

Probabilistic analysis of algorithms for cost constrained minimum weighted combinatorial objects

Alan Frieze* and Tomasz Tkocz†
Carnegie Mellon University
Pittsburgh PA15213
U.S.A.

Abstract

We consider cost constrained versions of the minimum spanning tree problem and the assignment problem. We assume edge weights are independent copies of a continuous random variable Z that satisfies $F(x) = \mathbb{P}(Z \leq x) \approx x^\alpha$ as $x \rightarrow 0$, where $\alpha \geq 1$. Also, there are $r = O(1)$ budget constraints with edge costs chosen from the same distribution. We use Lagrangean duality to construct polynomial time algorithms that produce asymptotically optimal solutions. For the spanning tree problem, we allow $r > 1$, but for the assignment problem we can only analyse the case $r = 1$.

1 Introduction

Let X be a finite set and let the elements of X be given independent random weights $w(e), e \in X$ and random costs $c_i(e), i = 1, 2, \dots, r$ for $e \in X$. Suppose we are given cost budgets of $\mathbf{C} = (C_i, i = 1, 2, \dots, r)$ and we consider the following problem: let \mathcal{S} denote some collection of subsets of X . For a function $f : X \rightarrow \mathbb{R}$ and $S \subseteq X$ we let $f(S) = \sum_{e \in S} f(e)$. We consider the optimization problem:

$$Opt(\mathcal{S}, \mathbf{C}) : \text{Minimise } w(S) \text{ subject to } S \in \mathcal{S}_r = \{S \in \mathcal{S} : c_i(S) \leq C_i, i = 1, 2, \dots, r\},$$

and let

$$w^* = w^*(\mathcal{S}, \mathbf{C}) \text{ denote the minimum value in } Opt(\mathcal{S}, \mathbf{C}).$$

This all sounds pretty general, but here we will only consider $X = E(K_n)$ or $X = E(K_{n,n})$, the edge sets of the complete graph and bipartite graph respectively. We will either have \mathcal{S} be the set of spanning trees of K_n or the set of perfect matchings of K_n or $K_{n,n}$.

In previous papers [3], [5], [6] we focussed on giving high probability asymptotic estimates of w^* in the case of trees, matchings, shortest paths and Hamilton cycles. In this paper we concentrate on finding polynomial time algorithms that w.h.p. find feasible solutions to $Opt(\mathcal{S}, \mathbf{C})$ with weight $(1 + o(1))w^*$. We do this without knowing asymptotic estimates of w^* . Finding such remain as open questions.

*Research supported in part by NSF grant DMS1661063

†Research supported in part by NSF grant DMS1955175

Spanning Trees Problems like this have been the focus of much research in the worst-case. For example Goemans and Ravi [7] consider the spanning tree problem with a single cost constraint. They give a polynomial time algorithm that finds a $(1 + \varepsilon, 1)$ solution to this problem. Here an (α, β) solution T is one that satisfies $c_i(T) \leq \alpha C_i$, $i = 1, 2, \dots, r$ and $w(T) \leq \beta w^*$. When $\alpha = 1$ and $\beta = 1 + o(1)$ we say that we have an *asymptotically optimal* solution. In [5], Frieze and Tkocz consider the case where $r = 1$ and costs and weights are independent copies of the uniform $[0, 1]$ random variable U and give a polynomial time algorithm that w.h.p. finds an asymptotically optimal solution. When $r > 1$ Grandoni, Ravi, Singh and Zenklusen [9] consider a maximization version and give a $(1 + \varepsilon, 1)$ algorithm that runs in time $n^{O(r^2/\varepsilon)}$ time on an n -vertex graph.

We will assume that we are working with real numbers. This may seem unrealistic, but we can instead work on a discretised version where we keep a polynomial number of bits m . It is routine to modify the proof below and make it work when $m = Kn \log n$, for large K .

Our random weights and costs will be distributed as a continuous random variable Z where $F(x) = \mathbb{P}(Z \leq x) \approx x^\alpha$ as $x \rightarrow 0$, where $\alpha \geq 1$. Here $A \approx B$ is an abbreviation for $A = (1 + o(1))B$ as $n \rightarrow \infty$, assuming that $A = A(n), B = B(n)$.

In this paper we prove

Theorem 1. *Suppose that the edges of the complete graph K_n are all given independent copies of Z . Let \mathcal{T} denote the set of spanning trees of K_n . Suppose there are $r = O(1)$ cost constraints whose coefficients are also independent copies of Z or Z_d . Suppose also that $n \geq C_i = \omega n^{1-r/(\alpha(r+1))} \log n$, $i = 1, 2, \dots, r$ where $\omega \rightarrow \infty$. Then there is a polynomial time algorithm that w.h.p. finds an asymptotically optimal solution to $\text{Opt}(\mathcal{T}, \mathbf{C})$.*

Remark 1. *Using a result of Gupta, Lee and Li [8], we can replace K_n by an arbitrary dense regular graph G with minimum degree $\delta = \tilde{\Omega}(n^{1/2})$ and edge connectivity $\kappa \geq \delta$. For this case, we will require that $n \geq C_i = \omega n \delta^{-1/(\alpha(r+1))} \log n$.*

Matchings Berger, Bonifaci, Grandoni and Schäfer [1] consider the case where $\mathcal{S} = \mathcal{M}$, the set of matchings of a graph. They consider the maximization version and describe an $n^{O(1/\varepsilon)}$ time algorithm that provides a $(1, 1 + \varepsilon)$ solution for the case where $r = 1$. In this paper we prove

Theorem 2. *Suppose that the edges of the complete graph K_n or the complete bipartite graph $K_{n,n}$ are given independent copies of Z . Let \mathcal{M} denote the set of perfect matchings of the complete bipartite graph $K_{n,n}$. Suppose there a single cost constraint whose coefficients are also independent copies of Z . And suppose that the RHS $C_1 \gg n^{1/2}$. Then there is a polynomial time algorithm that w.h.p. finds an asymptotically optimal solution to $\text{Opt}(\mathcal{M}, \mathbf{C})$.*

2 Trees

We consider the dual problem $\text{Dual}(\mathcal{T})$:

$$\text{Maximise } \phi(\boldsymbol{\lambda}) \text{ over } \boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_r) \geq 0, \text{ where } \phi(\boldsymbol{\lambda}) = \min \left\{ w(T) + \sum_{i=1}^r \lambda_i (c_i(T) - C_i) : T \in \mathcal{T} \right\}. \quad (1)$$

We note that

$$\text{if } \boldsymbol{\lambda} \geq 0 \text{ and } T \text{ is feasible for } \text{Opt}(\mathcal{T}, \mathbf{C}) \text{ then } \phi(\boldsymbol{\lambda}) \leq w(T). \quad (2)$$

We will show that w.h.p.

$$\text{that if } \boldsymbol{\lambda}^* \text{ solves (1) and } T^* \text{ solves } \text{Opt}(\mathcal{T}, \mathcal{C}) \text{ then } \phi(\boldsymbol{\lambda}^*) \approx w(T^*). \quad (3)$$

We note that solving (1) is equivalent to solving the Linear Program $LP(\mathcal{T})$:

$$\text{Minimise } \sum_{e \in E_n} w(e)x(e) \text{ subject to } \mathbf{x} \in P_{\mathcal{T}}, \sum_{e \in E_n} c_i(e)x(e) \leq C_i, i = 1, \dots, r, \quad x(e) \geq 0, e \in E_n, \quad (4)$$

where $E_n = \binom{[n]}{2}$ and $P_{\mathcal{T}}$ is the convex hull of the incidence vectors of the set \mathcal{T} .

We also note that (3) implies that the relative integrality gap for the integer program $\text{Opt}(\mathcal{T}, \mathcal{C})$ is $(1 + o(1))$ w.h.p.

Next let

$$w_{\boldsymbol{\lambda}}(T) = w(T) + \sum_{i=1}^r \lambda_i c_i(T) \quad \text{for } T \in \mathcal{T}.$$

Let $\mathcal{T}_{\boldsymbol{\lambda}}$ denote the set of trees that minimise $w_{\boldsymbol{\lambda}}$ and $O_{\mathcal{T}}(\boldsymbol{\lambda}) = \{\mathbf{x}(T) : T \in \mathcal{T}_{\boldsymbol{\lambda}}\}$ denote the set of incidence vectors of the trees in $\mathcal{T}_{\boldsymbol{\lambda}}$.

Lemma 3. $|\mathcal{T}_{\boldsymbol{\lambda}}| \leq r + 1$ with probability one.

Proof. First assume that we are using U . Suppose that $\mathcal{T}_{\boldsymbol{\lambda}} = \{T_1, T_2, \dots, T_s\}$ where $s > r + 1$. Let C be the $s \times r$ matrix ($C_{i,j} = c_i(T_j)$). Let C_k be the matrix consisting of the first k rows of C and let \mathbf{w}_k be the column vector ($w(T_j), j = 1, 2, \dots, k$) and let $\mathbf{1}_k$ be the all ones vector of dimension k . Then we have $\phi(\boldsymbol{\lambda})\mathbf{1}_r = \mathbf{w}_r + C_r\boldsymbol{\lambda}$. Now C_r is non-singular with probability one and so $\boldsymbol{\lambda} = C_r^{-1}(\phi(\boldsymbol{\lambda})\mathbf{1}_r - \mathbf{w}_r)$. Rows $r + 1, r + 2$ of the equation $\phi(\boldsymbol{\lambda})\mathbf{1}_{r+2} = \mathbf{w}_{r+2} + C_{r+2}\boldsymbol{\lambda}$ give us two distinct expressions for $\phi(\boldsymbol{\lambda})$. By equating them we find a non-trivial algebraic expression involving $w(T_j), c_i(T_j), i = 1, 2, \dots, r, j = 1, 2, \dots, s$ and such an expression exists with probability zero. The expression implies an explicit value for w_{r+2} , given the other parameters. \square

Remark 2. If we only keep weights/costs to m bit accuracy then this claim has to be modified to be with probability $1 - 2^{-\Omega(m)}$. We can afford to use the union bound over all possible choices of r spanning trees.

The optimum solution to $\text{Opt}(\mathcal{T}, \mathcal{C})$ lies in the face of $P_{\mathcal{T}}$ generated by the incidence vectors of the trees in $\mathcal{T}_{\boldsymbol{\lambda}^*}$. They generate a face because they are the vertex solutions to a linear program. If F is a face of a polytope P and E is an edge of F then E is an edge of P . Now if T_1, T_2 give rise to adjacent vertices of the polytope $P_{\mathcal{T}}$ then $E(T_2) = (E(T_1) \setminus e) \cup \{f\}$ for edges e, f . It then follows from Lemma 3 that we have

Lemma 4. If T_1, T_2 minimise $w_{\boldsymbol{\lambda}}$ then $|E(T_1) \setminus E(T_2)| \leq r$, with probability one.

Next let w_{\max} denote the maximum weight of any edge in any of the trees in $\mathcal{T}_{\boldsymbol{\lambda}}$ and let c_{\max} denote the maximum of any of the costs of any of the edges of any of the trees in $\mathcal{T}_{\boldsymbol{\lambda}}$.

Lemma 5. With probability one, there exists j such that $w(T_j) \leq w^* + rw_{\max}$ and $c_i(T_j) \leq C_i + rc_{\max}$ for $j = 1, 2, \dots, r$.

Proof. Let $\boldsymbol{\lambda} = \boldsymbol{\lambda}^*$ solve $\text{Dual}(\mathcal{T})$ and \mathbf{x}^* solve $LP(\mathcal{T})$. Then we have that \mathbf{x}^* is a convex combination of $\{\mathbf{x}(T) : T \in O_{\mathcal{T}}(\boldsymbol{\lambda})\}$. It follows that there exist $T_0, T_1, \dots, T_r \in \mathcal{T}_{\boldsymbol{\lambda}}$ such that (i) $w(T_0) \leq w^*$ and (ii) $c_i(T_i) \leq C_i, i = 1, 2, \dots, r$. (If $c_i(T_j) > C_i, j \geq 1$ then $\phi(\boldsymbol{\lambda}^* + \boldsymbol{\epsilon}) > \phi(\boldsymbol{\lambda}^*)$ for a sufficiently small perturbation $\boldsymbol{\epsilon}$. This contradicts the fact that $\boldsymbol{\lambda}^*$ maximises f .) \square

It then follows from Lemma 5 that

$$w(T) \leq w^* \text{ and } c_i(T) \leq C_i + r, i = 1, 2, \dots, r. \quad (5)$$

This almost solves our problem, except that T is not guaranteed to be feasible. We show next that a small adjustment to T results in an asymptotically optimal feasible solution.

If $c_i(T) > C_i/2$ then T contains at least $\frac{C_i}{4-C_i/(n-1)}$ edges $e \in X_i = \{e \in T : c_i(e) \geq C_i/4n\}$. Delete $4rn/C_i$ edges of X_i from T , for each i , to create a forest F for which $w(F) < w^*$ and $c_i(F) \leq C_i - r, i = 1, 2, \dots, r$. Now observe that if Γ is the subgraph of K_n spanned by edges e for which

$$w(e) \leq \psi = F^{-1}(n^{-1/(r+1)} \log^{1/r} n) \approx n^{-1/(\alpha(r+1))} \log^{1/(\alpha r)} n, \quad c_i(e) \leq \psi, i = 1, 2, \dots, r$$

then Γ is distributed as $G_{n,p}$ where $p \approx n^{-1} \log^{(r+1)/r} n$. Thus Γ is connected w.h.p. and so we can add $4r^2n/C_{\min}$, ($C_{\min} = \min_{i=1}^r C_i$) edges from Γ to F to make a spanning tree \hat{T} . (The claim that $G_{n,p}$ is connected follows from Erdős and Rényi [2].) We have

$$w(\hat{T}) \leq w^* + \frac{4r^2n\psi}{C_{\min}} \text{ and } c_i(\hat{T}) \leq C_i - r + \frac{4r^2n\psi}{C_{\min}} < C_i, i = 1, 2, \dots, r,$$

as our assumption on the C_i implies that $\frac{n\psi}{C_{\min}} \rightarrow 0$.

Also, to find the s trees, we need only find one tree T and then consider all trees of the form $T + e - f$.

Remark 3. *The argument that leads to (5) is valid for an arbitrary matroid.*

Now consider the claim in Remark 1. Theorem 1.1 of [8] implies the following: let G be an arbitrary δ -regular graph as in Remark 1. If $\delta p - \log n \rightarrow \infty$ then w.h.p. G_p is connected. Here G_p is obtained from G by independently deleting edges with probability $1 - p$.

As just observed we can w.h.p. find a tree T satisfying (5) in polynomial time. If $c_i(T) > C_i/2$ then T contains at least $\frac{C_i}{4-C_i/(n-1)}$ edges $e \in X_i = \{e \in T : c_i(e) \geq C_i/4n\}$. Delete $4rn/C_i$ edges of X_i from T , for each i , to create a forest F for which $w(F) < w^*$ and $c_i(F) \leq C_i - r, i = 1, 2, \dots, r$. Now observe that if Γ is the subgraph of G spanned by edges e for which $w(e) \leq \eta = F^{-1}(\delta^{-1/(r+1)} \log^{1/r} n), c_i(e) \leq \eta, i = 1, 2, \dots, r$ then Γ is distributed as G_p where $p \approx \delta^{-1} \log^{(r+1)/r} n$. Thus Γ is connected w.h.p. and so we can add $4rn/C_i$ edges from Γ to F to make a spanning tree \hat{T} . We have

$$w(\hat{T}) \leq w^* + \frac{4^2n\eta}{C_i} \text{ and } c_i(\hat{T}) \leq C_i - r + \frac{4r^2n\eta}{C_i} < C_i, i = 1, 2, \dots, r,$$

as our assumption on the C_i implies that $\frac{n\eta}{C_i} \rightarrow 0$.

3 Matchings

We analyze the algorithm of [1], but we avoid the enumeration that gives a running time of $n^{O(1/\varepsilon)}$. We will only consider bipartite matchings. Our analysis only uses alternating paths and avoids the use blossoms and so the non-bipartite case is almost identical to the bipartite case.

We let $w_\lambda(M) = w(M) + \lambda(c_1(M) - C_1)$ for $M \in \mathcal{M}$ and $\lambda \geq 0$. We consider the dual problem $Dual(\mathcal{M}, C_1)$:

$$\text{Maximise : } \phi(\lambda), \lambda \geq 0 \text{ where } \phi(\lambda) = \min \{w_\lambda(M) : M \in \mathcal{M}\}. \quad (6)$$

We note that solving (6) is equivalent to solving the Linear program $LP(\mathcal{M})$:

Minimise $\sum_{i,j} w(i,j)x(i,j)$ subject to

$$\sum_{i=1}^n x(i,j) = 1, \forall j \in [n] \text{ and } \sum_{j=1}^n x(i,j) = 1, \forall i \in [n] \text{ and } \sum_{i,j} c_1(i,j)x_{i,j} \leq C_1.$$

$LP(\mathcal{M})$ is a relaxation of $Opt(\mathcal{M}, C_1)$ and so we should assume that its optimal solution is not integral. This would mean that the constraint $c_1(\mathbf{x}) \leq C_1$ is tight at the optimum. Here $\mathbf{x} = (x(i,j), i, j \in [n])$.

Let $\mathcal{M}^*(\lambda)$ denote the members of \mathcal{M} that minimise w_λ .

Lemma 6. $|\mathcal{M}^*(\lambda)| \leq 2$, with probability one.

Proof. First assume that we are using U . Suppose that there are three distinct members of \mathcal{M} that minimise w_λ . This implies that there are three distinct $M_i \in \mathcal{M}, i = 1, 2, 3$ such that $w(M_i) + \lambda c_1(M_i) = C$ where $C = \phi(\lambda) + \lambda C_1$. But this implies, after eliminating C, λ that $\frac{w(M_1) - w(M_2)}{c_1(M_2) - c_1(M_1)} = \frac{w(M_1) - w(M_3)}{c_1(M_3) - c_1(M_1)}$, an event of probability zero.

In the case where we use U_d , given M_1, M_2, M_3 , we see from the previous sentence that this probability is 2^{-m} . There are $(n!)^3$ choices for the perfect matchings and we can use the union bound. \square

So, there exist $M_1, M_2 \in \mathcal{M}$ that satisfy $w(M_1) + \lambda c_1(M_1) = w(M_2) + \lambda c_1(M_2) = \phi_k(\lambda)$.

Let $\mathbf{x}(M)$ denote the n^2 -dimensional $\{0, 1\}$ index vector of matching M and let $P_{\mathcal{M}}$ denote the convex hull of these incidence vectors. The optimum solution to $Opt(\mathcal{M}, C_1)$ lies in the line segment of $P_{\mathcal{M}}$ generated by the incidence vectors of the two matchings minimizing w_{λ^*} . So, if we know at least one of M_1, M_2 and we know the optimum solution to $LP(\mathcal{M})$ then we can construct the other matching. We can find one of M_1, M_2 if we know λ^* . We just have to solve the assignment problem with weights w_{λ^*} . Because $\phi(\lambda)$ is a concave function, we can find λ^* to within accuracy $2^{-poly(n)}$ by solving $poly(n)$ assignment problems. Alternatively, we can read off λ^* from the solution to the dual of $LP(\mathcal{M})$:

$$\text{Maximise } -C_1\lambda + \sum_{i=1}^n u_i + \sum_{j=1}^n v_j \text{ subject to } \lambda \geq 0 \text{ and } -c_1(i,j)\lambda + u_i + v_j \leq w(i,j), \forall i, j.$$

Or

$$\text{Maximise }_{\lambda \geq 0} \left(\text{maximum } \sum_{i=1}^n u_i + \sum_{j=1}^n v_j \text{ subject to } u_i + v_j \leq w_\lambda(i,j) \right).$$

Assume then that we know λ^*, M_1, M_2 . Now we cannot have $C_1 < \min \{c_1(M_1), c_1(M_2)\}$ else $\phi(\lambda^* + \varepsilon) > \phi(\lambda^*)$ for sufficiently small ε . This follows from Lemma 6. Assume then that $c_1(M_1) < C_1 < c_1(M_2)$. We have

$$w(M_1) + \lambda^*(c_1(M_1) - C_1) = w(M_2) + \lambda^*(c_1(M_2) - C_1) \leq w^*, \quad (7)$$

where the inequality come from weak duality.

Let $C = \{e_1, e_2, \dots, e_k\} = M_1 \oplus M_2$. Let $a_i = \delta(e_i)w_{\lambda^*}(e_i)$ where $\delta(e_i) = 1$ for $i \in M_2$ and -1 otherwise. Then

$$w_{\lambda^*}(M_1) - w_{\lambda^*}(M_2) = \sum_{i=1}^k a_i = 0$$

and so there exists ℓ such that

$$\sum_{j=1}^t a_{\ell+j} \leq 0 \text{ for } t = 0, 1, \dots, k-1. \quad (8)$$

This is the content of the gasoline lemma of Lovász [10], Problem 3.21.

For $t \geq 0$ let $X_t = M_1 \cup \{e_{\ell+j} : j \leq t, \ell+j \text{ is odd}\} \setminus \{e_{\ell+j} : j \leq t, \ell+j \text{ is even}\}$. Let $\tau = \max \{t : c(X_t) \leq C_1\}$. Then we must have $\ell + \tau$ even, because the a_{2i-1} are positive and the a_{2i} are negative. Note that $M = X_\tau \setminus \{e_1\}$ is a matching and $c(M) \leq c(X_\tau) \leq C_1$ and that $|M| = n-1$. Note that (8) implies that

$$w_{\lambda^*}(M) \leq w_{\lambda^*}(X_\tau) \leq w_{\lambda^*}(M_1).$$

Now if M^* solves $Opt(\mathcal{M}, C_1)$ then

$$w(M_1) = w_{\lambda^*}(M_1) - \lambda^* C_1 \leq w_{\lambda^*}(M^*) - \lambda^* C_1 \leq w_{\lambda^*}(M^*) - \lambda^* c_1(M^*) = w(M^*)$$

as the M_i minimize w_{λ^*} . So,

$$\begin{aligned} w(X_\tau) = w_{\lambda^*}(X_\tau) - \lambda^* c_1(X_\tau) &= w_{\lambda^*}(X_\tau) - \lambda^* C_1 + \lambda^*(C_1 - c_1(X_\tau)) \leq \\ &w_{\lambda^*}(M_1) - \lambda^* C_1 + \lambda^*(C_1 - c_1(X_\tau)) \leq w^* + \lambda^*(C_1 - c_1(X_\tau)), \end{aligned}$$

where the final inequality is from (7).

Let $f = x_{\ell+\tau+1}$. Then, the maximality of τ implies that $c_1(f) > C_1 - c_1(X_\tau) \geq 0$. So,

$$w(M) \leq w(X_\tau) \leq w^* + \lambda^* c_1(f) \leq w^* + \lambda^*. \quad (9)$$

Furthermore, by construction,

$$c(M) \leq C_1. \quad (10)$$

At this point, we need to do two things. The first is to bound λ^* and the second is to deal with the fact that $|M| = n-1$. The following lemma deals with λ^*

Lemma 7. *There is a constant $D = D(\alpha) > 0$ depending only on α such that*

$$\lambda^* \leq \frac{Dn^{2-1/\alpha}}{C_1^2}, \quad w.h.p.$$

To deal with the second we prove the following lemma:

Lemma 8. *Suppose that M is a matching of size at most $n-1$. Then w.h.p. there is an augmenting path that creates a matching M' with (i) $|M'| = |M|+1$, (ii) $w(M') \leq w(M) + 3n^{-1/3}$ and (iii) $c_1(M') \leq c_1(M) + 3n^{-1/3}$.*

It follows from (9), (10) and these two lemmas that we can w.h.p. find a perfect matching M' such that

$$w(M') \leq w(M^*) + \frac{Dn^{2-1/\alpha}}{C_1^2} + 3n^{-1/3} \text{ and } c_1(M') \leq c_1(M) + 3n^{-1/3}.$$

Now if $c_1(M) + 3n^{-1/3} \leq C_1$ then we are done. Otherwise, we remove an edge from M' of cost at least $3/n^{1/3}$. Such edges exist as we have assumed that $C_1 \gg n^{2/3}$. Applying Lemma 8 again we have a perfect matching M'' satisfying

$$w(M'') \leq w(M^*) + \frac{Dn^{2-1/\alpha}}{C_1^2} + 6n^{-1/3} \text{ and } c_1(M'') \leq C_1.$$

Note that [3], Theorem 3 implies that $w^* = \Omega(n^{1-1/\alpha})$ and this will complete the proof of Theorem 2, since we have assumed that $C_1 \gg n^{1/2}$.

3.1 Proof of Lemma 7

Proof. Here we assume that the weights and costs are i.i.d. copies of a continuous random variable X with $\mathbb{P}(X \leq t) \approx t^\alpha$ as $t \rightarrow 0$ so that the density of X , call it f , satisfies $f(x) \approx \alpha x^{\alpha-1}$ as $x \rightarrow 0$. For a fixed $\lambda > 0$, the density f_Z of $Z = X + \lambda X'$, where X' is an independent copy of X , satisfies

$$f_Z(x) = \int_0^x \lambda^{-1} f(\lambda^{-1}t) f(x-t) dt \approx \lambda^{-\alpha} \alpha^2 \int_0^x [t(x-t)]^{\alpha-1} dt = \lambda^{-\alpha} D_\alpha x^{2\alpha-1}, \quad D_\alpha = \alpha^2 \frac{\Gamma(a+1)^2}{\Gamma(2a+2)},$$

as $x \rightarrow 0$. Thus, $\mathbb{P}(Z \leq t) \approx \frac{D_\alpha}{2\alpha} \lambda^{-\alpha} t^{2\alpha}$, so by the results from [3] (see Theorem 3, the unconstrained case $r = 0$ and Section 6) applied to a rescaled version of Z (so that its CDF behaves like $t^{2\alpha}$), we have w.h.p.,

$$\phi(\lambda) + C_1 \lambda = \min_M w_\lambda(M) = \lambda^{1/2} (2\alpha/D_\alpha)^{1/(2\alpha)} \Theta(n^{1-1/(2\alpha)})$$

Let $\phi_A(\lambda) = A \lambda^{1/2} n^{1-1/(2\alpha)} - C_1 \lambda$ for constant $A > 0$. Thus w.h.p. for some constants $0 < A < B$ depending only of α , we have

$$\phi_A(\lambda) \leq \phi(\lambda) \leq \phi_B(\lambda).$$

Moreover, $\phi_X, X = A, B$ is maximised at $\lambda_X^* = \frac{X^2 n^{2-1/\alpha}}{4C_1^2}$ and then

$$\phi_X(\lambda) = \phi_X(\lambda_X^*) - C_1 \left(\sqrt{\lambda} - \sqrt{\lambda_X^*} \right)^2 \quad \text{where } \phi_X(\lambda_X^*) = \frac{X^2 n}{4C_1^2}. \quad (11)$$

If λ^* maximises ϕ then for $K = O(1)$ we have, using (11),

$$\phi(\lambda^*) \geq \phi(\lambda_A^*) = \frac{A^2 n^{2-1/\alpha}}{4C_1} > \phi_B(K^2 \lambda_B^*) = (1 - (K-1)^2) \frac{B^2 n^{2-1/\alpha}}{4C_1}.$$

It follows that $\lambda^* \leq K^2 \lambda_B$ where $(1 - (K-1)^2) B^2 = A^2/2$. □

3.2 Proof of Lemma 8

Proof. For this we consider the random bipartite graph H which consists of those edges e for which $w(e), c_1(e) \leq n^{-1/3}$. This is distributed as the random bipartite graph $G_{n,n,p}$ where $p = n^{-2/3}$. Suppose that $a \in B_1, b \in B_2$ are the vertices not covered by M . Next let A be the set of vertices in V_2 that can be reached by an alternating path of length at most five. We first observe that the minimum/maximum degree in H is at least $\approx np$ w.h.p. See for example Frieze and Karoński [4], Theorem 3.4. We show next that w.h.p.

$$S \subseteq V_1, n^{1/4} \leq |S| \leq n_0 = \frac{n^{2/3}}{\log n} \text{ implies } |N(S)| \geq np|S|/4. \quad (12)$$

Indeed, if $v \in V_2$ then $\mathbb{P}(v \in N(S)) = 1 - (1-p)^{|S|} \geq p|S|/2$. So,

$$\mathbb{P}(\neg(12)) \leq \sum_{s=n^{1/4}}^{n_0} \binom{n}{s} \mathbb{P}(\text{Bin}(n, ps/2) \leq nps/4) \leq \sum_{s=n^{1/4}}^{n_0} \left(\frac{ne}{s} \right)^s e^{-nps/16} = \sum_{s=n^{1/4}}^{n_0} \left(\frac{ne^{1-np/16}}{s} \right)^s = o(1).$$

Given the property in (12) we see that regardless of M , there are at least $np/2$ alternating paths of length two, ending in V_1 . Then there must be at least $np/2 \times np/4 - 2np \geq n^2 p^2 / 10 > n_0$ alternating paths of length four ending in V_1 . Finally we see that $|A| \geq n_0 np / 4 = \Omega(n / \log n)$ and similarly for $|B|$. We then observe that w.h.p. there is an edge of H connecting A and B . Indeed, the probability there is no such edge is at most $2^{2n} (1-p)^{\Omega(n^2 / \log^2 n)} = o(1)$. It follows that we can convert M to a perfect matching at an additional weight and cost of at most $3n^{-1/3}$. □

References

- [1] A. Berger, V. Bonifac, F. Grandoni and G. Schaefer, Budgeted matching and budgeted matroid intersection via the gasoline puzzle, *Lecture Notes in Computer Science* 5035 (2007) 273-287.
- [2] P. Erdős and A. Rényi, On random graphs I, *Publ. Math. Debrecen* 6 (1959) 290-297.
- [3] A.M. Frieze, W. Pegden, G. Sorkin and T. Tkocz, Minimum-weight combinatorial structures under random cost-constraints.
- [4] A.M. Frieze and M. Karoński, *Introduction to Random Graphs*, Cambridge University Press, 2015.
- [5] A.M. Frieze, and T. Tkocz, A randomly weighted minimum spanning tree with a random cost constraint.
- [6] A.M. Frieze, and T. Tkocz, A randomly weighted minimum arborescence with a random cost constraint
- [7] M. Goemans and R. Ravi, *The constrained minimum spanning tree problem*, Fifth Scandinavian Workshop on Algorithm Theory, LNCS 1097, Reykjavik, Iceland (1996) 66-75.
- [8] A. Gupta, E. Lee and J. Li, The Connectivity Threshold for Dense Graphs.
- [9] F. Grandoni, R. Ravi, M. Singh and R. Zenklusen, New approaches to multi-objective optimization, *Mathematical Programming* 146, (2014) 525-554.
- [10] L. Lovász, *Combinatorial Problems and Exercises*, AMS Chelsea Publishing, 2nd Edition, 2007.
- [11] D. Walkup, On the expected value of a random assignment problem, *SIAM Journal on Computing* 8 (1979) 440-442.