Prague Summer School: Discrepancy Theory Summer 2018

Lecture 5

Aleksandar Nikolov

1 Discrepancy and Rounding

Using the connection between discrepancy and rounding, in this lecture we will give an approximation algorithm for (a special case of) the bin packing problem. The approach will be to solve a linear programming relaxation of the problem, and then round the linear program solution using discrepancy. The main fact that enables this approach to work is the following lemma we saw back in the first lecture.

Lemma 1. *For any $m \times n$ matrix A, and any $w \in [0,1]^n$, there exists a $x \in \{0,1\}^n$ such that* $\|Ax - Aw\|_\infty \leq \mathrm{herdisc}(A)$.

Linear discrepancy in fact allows us to round any real vactor, and not just vectors with entries in $[0,1]$. This is captures by the following proposition:

Proposition 2. $\forall x \in \mathbb{R}^n_+, \exists y \in \mathbb{Z}^n_+$ *such that* $\|Ax - Ay\|_\infty \leq \mathrm{herdisc}(A)$

Proof. Let $\bar{x}_i = x_i - \lfloor x_i \rfloor$, then $\bar{x}_i \in [0,1]$, and there exists $\bar{y} \in \{0,1\}^n$ and $y = \bar{y} + \lfloor x \rfloor$ (where the floor is applied coordinate-wise) such that:

$$\begin{aligned} \|A\bar{x} - A\bar{y}\|_\infty &\leq \mathrm{herdisc}(A) \\ \|(A\bar{x} + A\lfloor x \rfloor) - (A\bar{y} + A\lfloor x \rfloor)\|_\infty &\leq \mathrm{herdisc}(A) \\ \|Ax - Ay\|_\infty &\leq \mathrm{herdisc}(A) \end{aligned}$$

□

2 Approximation to Bin Packing

The bin packing problem takes as input a set of n items of sizes $s_1, s_2, \ldots, s_n$, where each $s_i \in [0,1]$. The goal is to pack the items into the smallest number of bins, each of size 1. Let us use OPT to denote the optimal number of bins needed to pack all items. Bin packing is NP-hard, and remains NP-hard to decide whether we need 2 vs. 3 bins (i.e. whether OPT ≤ 2 or ≥ 3). Before getting started, we list the current known approximations:

- The First-Fit (arbitrary order) algorithm gives a 2-approximation.
- The First-Fit in decreasing order gives a 1.7 OPT + 1 [2].
- De La Vega and Lucker gave an asymptotic PTAS that uses $\leq (1+\epsilon)\mathrm{OPT} + 1$ bins, for any given $\epsilon \geq 0$.

- Karmarkar and Karp (KK) gave an algorithm that achieves $\le \text{OPT} + O(\log^2(\text{OPT}))$ many bins [3].
- Rothvoss used discrepancy to give an algorithm that uses $\le \text{OPT}+O(\log(\text{OPT}) \log\log(OPT))$ bins [4].
- Rothvoss and Hoberg refined the discrepancy approach to give an algorithm that uses $\le \text{OPT} + O(\log(\text{OPT}))$ bins [1].

For the special case when $s_i > \frac{1}{k+1}, \forall i$ (i.e. we can place at most k items per bin), and k is a constant, the KK result already gives $\text{OPT} + O(\log(\text{OPT}))$ bins. Hoberg and Rothvoss's result essentially reduces the general case to this case. It is open whether there exists an efficient algorithm that packs all items in $\text{OPT} + 1$ bins (!).

For the remainder of the lecture, we will show how to get the KK bound for $s_i > \frac{1}{4}, \forall i$ using discrepancy on the Gilmore Gomory LP relaxation, defined below. We will show that we can pack all items in $\text{OPT} + O(\log n)$ many bins. There is a standard reduction that then implies that we can do so with $\text{OPT} + O(\log \text{OPT})$ many bins.

For convenience, let us assume all the weights are sorted in decreasing order:

$$1 \ge s_1 \ge s_2 \ge \ldots \ge s_n > \frac{1}{4}$$

Let $s = (s_1, s_2, \ldots, s_n)$ denote the sorted vector, and let $\mathcal{P}_s$ be all possible ways to pack the items:

$$\mathcal{P}_s = \left\{ p \in \{0,1\}^n \;\; : \;\; p^T s = \sum p_i s_i \le 1 \right\}$$

Think of p as the indicator of a feasible set of items which we can assign to a single bin. Since we can only fit at most 3 items per bin, it follows that $|\mathcal{P}_s| = O(n^3)$. Now consider the following linear program, which is a relaxation of the bin packing problem:

$$\begin{aligned} \text{minimize} \quad & \sum x_p & & (\text{\# of bins}) \\ \text{subject to} \quad & \sum x_p \cdot p \ge 1 & & (\text{each item in at least a bin}) \\ & x_p \ge 0 \;\; \forall p \in \mathcal{P}_s \end{aligned}$$

If we denote by LP the value of this linear program, then, because every bin packing gives a solution to the program, we have $\text{LP} \le \text{OPT}$.

Let x be an optimal solution to the lienar program such that $|\{p : x_p > 0\}| \le n$. That such an optimal solution always exists follows from the theory of basic feasible solutions. Let B be the constraint matrix of the program, with columns restricted to p such that $x_p > 0$, and let us replace x with x restricted to these patterns p, as well. Given the item size constraint, notice that each column in B has at most 3 zeroes:

$$B = \begin{pmatrix} \vdots & & \vdots \\ p_1 & \ldots & p_n \\ \vdots & & \vdots \end{pmatrix} \qquad \{p_1 \ldots p_n\} = \{p : x_p > 0\}$$

$$Bx = 1$$

Let A be a matrix constructed as follows:

$$A = \begin{pmatrix} T_n B \\ 3 \dots 3 \end{pmatrix}$$

where T_n is the lower triangular $\{0,1\}$ matrix:

$$\begin{pmatrix} 1 & \dots & 0 \\ 1 & \ddots & 0 \\ 1 & \dots & 1 \end{pmatrix}$$

This means that each of the first n rows of A is of the form $A_{i*} = \sum_{j=1}^{i} B_{j*}$. Since the columns of B have at most 3 ones each, and all other entries are 0, it follows that $A \in \{0,1,2,3\}^{(n+1)\times n}$. Moreover, each column in A is monotone non-decreasing. Because $Bx = 1$, we have:

$$Ax = \begin{pmatrix} 1 \\ 2 \\ \vdots \\ n \\ 3 \cdot LP \end{pmatrix}$$

where $LP \leq OPT$. Assume $x_p \in [0,1)$, and let D denote herdisc(A). By Proposition 2,

$$\exists y \in \mathbb{Z}_+^n \text{ s.t. } \|Ax - Ay\|_\infty \leq D$$
$$|3 \cdot \mathbf{1}^T y - 3 \cdot \mathbf{1}^T x| \leq D$$

The second inequality implies $1^T y \leq \text{LP} + \frac{1}{3}D \leq \text{OPT} + \frac{1}{3}D$.

Claim 3. *We can pack all items in $D + \sum_p y_p$ many bins.*

Proof. We will use y_p copies of pattern p, and additional D copies of the pattern that contains a space for item 1. We will pack items according to these patterns, except that we will also allow ourselves to put a smaller item in the space reserved for a larger item. Since item 1 is the largest one[1], we can pack any item in a space reserved for it.

We construct a bipartite graph $G(V = \mathcal{U} \cup \mathcal{V}, E)$ as follows: Every vertex $u_i \in \mathcal{U}$ corresponds to an item i of size s_i. For item 1, we create $(By)_1 + D$ copies of v_1, which is the number of slots where item 1 can fit (or any item with size $\leq s_1$). For the remaining items, create $(By)_i$ copies of vertex v_i for each item i (i.e. the # of slots for items of size $\leq s_i$). For every $u_i \in \mathcal{U}$, add an edge

[1] recall that items are sorted in non-increasing order

$(u_i, v_j) \in E$ for all $j \leq i$. The goal is to compute a $\mathcal{U}$-perfect matching. Then, for any edge (u_i, v_j) of the matching, we will pack item i in any slot reserved for item j. Since the matching is perfect, we can pack all items.

The proof that a $\mathcal{U}$-perfect matching exists is based on Hall's theorem, which says:

Theorem 4. *$\mathcal{U}$-perfect matching exists if and only if for all $\mathcal{W} \subseteq \mathcal{U}, |N(\mathcal{W})| \geq |\mathcal{W}|$, where $N(\mathcal{W})$ is the set of neighbors of $\mathcal{W}$.*

Now we verify the condition of Hall's theorem. Assume $\mathcal{W} = \{u_1, \ldots, u_i\}$. This is without loss of generality, because $\forall j \leq i$, $N(u_j) \subseteq N(u_i)$ by construction. Therefore

$$\begin{aligned}|N(\mathcal{W})| &= \sum_{j=1}^{i} (By)_j + D \quad \text{for } j = 1 \\ &= (Ay)_i + D \\ &\geq (Ax)_i = i = |\mathcal{W}|\end{aligned}$$

Therefore a $\mathcal{U}$-perfect matching exists and all items can be packed, and the cost of the solution is at most $\leq$ LP $+ (1 + \frac{1}{3})D$. □

On the other hand, for all A with monotone non-decreasing columns and entries in $[k]$, we have $\text{herdisc}(A) = O(k \log n)$, thus:

$$\begin{aligned}D &= \text{herdisc}(A) \\ &\leq \text{herdisc}(A) \\ &= O(3 \log n) = O(\log n)\end{aligned}$$

Let us prove a somewhat weaker bound using an exercise from a previous lecture. Let $\pi_1, \ldots, \pi_k$ be three permutations on $\{1, \ldots, n\}$ and let $\mathcal{S}$ be sets of the type $\{\pi_i(1), \ldots, \pi_i(j)\}$, i.e. prefixes of the permutations. Recall that we used the partial coloring result of Lovett and Meka to show that such a set system has discrepancy $O(\sqrt{k} \log n)$ for any choice of $\pi_1, \ldots, \pi_k$. We use this result in the following proposition.

Proposition 5. *For any matrix A with monotone non-decreasing columns and entries in $[k]$ we have* $\text{herdisc}(A) = O(k^{3/2} \log n)$.

Proof. Let us bound the discrepancy of an arbitrary matrix A satisfying the condition of the proposition. Since this class of matrices is closed under taking submatrices, this automatically bounds the hereditary discrepancy as well.

We decompose A into k matrices where the entries are of the form:

$$A_{ij}^{(\ell)} = \begin{cases} 1 \text{ if } A_{ij} \geq \ell \\ 0 \text{ otherwise} \end{cases}$$

Therefore $A = A^{(1)} + \ldots + A^{(k)}$. Since A has monotone columns, thus each $A^{(\ell)}$ has monotone columns. Therefore the matrix

$$\tilde{A} = \begin{pmatrix} A^{(1)} \\ \vdots \\ A^{(k)} \end{pmatrix}$$

is the incidence matrix of a set system $\mathcal{S}$ consisting of the prefixes of k permutations, as we defined it above. Then $\mathrm{disc}(\mathcal{S}) = \mathrm{disc}(\tilde{A}) = O(\sqrt{k}\log n)$, i.e. there exists an $x \in \{-1,1\}^n$ such that $\|\tilde{A}x\|_\infty = O(\sqrt{k}\log n)$, or equivalently, $\|A^{(\ell)}x\|_\infty = O(\sqrt{k}\log n)$ for all ℓ. It follows that

$$\|Ax\|_\infty = \|A^{(1)}x + \ldots + A^{(\ell)}x\|_\infty \leq \|A^{(1)}x\|_\infty + \ldots + \|A^{(\ell)}x\|_\infty = O(k^{3/2}\log n),$$

as we wanted. □

This proposition and the claim we proved earlier mean that we can pack all items in number of bins bounded by

$$\sum_p y_p + D \leq \mathrm{OPT} + \frac{4}{3}D = \mathrm{OPT} + O(\log n).$$

References

[1] Rebecca Hoberg and Thomas Rothvoss. A logarithmic additive integrality gap for bin packing. *CoRR*, abs/1503.08796, 2015.

[2] David S Johnson. *Near-optimal bin packing algorithms*. PhD thesis, Massachusetts Institute of Technology, 1973.

[3] Narendra Karmarkar and Richard M Karp. An efficient approximation scheme for the one-dimensional bin-packing problem. In *Foundations of Computer Science, 1982. SFCS'08. 23rd Annual Symposium on*, pages 312–320. IEEE, 1982.

[4] Thomas Rothvoß. Approximating bin packing within o(log OPT * log log OPT) bins. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 20–29. IEEE Computer Society, 2013.