Bin Packing via Discrepancy of Permutations

Friedrich Eisenbrand* Dömötör Pálvölgyi† Thomas Rothvoß‡

Abstract

A well studied special case of *bin packing* is the *3-partition problem*, where n items of size $> \frac{1}{4}$ have to be packed in a minimum number of bins of capacity one. The famous *Karmarkar-Karp algorithm* transforms a fractional solution of a suitable LP relaxation for this problem into an integral solution that requires at most $O(\log n)$ additional bins.

The *three-permutations-conjecture* of Beck is the following. Given any 3 permutations on n symbols, one can color the symbols red and blue, such that in any interval of any of those permutations, the number of red and blue symbols differs only by a constant. Beck's conjecture is well known in the field of *discrepancy theory*.

We establish a surprising connection between bin packing and Beck's conjecture: If the latter holds true, then the additive integrality gap of the 3-partition linear programming relaxation is bounded by a constant.

1 Introduction

The *bin packing* problem is the following. Given n items of size $s_1, \dots, s_n \in [0,1]$ respectively, the goal is to pack these items in as few bins of capacity one as possible. Bin packing is a fundamental problem in Computer Science with numerous applications in theory and practice.

The development of heuristics for bin packing with better and better performance guarantee is an important success story in the field of *Approximation Algorithms*. Johnson [14, 15] has shown that the *First Fit* algorithm requires at most $1.7 \cdot OPT + 1$ bins and that *First Fit Decreasing* yields a solution with $\frac{11}{9}OPT + 4$ bins (see [7] for a tight bound of $\frac{11}{9}OPT + \frac{6}{9}$). An important step forward was made by Fernandez de la Vega and Luecker [9] who provided an asymptotic polynomial time approximation scheme for bin packing. The *rounding technique* that is introduced in their paper has been very influential in the design of PTAS's for many other difficult combinatorial optimization problems.

In 1982, Karmarkar and Karp [16] proposed an approximation algorithm for bin packing that can be analyzed to yield a solution using at most $OPT + O(\log^2 n)$ bins. This seminal procedure is based on the *Gilmore Gomory LP relaxation* [11, 8]:

$$\text{(LP)} \qquad \begin{array}{lcl} \min \sum_{p \in \mathcal{P}} x_p & & \\ \sum_{p \in \mathcal{P}} p \cdot x_p & \geq & \mathbf{1} \\ x_p & \geq & 0 \quad \forall p \in \mathcal{P} \end{array}$$

Here $\mathbf{1} = (1, \dots, 1)^T$ denotes the all ones vector and $\mathcal{P} = \{p \in \{0,1\}^n \colon s^T p \leq 1\}$ is the set of all feasible *patterns*, i.e. every vector in $\mathcal{P}$ denotes a feasible way to pack one bin. Let OPT and OPT_f be the value of the best integer and fractional solution respectively. The linear program (LP) has an exponential number of variables but still one can compute a basic solution x with $\mathbf{1}^T x \leq OPT_f + \delta$ in time polynomial in n and $1/\delta$ [16] using the Grötschel-Lovász-Schrijver variant of the Ellipsoid method [12].

The procedure of Karmarkar and Karp [16] yields an *additive integrality gap* of $O(\log^2 n)$, i.e. $OPT \leq OPT_f + O(\log^2 n)$, see also [24]. This corresponds to an asymptotic FPTAS[1] for bin packing. The authors in [19] conjecture that even $OPT \leq \lceil OPT_f \rceil + 1$ holds and this even if one replaces the right-hand-side $\mathbf{1}$ by any other positive integral vector b. This *Modified Integer Round-up Conjecture* was proven by Sebő and Shmonin [20] if the number of different item sizes is at most 7. We would like to mention that Jansen and Solis-Oba [13] recently provided an $OPT + 1$ approximation-algorithm for bin packing if the number of item sizes is fixed.

Much of the hardness of bin packing seems to appear already in the special case of *3-partition*, where $3n$ items of size $\frac{1}{4} < s_i < \frac{1}{2}$ with $\sum_{i=1}^{3n} s_i = n$ have to be packed. It is strongly **NP**-hard to distinguish between $OPT \leq n$ and $OPT \geq n + 1$ [10]. No stronger hardness result is known for general bin packing. A closer look into [16] reveals that, with the restriction $s_i > \frac{1}{4}$, the Karmarkar-Karp algorithm uses $OPT_f + O(\log n)$ bins[2].

*EPFL, Lausanne, Switzerland, friedrich.eisenbrand@epfl.ch. Supported by the Swiss National Science Foundation (SNF).

†Eötvös Loránd University (ELTE), Budapest, Hungary, dom@cs.elte.hu

‡EPFL, Lausanne, Switzerland, thomas.rothvoss@epfl.ch. Supported by the German Research Foundation (DFG) within the Priority Program 1307 "Algorithm Engineering".

[1] An asymptotic fully polynomial time approximation scheme (AFPTAS) is an approximation algorithm that produces solutions of cost at most $(1+\varepsilon)OPT + p(1/\varepsilon)$ in time polynomial in n and $1/\varepsilon$, where also p must be a polynomial.

[2] The *geometric grouping* procedure (Lemma 5 in [16]) discards

Discrepancy theory. Let $[n] := \{1, \ldots, n\}$ and consider a set system $\mathcal{S} \subseteq 2^{[n]}$ over the ground set $[n]$. A *coloring* is a mapping $\chi : [n] \to \{\pm 1\}$. In *discrepancy theory*, one aims at finding colorings for which the difference of "red" and "blue" elements in all sets is as small as possible. Formally, the *discrepancy* of a set system $\mathcal{S}$ is defined as

$$\mathrm{disc}(\mathcal{S}) = \min_{\chi:[n]\to\{\pm 1\}} \max_{S\in\mathcal{S}} |\chi(S)|.$$

where $\chi(S) = \sum_{i\in S} \chi(i)$. A random coloring provides an easy bound of $\mathrm{disc}(\mathcal{S}) \leq O(\sqrt{n \log |\mathcal{S}|})$ [18]. The famous "*Six Standard Deviations suffice*" result of Spencer [21] improves this to $\mathrm{disc}(\mathcal{S}) \leq O(\sqrt{n \log(2|\mathcal{S}|/n)})$.

If every element appears in at most t sets, then the *Beck-Fiala Theorem* [3] yields $\mathrm{disc}(\mathcal{S}) < 2t$. The same authors conjecture that in fact $\mathrm{disc}(\mathcal{S}) = O(\sqrt{t})$. Srinivasan [23] gave a $O(\sqrt{t} \log n)$ bound, which was improved by Banaszczyk [1] to $O(\sqrt{t \log n})$. Many such discrepancy proofs are purely existential, for instance due to the use of the pigeonhole principle. In a very recent breakthrough Bansal [2] showed how to obtain the desired colorings for the Spencer [21] and Srinivasan [23] bounds by considering a random walk, guided by the solution of a semidefinite program.

The following conjecture is coined *three-permutations-conjecture* or simply *Beck's conjecture* (see Problem 1.9 in [4]):

> Given any 3 permutations on n symbols, one can color the symbols with red and blue, such that in every interval of every of those permutations, the number of red and blue symbols differs by $O(1)$.

A set of permutations $\pi_1, \ldots, \pi_k : [n] \to [n]$ induces a set-system[3]

$$\mathcal{S} = \{\{\pi_i(1), \ldots, \pi_i(j)\} : j = 1, \ldots, n;\ i = 1, \ldots, k\}.$$

If we denote the maximum discrepancy of such a set-system induced by k permutations over n symbols as $D_k^{\mathrm{perm}}(n)$, then Beck's conjecture can be rephrased as $D_3^{\mathrm{perm}}(n) = O(1)$.

So far the best known bound on $D_3^{\mathrm{perm}}(n)$ is $O(\log n)$ and more generally $D_k^{\mathrm{perm}}(n)$ can be bounded by $O(k \log n)$ [5] and by $O(\sqrt{k} \log n)$ [23, 22] using the so-called *entropy method*.

items of size $O(\log \frac{1}{s_{\min}})$, where $s_{\min}$ denotes the size of the smallest item. The geometric grouping is applied $O(\log n)$ times in the Karmarkar-Karp algorithm. The claim follows by using that $s_{\min} > \frac{1}{4}$ for 3-partition.

[3] We only consider intervals of permutations that start from the first element. Since any interval is the difference of two such prefixes, this changes the discrepancy by a factor of at most 2.

Our contribution. The main result of this paper is a proof of the following theorem.

THEOREM 1.1. *If Beck's conjecture holds, then the integrality gap of the linear program* (LP) *restricted to 3-partition instances is bounded by an additive constant.*

This result is constructive in the following sense. If one can find a constant discrepancy coloring for any three permutations in polynomial time, then there is an $OPT + c$ approximation algorithm for 3-partition for a constant c.

The proof of Theorem 1.1 itself is via two steps.

i) We show that the additive integrality gap of (LP) is at most twice the maximum *linear discrepancy* of a k-*monotone matrix* if all item sizes are larger than $1/(k+1)$ (Section 3). This step is based on matching techniques and *Hall's theorem.*

ii) We then show that the linear discrepancy of a k-monotone matrix is at most k times the discrepancy of k permutations (Section 4.1). This result uses a theorem of Lovász, Spencer and Vesztergombi.

The theorem then follows by setting k equal to 3 in the above steps. Furthermore, we show that the discrepancy of k permutations is at most 4 times the linear discrepancy of a k-monotone matrix. And finally, we provide a $5k \cdot \log_2(2\min\{m, n\})$ upper bound on the linear discrepancy of a k-monotone $n \times m$-matrix.

2 Preliminaries

We first review some further necessary preliminaries on discrepancy theory. We refer to [18] for further details.

If A is a matrix, then we denote the ith row of A by A_i and the jth entry in the ith row by A_{ij}. The notation of discrepancy can be naturally extended to real matrices $A \in \mathbb{R}^{m\times n}$ as

$$\mathrm{disc}(A) := \min_{x\in\{0,1\}^n} \|A(x - 1/2 \cdot \mathbf{1})\|_\infty,$$

see, e.g. [18]. Note that if A is the incidence matrix of a set system $\mathcal{S}$ (i.e. each row of A corresponds to the characteristic vector of a set $S \in \mathcal{S}$), then $\mathrm{disc}(A) = \frac{1}{2}\mathrm{disc}(\mathcal{S})$, hence this notation is consistent — apart from the $\frac{1}{2}$ factor.

The *linear discrepancy* of a matrix $A \in \mathbb{R}^{m\times n}$ is defined as

$$\mathrm{lindisc}(A) := \max_{y\in[0,1]^n} \min_{x\in\{0,1\}^n} \|Ax - Ay\|_\infty.$$

This value can be also described by a two player game. The first player chooses a fractional vector y, then the second player chooses a 0/1 vector x. The goal of the first player is to maximize, of the second to minimize $\|Ax-Ay\|_\infty$. The inequality $\mathrm{disc}(A) \le \mathrm{lindisc}(A)$ holds by choosing $y := (1/2,\dots,1/2)$. One more notion of defining the "complexity" of a set system or a matrix is that of the *hereditary discrepancy*:

$$\mathrm{herdisc}(A) := \max_{B \text{ submatrix of } A} \mathrm{disc}(B).$$

Notice that one can assume that B is formed by choosing a subset of the columns of A. This parameter is obviously at least $\mathrm{disc}(A)$ since we can choose $B := A$ and in [17] even an upper bound for $\mathrm{lindisc}(A)$ is proved (see again [18] for a recent description).

THEOREM 2.1. (LOVÁSZ, SPENCER, VESZTERGOMBI) *For $A \in \mathbb{R}^{m\times n}$ one has*

$$\mathit{lindisc}(A) \le 2 \cdot \mathit{herdisc}(A).$$

3 Bounding the gap via the discrepancy of monotone matrices

A matrix A is called *k-monotone* if all its column vectors have non-decreasing entries from $0,\dots,k$. In other words $A \in \{0,\dots,k\}^{m\times n}$ and $A_{1j} \le \dots \le A_{mj}$ for any column j. We denote the maximum linear discrepancy of such matrices by

$$D_k^{\mathrm{mon}}(n) := \max_{\substack{A\in\mathbb{Z}^{m\times n}\\ k\text{-monotone}}} \mathrm{lindisc}(A).$$

The next theorem establishes step i) mentioned in the introduction.

THEOREM 3.1. *Consider the linear program* (LP) *and suppose that the item sizes satisfy $s_1,\dots,s_n > \frac{1}{k+1}$. Then*

$$OPT \le OPT_f + \left(1+\frac{1}{k}\right) D_k^{mon}(n).$$

Proof. Assume that the item sizes are sorted such that $s_1 \ge \dots \ge s_n$. Let y be any optimum basic solution of (LP) and let $p_1,\dots,p_m$ be the list of patterns. Since y is a basic solution, its support satisfies $|\{i\colon y_i > 0\}| \le n$. Hence by deleting unused patterns, we may assume [4] that $m = n$.

We define $B = (p_1,\dots,p_n) \in \{0,1\}^{n\times n}$ as the matrix composed of the patterns as column vectors. Clearly $By = \mathbf{1}$. Let A be the matrix that is defined by $A_i := \sum_{j=1}^{i} B_j$, again A_i denotes the ith row of A. In other words, A_{ij} denotes the number of items of types $1,\dots,i$ in pattern p_j. Since $By = \mathbf{1}$ we have $Ay = (1,2,3,\dots,n)^T$. Each column of A is monotone. Furthermore, since no pattern contains more than k items one has $A_{ij} \in \{0,\dots,k\}$, thus A is k-monotone.

[4] In case that there are less than n patterns, we add empty patterns.

We attach a row $A_{n+1} := (k,\dots,k)$ as the new last row of A. Clearly A remains k-monotone. There exists a vector $x \in \{0,1\}^n$ with

$$\|Ax - Ay\|_\infty \le \mathrm{lindisc}(A) \le D_k^{\mathrm{mon}}(n).$$

We *buy* x_i times pattern p_i and $D_k^{\mathrm{mon}}(n)$ times the pattern that only contains the largest item of size s_1.

It remains to show: (1) this yields a feasible solution; (2) the number of patterns does not exceed the claimed bound of $OPT_f + (1+\frac{1}{k}) \cdot D_k^{\mathrm{mon}}(n)$.

For the latter claim, recall that the constraint emerging from row $A_{n+1} = (k,\dots,k)$ together with $\sum_{i=1}^{n} y_i = OPT_f$ provides

$$k\sum_{i=1}^{n} x_i \le k \cdot \sum_{i=1}^{n} y_i + D_k^{\mathrm{mon}}(n) = k \cdot OPT_f + D_k^{\mathrm{mon}}(n).$$

We use this to upper bound the number of opened bins by

$$\sum_{i=1}^{n} x_i + D_k^{\mathrm{mon}}(n) \le OPT_f + \left(1+\frac{1}{k}\right) \cdot D_k^{\mathrm{mon}}(n).$$

It remains to prove that our integral solution is feasible. To be more precise, we need to show that any item i can be assigned to a space reserved for an item of size s_i or larger.

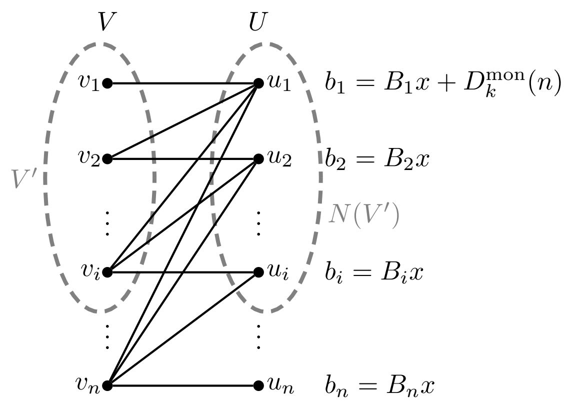

Figure 1: The bipartite graph in the proof of Theorem 3.1

To this end, consider a bipartite graph with nodes $V = \{v_1,\dots,v_n\}$ on the left, representing the items. The nodes on the right are the set $U = \{u_1,\dots,u_n\}$, where each u_i is attributed with a multiplicity b_i representing the number of times that we reserve space for items of size s_i in our solution, see Figure 1. Recall that

$$b_i = \begin{cases} B_i x + D_k^{\mathrm{mon}}(n) & \text{if } i = 1 \\ B_i x & \text{otherwise} \end{cases}.$$

We insert an edge (v_i, u_j) for all $i \geq j$. The meaning of this edge is the following. One can assign item i into the space which is reserved for item j since $s_i \leq s_j$. We claim that there exists a V-perfect matching, respecting the multiplicities of U. By *Hall's Theorem*, see, e.g. [6], it suffices to show for any subset $V' \subseteq V$ that the multiplicities of the nodes in $N(V')$ (the neighborhood of V') are at least $|V'|$. Observe that $N(v_i) \subseteq N(v_{i+1})$, hence it suffices to prove the claim for sets of the form $V' = \{1, \ldots, i\}$. For such a V' one has

$$\begin{aligned} \sum_{u_j \in N(V')} b_j &= D_k^{\mathrm{mon}}(n) + \sum_{j=1}^{i} B_j x \\ &= D_k^{\mathrm{mon}}(n) + A_i x \geq A_i y = i \end{aligned}$$

and the claim follows.

4 Bounding the discrepancy of monotone matrices by the discrepancy of permutations

In this section, we show that the linear discrepancy of k-monotone matrices is essentially bounded by the discrepancy of k permutations. This corresponds to step ii) in the proof of the main theorem. By Theorem 2.1 it suffices to bound the discrepancy of k-monotone matrices by the discrepancy of k permutations times a suitable factor.

We first explain how one can associate a permutation to a 1-monotone matrix. Suppose that $B \in \{0,1\}^{m \times n}$ is a 1-monotone matrix. If B^j denotes the j-th column of B, then the permutation π that we associate with B is the (not necessarily unique) permutation that satisfies $B^{\pi(1)} \geq B^{\pi(2)} \geq \cdots \geq B^{\pi(n)}$ where $u \geq v$ for vectors $u, v \in \mathbb{R}^m$ if $u_i \geq v_i$ for all $1 \leq i \leq m$. On the other hand the matrix B (potentially plus some extra rows and after merging identical rows) gives the incidence matrix of the set-system induced by π.

A k-monotone matrix B can be decomposed into a sum of 1-monotone matrices $B^1, \ldots, B^k$. Then any B^ℓ naturally corresponds to a permutation π_ℓ of the columns as we explained above. A low-discrepancy coloring of these permutations yields a coloring that has low discrepancy for any B^ℓ and hence also for B, as we show now in detail.

Theorem 4.1. *For any $k, n \in \mathbb{N}$, one has $D_k^{mon}(n) \leq k \cdot D_k^{perm}(n)$.*

Proof. Consider any k-monotone matrix $A \in \mathbb{Z}^{m \times n}$. By virtue of Theorem 2.1, there is a $m \times n'$ submatrix, B, of A such that $\mathrm{lindisc}(A) \leq 2 \cdot \mathrm{disc}(B)$, thus it suffices to show that $\mathrm{disc}(B) \leq \frac{k}{2} \cdot D_k^{\mathrm{perm}}(n)$. Of course, B itself is again k-monotone.

Let B^ℓ also be a $m \times n'$ matrix, defined by

$$B_{ij}^\ell := \begin{cases} 1 & \text{if } B_{ij} \geq \ell \\ 0, & \text{otherwise.} \end{cases}$$

The matrices B^ℓ are 1-monotone, and the matrix B decomposes into $B = B^1 + \ldots + B^k$. As mentioned above, for any ℓ, there is a (not necessarily unique) permutation π_ℓ on $[n']$ such that $B^{\ell,\pi_\ell(1)} \geq B^{\ell,\pi_\ell(2)} \geq \ldots \geq B^{\ell,\pi_\ell(n')}$, where $B^{\ell,j}$ denotes the jth column of B^ℓ. Observe that the row vector B_i^ℓ is the characteristic vector of the set $\{\pi_\ell(1), \ldots, \pi_\ell(j)\}$, where j denotes the number of ones in B_i^ℓ.

Let $\chi : [n'] \to \{\pm 1\}$ be the coloring that has discrepancy at most $D_k^{\mathrm{perm}}(n)$ with respect to all permutations $\pi_1, \ldots, \pi_k$. In particular $|B_i^\ell \chi| \leq D_k^{\mathrm{perm}}(n)$, when interpreting χ as a ± 1 vector. Then by the triangle inequality

$$\mathrm{disc}(B) \leq \frac{1}{2} \|B\chi\|_\infty \leq \frac{1}{2} \sum_{\ell=1}^{k} \|B^\ell \chi\|_\infty \leq \frac{k}{2} D_k^{\mathrm{perm}}(n).$$

Combining Theorem 3.1 and Theorem 4.1, we conclude

Corollary 4.1. *Given any bin packing instance with n items of size bigger than $\frac{1}{k+1}$ one has*

$$OPT \leq OPT_f + 2k \cdot D_k^{perm}(n).$$

In particular, this proves Theorem 1.1, our main result.

Bounding the discrepancy of permutations in terms of the discrepancy of monotone matrices. In addition we would like to note that the discrepancy of permutations can be also bounded by the discrepancy of k-monotone matrices as follows.

Theorem 4.2. *For any $k, n \in \mathbb{N}$, one has $D_k^{perm}(n) \leq 4 \cdot D_k^{mon}(n)$.*

Proof. We will show that for any permutations $\pi_1, \ldots, \pi_k$ on $[n]$, there is a $kn \times n$ k-monotone matrix C with $\mathrm{disc}(\pi_1, \ldots, \pi_k) \leq 4 \cdot \mathrm{disc}(C)$. Let $\Sigma \in \{1, \ldots, n\}^{kn}$ be the string which we obtain by concatenating the k permutations. That means $\Sigma = (\pi_1(1), \ldots, \pi_1(n), \ldots, \pi_k(1), \ldots, \pi_k(n))$. Let C the matrix where C_{ij} is the number of appearances of $j \in \{1, \ldots, n\}$ among the first $i \in \{1, \ldots, kn\}$ entries of Σ. By definition, C is k-monotone, in fact it is the "same" k-monotone matrix as in the previous proof.

Choose $y := (\frac{1}{2}, \dots, \frac{1}{2})$ to have $Cy = (\frac{1}{2}, 1, \dots, \frac{kn}{2})$. Let $x \in \{0,1\}^n$ be a vector with $\|Cx - Cy\|_\infty \le \text{disc}(C)$. Consider the coloring $\chi : [n] \to \{\pm 1\}$ with $\chi(j) := 1$ if $x_j = 1$ and $\chi(j) := -1$ if $x_j = 0$. We claim that the discrepancy of this coloring is bounded by $4 \cdot \text{disc}(C)$ for all k permutations. Consider any prefix $S := \{\pi_i(1), \dots, \pi_i(\ell)\}$. Let $r = C_{(i-1)n+\ell} \in \{i-1, i\}^n$ be the row of C that corresponds to this prefix. With these notations we have

$$\begin{aligned} |\chi(S)| &\le |(r - (i-1)\mathbf{1}) \cdot (2x - 2y)| \\ &\le 2 \cdot \big(\underbrace{|r(x-y)|}_{\le \text{disc}(C)} + \underbrace{|k \cdot \mathbf{1}(x-y)|}_{\le \text{disc}(C)} \big) \\ &\le 4 \cdot \text{disc}(C). \end{aligned}$$

Here the inequality $|(k \cdot \mathbf{1}) \cdot (x - y)| \le \text{disc}(C)$ comes from the fact that $k \cdot \mathbf{1} = (k, \dots, k)$ is the last row of C.

From Theorem 4.1 and Theorem 4.2, we obtain that the following conjecture is equivalent to the Three-Permutations-Conjecture.

Conjecture 4.1. *There exists a constant $c > 0$, such that for any 3-monotone matrix $A \in \mathbb{Z}^{m \times n}$ one has*

$$\textit{lindisc}(A) \le c.$$

5 A bound on the discrepancy of monotone matrices

Finally, we want to provide a non-trivial upper bound on the linear discrepancy of k-monotone matrices. The result of Spencer, Srinivasan and Tetali [22, 23] together with Theorem 4.1 yields a bound of $D_k^{\text{mon}}(n) = O(k^{3/2} \log n)$. This bound can be reduced by a direct proof that shares some similarities with that of Bohus [5]. Note that $D_k^{\text{mon}}(n) \ge k/2$, as the k-monotone 1×1 matrix $A = (k)$ together with target vector $y = (1/2)$ witnesses.

Theorem 5.1. *Consider any k-monotone matrix $A \in \mathbb{Z}^{n \times m}$. Then*

$$\textit{lindisc}(A) \le 5k \cdot \log_2(2 \min\{n, m\}).$$

Proof. If $n = m = 1$, $\text{lindisc}(A) \le \frac{k}{2}$, hence the claim is true. Let $y \in [0,1]^m$ by any vector. We can remove all columns i with $y_i = 0$ or $y_i = 1$ and then apply induction (on the size of the matrix). Next, if $m > n$, i.e. the number of columns is bigger then the number of constraints, then y is not a basic solution of the system

$$\begin{aligned} Ay &= b \\ 0 \le y_i &\le 1 \quad \forall i = 1, \dots, m. \end{aligned}$$

We replace y by a basic solution y' and apply induction (since y' has some integer entries and $Ay = Ay'$).

Finally it remains to consider the case $m \le n$. Let $a_1, \dots, a_n$ be the rows of A and let $d(j) := \|a_{j+1} - a_j\|_1$ for $j = 1, \dots, n-1$, i.e. $d(j)$ gives the cumulated differences between the jth and the $(j+1)$th row. Since the columns are k-monotone, each column contributes at most k to the sum $\sum_{j=1}^{n-1} d(j)$. Thus

$$\sum_{j=1}^{n-1} d(j) \le mk \le nk.$$

By the pigeonhole principle at least $n/2$ many rows j have $d(j) \le 2k$. Take any second of these rows and we obtain a set $J \subseteq \{1, \dots, n-1\}$ of size $|J| \ge n/4$ such that for every $j \in J$ one has $d(j) \le 2k$ and $(j+1) \notin J$. Let $A'y = b'$ be the subsystem of $n' \le \frac{3}{4}n$ many equations, which we obtain by deleting the rows in J from $Ay = b$. We apply induction to this system and obtain an $x \in \{0,1\}^m$ with

$$\begin{aligned} \|A'x - A'y\|_\infty &\le 5k \cdot \log_2(2n') \\ &\le 5k \log_2\left(2 \cdot \frac{3}{4} n\right) \\ &\le 5k \log_2(2n) - 5k \log_2\left(\frac{4}{3}\right) \\ &\le 5k \log_2(2n) - 2k. \end{aligned}$$

Now consider any $j \in \{1, \dots, n\}$. If $j \notin J$, then row j still appeared in $A'y = b'$, hence $|a_j^T x - a_j^T y| \le 5k \log_2(2n) - 2k$. Now suppose $j \in J$. We remember that $j + 1 \notin J$, thus $|a_{j+1}^T (x - y)| \le 5k \log_2(2n) - 2k$. But then using the triangle inequality

$$\begin{aligned} |a_j^T x - a_j^T y| &\le \underbrace{|(a_{j+1} - a_j)^T (x - y)|}_{\le d(j) \le 2k} + \underbrace{|a_{j+1}^T (x - y)|}_{\le 5k \log_2(2n) - 2k} \\ &\le 5k \cdot \log(2n). \end{aligned}$$

References

[1] W. Banaszczyk. Balancing vectors and Gaussian measures of n-dimensional convex bodies. *Random Structures Algorithms*, 12(4):351–360, 1998.

[2] N. Bansal. Constructive algorithms for discrepancy minimization. *CoRR*, abs/1002.2259, 2010. informal publication.

[3] J. Beck and T. Fiala. "Integer-making" theorems. *Discrete Appl. Math.*, 3(1):1–8, 1981.

[4] J. Beck and V. Sós. Discrepancy theory. In *Handbook of combinatorics, Vol. 1, 2*, pages 1405–1446. Elsevier, Amsterdam, 1995.

[5] G. Bohus. On the discrepancy of 3 permutations. *Random Structures Algorithms*, 1(2):215–220, 1990.

[6] R. Diestel. *Graph Theory*, volume 173 of *Graduate Texts in Mathematics*. Springer-Verlag, 1997.

[7] G. Dósa. The tight bound of first fit decreasing bin-packing algorithm is FFD(I) <= 11/9OPT(I) + 6/9. In Bo Chen, Mike Paterson, and Guochuan Zhang, editors, *Combinatorics, Algorithms, Probabilistic and Experimental Methodologies, First International Symposium, ESCAPE 2007, Hangzhou, China, April 7-9, 2007, Revised Selected Papers*, volume 4614 of *Lecture Notes in Computer Science*, pages 1–11. Springer, 2007.

[8] K. Eisemann. The trim problem. *Management Science*, 3(3):279–284, 1957.

[9] W. Fernandez de la Vega and G. S. Lueker. Bin packing can be solved within $1 + \varepsilon$ in linear time. *Combinatorica*, 1(4):349–355, 1981.

[10] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, New York, New York, 1979.

[11] P. C. Gilmore and R. E. Gomory. A linear programming approach to the cutting-stock problem. *Operations Research*, 9:849–859, 1961.

[12] M. Grötschel, L. Lovász, and A. Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2):169–197, 1981.

[13] K. Jansen and R. Solis-Oba. An $opt + 1$ algorithm for the cutting stock problem with constant number of object lengths. In Friedrich Eisenbrand and F. Bruce Shepherd, editors, *Proceedings of the 14th International Conference on Integer Programming and Combinatorial Optimization, IPCO 2010*, pages 438–449, 2010.

[14] D. S. Johnson. *Near-optimal bin packing algorithms*. PhD thesis, MIT, Cambridge, MA, 1973.

[15] D. S. Johnson, A. Demers, J. D. Ullman, M. R. Garey, and R. L. Graham. Worst-case performance bounds for simple one-dimensional packing algorithms. *SIAM Journal on Computing*, 3(4):299–325, 1974.

[16] N. Karmarkar and R. M. Karp. An efficient approximation scheme for the one-dimensional bin-packing problem. In *23rd annual symposium on foundations of computer science (Chicago, Ill., 1982)*, pages 312–320. IEEE, New York, 1982.

[17] L. Lovász, J. Spencer, and K. Vesztergombi. Discrepancy of set-systems and matrices. *European J. Combin.*, 7(2):151–160, 1986.

[18] J. Matoušek. *Geometric discrepancy*, volume 18 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 1999. An illustrated guide.

[19] G. Scheithauer and J. Terno. Theoretical investigations on the modified integer round-up property for the one-dimensional cutting stock problem. *Operations Research Letters*, 20(2):93 – 100, 1997.

[20] A. Sebő and G. Shmonin. Proof of the modified integer round-up conjecture for bin packing in dimension 7. Personal communication, 2009.

[21] J. Spencer. Six standard deviations suffice. *Transactions of the American Mathematical Society*, 289(2):679–706, 1985.

[22] J. H. Spencer, A. Srinivasan, and P. Tetali. The discrepancy of permutation families. Unpublished manuscript.

[23] A. Srinivasan. Improving the discrepancy bound for sparse matrices: Better approximations for sparse lattice approximation problems. In *Proceedings of the 8th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA'97 (New Orleans, Louisiana, January 5-7, 1997)*, pages 692–701, Philadelphia, PA, 1997. ACM SIGACT, SIAM, Society for Industrial and Applied Mathematics.

[24] D. Williamson. Lecture notes on approximation algorithms, Fall 1998. IBM Research Report, 1998. `http://legacy.orie.cornell.edu/~dpw/cornell.ps` .