


A Randomized Rounding Approach for DAG Edge Deletion

Sina Kalantarzadeh 

University of Waterloo, Canada

Nathan Klein 

Boston University, MA, USA

Victor Reis 

Microsoft Research, Seattle, WA, USA

Abstract

In the DAG Edge Deletion problem, we are given an edge-weighted directed acyclic graph and a parameter k , and the goal is to delete the minimum weight set of edges so that the resulting graph has no paths of length k . This problem, which has applications to scheduling, was introduced in 2015 by Kenkre, Pandit, Purohit, and Saket. They gave a k -approximation and showed that it is UGC-Hard to approximate better than $\lfloor 0.5k \rfloor$ for any constant $k \geq 4$ using a work of Svensson from 2012. The approximation ratio was improved to $\frac{2}{3}(k+1)$ by Klein and Wexler in 2016.

In this work, we introduce a randomized rounding framework based on distributions over vertex labels in $[0, 1]$. The most natural distribution is to sample labels independently from the uniform distribution over $[0, 1]$. We show this leads to a $(2 - \sqrt{2})(k+1) \approx 0.585(k+1)$ -approximation. By using a modified (but still independent) label distribution, we obtain a $0.549(k+1)$ -approximation for the problem, as well as show that no independent distribution over labels can improve our analysis to below $0.542(k+1)$. Finally, we show a $0.5(k+1)$ -approximation for bipartite graphs and for instances with structured LP solutions. Whether this ratio can be obtained in general is open.

2012 ACM Subject Classification Theory of computation \rightarrow Approximation algorithms analysis; Mathematics of computing \rightarrow Approximation algorithms

Keywords and phrases Approximation Algorithms, Randomized Algorithms, Linear Programming, Graph Algorithms, Scheduling

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2025.18

Category APPROX

Related Version *Full Version:* <https://arxiv.org/abs/2507.07943>

1 Introduction

Given an edge-weighted directed acyclic graph $G = (V, A)$ and a parameter k , the DAG Edge Deletion problem of k , or $\text{DED}(k)$, is to delete a minimum weight set of edges so that the resulting graph has no paths with k edges. $\text{DED}(k)$ was first introduced by Kenkre, Pandit, Purohit, and Saket [9], who gave a simple k -approximation and showed it was UGC-Hard to approximate better than $\lfloor \frac{k}{2} \rfloor$ for any constant k using a result on the vertex deletion version of the problem by Svensson [17]. Kenkre et al. also studied the maximization version of $\text{DED}(k)$, known as Max k -Ordering, which can be studied on general directed graphs. In Max k -Ordering, we want to *keep* as many edges as possible so that there are no paths of length k remaining. They gave a 2-approximation for Max k -Ordering. Max k -Ordering and $\text{DED}(k)$ are NP-Hard for $k \geq 2$ [13, 12].

For $\text{DED}(k)$, it is important that the input is a DAG. Otherwise, it becomes impossible to approximate when k depends on n (which it may), as it can encode the Directed Hamiltonian Path problem, and requires prohibitive running time when k is large.¹ However, on a DAG,

¹ Although, somewhat surprisingly, one can find a path of length k in time $2^k \cdot \text{poly}(n, k)$ [19].



finding paths of arbitrary length is polynomial time solvable using a dynamic program. For Max k -Ordering, it is not necessary to assume the input is a DAG since one can still obtain an approximate solution by comparing to an optimal solution that may keep all the edges. In this work we will focus on the regime in which our running time cannot depend exponentially on k and k may depend on n .

DED(k) is identical to the following problem: find a partition V_1, \dots, V_k of G such that the weight of edges that do not go from a smaller indexed component to a larger one is minimized. This was observed by [9], and follows from the fact that a DAG with no paths of length k has a topological ordering with at most k layers. This gives a natural scheduling (or parallel computing) application for DED(k): if the vertices represent tasks and the edges represent constraints that task i must occur before task j , DED(k) is equivalent to finding the minimum weight set of constraints to violate such that all the tasks can be completed within k serial steps. DED(k) is also motivated by a problem in VSLI design, similar to the DAG vertex deletion version (DVD(k)) introduced by Paik [15] in the 90s. For DVD(k), we are given a circuit as input and the goal is to upgrade the nodes of the circuit so that no electrical signal has to travel through k consecutive un-upgraded nodes. DED(k) can be motivated identically, except here we upgrade the wires instead of the nodes. DVD(k) also has been studied for its applications to secure multi-party computation, see for example [4].

After the work of Kenkre et al. introducing DED(k), giving a linear program for it, and showing a k -approximation, Klein and Wexler [12] showed a $\frac{2}{3}(k+1)$ -approximation using a combinatorial rounding approach, as well as improved results for small values of k . They also showed an integrality gap of at least $\frac{k+1}{2}$ for the natural LP using a construction of Alon, Bollobás, Gyárfás, Lehel, and Scott [2]. In Section 2, we show the following improved approximation ratio and integrality gap:

► **Theorem 1.** *There is a randomized approximation algorithm for DED(k) with approximation ratio $0.549(k+1)$. In particular, given a solution x to the natural LP (1), there is a randomized algorithm producing a solution of expected cost at most $0.549(k+1) \cdot c(x)$, where $c(x)$ is the cost of the LP. Therefore, the integrality gap of LP (1) is at most $0.549(k+1)$.*

Note that it is easy to derandomize our algorithm up to an arbitrary loss in precision using the method of condition expectation. We will not discuss this as it is standard.

Our algorithm rounds using a distribution over vertex labels in $[0, 1]$. The most natural distribution to pick is uniform over the interval $[0, 1]$, which leads to an approximation ratio of $(2 - \sqrt{2})(k+1) \approx 0.585(k+1)$. The above theorem, giving a factor of $0.549(k+1)$, is obtained using a distribution which has more mass near 0 and 1. A natural question is what the optimal distribution over labels is. While we are unable to answer this question precisely, in Section 3, we show that there is no independent distribution obtaining better than $0.542(k+1)$ using our framework. The lower bound has some similarity to recent work in additive combinatorics on Sidon sets, where one needs a lower bound on the supremum of the probability density function of autoconvolutions [5, 14].

We also show a $0.5(k+1)$ -approximation for *bipartite* DAGs using a correlated labeling distribution. As the UGC-Hardness result of Kenkre et al. [9] holds for bipartite graphs, this rounding is optimal up to constants under the UGC. Finally, we show a $0.5(k+1)$ -approximation for structured LP solutions using an independent label distribution. In particular, if x is an optimal LP solution for an instance and there exists some $c > 0$ so that $x_e \in \{0, c\}$ for all $e \in A$, then we show a randomized $0.5(k+1)$ -approximation. (In reality, we show something slightly stronger than this, see Section 4 for more details.) This rounding is optimal, as the integrality gap example in [12] has this property.

1.1 Hardness Results and Other Related Work

In 2012, Svensson [17] gave lower bounds for DAG Vertex Deletion. In particular, he showed that for any constant k , it is hard to approximate $DVD(k)$ with a ratio better than k assuming the UGC. A k -approximation for this problem is easy to obtain, so the approximability of $DVD(k)$ is completely understood (assuming the UGC) for all constants k . When k may depend on n , however, the problem is still not fully understood, even under the UGC, as the lower bound no longer holds. The vertex deletion version has also been studied on undirected graphs for small k [3, 18]. As for NP-Hardness, there is an approximation preserving reduction from Vertex Cover to $DVD(2)$ and to $DED(4)$, so $DVD(k)$ is NP-Hard to approximate better than $\sqrt{2}$ for $k \geq 2$, as is $DED(k)$ for $k \geq 4$. NP-Hardness of $\sqrt{2}$ for Vertex Cover was shown by Khot, Minzer, and Safra [11] in 2017, improving upon the classical result of Dinur [8].

Notice that the NP-Hardness result works for *any* $k \geq 2$ for $DVD(k)$ and any $k \geq 4$ for $DED(k)$, where k may depend on n , as given hardness for parameter c , one can add paths of length $k - c$ with infinite weight vertices or edges to every vertex. This reduction works for the UGC-Hardness results as well, meaning that for any constant k and any $k' \geq k$ (where k' may depend on n), $DED(k')$ is still UGC-Hard to approximate better than $\lfloor \frac{k}{2} \rfloor$ (or k for $DVD(k')$). In other words, both problems do not become easier as k grows.

As noted in [9], $DED(k)$ is related to the well-studied Maximum Acyclic Subgraph Problem (MAS). Here we are given a directed graph (which is not a DAG) and asked to find an acyclic subgraph with as many edges as possible. Note that the maximization version of $DED(k)$, Max k -Ordering, is a generalization of MAS, recovering this problem when k is set to n . Guruswami, Manokaran, and Raghavendra showed that MAS is UGC-Hard to approximate better than $2 - \epsilon$ for any $\epsilon > 0$ [7], and Kenkre et al. [9] note that this implies UGC-Hardness of $2 - o_k(1)$ for Max k -Ordering. Also related is the Restricted Maximum Acyclic Subgraph (RMAS) problem, where each vertex has a finite set of possible integer labels, and the goal is to pick a valid labeling that maximizes the number of edges going from a smaller to a larger label in the assignment. This problem was introduced by Khandekar, Kimbrel, Makarychev, and Sviridenko [10] and its approximability was studied by Grandoni, Kociumaka, and Włodarczyk [6], who gave a $2\sqrt{2}$ -approximation. RMAS clearly generalizes MAS. To the best of our knowledge, a minimization version of RMAS on DAGs has not been studied.

A more general setting is the k -hypergraph vertex cover problem for which a simple rounding gives a k -approximation. For the setting where the underlying hypergraph is also k -partite, a theorem of Lovász (see [1]) gives a $0.5k$ -approximation.

1.2 Formulation as a CSP

$DED(k)$ can easily be stated as a constraint satisfaction problem. In particular, one can create a variable x_v for each vertex $v \in V$ and allow $x_v \in \{1, 2, \dots, k\}$. Then, we have a set of constraints given by the edges: for each $e = (u, v)$ we create a constraint that $x_u < x_v$. Max k -Ordering is then the problem of maximizing the number of satisfied constraints, and $DED(k)$ is the problem of minimizing the number of unsatisfied constraints. So, both problems fall under the purview of Raghavendra's UGC-optimal algorithm for CSPs [16] *whenever k is an absolute constant*. However, the approximation ratio of this algorithm for constant k is not known, and the runtime of the rounding step depends exponentially on k . (In addition, for $DED(k)$, the objective function is negative, and there is an additive error in Raghavendra's algorithm which could produce issues for this approach.) So, while this is an important view on the problem, it does not give satisfactory guarantees as $k \rightarrow \infty$ or when k depends on n .

1.3 Organization of the Paper

In Section 2, we show our improved approximation algorithm. Our approach is to assign labels $\ell(v) \in [0, 1]$ to each vertex v independently and cut edges $e = (u, v)$ for which $\ell(v) - \ell(u)$ is small compared to x_e . By choosing an appropriate definition of small, the feasibility of the LP solution then implies cutting these edges results in a feasible solution to DED(k). We first show that drawing labels $\ell(v) \sim \text{Uniform}(0, 1)$ results in a $(2 - \sqrt{2})(k + 1) \approx 0.585(k + 1)$ -approximation, and then show a modified distribution that results in an algorithm with ratio slightly below $0.549(k + 1)$.

In Section 3, we show that no independent labeling distribution can improve the approximation ratio of our algorithm to below $0.542(k + 1)$, at least within our analysis framework, which bounds the probability each edge is cut by αx_e to obtain an approximation ratio of α .

Finally, in Section 4, we give our two additional results showing $0.5(k + 1)$ -approximation algorithms for bipartite DAGs and instances with structured LP solutions.

2 Improved Approximation Algorithm

The following is the LP from Kenkre et al. [9] which they used to achieve a k -approximation for edge-weighted DAGs, where c_e gives the cost of the edge e and x_e denotes the relaxation of the decision variable on whether edge e is cut.

$$\begin{aligned}
 \min \quad & \sum_{e \in A} c_e x_e, \\
 \text{s.t.} \quad & \sum_{e \in P} x_e \geq 1, \quad \forall \text{ paths } P \text{ of length } k, \\
 & x_e \geq 0, \quad \forall e \in A.
 \end{aligned} \tag{1}$$

Although there are an exponential number of constraints, a separation oracle can be implemented in polynomial time, as the input graph is a DAG. Therefore by the ellipsoid method, this LP can be solved in polynomial time. Note we will use $c(x)$ to denote $\sum_{e \in A} c_e x_e$ and for a set of edges $S \subseteq A$, $c(S) = \sum_{e \in S} c_e$.

2.1 Algorithmic Framework

Our algorithm is as follows, and depends on a *label distribution* $\mu : [0, 1]^V \rightarrow \mathbb{R}_{\geq 0}$ over label sets $\ell \in [0, 1]^V$. We will use $\ell(v)$ to denote the label of vertex v .

First we solve (1) to obtain an optimal solution x . Then we draw a label set ℓ according to μ , and delete all edges $e = (u, v)$ with $\ell(v) - \ell(u) \leq (k + 1)x_e - 1$.

We will focus on upper and lower bounds for distributions in which the label of every vertex is drawn independently from the same fixed distribution over $[0, 1]$. However, in Section 4 we will use a distribution that is not independent. Before analyzing the performance of two label distributions, we show now that it results in a feasible solution.

► **Lemma 2.** *For any solution x to (1) and any set of labels $\ell \in [0, 1]^V$, deleting all the edges with $\ell(v) - \ell(u) \leq (k + 1)x_e - 1$ results in a feasible solution.*

Proof. Let P be a path of length k beginning at vertex u and ending at v . By way of contradiction, suppose no edges in the path are deleted. Then:

$$\begin{aligned}
 \ell(v) &> \ell(u) + \sum_{e \in P} [(k+1)x_e - 1] \\
 &\geq \sum_{e \in P} [(k+1)x_e - 1] && \text{Since } \ell(u) \geq 0 \\
 &= \sum_{e \in P} (k+1)x_e - k && \text{Since } P \text{ is of length } k \\
 &\geq (k+1) - k && \text{Since } x \text{ is a feasible solution} \\
 &= 1.
 \end{aligned}$$

This gives a contradiction because $\ell(v) \in [0, 1]$. Note the first inequality is strict because we keep an edge $e = (a, b)$ only if $\ell(b) - \ell(a) > (k+1)x_e - 1$. ◀

To bound the approximation ratio of our algorithm, we will use the following.

► **Lemma 3.** *Let μ be a labeling distribution and $\alpha \geq \frac{1}{2}$. If for any edge $e = (u, v)$ and any $t \in [-1, 1]$ we have*

$$\frac{\mathbb{P}[\ell(v) - \ell(u) \leq t]}{t + 1} \leq \alpha,$$

then using μ in the above algorithm produces a solution of expected cost at most $\alpha(k+1) \cdot c(x)$ and results in a randomized $\alpha(k+1)$ -approximation.

Proof. Let $e = (u, v)$ be an edge. We will show that

$$\mathbb{P}[\ell(u) - \ell(v) \leq (k+1)x_e - 1] \leq (k+1)\alpha \cdot x_e. \quad (2)$$

So long as we have this, the claim follows, because if S is the set of edges cut by the algorithm, then:

$$\begin{aligned}
 \mathbb{E}[c(S)] &= \sum_{e=(u,v) \in A} \mathbb{P}[\ell(u) - \ell(v) \leq (k+1)x_e - 1] c_e \\
 &\leq \alpha(k+1) \sum_{e \in A} c_e x_e = \alpha(k+1) \cdot c(x).
 \end{aligned}$$

To obtain (2), set $t = (k+1)x_e - 1$. Clearly, $t \geq -1$. If $t \geq 1$, then $x_e \geq \frac{2}{k+1}$, but then (2) is trivially satisfied. So, $t = (k+1)x_e - 1 \in [-1, 1]$ and by the assumption,

$$\frac{\mathbb{P}[\ell(v) - \ell(u) \leq (k+1)x_e - 1]}{(k+1)x_e} \leq \alpha.$$

Rearranging gives (2). ◀

2.2 Uniform Distribution

Here we consider the most natural label distribution: for every vertex, we simply assign it a label uniformly at random from the interval $[0, 1]$, independent of all other vertices. First we show the following lemma:

► **Lemma 4.** Consider the labeling distribution μ which independently assigns every vertex a label uniformly at random from $[0, 1]$. Then, for all $u, v \in V$, $t \in [-1, 1]$,

$$\frac{\mathbb{P}[\ell(v) - \ell(u) \leq t]}{t + 1} \leq 2 - \sqrt{2}.$$

Proof. Let $\ell(u), \ell(v)$ be draws from $\text{Uniform}(0, 1)$. Then, the CDF $F(x)$ of $\ell(v) - \ell(u)$ is as follows:

$$F(t) = \begin{cases} \frac{1}{2}t^2 + t + \frac{1}{2} & -1 \leq t < 0 \\ -\frac{1}{2}t^2 + t + \frac{1}{2} & 0 \leq t \leq 1 \end{cases}$$

If $-1 \leq t < 0$, then

$$\frac{F(t)}{t + 1} = \frac{(t + 1)(\frac{1}{2}t + \frac{1}{2})}{t + 1} = \frac{1}{2}(t + 1) \leq \frac{1}{2}.$$

If $0 \leq t \leq 1$, then

$$\frac{F(t)}{t + 1} = \frac{-\frac{1}{2}t^2 + t + \frac{1}{2}}{t + 1}.$$

Computing the derivative of this expression with respect to t , we obtain $\frac{-\frac{1}{2}t^2 - t + \frac{1}{2}}{(t + 1)^2}$, and setting this to 0, we know that the maximizer must be either $-1 - \sqrt{2}$, $-1 + \sqrt{2}$, or an endpoint $t = 0$ or $t = 1$. Of the two roots, only $-1 + \sqrt{2}$ is in the range $t \in [0, 1]$. When evaluated this is $2 - \sqrt{2}$. Evaluating at $t = 0$ or $t = 1$ we obtain $\frac{1}{2}$, so the maximizer is $2 - \sqrt{2}$. ◀

Using Lemma 3, we obtain the following as a corollary:

► **Theorem 5.** Assigning $\ell(v)$ to each vertex v independently and uniformly from $[0, 1]$ results in a solution with expected cost at most $(2 - \sqrt{2})(k + 1) \cdot c(x) \approx 0.585(k + 1) \cdot c(x)$. By using an optimal solution x , we obtain a $(2 - \sqrt{2})(k + 1) \approx 0.585(k + 1)$ -approximation.


2.3 Improving Upon Uniform

While using the uniform distribution is a natural candidate, it is (perhaps somewhat surprisingly) not optimal. If we wanted to obtain an $0.5(k + 1)$ approximation using our framework, it is not difficult to see that the CDF of $\ell(v) - \ell(u)$ must be equal to $\text{Uniform}(-1, 1)$ for all u, v with an edge between them with $x_e > 0$. If this were obtainable, we would have:

$$\mathbb{P}[\ell(v) - \ell(u) \leq (k + 1)x_e - 1] = \frac{k + 1}{2}x_e,$$

for all values of $x_e \in [0, \frac{2}{k+1}]$. (Edges of value greater than $\frac{2}{k+1}$ are cut with probability 1 for any label distribution.)

So, in some sense, **our goal is to design an independent distribution over labels so that the CDF of $\ell(v) - \ell(u)$ is “closer” to $\text{Uniform}(-1, 1)$** . Using numerical verification and exhaustive search, we arrived at the probability distribution \mathcal{D} over $[0, 1]$ with density $p(x) = \frac{2}{3} + \frac{23}{3}(1 - 2x)^{22}$. If X, Y are sampled independently from \mathcal{D} , we obtain a CDF for which $X - Y$ is closer to $\text{Uniform}(-1, 1)$ in the relevant respect. We show that using this distribution yields a better approximation ratio than uniform labels. Due to the complexity of the resulting expressions, we use numerical verification. Figure 1 shows the comparison of the densities of \mathcal{D} and $\text{Uniform}(-1, 1)$.



■ **Figure 1** Densities of \mathcal{D} and $\text{Uniform}(-1, 1)$.

► **Theorem 6.** *Assigning labels ℓ to all vertices from the interval $[0, 1]$ independently according to the density $p(x) = \frac{2}{3} + \frac{23}{3}(1 - 2x)^{22}$ results in a solution with expected cost below $0.549(k + 1) \cdot c(x)$. This results in a randomized $< 0.549(k + 1)$ -approximation for $\text{DED}(k)$.*

Proof. As in the uniform case, to obtain the theorem it is sufficient to prove the following claim and apply Lemma 3.

▷ **Claim 7.** Suppose X, Y are independent random variables sampled from \mathcal{D} , then $\frac{\mathbb{P}[X - Y \leq t]}{t + 1} < 0.549$ for all $t \in [-1, 1]$.

Proof. This is verified by Mathematica. See Section A for the code used and an execution. The maximum of $\frac{\mathbb{P}[X - Y \leq t]}{t + 1}$ occurs at approximately $t = 0.27$ with value slightly below 0.5482. See Figure 3. ◀

3 Lower Bound

An interesting problem is to find the function that minimizes the approximation ratio when using an *independent* distribution over labels:

► **Problem 8.** Let \mathcal{P} denote the set of probability distributions supported on $[0, 1]$. Find

$$\alpha := \inf_{D \in \mathcal{P}} \sup_{t \in [-1, 1]} \frac{\mathbb{P}_{\ell(u), \ell(v) \sim D}[\ell(u) - \ell(v) \leq t]}{t + 1}.$$

While this problem does not capture the possibility of correlated label distributions (which we exploit in the bipartite case in Section 4), we find it quite interesting on its own. While some related questions have been studied in additive combinatorics, for example by Cloninger and Steinerberger [5] and Matolcsi and Vinuesa [14], we were unable to find a resolution of this question in the literature. We have been able to bound:

$$0.542 < \alpha < 0.549,$$

where the upper bound is from Section 2. Here, we show a complete proof of a lower bound of 0.539.

18:8 A Randomized Rounding Approach for DAG Edge Deletion

► **Lemma 9.** $\alpha > 0.539$.

Proof. Let D be a distribution supported on $[0, 1]$ and X, Y be independent samples from D . Suppose that

$$\mathbb{P}[X - Y \leq x] \leq \alpha(x + 1) \quad \text{for all } x \in [-1, 1].$$

Denote $Z := X - Y$. First note that symmetry of Z gives the complementary lower bound

$$\mathbb{P}[Z \leq x] = \mathbb{P}[Z \geq -x] = 1 - \mathbb{P}[Z \leq -x] \geq 1 - \alpha(1 - x), \quad \text{for all } x \in [-1, 1].$$

These inequalities imply

$$-(\alpha - \tfrac{1}{2})(1 - x) \leq \mathbb{P}[Z \leq x] - \frac{x + 1}{2} \leq (\alpha - \tfrac{1}{2})(1 + x) \quad \text{for all } x \in [-1, 1]. \quad (\star)$$

Fix some $t \in (\pi, 2\pi)$ that we choose later (so $\sin t < 0$) and observe that

$$\begin{aligned} \mathbb{E}[\cos(tZ)] &= \mathbb{E}[\cos(t(X - Y))] \\ &= \mathbb{E}[\cos(tX)]\mathbb{E}[\cos(tY)] + \mathbb{E}[\sin(tX)]\mathbb{E}[\sin(tY)] \\ &= \mathbb{E}[\cos(tX)]^2 + \mathbb{E}[\sin(tX)]^2 \geq 0. \end{aligned}$$

Integration by parts gives

$$\mathbb{E}[\cos(tZ)] = \cos t + t \int_{-1}^1 \sin(tx) \mathbb{P}[Z \leq x] dx.$$

For $U \sim \text{Uniform}(-1, 1)$ (whose CDF is $\frac{x+1}{2}$), we have $\mathbb{E}_{U \sim \text{Uniform}(-1, 1)}[\cos(tU)] = \frac{\sin t}{t}$, therefore

$$\frac{\sin t}{t} = \cos t + t \int_{-1}^1 \sin(tx) \cdot \left(\frac{x + 1}{2} \right) dx.$$

Using $\mathbb{E}[\cos(tZ)] \geq 0$ and subtracting the two formulas above yields

$$-\frac{\sin t}{t} \leq \mathbb{E}[\cos(tZ)] - \frac{\sin t}{t} = t \int_{-1}^1 \sin(tx) \left(\mathbb{P}[Z \leq x] - \frac{x + 1}{2} \right) dx.$$

By symmetry, we have

$$\int_{-1}^1 \sin(tx) \left(\mathbb{P}[Z \leq x] - \frac{x + 1}{2} \right) dx = 2 \int_0^1 \sin(tx) \left(\mathbb{P}[Z \leq x] - \frac{x + 1}{2} \right) dx.$$

Since $\sin(tx)$ has only one root π/t in $(0, 1)$, we may use (\star) to upper bound the integral

$$\frac{1}{\alpha - \frac{1}{2}} \int_0^1 \sin(tx) \left(\mathbb{P}[Z \leq x] - \frac{x + 1}{2} \right) dx \leq \int_0^{\pi/t} \sin(tx)(1 + x) dx - \int_{\pi/t}^1 \sin(tx)(1 - x) dx,$$

which evaluates to $\frac{3t + \sin t}{t^2}$. It follows that

$$\alpha \geq \frac{1}{2} + \sup_{t \in (\pi, 2\pi)} \frac{-\sin t}{6t + 2\sin t} \geq \frac{1}{2} + \frac{-\sin 4.5}{27 + 2\sin 4.5} > 0.539. \quad \blacktriangleleft$$

We remark that using the function $\cos(4.4X) + 0.29\cos(8.9X)$ instead of $\cos(4.5X)$ in the above proof shows a slightly stronger lower bound $\alpha > 0.542$.

From the Lemma 9 it follows that every distribution $D \in \mathcal{P}$ has a corresponding CDF of $D - D$ which significantly exceeds the CDF of $\text{Uniform}(-1, 1)$ at some point. Therefore we cannot attain a ratio of $\frac{1}{2}$. However, this may still be possible for correlated label distributions, and indeed, depending on the structure of the graph such distributions *can* exist. In the next section we show they can be constructed easily for bipartite graphs.

4 Optimal Rounding for Special Cases

In this section, we show how to obtain our desired bound of $0.5(k+1)$ for two special cases. The result for bipartite graphs is tight up to constants, while the second result is tight as there is a matching integrality gap on a point with $x_e = \frac{1}{k}$ for all edges.

4.1 Bipartite Graphs

Here we show that for any bipartite DAG there exists a distribution over labels so that the randomized algorithm from above achieves an approximation ratio of $0.5(k+1)$. Notably, the instances from Kenkre et al. used to obtain UGC hardness of $\lfloor 0.5k \rfloor$ are bipartite, so this is tight up to constants assuming the UGC.

► **Theorem 10.** *Let $G = (V, A)$ be bipartite with bipartition (A', B') and x be a feasible solution to (1). Then, there is an algorithm that produces a DED(k) solution with expected cost at most $(0.5)(k+1) \cdot c(x)$.*

Proof. Our algorithm is identical to before, except now we produce a correlated distribution over labels. To produce the labeling $\ell : V \rightarrow [0, 1]$, we first sample $y \in [0, 1]$ uniformly. Then, with probability $\frac{1}{2}$, assign $\ell(v) = 0$ for all $v \in A'$ and $\ell(v) = y$ for all $v \in B'$. Otherwise, assign $\ell(v) = y$ for all $v \in A'$ and $\ell(v) = 0$ for all $v \in B'$.

For every edge $a = (u, v) \in A$, the distribution over $\ell(v) - \ell(u)$ is now uniform over $[-1, 1]$, as with probability $\frac{1}{2}$ we obtain a uniformly random number between 0 and 1 and otherwise a uniformly random number between -1 and 0. Therefore, if we now run the same algorithm as before, we have:

$$\begin{aligned} \mathbb{P}[a \text{ is cut}] &= \mathbb{P}[\ell(v) - \ell(u) \leq (k+1)x_e - 1] \\ &= \frac{(k+1)x_e - 1 + 1}{2} && \text{As } \ell(v) - \ell(u) \sim \text{Uniform}(-1, 1) \\ &= 0.5(k+1)x_e, \end{aligned}$$

as desired. ◀

It is unclear whether correlated distributions can be used to improve the approximation ratio in the general case.

4.2 Instances with Structured LP Solutions

Here we show that an independent distribution over labels exists when we have a structured LP solution. It implies, for example, that if our LP solution has the property that $x_e \in \{0, c\}$ for all $e \in E$ for some fixed value c , we can obtain a $0.5(k+1)$ -approximation, but is slightly more general (for example, we can obtain the same ratio if all edges have value at least $\frac{1.5}{k+1}$).

► **Theorem 11.** *If x is a feasible solution to (1), and for every edge e with $x_e > 0$ we have:*

$$\frac{1 + \frac{1}{r+1}}{k+1} \leq x_e < \frac{1 + \frac{1}{r}}{k+1},$$

for some integer $r \geq 1$, then there is an algorithm that produces a solution of expected cost at most $\frac{k+1}{2}c(x)$. In the case that $r = 1$, the upper bound is not necessary as we may simply cut all edges with value at least $\frac{2}{k+1}$ and round the remaining edges.

Proof. Let μ_r be the uniform distribution over labels in the set $\{0, \frac{1}{r}, \frac{2}{r}, \dots, 1\}$. We claim applying our framework by selecting each label independently from μ_r proves the theorem. We cut an edge $e = (u, v)$ if

$$\ell(v) - \ell(u) \leq (k+1)x_e - 1 < \frac{1 + \frac{1}{r}}{k+1}(k+1) - 1 = \frac{1}{r}.$$

Therefore, since $\ell(u) - \ell(v)$ is either negative, 0, or at least $\frac{1}{r}$, we cut an edge exactly when $\ell(v) \leq \ell(u)$. So, the probability we cut any edge with $x_e > 0$ is:

$$\begin{aligned} \mathbb{P}[\ell(v) \leq \ell(u)] &= 1 - \mathbb{P}[\ell(v) > \ell(u)] \\ &= 1 - \sum_{i=0}^r \mathbb{P}\left[\ell(v) = \frac{i}{r}\right] \cdot \mathbb{P}\left[\ell(u) < \frac{i}{r}\right] \\ &= 1 - \sum_{i=0}^r \frac{1}{r+1} \cdot \frac{i}{r+1} = 1 - \frac{r}{2(r+1)} = \frac{r+2}{2(r+1)}. \end{aligned}$$

Now, the probability we cut an edge compared to x_e (our approximation ratio) is:

$$\frac{r+2}{2(r+1)} \cdot \frac{1}{x_e} \leq \frac{r+2}{2(r+1)} \cdot \frac{k+1}{1 + \frac{1}{r+1}} = \frac{k+1}{2}.$$

As desired. ◀

5 Conclusion

The main open question remaining is whether it is possible to obtain an algorithm with approximation ratio $0.5(k+1)$, or whether this is UGC-Hard. We are unaware of any instances with integrality gap above $0.5(k+1)$, so improving this bound is also of interest. Another open problem is whether it is NP-Hard to obtain an $O(1)$ approximation for $\text{DED}(k)$.

References

- 1 Ron Aharoni, Ron Holzman, and Michael Krivelevich. On a theorem of Lovász on covers in r -partite hypergraphs. *Combinatorica*, 16(2):149–174, 1996. doi:10.1007/BF01844843.
- 2 Noga Alon, Béla Bollobás, András Gyárfás, Jenő Lehel, and Alex Scott. Maximum directed cuts in acyclic digraphs. *Journal of Graph Theory*, 55(1):1–13, 2007. doi:10.1002/jgt.20215.
- 3 Boštjan Brešar, František Kardoš, Ján Katrenič, and Gabriel Semanišin. Minimum γ -path vertex cover. *Discrete Applied Mathematics*, 159(12):1189–1195, 2011. doi:10.1016/j.dam.2011.04.008.
- 4 Pierre Charbit, Geoffroy Couteau, Pierre Meyer, and Reza Naserasr. A note on low-communication secure multiparty computation via circuit depth-reduction. In Elette Boyle and Mohammad Mahmoody, editors, *Theory of Cryptography*, pages 167–199, Cham, 2025. Springer Nature Switzerland.
- 5 Alexander Cloninger and Stefan Steinerberger. On suprema of autoconvolutions with an application to sidon sets. *Proceedings of the American Mathematical Society*, 145(8):3191–3200, 2017. doi:10.1090/proc/13690.
- 6 Fabrizio Grandoni, Tomasz Kociumaka, and Michał Włodarczyk. An lp-rounding 22-approximation for restricted maximum acyclic subgraph. *Information Processing Letters*, 115(2):182–185, 2015. doi:10.1016/j.ipl.2014.09.008.
- 7 Venkatesan Guruswami, Rajsekar Manokaran, and Prasad Raghavendra. Beating the random ordering is hard: Inapproximability of maximum acyclic subgraph. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 573–582, 2008. doi:10.1109/F0CS.2008.51.

- 8 Samuel Safra Irit Dinur. On the hardness of approximating minimum vertex cover. *Annals of Mathematics*, 162(1):439–485, 2005. URL: <http://www.jstor.org/stable/3597377>.
- 9 Sreyash Kenkre, Vinayaka Pandit, Manish Purohit, and Rishi Saket. On the approximability of digraph ordering. *Algorithmica*, 78(4):1182–1205, August 2017. doi:10.1007/s00453-016-0227-7.
- 10 Rohit Khandekar, Tracy Kimbrel, Konstantin Makarychev, and Maxim Sviridenko. On hardness of pricing items for single-minded bidders. In Irit Dinur, Klaus Jansen, Joseph Naor, and José Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 202–216, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg. doi:10.1007/978-3-642-03685-9_16.
- 11 Subhash Khot, Dor Minzer, and Muli Safra. On independent sets, 2-to-2 games, and grassmann graphs. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017*, pages 576–589, New York, NY, USA, 2017. Association for Computing Machinery. doi:10.1145/3055399.3055432.
- 12 Nathan Klein and Tom Wexler. On the approximability of dag edge deletion, 2015.
- 13 Michael Lampis, Georgia Kaouri, and Valia Mitsou. On the algorithmic effectiveness of digraph decompositions and complexity measures. In Seok-Hee Hong, Hiroshi Nagamochi, and Takuro Fukunaga, editors, *Algorithms and Computation*, volume 5369 of *Lecture Notes in Computer Science*, pages 220–231. Springer Berlin Heidelberg, 2008. doi:10.1007/978-3-540-92182-0_22.
- 14 Máté Matolcsi and Carlos Vinuesa. Improved bounds on the supremum of autoconvolutions. *Journal of Mathematical Analysis and Applications*, 372(2):439–447.e2, 2010. doi:10.1016/j.jmaa.2010.07.030.
- 15 D. Paik, S. Reddy, and S. Sahni. Deleting vertices to bound path length. *IEEE Transactions on Computers*, 43(9):1091–1096, September 1994. doi:10.1109/12.312117.
- 16 Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing, STOC '08*, pages 245–254, New York, NY, USA, 2008. Association for Computing Machinery. doi:10.1145/1374376.1374414.
- 17 Ola Svensson. Hardness of vertex deletion and project scheduling. In Anupam Gupta, Klaus Jansen, José Rolim, and Rocco Servedio, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, volume 7408 of *Lecture Notes in Computer Science*, pages 301–312. Springer Berlin Heidelberg, 2012. doi:10.1007/978-3-642-32512-0_26.
- 18 Jianhua Tu and Wenli Zhou. A factor 2 approximation algorithm for the vertex cover problem. *Information Processing Letters*, 111(14):683–686, 2011. doi:10.1016/j.ipl.2011.04.009.
- 19 Ryan Williams. Finding paths of length k in $O(2^k)$ time. *Inf. Process. Lett.*, 109(6):315–318, February 2009. doi:10.1016/j.ipl.2008.11.004.

18:12 A Randomized Rounding Approach for DAG Edge Deletion

A Automated Verification

Here we include a copy of a high precision run in Mathematica showing a worst case ratio of slightly below 0.5482.

```
In[1]:= P = ProbabilityDistribution[2/3 + (23/3) * (1 - 2 x)^(22), {x, 0, 1}]
T = TransformedDistribution[y - z, {y ≈ P, z ≈ P}]

Out[1]= ProbabilityDistribution[ $\frac{2}{3} + \frac{23}{3} (1 - 2 x)^{22}$ , {x, 0, 1}]

Out[2]= TransformedDistribution[y - z, {y ≈ ProbabilityDistribution[ $\frac{2}{3} + \frac{23}{3} (1 - 2 x)^{22}$ , {x, 0, 1}], z ≈ ProbabilityDistribution[ $\frac{2}{3} + \frac{23}{3} (1 - 2 x)^{22}$ , {x, 0, 1}]}]}

In[3]:= NMaximize[
  {CDF[T, t]/(t+1), -1 ≤ t ≤ 1},
  t,
  WorkingPrecision → 50,
  PrecisionGoal → 30,
  AccuracyGoal → 30
]

Out[3]= {0.54819688281860859640147603001950034519824400446040, {t → 0.26655712446157772938814603176240996931271738299061}}
```

■ **Figure 2** Execution of upper bound code.

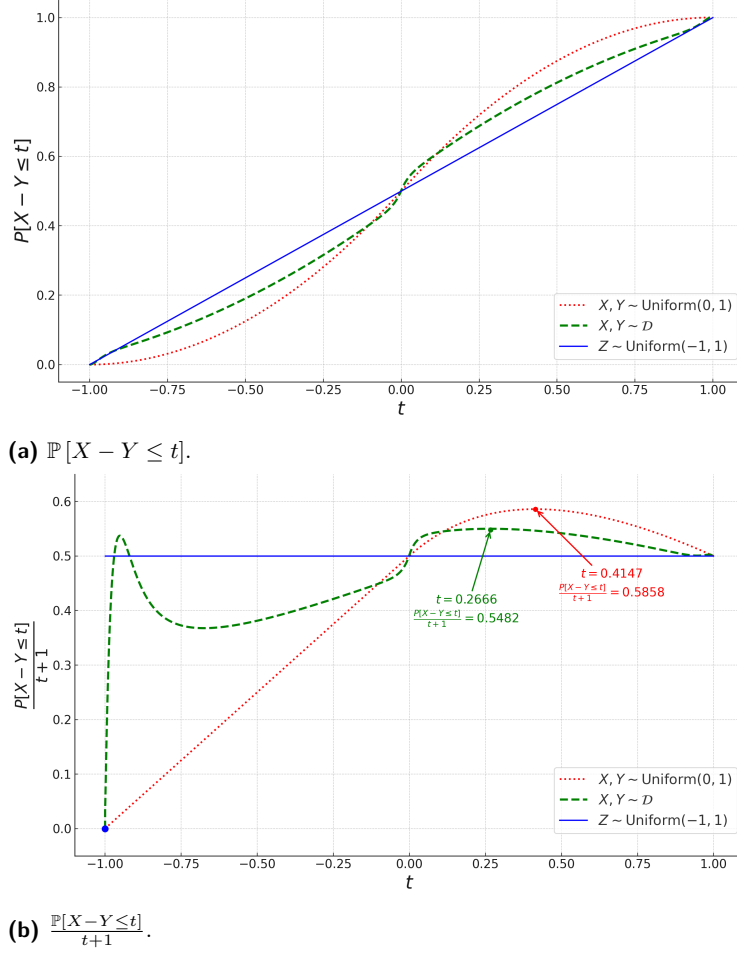
B Figure 3

Figure 3 Figure 3a and Figure 3b compare the behavior of the difference $X - Y$, where X, Y are sampled independently from two candidate label distributions: $\text{Uniform}(0, 1)$ (red), and the refined distribution \mathcal{D} with density $p(x) = \frac{2}{3} + \frac{23}{3}(1 - 2x)^{22}$ (green). As a reference, we also include the random variable $Z \sim \text{Uniform}(-1, 1)$ (blue), which represents a target distribution that would yield a 0.5-approximation if realizable. While Section 3 proves that no independent label distribution over $[0, 1]$ can make $X - Y \sim \text{Uniform}(-1, 1)$, the behavior of Z remains a useful benchmark. In both subfigures, the distribution \mathcal{D} yields a difference distribution that more closely mimics Z than $\text{Uniform}(0, 1)$ does. In particular, as established in Lemma 4 and Claim 7, the maximum of $\frac{\mathbb{P}[X - Y \leq t]}{t + 1}$ occurs at $t = 2 - \sqrt{2} \approx 0.4147$ for $\text{Uniform}(0, 1)$ and at $t = 0.2666$ for \mathcal{D} , with corresponding values 0.5858 and 0.5482, respectively.