_w(\vec{v}, \vec{x}) d\vec{x} \\ \overset{(a)}{=} \quad & 2(m-j+1)\binom{n-i+j}{2} \int_0^U \int_0^\infty \cdots \int e^{-(s_v + (n-i+j)x_a + u_1 + \ldots + u_{n-i+j-2})} \, du_1 \ldots du_{n-i+j-2} \, dx_a \\ = \quad & (m-j+1)(n-i+j-1) e^{-s_v} \left(1 - e^{-(n-i+j)U}\right). \end{aligned} \tag{16}$$

The factor $\binom{n-i+j}{2}$ at equality (a) comes from the possible choices for a, b from the set $\{1, \ldots, n-i+j\}$, the factor $(m-j+1)$ comes from the row choices available to e_1 as in Section 4, and the factor 2 corresponds to the partition, $\mathcal{S}_1$ or $\mathcal{S}_2$ (defined likewise), that P_{m-j+1} belongs to.

Similarly, on $\mathcal{R}_2$, we have that $x_a = U$ and we shall set $u_l = x_l - x_a$ for $l \neq a$ to obtain

$$
\begin{aligned}
&\int_{\mathcal{R}_2} f_w(\vec{v}, \vec{x}) d\vec{x} \\
\overset{(b)}{=}\quad & \frac{1}{2}\left[2(m-j+1)(n-i+j-1)\int_0^{\infty}\cdots\int e^{-(s_v+(n-i+j)U+u_1+\ldots+u_{n-i+j-1})}\,du_1\ldots du_{n-i+j-1}\right] \\
=\quad & (m-j+1)(n-i+j-1)e^{-s_v}e^{-(n-i+j)U}. \qquad (17)
\end{aligned}
$$

In equality (b) above, the factor $(n-i+j-1)$ comes from the choice of positions available to x_a (note that x_a cannot occur in same column as the entry d_o which was defined in Lemma 10). The factor $(m-j+1)$ comes from the row choices available to e_1, and the factor 2 is due to the partition, $\mathcal{S}_1$ or $\mathcal{S}_2$, that P_{m-j} belongs to. Finally, the factor $\frac{1}{2}$ comes from the fact that on $\mathcal{R}_2$, e_1 and e_2 occur on different rows. Therefore, P_{m-j+1} will map to the desired P_{m-j} with probability $\frac{1}{2}$.

Putting (16) and (17) together, we obtain

$$f_v(\vec{v}) = e^{-s_v} = (e^{-(p_{1,1}+\ldots+p_{m-j,n})} \times (m-j+1)(n-i+j-1)e^{-(m-j+1)(n-i+j-1)\theta} \times e^{-(r_1+\ldots+r_{i-j})}.$$

We summarize the above in the following lemma.

Lemma 14 *The following hold:*
(i) P_{m-j} *consists of i.i.d.* $\exp(1)$ *variables.*
(ii) $\theta = \bar{T}_2^{m-j+1} - \bar{T}_1^{m-j+1}$ *is an* $\exp((m-j+1)(n-i+j-1))$ *random variable.*
(iii) $\vec{r}$ *consists of i.i.d.* $\exp(1)$ *variables and* $m-i$ *zeroes.*
(iv) P_{m-j}*,* $(\bar{T}_2^{m-j} - \bar{T}_1^{m-j})$*, and* $\vec{r}$ *are independent.*

This completes Step 2 of the induction. ■

From Claim 2, we have that $\bar{T}_2^{m-j+1} - \bar{T}_1^{m-j+1} = V_2^{i-j+1,m-j+1} - V_1^{i-j+1,m-j+1}$, and from the inductive hypothesis we have $V_2^{i-j+1,m-j+1} - V_1^{i-j+1,m-j+1} = V_{j+1}^i - V_j^i$. Hence we have the following corollary.

Corollary 6 $V_{j+1}^i - V_j^i \sim \exp((m-j+1)(n-i+j-1))$ *for* $j = 1, 2, .., i-1$*.*

To complete the proof of Conjecture 2 we need to compute the distribution of the two increments $V_{i+1}^i - V_i^i$ and $V_1^{i+1} - V_{i+1}^i$. At the last step of the induction, i.e. $j = i-1$, we have a matrix P_{m-i+1} consisting of i.i.d. exp(1) random variables and satisfying the following properties: $V_2^{1,m-i+1} - V_1^{1,m-i+1} = V_{i+1}^i - V_i^i$ and $V_1^{2,m-i+1} - V_2^{1,m-i+1} = V_1^{i+1} - V_{i+1}^i$. The following lemma completes the proof of Conjecture 2.

Lemma 15 *The following identities hold:*

(i) $V_2^{1,m-i+1} - V_1^{1,m-i+1} \sim \exp((m-i+1)(n-1))$*.*

(ii) $V_1^{2,m-i+1} - V_2^{1,m-i+1} \sim \exp((m-i)n)$*.*

Proof This can be easily deduced from the memoryless property of the exponential distribution; equally, one can refer to Lemma 1 in [Na 02] for the argument. (Remark: There is a row and column interchange in the definitions of the *V-matchings* in [Na 02].) ■

Thus, we have fully established Conjecture 2 and obtain

Theorem 3

$$F(k, m, n) = \sum_{i,j \geq 0, i+j<k} \frac{1}{(m-i)(n-j)}.$$

Proof

$$\begin{aligned} V_1^k &= (V_1^k - V_k^{k-1}) + (V_k^{k-1} - V_{k-1}^{k-1}) + \cdots + (V_2^1 - V_1^1) + V_1^1 \\ \Rightarrow \quad \mathbb{E}(V_1^k) &= \sum_{i,j \geq 0, i+j<k} \frac{1}{(m-i)(n-j)}, \\ \text{and} \quad F(k,m,n) &\triangleq \mathbb{E}(V_1^k). \end{aligned}$$

■

This gives an alternate proof to Parisi's conjecture since [CS 99] shows that $E_n = F(n,n,n) = \sum_{i=1}^{n} \frac{1}{i^2}$.

7 Conclusions and further work

This paper provides a proof of the conjectures by Parisi [Pa 98] and Coppersmith-Sorkin [CS 99]. However, there are still some interesting properties of matchings in matrices with i.i.d. exp(1) entries that remain unresolved. Of particular interest is the following conjecture: Let Q be an $(n-1) \times n$ matrix of i.i.d. *exp*(1) entries and let T_i^Q denote its *T-matchings*. Let Υ denote the set of all placements of the row-wise minimum entries of Q; for example, all the row-wise minima in the same column, all in distinct columns, etc. Consider any fixed placement of the row minima $\xi \in \Upsilon$. Then we make the following conjecture.

Conjecture 3 *Conditioned on a particular placement, ξ, of the row-wise minima of Q, $T_{j+1}^Q - T_j^Q \sim \exp(j(n-j))$ for $j = 1, \ldots, n-1$. Furthermore, these increments are independent of each other.*

Clearly, if we average over all $\xi \in \Upsilon$ then we recover Conjecture 1. It is simple to prove Conjecture 3 in the case when ξ is the placement that corresponds to all row minima being in distinct columns. Therefore, if a simple argument could be advanced for showing that the distribution of $\{T_{j+1}^Q - T_j^Q\}$ is invariant with respect to the placement of the row minima, we will prove Conjecture 3 and hence obtain another proof of Parisi's conjecture.

We have verified Conjecture 3 for $n \leq 4$. We believe that similar conjectures exist for $m \times n$ matrices ($m < n$) as well, i.e. distributional properties of the increments of the *T-matchings* ought to be independent of the row minima placement.

As Aldous [Al 01] has noted, [Do 00] contains an incomplete proof to Parisi's conjecture. We believe that some of the claims in [Do 00] can be justified using a generalization of Conjecture 3. It is also interesting to note that we do not use the exponentiality of the variables until we "integrate out the marginals". So it might be possible to say something in cases where the entries are not exponential or are exponentials with unequal rates as in [BCR 02].

References

[Al 92] D. J. Aldous. Asymptotics in the random assignment problem. *Probab. Th. Rel. Fields*, 93 (1992) 507-534.

[Al 01] D. J. Aldous. The $\zeta(2)$ Limit in the Random Assignment Problem. *Random Structures and Algorithms* 18 (2001) 381-418.

[AS 02] S. E. Alm and G. Sorkin. Exact expectations and distributions for the random assignment problem. *Combinatorics, Probability, and Computing*, vol 11 no.3 (2002) 217-248.

[BCR 02] M. W. Buck, C. S. Chan and D. P. Robbins. On the Expected Value of the Minimum Assignment. *Random Structures and Algorithms* 21 (2002) 33-58.

[BKMP 86] Brunetti, Krauth, Mézard and Parisi (Unpublished).

[CS 99] D. Coppersmith and G. Sorkin. Constructive bounds and exact expectations for the random assignment problem. *Random Structures and Algorithms*, 15 (1999) 113-144.

[CS 02] D. Coppersmith and G. Sorkin. On the Expected Incremental Cost of a Minimum Assignment. In *Contemporary Combinatorics* (B. Bollobas, ed.), *Bolyai Society Mathematical Studies*, vol 10, chapter 8, Springer, 2002

[Do 00] V. S. Dotsenko. Exact solution of the random bipartite matching model. *J. Phys. A*, vol 33, (2000) 2015-2030.

[GK 93] M. X. Goemans and M. S. Kodialam. A lower bound on the expected cost of an optimal assignment. *Math. Oper. Res.*, 18 (1993) 267-274.

[Ka 84] R. M. Karp. An upper bound on the expected cost of an optimal assignment. Technical Report, Computer Science Division, University of California, Berkeley CA, 1984.

[Ka 87] R. M. Karp. An upper bound on the expected cost of an optimal assignment. In *Discrete Algorithms and Complexity: Proceedings of the Japan-U.S Joint Seminar*, pages 1-4, Academic Press, 1987.

[KKV 94] R. M. Karp, A. H. G Rinnooy Kan and R. Vohra. Average case analysis of a heuristic for the assignment problem. *Mathematics of Operations Research*, 19(3) (1994) 513-522.

[La 93] A. Lazarus. Certain expected values in the random assignment problem. *Oper. Res. Letters*, 14 (1993) 207-214.

[LW 00] S. Linusson and J. Wästlund. A Generalization of the Random Assignment Problem. http://www.mai.liu.se/˜ svlin/

[LW 03] S. Linusson and J. Wästlund. A Proof of Parisi's Conjecture on the Random Assignment Problem. http://www.mai.liu.se/˜ svlin/

[MP 87] M. Mézard and G. Parisi. On the solution of the random link matching problem. *J. Physique*, 48 (1987) 1451-1459.

[MPV 87] M. Mézard, G. Parisi and M. A. Virasoro. *Spin Glass Theory and Beyond*. World Scientific, Singapore, 1987.

[Na 02] C.Nair. Towards the Resolution of Coppersmith-Sorkin Conjectures. *Proceedings of the 40th Annual Allerton Conference on Communication, Control and Computing*, 2002.

[Ol 92] B. Olin. *Asymptotic Properties of Random Assignment Problems*. PhD thesis, Kungl Tekniska Högskolan, Stockholm, Sweden, 1992.

[Pa 98] G. Parisi. A conjecture on random bipartite matching. Technical Report cond-mat/9801176, XXX LANL Archive, 1998

[SP 02] M. Sharma and B. Prabhakar. On Parisi's Conjecture for the Finite Random Assignment Problem. *Proceedings of the 40th Annual Allerton Conference on Communication, Control and Computing* 2002.

[SP 03] M. Sharma and B. Prabhakar. On Parisi's Conjecture for the Finite Random Assignment Problem. *To be submitted.*

[St 97] J. Michael Steele. *Probability Theory and Combinatorial Optimization*. Number 69 in CBMS-NSF Regional Conference Series in Applied Mathematics. SIAM, Philadelphia PA, 1997.

[Wa 79] D. W. Walkup. On the expected value of a random assignment problem. *SIAM J. Comput.*, 8 (1979) 440-442.

Appendix: Proofs of the combinatorial lemmas

To execute some of the proofs we will use the alternate representation of an arbitrary $k \times l$ matrix M as a complete bipartite graph, with k vertices on the left and l vertices on the right corresponding to the rows and columns of M, respectively. The edges are assigned weights m_{ij} with the obvious numbering.

Proof of Lemma 1 We represent the matrix M as a complete bipartite graph $G_{n,n+l}$. Choose an arbitrary column $\mathbf{c} \in K_1$ and focus on the subgraph consisting of only those edges that participate in T_1^n (as defined in Section 1) and $\mathcal{M}_{\mathbf{c}}$. Such a subgraph is depicted in Figure 1 where the bold edges belong to T_1^n and the dashed edges correspond to $\mathcal{M}_1$, the smallest matching in the matrix obtained by deleting column 1.

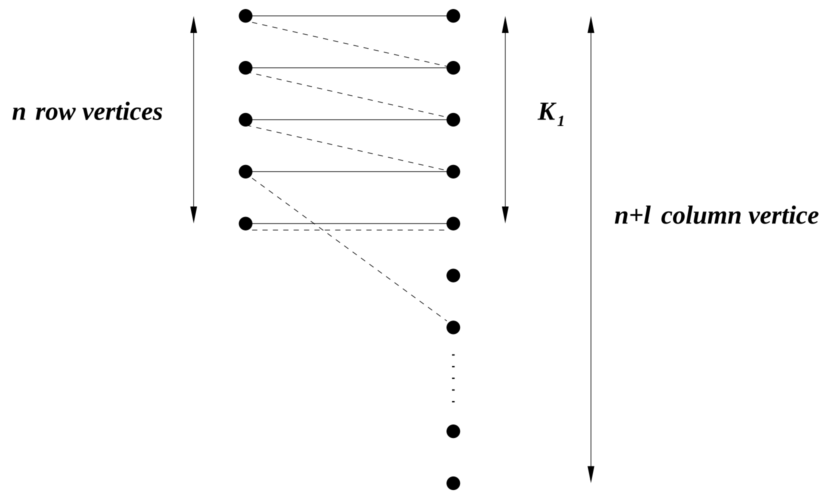

Figure 1: Subgraph depicting an even-length path and a 2-cycle

In general, a subgraph of a bipartite graph can consist of the following components: cycles, and paths of even or odd lengths. We claim that it is impossible for the subgraph induced by the edges of T_1^n and $\mathcal{M}_1$ to have cycles of length greater than two, or paths of odd length.

A cycle of length greater than two is impossible because it would correspond to two different sub-matchings being chosen by T_1^n and $\mathcal{M}_1$ on a common subset of rows and columns. This would contradict the minimality of either T_1^n or of $\mathcal{M}_1$.

An odd-length path is not possible because every vertex on the left has degree 2 and therefore any path that enters it will be able to leave. Thus, any path will have to be of even length. (Note that we can have cycles of length two because they correspond to a common entry being chosen by both matchings.)

Next we show that the only other component that can be present is a path of even length whose degree-one vertices are on the right. Every node on the left has degree 2 and hence even paths with two degree-1 nodes on the left is not possible. Suppose there are two or more paths of even length. Consider any two of them: at least one of them will not be incident on column 1. Now the bold edges along this path are smaller in weight than the dashed edges by the minimality of T_1^n. Thus we can append these bold edges to the dashed edges not on this path to obtain a new matching $\mathcal{M}_1'$ which would be smaller than $\mathcal{M}_1$. This contradicts the minimality of $\mathcal{M}_1$ amongst all matchings that do not use column 1.

Having deduced that this subgraph only consists of 2-cycles and one even-length path, we can now prove Lemma 1. A 2-cycle doesn't violate Lemma 1 because it corresponds to a right-side vertex (or column) in K_1 being used by $\mathcal{M}_1$. Consider the even-length path, and observe that all but one of the dashed edges along the even-length path are incident on columns in K_1. This is because dashed edges alternate with bold edges and the latter are incident only on columns in K_1. This leaves exactly one dashed edge (the last one in the path) being incident on a column outside K_1. ■

Proof of Lemma 2 First note that we only need to prove the result for $l \leq n$ because T_1^n can use at most n columns outside those in K.

The proof is by induction on matrices of size $n \times (n+j)$, each containing the columns K. The lemma is trivially true for $j = 1$. Suppose that it is true for matrices of size $n \times (n+j)$, $j = 1, \ldots, l-1$. If the statement is not true for $j = l$, then we show that Lemma 1 is contradicted.

Assume therefore that T_1^n consists of elements from more than one column outside K. Observe that the number of columns that T_1^n uses outside K cannot be less than l because that would violate the induction hypothesis for some lower j. Hence assume that T_1^n uses all l columns outside K.

Now choose any one of these l columns and find the smallest matching that does not take an element from it. We know from the induction hypothesis that this matching is $\mathcal{M}$. But since $\mathcal{M}$ and T_1^n have at least l columns which are not common, we arrive at a contradiction with Lemma 1. ■

Proof of Lemma 3 As before, we represent the augmented matrix M_{ext} as a complete bipartite graph $G_{n+r,n+l}$ and focus on the subgraph (see Figure 2) consisting of only those edges that are part of $\mathcal{M}$ (bold edges) and $\mathcal{M}'$ (dashed edges).

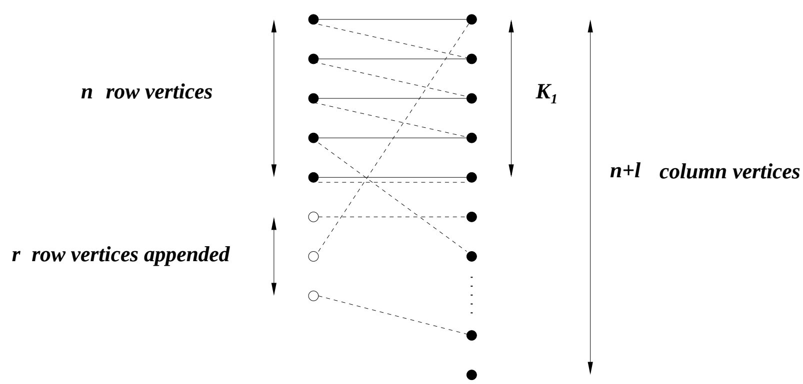

Figure 2: Subgraph depicting odd-length paths and a 2-cycle

We proceed by eliminating the possibilities for components of this subgraph. As in Lemma 1, a cycle of length greater than 2 is not possible, but 2-cycles (common edges) are possible.

Next we show that paths of even length cannot exist. There are two types of even-length paths, those with degree-1 vertices on the left and those with degree-1 vertices on the right. Since every vertex on the left in the original matrix has degree 2, even-length paths that have degree 1 on the left are not possible.

The other type of even-length path where the degree-1 vertices are on the right has the property that the solid and dashed edges use the same vertices on the left (i.e. same set of rows). Now, we have two matchings on the same set of rows and therefore by choosing the lighter one, we can contradict the minimality of $\mathcal{M}$ or $\mathcal{M}'$.

We now consider odd-length paths. Again, since every vertex on the left in the original matrix has degree 2, this implies that the only type of odd-length paths that are possible are those for which the number of edges from M_{ext} is one more than the number of edges from M. But in such an odd-length path, the vertices on the right (columns) used by the original matrix are also used by the augmented matrix.

The only other component left to consider are 2-cycles (or common edges). Here we have the same column used by both matchings. This completes the proof that every column used by $\mathcal{M}$ is also used by $\mathcal{M}'$. ■

Proof of Lemma 4 The proof is by induction on i. By definition, T_2^{n-j} is the smallest matching of size $n-j$ in L_{n-j} that contains exactly $n-j-1$ elements from K_1. Therefore T_2^{n-j} avoids one column in K_1, say i_o. This implies that $T_2^{n-j} \geq S_{i_o}^{n-j}$. Now from Lemma 1 we know that the smallest matching that avoids any column $i \in K_1$ contains just one element outside K_1. This implies that for all $i \in K_1$, S_i^{n-j} contains exactly $n-j-1$ elements from K_1. Thus $T_2^{n-j} \leq S_i^{n-j}, \ \ i = 1, \ldots, n-j$. In particular, $T_2^{n-j} \leq S_{i_o}^{n-j}$ and this yields the following inequality: $T_2^{n-j} = S_{i_o}^{n-j} \leq S_i^{n-j}, \ \ i = 1, \ldots, n-j$. Again, by definition, $R_1^{n-j} = S_{i_o}^{n-j}$, the smallest of the S_i^{n-j}'s. Therefore, $T_2^{n-j} = R_1^{n-j}$ and the result holds for $i = 1$.

Now assume $T_{i+1}^{n-j} = R_i^{n-j}, \ i = 1, ..., k-1$ for some $k \leq n-j$. Consider the columns present in $K_{12..k}$. These columns were by definition present in all of $T_1^{n-j}, T_2^{n-j}, ..., T_k^{n-j}$ and hence, by the induction hypothesis, in all of $R_1^{n-j}, R_2^{n-j}, ..., R_{k-1}^{n-j}$. Note that there are exactly $k-1$ columns that are present in K_1 but not in $K_{12..k}$. Therefore, $R_i^{n-j}, \ \ i = 1, ..., k-1$ is the ordered sequence of weights of the smallest matchings of size $n-j$ when one of these columns is removed at a time.

This implies that R_k^{n-j} is the weight of the smallest matching that avoids a column in $K_{12..k}$. By definition, T_{k+1}^{n-j} is the smallest matching that uses all but one column in $K_{12..k}$. Therefore, T_{k+1}^{n-j} avoids one column in $K_{12..k}$ and we obtain $T_{k+1}^{n-j} \geq R_k^{n-j}$. From Lemma 1 we know that all S_i^{n-j}'s contain exactly $n-j-1$ elements from K_1. This implies in particular that R_k^{n-j} contains all but one column in $K_{12..k}$ since it avoids a column in $K_{12..k}$. Clearly R_k^{n-j} was a possible choice for T_{k+1}^{n-j}. Hence $T_{k+1}^{n-j} \leq R_k^{n-j}$.

Therefore, we have $T_{k+1}^{n-j} = R_k^{n-j}$ and this completes the induction step. ■

Proof of Corollary 1 From Lemma 1 we know that for all i, S_i^{n-j} contains exactly $n-j-1$ elements from K_1. From this fact and the equivalence to R_i^{n-j}'s established in Lemma 4, the corollary follows. ■

Proof of Corollary 2 This follows easily by noticing that in the proof of the induction step in Lemma 4, $R_k^{n-j}, \ k = 1, \ldots, n-j$ is the weight of the smallest matching that avoids exactly one column in $K_{12...k}$. But $T_{k+1}^{n-j} = R_k^{n-j}, \ k = 1, \ldots, n-j$ and we are done. ■

Proof of Lemma 6 For the proof we will augment the given matrix by an $(l-1) \times (n+l)$ matrix with all elements equal to 0 as in Figure 3 and call the augmented matrix M_{aug}. In M_{aug}, mark all the elements that participate in any of the minimum permutations in the $n+l$ square sub-matrices obtained by deleting one column at a time. Denote by S_i the weight of the smallest matching in the square sub-matrix of M_{aug} resulting from the deletion of column i. Denote by $\{T_i\}_1^n$ the ordered sequence of S_i's. From Lemma 5, exactly two elements per row will be marked in M_{aug} by the elements of $\{T_i\}_1^n$.

Now since the lower part of M_{aug} always contributes zero towards every matching in the original matrix, we note that $T_1 = T_1^n$. In fact, the deletion of any column i in M_{aug} that is not in K_1 will lead to $T_i = T_1^n$.

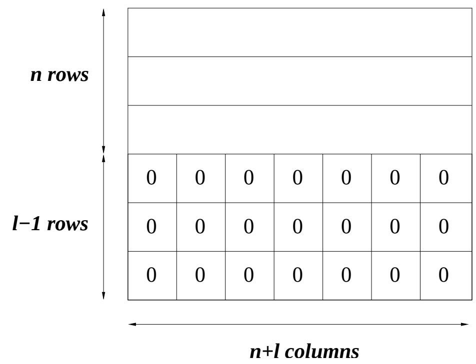

Figure 3: The augmented matrix

When we start deleting columns in K_1 we obtain the higher T_i's for M_{aug}. By the equivalent characterization of the sequence $T_1^n, T_2^n, \ldots, T_{n+1}^n$ described in Lemma 4, we notice that the elements that participate in these T_i^n's are precisely those that are marked in the upper n rows in M_{aug}. Since two dots per row are marked in M_{aug}, we obtain our result. ∎