Advanced Approximation Algorithms **(CMU 18-854B, Spring 2008)**

Lecture 8: Hardness of Min-Ek-Hypergraph-Independent-Set

February 7 2008

Lecturer: Ryan O'Donnell *Scribe: Aaron Roth*

1 Introduction

In this lecture, we start to prove a hardness result for Max-Ek-Independent-Set. The proof will be completed in the next lecture. We begin by defining the problem, and then devote a section to developing intuition for the reduction. Finally, we give the hardness reduction and prove completeness (but defer soundness to next lecture).

1.1 The Problem

First we consider a closely related problem, Min-Ek-Vertex-Cover.

Definition 1.1. For any $k \geq 2$, an instance of *weighted Min-Ek-Vertex Cover* is defined by a k-uniform hypergraph $G = (V, E)$, with weights $w(v)$ on each vertex $v \in V$ such that $\sum_v w(v) = 1$. A valid solution is a vertex cover $\mathcal{C} \subseteq V$ such that for all hyperedges $e \in E$, $\mathcal{C} \cap e \neq \emptyset$. The objective is to minimize the weight of the vertex cover: $\sum_{v \in \mathcal{C}} w(v)$.

Note that for $k = 2$, this is the standard weighted vertex cover problem.

On homework 1, we saw several simple algorithms that guarantee factor-k approximations to Min-Ek-Vertex-Cover. Given this, we may ask: "What would our dream hardness result be for this problem?" Naively, we may respond: "$k - \epsilon$ hardness for every constant $\epsilon > 0$!", but we must dream bigger! Perhaps we could prove $1/k + \delta$ vs. $1 - \epsilon$ hardness for any positive constants ϵ, δ. This would imply $k - \epsilon$ hardness, and would be stronger than (say) $1/k^2 + \delta$ vs. $1/k + \epsilon$ hardness, which would also imply $k - \epsilon$ hardness, because the gap instance is larger. Note that the stronger result can also be scaled to imply the weaker result.

Throughout the lecture, we will find it helpful to think of $1-$ the weight of the cover, rather than the cover itself. We recall that this quantity is always non-negative, since as defined above, the weights sum to 1.

As is the case when $k = 2$, vertex covers are closely related to independent sets:

Definition 1.2. Given a k-uniform hypergraph $G = (V, E)$, $\mathcal{I} \subseteq V$ is an *independent set* if and only if $V \setminus \mathcal{I}$ is a vertex cover. In words, an independent set is a collection of vertices $\mathcal{I}$ such that $\mathcal{I}$ does not include all of the vertices in any edge e.

This naturally leads to the definition of our main problem of instance:

Definition 1.3. For any $k \geq 2$ an instance of *Max-Ek-Independent-Set* is given by a k-uniform hypergraph, and a valid solution is an independent set $\mathcal{I}$. The objective is to maximize the weight of the independent set: $\sum_{v \in \mathcal{I}} w(v)$.

Note that an optimal solution to Min-Ek-Vertex-Cover also gives an optimal solution to Max-Ek-Independent-Set, and our dream hardness result translates to $1 - 1/k - \delta$ vs. ϵ hardness for any constants $\epsilon, \delta > 0$. Such a result would imply that the problem could not be approximated to *any* constant factor, and is particularly strong because of the "location of the gap". (i.e. even in the case when almost the entire graph is an independent set, we can't efficiently find even a tiny set).

This isn't quite known, but something close is:

Theorem 1.4 (Dinur et al. [1]). *The* $1 - 1/(k-1) - \delta$ *vs.* ϵ *decision problem for Max-Ek-Independent-Set is NP hard.*

This result is *almost* as good, but gives no guarantee for $k = 2$. In this lecture, we will give something even easier: $1 - 2/k - \delta$ vs. ϵ hardness. We will again reduce from label cover.

2 Intuition

In this section, we construct an instance that differs from, but should give intuition for, the instance we will construct from Label-Cover in the hardness proof.

Recall the gadget we used to prove hardness for max-coverage: A set of elements $\{0,1\}^{|K|}$ which we think of as bit-strings, and sets $S_{a,b} = \{x : x_a = b\}$. This gadget has the nice property that it has $|K|$ 'good' solutions of cardinality 2, and any other solution must be substantially worse.

2.1 A Gadget/Gap Instance

We want to construct a similar instance for Max-Ek-Independent-Set: One that has K good solutions, such that any other solution is substantially worse. We start by designing an instance with L good solutions[1]: our vertex set will be $\{0,1\}^L$. We want our good solutions to be the sets $\mathcal{D}_\alpha = \{x : x_\alpha = 1\}$. We refer to these as "dictator solutions" since inclusion into the set $\mathcal{D}_\alpha$ is dictated by the α coordinate.

We now must construct hyperedges to guarantee that the dictator solutions are (the only) large independent sets. We take a natural approach: we include *all* possible size k hyperedges that are consistent with the $\mathcal{D}_\alpha$'s being independent sets! For example, for $k = 4$, figure 2.1 gives a set of elements that would *not* form an edge in our construction – such an edge wouldn't be consistent with $\mathcal{D}_4$ being an independent set! Edges in our construction would include every 4-tuple of strings such that every column had at least one 0.

Definition 2.1. Distinct strings $(x_1, \ldots, x_k)$ form a hyperedge if and only if for all α, at least among $(x_1)_\alpha, \ldots, (x_k)_\alpha$ takes value 0.

[1] Note our foresight in switching to L from K

```
100101
010111
001100
101110
```

Figure 1: Not a hyperedge in our construction for $k = 4$.

We may alternatively identify strings $x \in \{0,1\}^L$ with subsets $A \subseteq [L]$ in the natural way, and will treat them in whichever way is more convenient throughout the rest of the lecture. Don't forget these strings/sets are also vertices!

Treating vertices as sets, we may rephrase our definition:

Definition 2.2. Distinct sets $(A_1, \ldots, A_k)$ form a hyperedge if and only if $A_1 \cap \ldots \cap A_k = \emptyset$.

So far so good. But thus far, our good solutions $\mathcal{D}_\alpha$ include half the strings! We want them to include a $1 - 2/k$ fraction... We could use k-ary (rather than binary) strings, but instead we add weights to the vertices.

Definition 2.3. Given $0 \leq p \leq 1$ the *p-biased weight* of a string $x \in \{0,1\}^d$ is:

$$w(x) = \prod_{\alpha \in [L]} p^{x_\alpha}(1-p)^{1-x_\alpha}.$$

In words, it is the probability you would choose the string if you were to select each coordinate i.i.d., choosing a 1 with probability p. Observe that the weights sum to 1: this defines a probability distribution over strings.

We now have weights for elements... What about for sets?

Definition 2.4. For any predicate P, if $\mathcal{I}$ is a collection of strings/sets that satisfy predicate P: $\mathcal{I} = \{x : P(x)\}$ then we say that the weight of $\mathcal{I}$ is $w(\mathcal{I})$, the probability of choosing a string string x such that $P(x)$ holds.

Since a single coordinate determines whether string x is in the dictator set $\mathcal{D}_\alpha$, we have $w(\mathcal{D}_\alpha) = p$. Why don't we set $p = 1 - 1/k$, throwing caution to the wind![2]

What we need to argue now is that other independent sets (that aren't defined on the basis of a single coordinate) have small weight. What can we say about other independent sets?

From the definition, $\mathcal{F} \subseteq \{0,1\}^L$ is an independent set if and only if, for all sets $A_1, \ldots, A_k$ such that $\bigcap_{i=1}^k A_i = \emptyset$, $A_i \notin \mathcal{F}$ for at least one $i \in [L]$ (since these sets form a hyperedge). Alternatively we know that for all distinct $A_1, \ldots, A_k \in \mathcal{F}$, $\bigcap_{i=1}^k A_i \neq \emptyset$. In other words[3], $\mathcal{F}$ is an independent set if and only if it is k-wise 1-intersecting.

[2] we are keeping the dream alive!

[3] words from homework 2

The idea is that if $\mathcal{F}$ is an independent set with non-negligible weight, it should "suggest" a small number of coordinates that in the reduction will correspond to labels in a Label-Cover instance. To 'decode', we will take the $A_1, \dots, A_k \in \mathcal{F}$ with the samllest intersection. Since $\mathcal{F}$ is an independent set, we are guaranteed that this intersection will have cardinality ≥ 1.

The 'bad case' would be when all k-tuples of sets in $\mathcal{F}$ have large intersection. We recall terminology from homework 2 to refer to such bad sets:

Definition 2.5. A collection $\mathcal{F}$ is *k-wise t-intersecting* if for all $A_1, \dots, A_k \in \mathcal{F}$, $|\bigcap_{i=1}^k A_i| \geq t$.

What we hope is that if a collection $\mathcal{F}$ is k-wise t-intersecting for large t, it will have to contain many 1's (and so have tiny weight). We will need to prove a theorem of the following form:

Theorem 2.6. *Let constants $\delta, \epsilon \geq 0$. Suppose $\mathcal{F} \subseteq \{0,1\}^L$ has p-biased weight $\geq \epsilon$ where $p = 1 - 1/k - \delta$. Then $\mathcal{F}$ cannot be[4] k-wise t-intersecting when $t \geq O(1/\delta^2 \log(1/(\epsilon\delta k)))$.*

Note that in the above dream theorem, t is a constant that depends only on ϵ and δ.[5]

Enough for intuition. Let's begin the reduction.

3 The Reduction (And Completeness)

Theorem 3.1. *For all constant even valued $k \geq 4$, and all constants $\epsilon, \delta > 0$, the Weighted Max-Ek-Independent-Set $1 - 2/k - \delta$ vs. ϵ decision problem is NP-hard.*

Proof of completeness. We reduce from Label-Cover(K, L). Given a Label-Cover instance $\mathcal{G}$ defined over the bipartite graph $G = (U \cup V, E)$, we create a set of vertices $V \times \{0,1\}^{|L|}$ in Max-Ek-Independent-Set instance $\mathcal{H}$. (So we have $|V| \cdot 2^{|L|}$ vertices). We assign p-biased weights on the vertices with $p = 1 - 2/k - \delta$.[6] For every pair of edges $(u, v), (u, v')$ in $\mathcal{G}$ that share the same endpoint in U, we create a hyperedge consisting of k sets $\{A_1^{(v)}, \dots, A_{k/2}^{(v)}, B_1^{(v')}, \dots, B_{k/2}^{(v')}\}$ if and only if the sets $\pi_{v \to u}(\cap A_i^{(v)})$ and $\pi_{v' \to u}(\cap B_i^{(v')})$ are disjoint (recall that the intersection of strings gives a set of coordinates, which correspond to labels in the Label-Cover instance).

We prove completeness, and leave soundness for the next lecture: Suppose $\mathsf{Opt}(\mathcal{G}) = 1$. Then there exists a function $f : V \to L, U \to K$ satisfying all constraints in $\mathcal{G}$. We must show $\mathsf{Opt}(\mathcal{H}) \geq p$.

For each $v \in V$, choose $\mathcal{D}_{f(v)}$, which has weight $p \cdot |V|$ (what we want!) Why must this be an independent set? Suppose it were not, and contained every vertex in some hyperedge $\{A_1^{(v)}, \dots, A_{k/2}^{(v)}, B_1^{(v')}, \dots, B_{k/2}^{(v')}\}$. We know $f(v) \in A_1^{(v)} \cap \dots \cap A_{k/2}^{(v)}$, and $f(v') \in B_1^{(v')} \cap \dots \cap B_{k/2}^{(v')}$, but this implies that $\pi_{v \to u}(f(v)) \in \pi_{v \to u}(\cap A_i^{(v)})$ and $\pi_{v' \to u}(f(v')) \in \pi_{v' \to u}(\cap B_i^{(v')})$. Since we know $\pi_{v' \to u}(f(v')) = \pi_{v \to u}(f(v)) = f(u)$, this can't have been a hyperedge, contradicting our assumption! □

[4] is too *big* to be

[5] Well, okay, also on k, but t gets *smaller* as k grows!

[6] the careful observer will notice that our weights actually sum to $|V|$, not 1, but the careful observer may simply divide by $|V|$ if she wishes.

References

[1] I. Dinur, V. Guruswami, S. Khot, and O. Regev. A New Multilayered PCP and the Hardness of Hypergraph Vertex Cover. *SIAM Journal on Computing*, 34:1129, 2005.