Edge cover and polymatroid flow problems

Martin Hessler
Dept of Mathematics
Linköping University
SE-581 83 Linköping, Sweden
mahes@mai.liu.se

Johan Wästlund
Dept of Mathematical Sciences
Chalmers University of Technology
SE-412 96 Gothenburg, Sweden
wastlund@chalmers.se

November 24, 2010
Mathematics Subject Classification: Primary: 60C05, 90C27, 90C35.

Abstract

In an n by n complete bipartite graph with independent exponentially distributed edge costs, we ask for the minimum total cost of a set of edges of which each vertex is incident to at least one. This so-called minimum edge cover problem is a relaxation of perfect matching.

We show that the large n limit cost of the minimum edge cover is $W(1)^2+2W(1) \approx 1.456$, where W is the Lambert W-function. In particular this means that the minimum edge cover is essentially cheaper than the minimum perfect matching, whose limit cost is $\pi^2/6 \approx 1.645$.

We obtain this result through a generalization of the perfect matching problem to a setting where we impose a (poly)matroid structure on the two vertex-sets of the graph, and ask for an edge set of prescribed size connecting independent sets.

1 Introduction

In 1998 G. Parisi published the two page note [10], with a conjectured exact formula for the so-called *random assignment problem* with exponential edge costs. The conjecture was that if the edges of a complete n by n bipartite graph are given independent random costs of mean 1 exponential distribution, then the minimum total cost of a perfect matching (or *assignment*) has expectation

$$1 + \frac{1}{4} + \frac{1}{9} + \cdots + \frac{1}{n^2}. \tag{1}$$

The asymptotic $\pi^2/6$-limit (conjectured already in [7]) was established in [1, 2], and (1) was subsequently proved and generalized in different directions, see [3, 6, 8, 9, 12, 15].

In the current paper we investigate the *minimum edge cover problem*. The random model is the same, an n by n bipartite graph whose edges are assigned independent exponentially distributed costs. An *edge cover* is a set of edges that covers all vertices, that is, each vertex is incident to at least one edge in the set. The minimum edge cover problem asks for the edge cover of minimum total cost.

A perfect matching is an edge cover, and the minimum cost perfect matching is therefore one of the solutions competing for the minimum edge cover. Therefore (1) is an upper bound on the expected cost of the minimum edge cover. Perfect matchings clearly minimize the number of edges in an edge cover, and therefore one would expect the minimum perfect matching to be a fairly good edge cover. But for $n \geq 3$, it may happen that an edge cover with more than n edges is cheaper than the minimum perfect matching (see Figure 1).

Such situations clearly have positive probability, and therefore the expected cost of the minimum edge cover is strictly smaller than (1) as soon as $n \geq 3$. The question arises whether for large n the minimum edge cover is considerably cheaper than the minimum perfect matching, or if the large n limit cost is still $\pi^2/6$. A related question is whether the number of edges in the minimum edge cover is in general close to n, or if it is more likely close to cn for some $c > 1$. Here c would be the average vertex degree in the minimum edge cover, in other words the average number of edges incident to a particular vertex.

We will answer these questions by establishing the following two theorems.

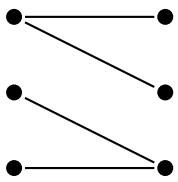

Figure 1: The drawn edges constitute an edge cover that doesn't contain a perfect matching as a subset. If these edges are relatively cheap and all other edges are considerably more expensive, then this will be the minimum edge cover.

Theorem 1.1. *As $n \to \infty$, the cost of the minimum edge cover converges in expectation and in probability to the number*

$$\min_{x \in \mathbf{R}} (x^2 + 2e^{-x}) \approx 1.4559.$$

The number x that minimizes $x^2 + 2e^{-x}$ is the solution to the equation $x = e^{-x}$, which is $W(1) \approx 0.5671$, where W is the Lambert W-function.

Theorem 1.2. *As $n \to \infty$, the average vertex degree in the minimum edge cover converges in expectation and in probability to* $2W(1) \approx 1.1343$.

Our results on the edge cover problem will be established as applications of results on another class of optimization problems (to which paradoxically the edge cover problem does not belong). We therefore start from the seemingly unrelated subject of matroid theory, and later return to the edge cover problem.

Matroids and polymatroids are introduced in Section 2. In Sections 3–4 we describe certain combinatorial optimization problems related to polymatroids. Random models of these *flow problems* are analyzed in Sections 5–9. In Section 10 we discuss a simple but artificial example of a random polymatroid flow problem which not covered by earlier generalizations of (1). In particular it turns out that the edge cover problem with specified number of edges belongs to this class of problems. This allows us in Section 11 to prove Theorems 1.1 and 1.2, roughly speaking by optimizing the number of edges to minimize the limit cost. Finally in Section 12 we discuss some further conjectures.

2 Matroids and polymatroids

In this section we introduce basic concepts and notation for matroids and polymatroids. For a more thorough introduction to matroid theory from the perspective of combinatorial optimization we refer to [5].

A *matroid* on the (finite) ground set A is given by declaring certain subsets of A to be *independent*. The matroid concept can be defined by various "axiom systems". One of the simplest is the following:

1. A subset of an independent set is independent.
2. For every subset $S \subseteq A$, all maximal (with respect to inclusion) independent subsets of S have the same number of elements.

The size of a maximal independent subset of S is called the *rank* of S. The axioms (1) and (2) together imply that there is a unique largest superset of S that has the same rank as S. This set is called the *span* (or *closure*) of S and is denoted $\sigma(S)$. A set is *closed* if it is equal to its own span. Closed sets are also called *subspaces* or *flats* in the literature.

A trivial example of a matroid is a so-called *discrete* matroid where every set is independent. As will become clear, this corresponds to the assignment problem.

The matroid concept was generalized to *polymatroids* by J. Edmonds [4] in 1970. Here we consider only so-called *integral* polymatroids, see [11]. An integral polymatroid on the ground set A is defined by declaring certain multisets of elements of A as independent. A finite multiset S on the ground set A is formally a function $A \to \mathbb{N}$, where $S(a)$ is the multiplicity of a in S. In order to make our algebra of multisets closed under arbitrary unions we allow for infinite multiplicities, and therefore regard multisets as functions $A \to \mathbb{N} \cup \{\infty\}$. If S and T are multisets on A, then we say that S is a subset of T if $S(a) \leq T(a)$ for every $a \in A$. The *size* $|S|$ of S is defined as

$$|S| = \sum_{a \in A} S(a),$$

in other words the number of elements of S counted with multiplicity. The union $S \cup T$ is given by $(S \cup T)(a) = \max(S(a), T(a))$, and this generalizes in the obvious way to infinite unions. We will use the notation $S + a$ to denote the multiset obtained from S by increasing the mutliplicity of a by 1.

An integral polymatroid structure on A is given by declaring certain multisets on A as independent, under the requirement that:

(1′) A multiset is independent if and only if all its finite subsets are independent.

(2′) For every multiset S on A, all maximal (with respect to inclusion) independent subsets of S have the same size.

The *rank* of a multiset is defined in the obvious way as the size of a maximal independent subset. The definition of the *span* of a multiset as the union of all supersets of the same rank also carries over, but notice that the span of a finite multiset can be infinite.

The polymatroid concept introduced by Edmonds [4] is a further generalization to functions $A \to \mathbb{R}_+$, so that a polymatroid can be regarded as a special type of polyhedron in $|A|$ dimensions. In this article, we consider only integral polymatroids, and in the following, we therefore refer to them simply as polymatroids.

The discrete matroid can be generalized to a polymatroid where a multiset is independent if and only if it is a subset of a given multiset C. This corresponds to the class of problems treated in Section 12 of [15], where C describes what [15] refers to as the *capacities* of the vertices.

3 Polymatroid flow problems

We let $\mathbf{A}$ and $\mathbf{B}$ be polymatroids on the ground sets A and B respectively (we have to distinguish in notation between the polymatroid and its ground set since we later consider different polymatroids on the same ground set). We construct an infinite bipartite graph which we simply denote (A, B), with vertex set $A \cup B$. For every pair $(a, b) \in A \times B$, there is a countably infinite sequence of edges $e_1(a, b), e_2(a, b), e_3(a, b), \ldots$ connecting a and b. We let $E = E(A, B)$ denote the set of all these edges.

If $F \subseteq E$ is a set of edges, then the *projections* F_A and F_B of F on A and B respectively are the multisets on A and B that record the number of edges of F incident to each vertex. If F_A and F_B are both independent in $\mathbf{A}$ and $\mathbf{B}$ respectively, then we say that F is a *flow* (with respect to $\mathbf{A}$ and $\mathbf{B}$). A flow consisting of k edges is called a k-flow.

A *cost function* is a function $c : E \to \mathbb{R}$ satisfying

$$0 \leq c(e_1(a, b)) \leq c(e_2(a, b)) \leq c(e_3(a, b)) \leq \ldots$$

for every $a \in A$ and $b \in B$. If a cost function is given, and k is a number such that there exist multisets of rank k in both $\mathbf{A}$ and $\mathbf{B}$, then we can ask for the minimum cost k-flow, that is, the set $F \subseteq E$ that minimizes

$$c(F) = \sum_{e \in F} c(e)$$

under the constraint that F is a k-flow. As is shown in the next section, this combinatorial optimization problem, the *polymatroid flow problem*, can be regarded as a special case of the *weighted matroid intersection problem* [5], for which polynomial time algorithms have been known since the 1970's.

Here it is worth making a couple of remarks on the relation of our new results to earlier work. The papers [6, 8, 12] considered only matching problems, which can be regarded as the special case corresponding to polymatroids where the independent multisets are precisely those that do not have any repeated elements (discrete matroids). In [15] this is generalized to a setting which could, in the current perspective, be called *fixed capacity polymatroids*. These are polymatroids where the independent multisets are precisely the subsets of a given multiset. In other words, each element of the ground set has a given *capacity*, and a multiset is independent if no element has multiplicity exceeding its capacity (in [15] the capacities were assumed to be finite but this is not necessary).

From a matroid-theoretic point of view, these polymatroids have the very special property that each subspace has a unique basis. This in turn implies that the so-called *nesting property* holds at vertex level, that is, the degree of a vertex in the minimum $(r+1)$-flow is at least equal to its degree in the minimum r-flow. It is easy to see that this need not hold for arbitrary polymatroid flow problems. Instead, as we shall see in the next section, the nesting property holds only at subspace level: The span of the projection of the minimum r-flow is a subset of the span of the projection of the minimum $(r+1)$-flow, although the projections themselves may be completely disjoint.

4 Combinatorial properties of the polymatroid flow problem

In this section we establish the necessary combinatorial properties of the polymatroid flow problem. These properties are direct generalizations of

the results of Section 12 of [15]. The result which is new and essentially stronger in the current paper is Theorem 4.1, while Theorems 4.2 and 5.1 are corollaries derived in the same way as in [15].

The polymatroid flow problem can be regarded as a special case of the weighted matroid intersection problem, by treating the edge set E as the ground set of two matroids $M_{\mathbf{A}}$ and $M_{\mathbf{B}}$, where a set of edges is independent in $M_{\mathbf{A}}$ and $M_{\mathbf{B}}$ if the projections on A and B are independent in $\mathbf{A}$ and $\mathbf{B}$ respectively.

Although most of the ingredients of the proof of the following theorem are present in [5], we have not been able to find it in the literature. The theorem is valid when $M_{\mathbf{A}}$ and $M_{\mathbf{B}}$ are arbitrary matroids on the same ground set E even if they do not originate in two polymatroids as in our setting. We assume that each element of E is assigned a real cost (which we may, without loss of generality, assume to be nonnegative). Let $\sigma_{\mathbf{A}}$ and $\sigma_{\mathbf{B}}$ be the closure (span) operators with respect to $M_{\mathbf{A}}$ and $M_{\mathbf{B}}$ respectively. A subset of E which is independent with respect to both $M_{\mathbf{A}}$ and $M_{\mathbf{B}}$ is called a *flow*.

A *circuit* (with respect to a given matroid) is a minimal dependent set, in other words a set which is not independent but where the removal of any element would give an independent set.

Theorem 4.1. *Suppose that $F \subseteq E$ is an r-flow which is not of minimum cost. Then there is an r-flow F' of smaller cost than F which contains at most one element outside $\sigma_{\mathbf{A}}(F)$.*

Proof. We let G be a minimum cost r-flow, and let $H = F \triangle G = (G - F) \cup (F - G)$, in other words the symmetric difference of F and G. We construct a bipartite graph with vertex sets $G - F$ and $F - G$ and edges labeled A and B as follows. For each element e of $G - F$, if $F + e$ contains an $M_{\mathbf{A}}$-circuit, then there are edges labeled A from e to every element of this circuit which is in $F - G$. If $F + e$ contains an $M_{\mathbf{B}}$-circuit, then there are edges labeled B from e to every element of this circuit which is in $F - G$.

A subset of H is *balanced* if it contains equally many elements from F and G. A balanced subset U of H is *complete* if there is a matching of all elements in $U \cap G \cap \sigma_A(F)$ to elements in $U \cap F$ via edges labeled A, and a matching of all elements in $U \cap G \cap \sigma_B(F)$ to elements in $U \cap F$ via edges labeled B. Moreover, we say that U is *improving*, if the total cost of $U \cap G$ is equal to or smaller than the cost of $U \cap F$.

If U is an arbitrary subset of $\sigma_{\mathbf{A}}(F) \cap (G - F)$, then since U cannot be spanned in $M_{\mathbf{A}}$ by fewer than $|U|$ elements, there are at least $|U|$ elements in

$F - G$ that have edges labeled A to U. Hence by the criterion of P. Hall, the edges labeled A contain a matching of all the elements of $\sigma_{\mathbf{A}}(F) \cap (G - F)$ to the elements of $F - G$. Similarly there is a matching of all the elements of $\sigma_{\mathbf{B}}(F) \cap (G - F)$ to $F - G$ via edges labeled B.

If we choose such matchings, then they will split H into a number of paths and cycles. The paths that have one more element from one of F and G than from the other can be paired so that H is split into a number of balanced subsets that for the purpose of this proof we call *pseudocycles*. A pseudo-cycle (with respect to a particular choice of matchings) is a set of nodes that constitutes an alternating cycle except that possibly one edge labeled A and/or one edge labeled B are missing.

Since H can be partitioned into pseudo-cycles, there must be an improving pseudo-cycle U. Since a pseudo-cycle contains at most one element outside $\sigma_A(F)$, it would be sufficient to prove that $F \triangle U$ is a flow. Unfortunately this need not always be the case. However, we will show that if U is an improving pseudo-cycle which is minimal in the sense that no proper subset of U is an improving pseudo-cycle under any (possibly different) choice of matchings, then $F \triangle U$ is a flow. We therefore assume that U is minimal in this sense.

Suppose first that the elements of U involved in the A-matching can be ordered $(f_1, g_1), \ldots, (f_n, g_n)$ (that is, the A-matching pairs $f_i \in F$ with $g_i \in G$) in such a way that there is no edge (f_i, g_j) labeled A for any $i < j$. In this case, we start from F, and for $i = 1, \ldots, n$ in turn add the element g_i and then delete f_i. In each step, the circuit created when adding g_i will be the same as the circuit of $F + g_i$. In particular, the circuit created when adding g_i will contain f_i, so that the deletion of f_i restores independence with respect to A. Possibly U also contains one element outside $\sigma_{\mathbf{A}}(F)$, and one element of $F - G$ which is not used in the A-matching. Adding the extra element outside $\sigma_{\mathbf{A}}(F)$ and deleting the superfluous element of $F - G$ will obviously maintain independence with respect to A.

If such an ordering is impossible, then the reason must be that there is an obstruction in the form of a set of pairs $(f_1, g_1), \ldots, (f_m, g_m)$ of the A-matching so that there are edges labeled A from f_i to g_{i+1} and also from f_m to g_1. We want to show that this contradicts the minimality of U.

We note that the elements of U have a natural cyclic ordering (unrelated to the obstruction) which is obtained by adding the possibly missing edges in the A- and B-matchings.

For each "extra" edge (f_i, g_j), $i \neq j$ of the obstruction, we associate a

smaller pseudo-cycle. We take the edge (f_i, g_j) together with the edges of the A-matching except (f_i, g_i) and (f_j, g_j). Altering the A-matching in this way splits U into at least two components. We now discard the components (possibly the same) that contain g_i and f_j. The remaining elements of U constitute a pseudo-cycle associated with the extra edge (f_i, g_j).

In this way, the obstruction $(f_1, g_1), \ldots, (f_m, g_m)$ gives rise to m smaller pseudo-cycles $U_1, \ldots, U_m$, see Figure 2. Now we observe that for any particular element in U, the number of pseudo-cycles $U_1, \ldots, U_m$ that it belongs to is equal to the number of times that the cyclic sequence $f_1, \ldots, f_m, f_1$ "winds" around the natural cyclic ordering of U, if each extra edge (f_i, g_j) is considered as a walk in the direction such that edges labeled A are traversed from G to F, and edges labeled B are traversed from F to G. In particular all elements of U belong to the same number of sets $U_1, \ldots, U_m$. It follows that at least one of $U_1, \ldots, U_m$ must be improving, contradicting the minimality hypothesis.

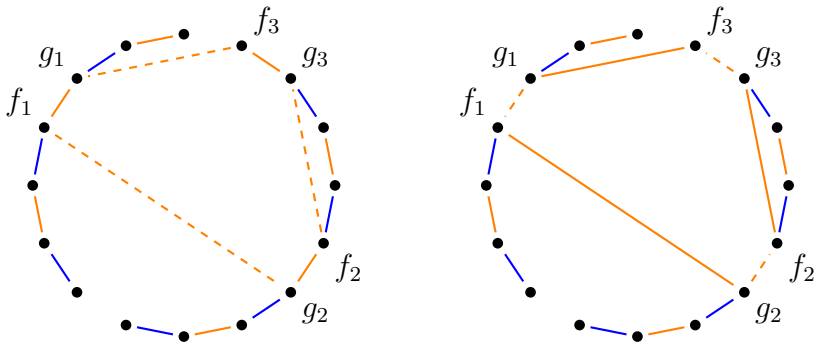

Figure 2: The pseudo-cycle splitting process. Left: A single pseudo-cycle before splitting. The dashed edges contitute the "obstruction". Right: Three pseudo-cycles after splitting. Edges labeled A are drawn in Amber, and edges labeled B are drawn in Blue.

We have shown that the minimality assumption on U implies that $F \triangle U$ is independent with respect to A. By the same argument, it follows that it is independent also with respect to B, and thereby it is a flow. This completes the proof. □

If $b \in B$, then we define the *contraction* $\mathbf{B}/b$ as the polymatroid on the ground set B where a multiset S is independent with respect to $\mathbf{B}/b$ iff $S + b$

is independent in $\mathbf{B}$. The following is a corollary to Theorem 4.1. To simplify the statement we assume that the edge costs are *generic* in the sense that no two different flows have the same cost. In the random model introduced in the next section, the edge costs are independent random variables of continuous distribution, which implies that genericity holds with probability 1.

Theorem 4.2. *Let $b \in B$. Let F be the minimum r-flow in $(\mathbf{A}, \mathbf{B}/b)$, and let G be the minimum $(r+1)$-flow in $(\mathbf{A}, \mathbf{B})$. Then G contains exactly one edge outside $\sigma_{\mathbf{A}}(F)$.*

Proof. In order to apply Theorem 4.1, we introduce an auxiliary element $a^\star$, and let $\mathbf{A}^\star$ be the polymatroid on the ground set $A \cup a^\star$ where a multiset S is independent in $\mathbf{A}^\star$ iff S restrcted to A is independent in $\mathbf{A}$. We let the first edge $(a^\star, b)$ have non-negative cost x, and let all other edges from $a^\star$ have infinite cost. If we put $x = 0$, then the minimum $(r+1)$-flow in $(\mathbf{A}^\star, \mathbf{B})$ consists of the edge $(a^\star, b)$ together with F. If we increase the value of x, then at some point the minimum $(r+1)$-flow in $(\mathbf{A}^\star, \mathbf{B})$ changes to G. If we let x have a value just above this point, so that the minimum $(r+1)$-flow in $(\mathbf{A}^\star, \mathbf{B})$ is G, but no other $(r+1)$-flow in $(\mathbf{A}^\star, \mathbf{B})$ has smaller cost than $F + (a^\star, b)$, then it follows from Theorem 4.1 that G contains exactly one edge outside $\sigma_{\mathbf{A}}(F)$. □

5 The random polymatroid flow problem

Suppose that each element of A and B is given a non-negative *weight*, and the weight of an element a is denoted $w(a)$. A random cost function is chosen by letting the cost $c_i(a, b) = c(e_i(a, b))$ be the ith point of a Poisson point process of rate $w(a)w(b)$ on the positive real numbers. The Poisson processes for different pairs of vertices are all independent. This is the bipartite version of the *friendly model* [15]. We let the random variable

$$C_k(\mathbf{A}, \mathbf{B})$$

denote the minimum cost of a k-flow.

The aim of this article is to obtain methods for computing the distribution of $C_k(\mathbf{A}, \mathbf{B})$. The following is a key theorem that in principle summarizes what we need to know about the random polymatroid flow problem. The characterization of the distribution of $C_k(\mathbf{A}, \mathbf{B})$ in terms of the urn processes on $\mathbf{A}$ and $\mathbf{B}$ described in Section 9 is then deduced by calculus.

As before, let F be the minimum r-flow in $(\mathbf{A}, \mathbf{B}/b)$ and let G be the minimum $(r+1)$-flow in $(\mathbf{A}, \mathbf{B})$. Moreover let a_G be the element of A incident to the unique edge of G which is not in $\sigma_{\mathbf{A}}(F)$.

Theorem 5.1 (Independence theorem)**.** *If we condition on $\sigma_{\mathbf{A}}(F)$ and the cost of F, then the cost of G is independent of a_G, and a_G is distributed on the set*

$$\{a \in A : \sigma_{\mathbf{A}}(F)(a) < \infty\} = \{a \in A : \mathrm{rank}_{\mathbf{A}}(F + a) = r + 1\}$$

with probabilities proportional to the weights.

Proof. We condition on (1) the costs of all edges in $\sigma_{\mathbf{A}}(F)$, and (2) for each $b \in B$, the minimum cost of all edges to b which are not in $\sigma_{\mathbf{A}}(F)$. By Theorem 4.1, we have thereby conditioned on the cost of G. By the memorylessness property of the Poisson process, the endpoint a_G of the edge of G which is not in $\sigma_{\mathbf{A}}(F)$ is still distributed, among the vertices in $\{a \in A : \sigma(F_{\mathbf{A}})(a) < \infty\}$, with probabilities proportional to the weights. □

6 The two-dimensional urn process

The two-dimensional urn process is a different random process which is governed by the same underlying parameters as the random flow problem. Each vertex a is drawn from an urn (and then put back) at times given by a Poisson process of rate $w(a)$ on the positive real numbers. For $x \geq 0$, we let A_x be the multiset that records, for each $a \in A$, the number of times that a has been drawn in the interval $[0, x]$, and similarly B_y records the number of times that $b \in B$ is drawn in the interval $[0, y]$.

For $0 \leq k \leq \min(\mathrm{rank}(\mathbf{A}), \mathrm{rank}(\mathbf{B}))$, we let R_k be the region consisting of all (x, y) in the positive quadrant for which

$$\mathrm{rank}(A_x) + \mathrm{rank}(B_y) < k.$$

This definition of R_k is a straighforward generalization of the one in [15]. We will prove that

$$EC_k(\mathbf{A}, \mathbf{B}) = E(\mathrm{area}(R_k)). \tag{2}$$

We obtain this as a special case of a more general theorem, from which it also follows that

$$\mathrm{var}(C_k(\mathbf{A}, \mathbf{B})) \leq \mathrm{var}(\mathrm{area}(R_k)). \tag{3}$$

7 The normalized limit measure

We use a method that has been used in [13, 14, 15]. We extend the set A by introducing an auxiliary special element $a^\star$ in the same way as in the proof of Theorem 4.2. A multiset S on $A \cup a^\star$ independent with respect to $\mathbf{A}^\star$ iff the restriction of S to A is independent with respect to $\mathbf{A}$.

The idea is to let the weight $w^\star = w(a^\star)$ of $a^\star$ tend to zero. When $w^\star \to 0$, the edges from $a^\star$ become expensive. Somewhat surprisingly, crucial information can be obtained by studying events involving the inclusion of an edge from $a^\star$ in the minimum k-flow. The probability of this is proportional to $w^\star$, and hence tends to zero as $w^\star \to 0$. What is interesting is the constant of proportionality. It is therefore convenient to introduce the *normalized limit measure* $E^\star$. Formally we can regard the parameter $w^\star$ as governing a family of probability measures on the same probability space. If ϕ is a random variable defined on the probability space of cost functions, then we let

$$E^\star(\phi) = \lim_{w^\star \to 0} \frac{1}{w^\star} E\left[\phi\right].$$

Obviously $E^\star$ can be defined on events by identifying an event with its indicator variable. We can informally regard $E^\star$ as a measure, the *normalized limit measure*. This measure can be realized by noting that the probability measure defined by the exponential distribution with rate $w^\star$, scaled up by a factor $1/w^\star$, converges to the Lebesgue measure on the positive real numbers as $w^\star \to 0$.

Hence the normalized limit measure can be obtained as follows: We first define the measure space of cost functions on the edges from the auxiliary vertex $a^\star$. This measure space is the set of all assignments of costs to these edges such that exactly one edge from $a^\star$ has nonnegative real cost, and the remaining edges from $a^\star$ have cost $+\infty$. For each $b \in B$, the cost functions for which the first edge $e_1(a^\star, b)$ has finite cost are measured by $w(b)$ times Lebesgue measure on the positive reals. The measure space of all cost functions for the edges from $a^\star$ is the disjoint union of these spaces over all $b \in B$. The measure $E^\star$ is then the product of this space with the probability space of cost functions on the ordinary edges given by the independent Poisson processes as described earlier.

8 A recursive formula

In this section we prove the analogues of Lemma 12.4 and Proposition 12.5 of [15]. This gives a recursive formula for all moments of the cost of the polymatroid flow problem. The calculation is in principle the same as in Section 12 of [15], but there are some differences in the details related to the remarks at the end of Section 6, and we therefore outline the steps of the calculation even though it means repeating some of the arguments of [15].

If T is a subspace with respect to $\mathbf{A}$, then we let $I_k(T, \mathbf{A}, \mathbf{B})$ be the indicator variable for the event that T is a subset of the span of the projection on A of the minimum k-flow with respect to $(\mathbf{A}, \mathbf{B})$.

Lemma 8.1. *Let N be a positive integer, and let S be a rank $k-1$ subspace of $\mathbf{A}$. For every $b \in B$, let I_b be the indicator variable for the event that the minimum k-flow with respect to $(\mathbf{A}^\star, \mathbf{B})$ contains an edge $(a^\star, b)$ and that the projection on A of the remaining edges spans S. Then for every $b \in B$,*

$$E\left[C_k(\mathbf{A}, \mathbf{B})^N \cdot I_{k-1}(S, \mathbf{A}, \mathbf{B}/b)\right] =$$
$$E\left[C_{k-1}(\mathbf{A}, \mathbf{B}/b)^N \cdot I_{k-1}(S, \mathbf{A}, \mathbf{B}/b)\right] + \frac{N}{w(b)} E^\star\left(C_k(\mathbf{A}^\star, \mathbf{B})^{N-1} \cdot I_b\right). \quad (4)$$

Proof. We compute $N/w(b) \cdot E^\star\left(C_k(\mathbf{A}^\star, \mathbf{B})^{N-1} \cdot I_b\right)$ by integrating over the cost, which we denote by t, of the first edge between $a^\star$ and b. We therefore condition on the costs of all other edges.

The density of t is $w^\star e^{-w^\star t}$, and we therefore get the normalized limit by dividing by $w^\star$ and instead computing the integral with the density $e^{-w^\star t}$. For every t this tends to 1 from below as $w^\star \to 0$, and by the principle of dominated convergence, we can interchange the limits and compute the integral using the density 1. This is the same thing as using the normalized limit measure.

We have

$$\frac{d}{dt}\left(C_k(\mathbf{A}^\star, \mathbf{B})^N \cdot I_{k-1}(S, \mathbf{A}, \mathbf{B}/b)\right) = N \cdot C_k(\mathbf{A}^\star, B)^{N-1} \cdot I_b.$$

According to the normalized limit measure, since we are conditioning on the edge $(a^\star, b)$ being the one with finite cost, $E^\star$ is just $w(b)$ times Lebesgue measure. Hence the statement follows by the fundamental theorem of calculus, since if we put $t = \infty$ we get

$$C_k(\mathbf{A}^\star, \mathbf{B})^N \cdot I_{k-1}(S, \mathbf{A}, \mathbf{B}/b) = C_k(\mathbf{A}, \mathbf{B})^N \cdot I_{k-1}(S, \mathbf{A}, \mathbf{B}/b),$$

while if we put $t = 0$, we get

$$C_k(\mathbf{A}^\star, \mathbf{B})^N \cdot I_{k-1}(S, \mathbf{A}, \mathbf{B}/b) = C_{k-1}(\mathbf{A}, \mathbf{B}/b)^N \cdot I_{k-1}(S, \mathbf{A}, \mathbf{B}/b).$$

□

If T is a subspace of a given polymatroid, and $S \subseteq T$, then we let $T - S$ denote the set of elements a of the ground set which have the property that $S + a$ is a subset of T and has higher rank than S. We write $\mathbf{B} - 0$ to denote the set of $b \in B$ for which $\{b\}$ has rank 1.

Theorem 8.2. *Let T be a rank k subspace of $\mathbf{A}$. Then*

$$\begin{aligned} &E\left[C_k(\mathbf{A}, \mathbf{B})^N \cdot I_k(T, \mathbf{A}, \mathbf{B})\right] = \\ &\sum_{\substack{S \subseteq T \\ \text{rank}(S) = k-1}} \frac{w(T-S)}{w(\mathbf{A}-S)} \cdot \sum_{b \in \mathbf{B}-0} \frac{w(b)}{w(\mathbf{B}-0)} E\left[C_{k-1}(\mathbf{A}, \mathbf{B}/b)^N \cdot I_{k-1}(S, \mathbf{A}, \mathbf{B}/b)\right] \\ &+ \frac{N}{w(\mathbf{B}-0)} \cdot \sum_{\substack{S \subseteq T \\ \text{rank}(S) = k-1}} \frac{w(T-S)}{w(\mathbf{A}-S)} \cdot E^\star\left(C_k(\mathbf{A}^\star, \mathbf{B})^{N-1} \cdot I_k(S + a^\star, \mathbf{A}^\star, \mathbf{B})\right), \end{aligned} \tag{5}$$

where the summations are taken over all S which are rank $k-1$ subspaces of T.

Proof. We multiply both sides of (4) by $w(b)$ and sum over all $b \in \mathbf{B} - 0$. This way we obtain

$$\begin{aligned} \sum_{b \in \mathbf{B}-0} w(b) \cdot E\left[C_k(\mathbf{A}, \mathbf{B})^N \cdot I_{k-1}(S, \mathbf{A}, \mathbf{B}/b)\right] = \\ \sum_{b \in \mathbf{B}-0} w(b) \cdot E\left[C_{k-1}(\mathbf{A}, \mathbf{B}/b)^N \cdot I_{k-1}(S, \mathbf{A}, \mathbf{B}/b)\right] \\ + N \cdot E^\star\left(C_k(\mathbf{A}^\star, \mathbf{B})^{N-1} \cdot I_k(S + a^\star, \mathbf{A}^\star, \mathbf{B})\right). \end{aligned} \tag{6}$$

We now use Theorem 5.1. Suppose that T is a rank k subspace of $\mathbf{A}$, and that $b \in \mathbf{B} - 0$. If the projection on A of the minimum k-flow with respect to $(\mathbf{A}, \mathbf{B})$ spans T, then the projection of the minimum $(k-1)$-flow in $(\mathbf{A}, \mathbf{B}/b)$ must span a rank $k-1$ subspace of T. By summing over the possible subspaces, we obtain

$$E\left[C_k(\mathbf{A},\mathbf{B})^N \cdot I_k(T,\mathbf{A},\mathbf{B})\right] = \sum_{\substack{S\subseteq T \\ \operatorname{rank}(S)=k-1}} \frac{w(T-S)}{w(\mathbf{A}-S)} \cdot E\left[C_k(\mathbf{A},\mathbf{B})^N \cdot I_{k-1}(S,\mathbf{A},\mathbf{B}/b)\right]. \quad (7)$$

Now we choose b randomly according to the weights, in other words, we multiply (7) by $w(b)/w(\mathbf{B}-0)$ and sum over all $b \in \mathbf{B}-0$. This leaves the left hand side intact, and we get

$$E\left[C_k(\mathbf{A},\mathbf{B})^N \cdot I_k(T,\mathbf{A},\mathbf{B})\right] = \sum_{\substack{S\subseteq T \\ \operatorname{rank}(S)=k-1}} \frac{w(T-S)}{w(\mathbf{A}-S)} \cdot \sum_{b\in\mathbf{B}-0} \frac{w(b)}{w(\mathbf{B}-0)} \cdot E\left[C_k(\mathbf{A},\mathbf{B})^N \cdot I_{k-1}(S,\mathbf{A},\mathbf{B}/b)\right]. \quad (8)$$

Now we use equation (6) divided by $w(\mathbf{B}-0)$ to rewrite the right hand side of (8). This establishes the theorem.

□

9 The higher moments in terms of the urn process

The recursion in Theorem 8.2 is the same as Proposition 12.5 of [15]. Therefore the moments of $C_k(\mathbf{A},\mathbf{B})$ can be expressed in terms of the urn process in complete analogy with the results in [15]. We will not repeat the proof of this here since the proof in [15] goes through essentially word by word, but we describe the conclusion.

We have already described the urn process: The elements of A and B are drawn from an urn independently with rates given by their respective weights. In the *extended urn process* of degree N there are moreover N points $(x_1,y_1),\ldots,(x_N,y_N)$ which are measured according to Lebesgue measure on the positive quadrant. The extended urn process therefore cannot be treated as a random process. In analogy with the normalized limit measure we use the notation $E^\star$ for the measure of a set of outcomes of the extended urn process,

since somehow it corresponds to introducing N new vertices of infinitesimal weight on each side of the urn process, and normalizing the probabilities of events involving these points.

Define the rank of the variables as

$$\mathrm{rank}_A(x) = \mathrm{rank}(A_x) + |\{i : x_i \leq x\}|,$$

and similarly $\mathrm{rank}_B(y)$. Then we have in analogy with Theorem 12.6 of [15]:

Theorem 9.1.

$$E(C_k(\mathbf{A}, \mathbf{B})^N) = E^\star(\text{for all } i = 1, \ldots, N,\ \mathrm{rank}_A(x_i) + \mathrm{rank}_B(y_i) \leq k + N).$$

It is worth describing, for $N = 1$ and $N = 2$, what this means in terms of the region R_k defined in Section 6. It is quite easy to see that for $N = 1$, the condition in the right hand side of (9) says precisely that the point (x_1, y_1) lies inside R_k. This establishes (2).

For $N = 2$, the condition in (9) is satisfied whenever (x_1, y_1) and (x_2, y_2) both belong to R_{k-1}, and on the other hand cannot be satisfied unless both belong to R_k. This together with (9) implies that (3) holds. If both points belong to R_k but not both belong to R_{k-1}, then the condition is satisfied if the points lie in *decreasing position*, that is, if one of them has greater x-coordinate and the other has greater y-coordinate.

10 The Fano-matroid flow problem

A favourite example of a matroid is the Fano-matroid. In this section we calculate the expectation and variance of the cost of the somewhat artificial but simple matroid flow problem with two Fano matroids. The Fano matroid can be defined as the linear independence structure of the seven nonzero elements of a vector space of dimension 3 over the field $GF(2)$ of two elements. It can be illustrated as in Figure 3. Any set of two distinct elements is independent, and three elements are independent if and only if they are not one of the seven "lines", one of which must be drawn as a circle.

The bipartite graph whose vertex sets are the elements of the two Fano matroids is shown in Figure 4, together with one of the solutions to the 3-flow problem.

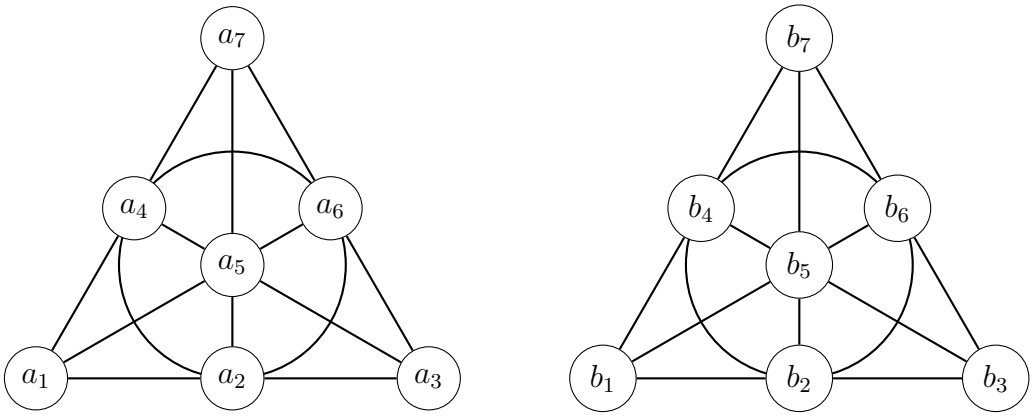

Figure 3: Two Fano matroids.

We assign exponential random variables of rate 1 to the edges and consider the case $k = 3$, that is, we ask for three edges such that the corresponding vertex sets are independent in the two matroids. We calculate the expectation and variance of the cost using Theorem 9.1.

Let w_0 be the time we have to wait until the first vertex is drawn in the urn process along the x-axis. This is a rate 7 exponential random variable. Then let w_1 be the time from w_0 until a second vertex is drawn. Since 6 vertices remain, w_1 is exponential of rate 6 and independent of w_0. Finally let w_2 be the time from the appearance of the second vertex until a vertex independent from the first two is drawn. Then w_2 is a rate 4 exponential variable and independent of w_0 and w_1.

The calculation of the first and second moments can be thought of as first placing (x_1, y_1) in a point in the positive quadrant with rank at most 2, counting only the points drawn in $\mathbf{A}$ and $\mathbf{B}$, and then for the calculation of the second moment placing (x_2, y_2) with rank at most 3, also taking into account the first point as depicted in Figure 5. By symmetry we could have started with the second point. We can therefore assume that the points are placed in one of two types of positions. For the first type the points are placed so that $x_1 < x_2$, $y_1 > y_2$ and with rank at most 2. For the second type the points are placed so that $x_1 < x_2$, $y_1 < y_2$ and with rank at most 1. In both these cases the rank is defined by counting only the points drawn in the urn-processes. The two urn-processes are independent (and therefore uncorrelated). Hence we get

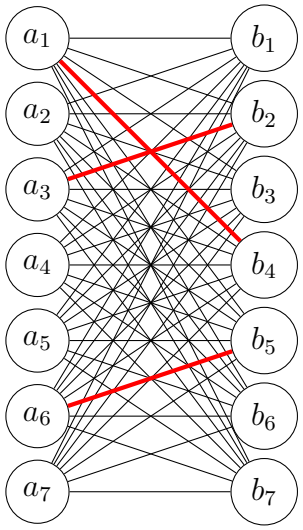

Figure 4: The complete bipartite graph whose vertex sets are the elements of the two Fano matroids. The red edges constitute one of the solutions to the 3-flow problem.

$$EC_3 = E(\text{Area}(R_3)) = E(w_0(w_0' + w_1' + w_2') + w_1(w_0' + w_1') + w_2 w_0') = \\ Ew_0(Ew_0' + Ew_1' + Ew_2') + Ew_1(Ew_0' + Ew_1') + Ew_2 Ew_0' = \\ \frac{1}{7} \cdot \left(\frac{1}{7} + \frac{1}{6} + \frac{1}{4}\right) + \frac{1}{6} \cdot \left(\frac{1}{7} + \frac{1}{6}\right) + \frac{1}{4} \cdot \frac{1}{7} = \frac{295}{1764}.$$

We can find the second moment using the index i_1 to denote the rank of the first point in the urn process on $\mathbf{A}$ and j_1 for the rank in the urn process on $\mathbf{B}$, and similarly i_2 and j_2 for the second point. We get

$$EC_3^2 = 2E\Big(\sum_{\substack{0 \le i_1 \le i_2 \\ 0 \le j_2 \le j_1 \\ i_1 + j_1 \le 2 \\ i_2 + j_2 \le 2}} w_{i_1} w_{i_2} w_{j_1} w_{j_2} + \sum_{\substack{0 \le i_1 \le i_2 \\ 0 \le j_1 \le j_2 \\ i_2 + j_2 \le 1}} w_{i_1} w_{i_2} w_{j_1} w_{j_2} \Big) = \frac{28589}{777924}.$$

The variance is therefore

$$\text{var}(C_3) = \frac{27331}{3111696}$$

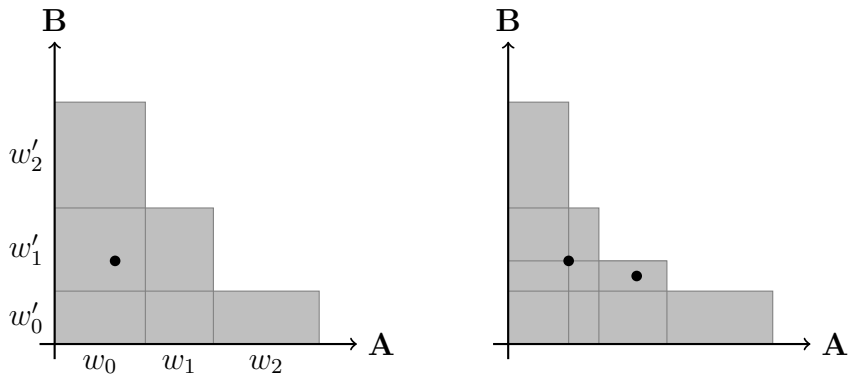

Figure 5: The Fano-matroid urn process. Left: The point (x_1, y_1) has to be within the shaded region. Right: Given the position of (x_1, y_1), the point (x_2, y_2) has to be within the shaded region.

and the standard deviation approximately 0.0937.

11 The minimum edge cover

Finally we return to the edge cover problem described in the introduction. We let the graph have n by n vertices of weight 1 (so that the cost of the cheapest edge between each pair of vertices has mean 1 exponential distribution). Recall that an edge cover is a set of edges such that each vertex is incident to at least one. The minimum edge cover obviously contains at most one edge between any pair of vertices, which means that when establishing Theorems 1.1 and 1.2 we may consider the friendly model with multiple edges as described in Section 5.

The minimum edge cover problem cannot be obtained as a polymatroid flow problem, since it is *non-uniform* in the sense that there are potentially optimal solutions with different numbers of edges. However, if we specify the number $k \geq n$ of edges, then the problem takes the form of a polymatroid flow problem.

The *k-edge cover problem* asks for an edge cover of exactly k edges. Clearly we must have $k \geq n$ in order for the problem to be solvable. The case $k = n$ is the assignment problem, corresponding to the polymatroid where a multiset is independent precisely if it has no repeated element.

Lemma 11.1. *For every k, the k-edge cover problem can be formulated as a polymatroid flow problem.*

Proof. The definition of independent multiset is obvious: A multiset of vertices is independent iff it is a subset of a multiset of size k that contains every element at least once. We have to verify that this class of multisets satisfies the axioms $(1')$ and $(2')$ for a polymatroid. Obviously $(1')$ holds. To verify $(2')$, let S be an arbitrary multiset. Clearly S is independent iff

$$|S| + \#\{a : S(a) = 0\} \leq k.$$

If on the other hand $|S| + \#\{a : S(a) = 0\} > k$, then a maximal independent subset S' of S can have $S'(a) = 0$ only if $S(a) = 0$, because if S' is independent and $S'(a) = 0$, then $S' + a$ is also independent. Therefore, assuming S is not independent, S' is a maximal independent subset of S if and only if $|S'| = k - \#\{a : S'(a) = 0\} = k - \#\{a : S(a) = 0\}$. This shows that all maximal independent subsets have the same size, which verifies $(2')$. □

The proof also shows that the rank of S is $k - \#\{a : S(a) = 0\}$ unless S is independent, in other words

$$\operatorname{rank}(S) = \min\left(|S|, k - \#\{a : S(a) = 0\}\right). \tag{9}$$

Hence we can find the expected cost of the minimum k-edge cover by studying the two-dimensional urn-process. We let k scale with n like $k = \alpha n + O(1)$ for a constant $\alpha \geq 1$, and determine the limit shape of the region R_k as a function of α. It can be shown using the techniques in [15] that limits can be interchanged in the sense that the limit cost is the same as the area of the limit shape.

Consider the urn process on one side of the graph. Vertices are drawn as the events of a rate 1 Poisson point process on the positive real numbers, and we want to determine the relative rank (rescaled by a factor n) of the process at time x. This is obtained by rescaling (9) by a factor n. The total number of vertices drawn up to time x (counting multiplicities and without regard to independence) is xn, and the number of vertices that have not yet been drawn even once is ne^{-x} (the random fluctuations are $o(n)$ for large n). Therefore the relative rank is given by

$$\min(x, \alpha - e^{-x}).$$

The limit region $R(\alpha)$ consists of the points (x, y) in the positive quadrant that satisfy

$$\min(x, \alpha - e^{-x}) + \min(y, \alpha - e^{-y}) \leq \alpha.$$

This is the same thing as the points that satisfy at least one of the four inequalities

$$x + y \leq \alpha \tag{10}$$
$$x \leq e^{-y} \tag{11}$$
$$y \leq e^{-x} \tag{12}$$
$$e^{-x} + e^{-y} \geq \alpha \tag{13}$$

Next consider the dynamics of $R(\alpha)$ as a function of α, see Figure 6. The two equations (11) and (12) are fixed in the sense that they do not depend on α. The curves where equality holds intersect in the point (t, t) where $t = W(1)$ is the solution to $t = e^{-t}$. The inequality (10) will add more points to $R(\alpha)$ only if it can hold when neither of (11) and (12) holds, and it is easy to see that this is the case only when the point (t, t) lies to the lower left of the line $x + y = \alpha$, in other words when $\alpha \geq 2t$.

The inequality (13) too adds points to $R(\alpha)$ only when the point (t, t) is to its lower left, in other words when $2e^{-t} \geq \alpha$, but this happens when $\alpha \leq 2t$.

We can therefore summarize the dynamics as follows: At $\alpha = 1$, $R(\alpha)$ is given by the boundary $e^{-x} + e^{-y} = 1$ (this is just the assignment problem). For $\alpha > 1$, $R(\alpha)$ has a fixed part consisting of the union of the two regions given by (11) and (12). The boundary of the fixed part has a cusp at the point (t, t). For α only slightly greater than 1, $R(\alpha)$ is the fixed part together with a region satisfying (13) which covers the cusp, and this region decreases as α increases. At $\alpha = 2t$, $R(\alpha)$ reaches a minimum consisting only of its fixed part, and here a phase transition occurs. For $\alpha > 2t$, $R(\alpha)$ consists of the fixed part together with the points below the straight line given by (10), and now it increases with α.

We let $g(\alpha)$ be the area of $R(\alpha)$. It follows from our analysis that $g(1) = \pi^2/6$, that g reaches a minimum at $g(2t) = t^2 + 2t$, and that $g(\alpha)$ is increasing for $\alpha \geq 2t$. Numerically the minimum is

$$g(1.134286581) \approx 1.455938093.$$

We notice that the minimum value can also be described as

$$\min_{x \in \mathbf{R}}(x^2 + 2e^{-x}),$$

since

$$\frac{d}{dx}(x^2 + 2e^{-x}) = 0$$

leads to $x = e^{-x}$.

The following theorem relies on making standard estimates of the behaviour of the urn process. What we need is first to allow for the interchange of limits concluding that the limit expected area is the same thing as the area of the limit region. Secondly we need to show that the variance of the cost tends to zero, which follows by showing that the variance of the area of R_n tends to zero. These estimates can be done in several ways, one of which is described in [15]. We do not repeat these arguments here.

Theorem 11.2. *For fixed α, the minimum cost of a vertex cover that contains exactly $\lfloor \alpha n \rfloor$ edges converges in probability to the area $g(\alpha)$ of $R(\alpha)$ as $n \to \infty$.*

The following observation now allows us to rigorously draw several conclusions about the (unrestricted size) edge cover problem:

Theorem 11.3. *For a fixed n and a fixed instance of the edge costs, the minimum cost C_k of an edge cover containing exactly k edges is a convex function of k.*

Proof. Consider two edge covers σ_1 and σ_2 of size $k-1$ and $k+1$ respectively. By a standard argument the symmetric difference of σ_1 and σ_2 can be split into alternating cycles and paths, and by switching a path of odd length whose first and last edges belong to σ_2, we obtain two edge covers of size k whose costs have the same sum as the costs of σ_1 and σ_2. One of them must have a cost which is at most the mean of the costs of σ_1 and σ_2. □

This rules out the possibility that, for a given instance of the graph, C_k is close to $g(k/n)$ for most k in (for instance) the range $n \leq k \leq 2n$ but far from it at a few exceptional values of k. More precisely:

Theorem 11.4.

$$\max_{n \leq k \leq 2n} |C_k - g(k/n)| \xrightarrow{\mathrm{p}} 0,$$

as $n \to \infty$.

From this we conclude that the minimum cost C_k of an edge cover without size constraint is given by the area of the fixed part of $R(\alpha)$:

Theorem 11.5. *As $n \to \infty$,*

$$\min_k C_k \xrightarrow{p} t^2 + 2t = \min_{x \in \mathbf{R}}(x^2 + 2e^{-x}) \approx 1.455938093.$$

Moreover, the average degree of a vertex in the solution converges to $2t \approx 1.134286581$.

A further consequence is that if $\alpha < 2t$ is fixed and $k = \lfloor \alpha n \rfloor$, then asymptotically almost surely, $C_{k-1} < C_k$, and consequently all components in the minimum edge cover of size k are stars.

It seems likely that as soon as $\alpha > 2t$, there will occur components in the solution that are not stars. Moreover, it also seems clear that at some higher value of α, there occurs a giant component in the solution. Here it is worth commenting on the behaviour of the limit cost $g(\alpha)$ at the phase transitions. One can show that $g(\alpha)$ is not analytic at the point $\alpha = 2t$. The easiest way to see this is probably by observing that the function h defined by $h(\alpha) = g(\alpha)$ for $\alpha > 2t$, and extended by analytic continuation to $1 \leq \alpha \leq 2t$, satisfies $h(1) = 3/2 \neq \pi^2/6$: When $\alpha = 1$, the line $x + y = \alpha$ connects the points $(0, 1)$ and $(1, 0)$, and with the natural interpretation of h in terms of areas, one finds that $h(1)$ is the sum of the areas of the regions $y \leq e^{-x}$ and $x \leq e^{-y}$ minus the area of the triangle $x+y \leq 1$. Since $g(1) = \pi^2/6$, g cannot be analytic.

At the point where there occur components that are not stars, the limit cost therefore clearly undergoes a phase transition. Here the structure of the solution changes in a way that can be observed locally. On the other hand at the later point where a giant component occurs, nothing in particular happens to the limit cost $g(\alpha)$. We speculate that this is part of a general pattern where phase transitions that are only visible on a global scale will not cause any dramatic effects on the cost of the optimum solution, but we currently cannot be more precise.

12 The outer-corner conjecture and the giant component

The set of outer corners of the region R_k describes a feasible solution to the optimization problem in an obvious way: Each outer corner corresponds to

the times at which a vertex is drawn in each of the two urn processes, and by putting an edge between all such pairs, we get an edge set, in this case an edge cover of the specified size.

The *outer-corner conjecture* is the conjecture that the probability distribution on feasible solutions defined in this way from the urn process is the same as the one obtained as the optimum solution in the graph with random edge costs. This conjecture has been verified in a number of special cases.

We show here that the outer-corner conjecture predicts exactly at which point the giant component occurs. Fix $\alpha > 2t$, and consider the solution obtained from the outer corners of R_k. We are interested in the expected size of the component containing a randomly chosen vertex, and in particular whether or not this expected size remains bounded as $n \to \infty$.

The break-point at which we stop accepting occurrences of vertices that have already been drawn lies at the intersection of the curves $x + y = \alpha$ and $y = e^{-x}$, in other words at the point x for which $x + e^{-x} = \alpha$. The outer corners that lie in the "tails" of R_k correspond to leafs in the k-edge cover, and therefore cannot contribute to a long path. The only edges that are relevant for the existence of a giant component are those that correspond to outer corners in the linear part of the boundary of R_k. The number of times that a vertex is drawn in the interval corresponding to the linear part of the boundary is Poisson distributed with mean equal to the length of this interval. Thus if we pick a vertex uniformly and explore its component in the solution, this is for large n approximately a Galton-Watson process with Poisson distributed offspring. The question of the existence of a giant component therefore reduces to the question whether the linear part of the boundary of R_k corresponds to an interval on the x-axis of length greater than 1.

The break-point in the urn process along the x-axis occurs at the x that solves $x + e^{-x} = \alpha$. The corresponding break-point in the urn process along the y-axis corresponds to the point x of the intersection of the line $x + y = \alpha$ with the curve $x = e^{-y}$, which occurs when x solves $x = e^{x-\alpha}$. The question is therefore for which α the difference between the solutions to these two equations is equal to 1. This occurs when $\alpha = 2W(1/e) + 1$, where W is the Lambert W-function defined as the inverse to the function We^W. Numerically this gives $\alpha \approx 1.556929086$.

References

[1] Aldous, David, *Asymptotics in the random assignment problem*, Probab. Theory Relat. Fields, **93** (1992) 507–534.

[2] Aldous, David, *The $\zeta(2)$ limit in the random assignment problem*, Random Structures & Algorithms **18:4** (2001), 381–418.

[3] Buck, Marshall W., Chan, Clara S. and Robbins, David P., *On the Expected Value of the Minimum Assignment*, Random Structures & Algorithms **21** (2002), no. 1, 33–58.

[4] Edmonds, J., *Submodular functions, matroids and certain polyhedra*, Proc. Int. Conf. on Combinatorics, Calgary 1970, Gordon and Breach, New York, 69–87.

[5] Lawler, Eugene, Combinatorial Optimization: Networks and Matroids, Holt, Rinehart and Winston 1976.

[6] Linusson, S. and Wästlund, J., *A proof of Parisi's conjecture on the random assignment problem*, Probab. Theory Relat. Fields **128** (2004), 419–440.

[7] Mézard, Marc and Parisi, Giorgio, *Replicas and optimization*, Journal de Physique Lettres **46** (1985), 771–778.

[8] Nair, Chandra, Prabhakar, Balaji and Sharma, Mayank, *Proofs of the Parisi and Coppersmith-Sorkin random assignment conjectures*, Random Structures & Algorithms **27:4** (2005), 413–444.

[9] Nair, Chandra, *Proofs of the Parisi and Coppersmith-Sorkin conjectures in the finite random assignment problem*, PhD thesis, Stanford University 2005.

[10] Parisi, Giorgio, *A conjecture on random bipartite matching*, arXiv:cond-mat/9801176, 1998.

[11] Welsh, D. J. A., *Matroid Theory*, Academic Press, London 1976.

[12] Wästlund, J., *A Proof of a Conjecture of Buck, Chan and Robbins on the Expected Value of the Minimum Assignment*, Random Structures & Algorithms **26:1–2** (2005) 237–251.

[13] Wästlund, J., *Random matching problems on the complete graph*, Electronic Communications in Probability **13** (2008), 258–265.

[14] Wästlund, J., *An easy proof of the zeta(2) limit in the assignment problem*, Electronic Communications in Probability, **14** (2009), 261–269.

[15] Wästlund, J., *The Mean Field Traveling Salesman and Related Problems*, Acta Mathematica **204:1** (2010), 91–150.

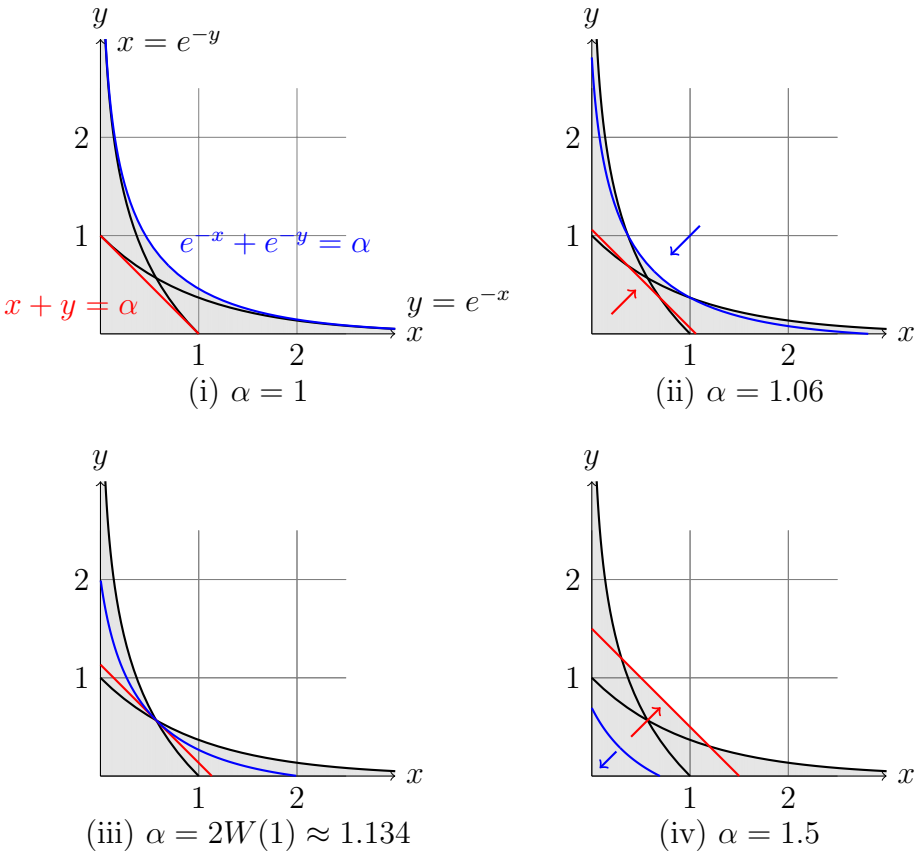

Figure 6: The dynamics of the shaded region $R(\alpha)$: (i) At $\alpha = 1$, the boundary of $R(\alpha)$ is given by the blue curve $e^{-x} + e^{-y} = 1$ and the area is $\pi^2/6$. (ii) The area of $R(\alpha)$ decreases with α as the curve $e^{-x}+e^{-y} = \alpha$ moves towards the origin. (iii) At $\alpha = 2W(1)$, the four curves meet in the point $(W(1), W(1))$ and the area of $R(\alpha)$ reaches its minimum. (iv) For larger α, the area of $R(\alpha)$ increases as the red curve $x + y = \alpha$ moves up and to the right.

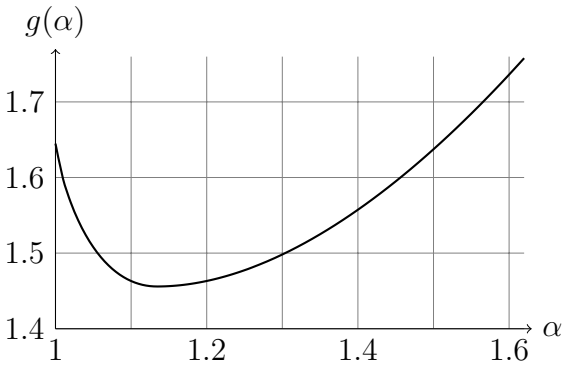

Figure 7: The limit cost function $g(\alpha)$.