

TUSNÁDY'S PROBLEM, THE TRANSFERENCE PRINCIPLE, AND NON-UNIFORM QMC SAMPLING

CHRISTOPH AISTLEITNER, DMITRIY BILYK, AND ALEKSANDAR NIKOLOV

ABSTRACT. It is well-known that for every $N \geq 1$ and $d \geq 1$ there exist point sets $x_1, \dots, x_N \in [0, 1]^d$ whose discrepancy with respect to the Lebesgue measure is of order at most $(\log N)^{d-1} N^{-1}$. In a more general setting, the first author proved together with Josef Dick that for any normalized measure μ on $[0, 1]^d$ there exist points x_1, \dots, x_N whose discrepancy with respect to μ is of order at most $(\log N)^{(3d+1)/2} N^{-1}$. The proof used methods from combinatorial mathematics, and in particular a result of Banaszczyk on balancings of vectors. In the present note we use a version of the so-called transference principle together with recent results on the discrepancy of red-blue colorings to show that for any μ there even exist points having discrepancy of order at most $(\log N)^{d-\frac{1}{2}} N^{-1}$, which is almost as good as the discrepancy bound in the case of the Lebesgue measure.

1. INTRODUCTION AND STATEMENT OF RESULTS

Many problems from applied mathematics require the calculation or estimation of the expected value of a function depending on several random variables; such problems are for example the calculation of the fair price of a financial derivative and the calculation of the expected loss of an insurance risk. The expected value $\mathbb{E}(g(Y^{(1)}, \dots, Y^{(d)}))$ can be written as

$$(1) \quad \int_{\mathbb{R}^d} g(y^{(1)}, \dots, y^{(d)}) d\nu(y^{(1)}, \dots, y^{(d)}),$$

where ν is an appropriate probability measure describing the joint distribution of the random variables $Y^{(1)}, \dots, Y^{(d)}$. Since a precise calculation of the value of such an integral is usually not possible, one looks for a numerical approximation instead. Two numerical methods for such problems are the *Monte Carlo method* (MC, using random sampling points) and the *Quasi-Monte Carlo method* (QMC, using cleverly chosen deterministic sampling points), where the QMC method is often preferred due to a faster convergence rate and deterministic error bounds. However, in the QMC literature it is often assumed that the problem asking for the value of (1) has at the outset already been transformed into the problem asking for the value of

$$(2) \quad \int_{[0,1]^d} f(x^{(1)}, \dots, x^{(d)}) dx^{(1)} \dots x^{(d)}.$$

Thus it is assumed that the integration domain is shrunk from \mathbb{R}^d to $[0, 1]^d$, and that the integration measure is changed from ν to the Lebesgue measure.

The first author is supported by the Austrian Science Fund (FWF), project Y-901.

The third author is supported by an NSERC Discovery Grant.

In principle, such a transformation always exists – this is similar to the way how general multivariate random sampling is reduced to sampling from the multivariate uniform distribution, using for example the Rosenblatt transform [28] (which employs sequential conditional inverse functions). The shrinking procedure is less critical, albeit also non-trivial¹; however, the change from a general measure to the uniform measure is highly problematic, in particular because QMC error bounds depend strongly on the regularity of the integrand. Note that the change from ν to the Lebesgue measure induces a transformation of the original function g to a new function f , and this transformation may totally ruin all “nice” properties of the initial function g such as the existence of derivatives or the property of having bounded variation. This is not the place to discuss this topic in detail; we just note that the main problem is not that each of $Y^{(1)}, \dots, Y^{(d)}$ may have a non-uniform marginal distribution, but rather that there may be a strong dependence in the joint distribution of these random variables (which by Sklar’s theorem may be encoded in a so-called *copula* – see [26] for details). We refer to [12] for a discussion of these issues from a practitioner’s point of view.

There are two possible ways to deal with the problems mentioned in the previous section, which can be seen as two sides of the same coin. Either one tries to transform the points of a classical QMC point set in such a way that they can be used for integration with respect to a different, general measure, and such that one obtains a bound for the integration error in terms of the discrepancy (with respect to the uniform measure) of the original sequence – this is the approach of Hlawka and Mück [18, 19], which has been taken up by several authors. On the other hand, one may try to sample QMC points directly in such a way that they have small discrepancy with respect to a given measure μ , and use error bounds which apply in this situation. This is the approach discussed in [1, 2], where it is shown that both key ingredients for QMC integration are given also in the setting of a general measure μ : there exist (almost) low-discrepancy point sets, and there exists a Koksma–Hlawka inequality giving bounds for the integration error in terms of the discrepancy (with respect to μ) of the sampling points and the degree of regularity of the integrand function. More precisely, let μ be a normalized Borel measure on $[0, 1]^d$, and let

$$(3) \quad D_N^*(x_1, \dots, x_N; \mu) = \sup_{A \in \mathcal{A}} \left| \frac{1}{N} \sum_{n=1}^N \mathbb{1}_A(x_n) - \mu(A) \right|$$

be the star-discrepancy of $x_1, \dots, x_N \in [0, 1]^d$ with respect to μ ; in this definition $\mathbb{1}_A$ denotes the indicator function of A , and the supremum is extended over the class \mathcal{A} of all axis-parallel boxes contained in $[0, 1]^d$ which have one vertex at the origin. Then in [1] it is shown that for arbitrary μ

¹One may wish to avoid this shrinking, and carry out QMC integration directly on \mathbb{R}^d instead. One way of doing so is to “lift” a set of sampling points from $[0, 1]^d$ to \mathbb{R}^d rather than shrinking the domain of the function. See for example [15] and the references there.

there exist points $x_1, \dots, x_N \in [0, 1]^d$ such that

$$(4) \quad D_N^*(x_1, \dots, x_N; \mu) \leq 63\sqrt{d} \frac{(2 + \log_2 N)^{(3d+1)/2}}{N}.$$

This improved an earlier result of Beck [8], where the exponent of the logarithmic term was $2d$. Note the amazing fact that the right-hand side of (4) does not depend on μ (whereas the choice of the points x_1, \dots, x_N clearly does). The estimate in (4) should be compared with corresponding results for the case of the uniform measure, where it is known that there exist point sets whose discrepancy is of order at most $(\log N)^{d-1}N^{-1}$ (so-called *low-discrepancy point sets*).²

A Koksma–Hlawka inequality for general measures was first proved by Götz [17] (see also [2]). For any points x_1, \dots, x_N , any normalized measure μ on $[0, 1]^d$ and any d -variate function f whose variation $\text{Var } f$ on $[0, 1]^d$ (in the sense of Hardy–Krause) is bounded, we have

$$(5) \quad \left| \int_{[0,1]^d} f(x) d\mu(x) - \frac{1}{N} \sum_{n=1}^N f(x_n) \right| \leq D_N^*(x_1, \dots, x_N; \mu) \text{Var } f.$$

This is a perfect analogue of the original Koksma–Hlawka inequality for the uniform measure. Combining (4) and (5) we see that in principle QMC integration is possible for the numerical approximation of integrals of the form

$$\int_{[0,1]^d} f(x^{(1)}, \dots, x^{(d)}) d\mu(x^{(1)} \dots x^{(d)}),$$

and that the convergence rate is almost as good as in the classical setting for the uniform measure. However, it should be noted that for the case of the uniform measure many explicit constructions of low-discrepancy point sets are known (see [16]), whereas the proof of (4) is a pure existence result and it is totally unclear how (and if) such point sets can be constructed with reasonable effort.

The purpose of the present note is to show that for any μ there actually exist points x_1, \dots, x_N whose discrepancy with respect to μ is of order at most $(\log N)^{d-\frac{1}{2}}N^{-1}$; this is quite remarkable, since it exceeds the corresponding bound for the case of the uniform measure only by a factor $(\log N)^{\frac{1}{2}}$.

Theorem 1. *For every $d \geq 1$ there exists a constant c_d (depending only on d) such that the following holds. For every $N \geq 2$ and every normalized Borel measure on $[0, 1]^d$ there exists points $x_1, \dots, x_N \in [0, 1]^d$ such that*

$$D_N^*(x_1, \dots, x_N; \mu) \leq c_d \frac{(\log N)^{d-\frac{1}{2}}}{N}.$$

The proof of this theorem uses a version of the so-called transference principle, which connects the combinatorial and geometric discrepancy, see Theorem 3. The novelty and the main observation of this paper lies in the

²While preparing the final version of this manuscript we learned that already in 1989 József Beck [7, Theorem 1.2] proved a version of (4) with the stronger upper bound $\mathcal{O}(N^{-1}(\log N)^{d+2})$. (However, the implied constant was not specified in his result.)

fact that this principle is still valid for general measures μ . This observation was previously made by Matoušek in [23], without providing any details, but otherwise it has been largely overlooked. In addition we shall use new upper bounds for Tusnády's problem due to the third author (for the discussion of Tusnády's problem, see §2.2). If one wants to refrain from the application of unpublished results, one can use Larsen's [20] upper bounds for Tusnády's problem instead, which yield Theorem 1 with the exponent $d + 1/2$ instead of $d - 1/2$ of the logarithmic term (see also [25]). An exposition of the connection between geometric and combinatorial discrepancies, together with the proof of Theorem 1, is given in the following section.

Before turning to combinatorial discrepancy, we want to make several remarks concerning Theorem 1. Firstly, whereas the conclusion of Theorem 1 is stated for finite point sets, one can use a well-known method to find an infinite sequence $(x_n)_{n \geq 1}$ whose discrepancy is of order at most $(\log N)^{d+1/2} N^{-1}$ for all $N \geq 1$. A proof can be modeled after the proof of Theorem 2 in [1].

Secondly, while the upper bound in Theorem 1 is already very close to the corresponding upper bound in the case of the classical discrepancy for the uniform measure, there is still a gap. One wonders whether this gap is a consequence of a deficiency of our method of proof, or whether the discrepancy bound in the case of general measures really has to be larger than that in the case of the uniform measure. In other words, we have the following open problem, which we state in a slightly sloppy formulation.

Open problem: Is the largest upper bound for the smallest possible discrepancy the one for the case of Lebesgue measure? In other words, is there any measure which is more difficult to approximate by finite atomic measures with equal weights than Lebesgue measure?

We think that this is a problem of significant theoretical interest. Note, however, that even in the classical case of the Lebesgue measure the problem asking for the minimal order of the discrepancy of point sets is famously still open³; while in the upper bounds for the best known constructions the logarithmic term has exponent $d - 1$, in the best known lower bound the exponent is $(d - 1)/2 + \varepsilon_d$ for small positive ε_d (see [11] for the latter result, and [10] for a survey). It is clear that some measures are much easier to approximate than Lebesgue measure – think of the measure having full mass at a single point, which can be trivially approximated with discrepancy zero. On the other hand, if the measure has a non-vanishing continuous component then one can carry over the orthogonal functions method of Roth [29] – this is done in [14]. However, the lower bounds for the discrepancy which one can obtain in this way are the same as those for the uniform case (and not larger ones). Intuitively, it is tempting to assume that the Lebesgue measure is essentially extremal – simply because intuition suggests that a measure

³This is known as the *Great Open Problem* of discrepancy theory, see, e.g., [9], page 8, and [24], page 178.

which is difficult to approximate should be “spread out everywhere”, and should be “equally large” throughout the cube.

2. COMBINATORIAL DISCREPANCY AND THE TRANSFERENCE PRINCIPLE

2.1. Combinatorial discrepancy. Let $V = \{1, \dots, n\}$ be a ground set and $\mathcal{C} = \{C_1, \dots, C_m\}$ be a system of subsets of V . Then the (combinatorial) discrepancy of \mathcal{C} is defined as

$$\text{disc } \mathcal{C} = \min_{y \in \{-1, 1\}^V} \text{disc}(\mathcal{C}, y),$$

where the vector y is called a (*red-blue*) *coloring* of V and

$$\text{disc}(\mathcal{C}, y) = \max_{i \in \{1, \dots, m\}} \left| \sum_{j \in C_i} y_j \right|$$

is the discrepancy of the coloring y . We may visualize the entries of the vector y as representing two different colors (usually red and blue). Then $\text{disc}(\mathcal{C}, y)$ is the maximal difference between the number of red and blue elements of V contained in a test set C_i , and $\text{disc } \mathcal{C}$ is the minimal value of $\text{disc}(\mathcal{C}, y)$ over all possible colorings y . For more information on this combinatorial notion of discrepancy, see [13, 24].

We will only be concerned with a geometric variation of this notion, i.e. the case when V is a point set in $[0, 1]^d$, and when the set system \mathcal{C} is the collection of all sets of the form $G \cap V$, where $G \in \mathcal{G}$ and \mathcal{G} is some geometric collection of subsets of $[0, 1]^d$: standard choices include, e.g., balls, convex sets, boxes (axis-parallel or arbitrarily rotated) etc. Let $\text{disc}(N, \mathcal{G}, d)$ denote the maximal possible discrepancy in this setting; that is, set

$$\text{disc}(N, \mathcal{G}, d) = \max_P (\text{disc } \mathcal{C}),$$

where the maximum is taken over all sets P of N points in $[0, 1]^d$ and where $\mathcal{C} = \mathcal{G}|_P = \{G \cap P : G \in \mathcal{G}\}$.

2.2. Tusnády’s problem. We denote by \mathcal{A} (as in the previous section) the class of all axis-parallel boxes having one vertex at the origin. The problem of finding sharp upper and lower bounds for $\text{disc}(N, \mathcal{A}, d)$ as a function of N and d is known as *Tusnády’s problem*. For the history and background of the problem, see [6, 23, 25].

We will use the following result, which has been recently announced by the third author [27]. A slightly weaker result [5] (with exponent d) by Bansal and Garg has been presented at MCQMC 2016 (and is also still unpublished). A yet weaker, but already published result (with exponent $d + 1/2$ instead of $d - 1/2$) is contained in [20].

Proposition 2. *For every $d \geq 1$ there exists a constant c_d (depending only on d) such that for all $N \geq 2$*

$$\text{disc}(N, \mathcal{A}, d) \leq c_d (\log N)^{d - \frac{1}{2}}.$$

Finally, we want to note that Tusnády’s problem is still unsolved; the known (and conjecturally optimal) lower bounds are of the order $(\log N)^{d-1}$, see [25]. As will become clear from the next subsection, further improvements of the upper bounds for Tusnády’s problem would directly imply improved upper bounds in Theorem 1. Note that actually Tusnády’s problem also falls within the framework of “discrepancy with respect to a general measure μ ”. Given N points, let μ be the discrete measure that assigns mass $1/N$ to each of these points. Then Tusnády’s problem asks for a set P of roughly $N/2$ points such that the discrepancy of P with respect to μ is small – under the additional requirement that the elements of P are chosen from the original set of N points. This additional requirement also explains why lower bounds for Tusnády’s problem do not imply lower bounds for the problem discussed in the present paper.

2.3. Transference principle. It is known that upper bounds for the combinatorial discrepancy of red-blue colorings can be turned into upper bound for the smallest possible discrepancy of a point set in the unit cube. This relation is called the *transference principle*.

For a system \mathcal{G} of measurable subsets of $[0, 1]^d$, its (unnormalized) geometric discrepancy is defined as

$$(6) \quad D_N(\mathcal{G}) = \inf_{P: \#P=N} \sup_{G \in \mathcal{G}} \left| \sum_{p \in P} \mathbb{1}_G(p) - N\lambda(G) \right|,$$

where the infimum is taken over all N -point sets $P \subset [0, 1]^d$ and $\lambda(G)$ is the Lebesgue measure of G . The *transference principle*, roughly speaking, says that (under some mild assumptions on the collection \mathcal{G}) **the geometric discrepancy is bounded above by the combinatorial discrepancy**, i.e. $D_N(\mathcal{G}) \lesssim \text{disc}(N, \mathcal{G}, d)$ with the symbol “ \lesssim ” interpreted loosely. Thus upper bounds on combinatorial discrepancy yield upper bounds for its geometric counterpart. This relation, in general, cannot be reversed: in the case when \mathcal{G} is the collection of all convex subsets of the unit cube, we have that $D_N(\mathcal{G})$ is of the order $N^{1-\frac{2}{d+1}}$, while $\text{disc}(N, \mathcal{G}, d)$ is of the order N as $N \rightarrow \infty$ (see, e.g., [24]).

In [24, 25] it is mentioned that M. Simmonovits attributes the idea of this principle to Vera T. Sós. It was used in the context of Tusnády’s problem by Beck [6] in 1981, and is stated in a rather general form in [22]. It can be found also in Matoušek’s book [24, p. 20] and in [25]. In all these instances it is formulated in a version which bounds the geometric discrepancy of point sets *with respect to the Lebesgue measure* (as defined in (6)) in terms of the combinatorial discrepancy of red-blue colorings.

However, upon examination of the proof, it turns out that the argument carries over to the case of the discrepancy with respect to an arbitrary measure (the only significant requirement is that the measure allows an ε -approximation for arbitrary ε ; see below). Similar to (3) and (6), we define

the geometric discrepancy of \mathcal{G} with respect to a Borel measure μ as

$$(7) \quad D_N(\mathcal{G}, \mu) = \inf_{P: \#P=N} D(P, \mathcal{G}, \mu),$$

where

$$(8) \quad D(P, \mathcal{G}, \mu) = \sup_{G \in \mathcal{G}} \left| \sum_{p \in P} \mathbb{1}_G(p) - N\mu(G) \right|.$$

Note that with this definition $D_N^*(x_1, \dots, x_N; \mu)$, as defined in (3), satisfies $D_N^*(x_1, \dots, x_N; \mu) = \frac{1}{N} D(\{x_1, \dots, x_N\}, \mathcal{A}, \mu)$. Below we state and prove the transference principle in a rather general form for arbitrary measures. The statement is similar to that in [24, 25] in the case of Lebesgue measure, and the proof follows along the lines of [24].

Theorem 3 (Transference principle for general measures). *Let μ be a Borel probability measure on $[0, 1]^d$ and let \mathcal{G} be a class of Borel subsets of $[0, 1]^d$ such that $[0, 1]^d \in \mathcal{G}$. Suppose that*

$$(9) \quad D_N(\mathcal{G}, \mu) = o(N) \quad \text{as } N \rightarrow \infty.$$

Assume furthermore that the combinatorial discrepancy of \mathcal{G} satisfies

$$(10) \quad \text{disc}(N, \mathcal{G}, d) \leq h(N)$$

for some function h with the property that $h(2N) \leq (2 - \delta)h(N)$ for some fixed $\delta > 0$. Then for every N

$$(11) \quad D_N(\mathcal{G}, \mu) = \mathcal{O}(h(N)),$$

i.e. there exist points $x_1, \dots, x_N \in [0, 1]^d$ so that $D(\{x_1, \dots, x_N\}, \mathcal{G}, \mu) \leq Ch(N)$.

Proof. Set $\varepsilon = h(N)/N$, and using (9) choose a positive integer k so large that there exists a set P_0 of $2^k N$ points in $[0, 1]^d$ so that $\frac{D(P_0, \mathcal{G}, \mu)}{2^k N} \leq \varepsilon$. By (11) we can find a red-blue coloring of the set P_0 with discrepancy at most $h(2^k N)$. The difference between the total number of red and blue points is also at most $h(2^k N)$, since the full unit cube is an element of our class of test sets. Without loss of generality we may assume that there are no more red than blue points (otherwise switch the roles of the red and the blue points). We keep all the red points and only so many of the blue points as to make sure that in total we have half the number of the original points, while we dispose of all the other blue points. Write P_1 for the new set. The cardinality of P_1 is $2^{k-1}N$. Furthermore,

$$\frac{D(P_1, \mathcal{G}, \mu)}{2^{k-1}N} \leq \varepsilon + \frac{h(2^k N)}{2^{k-1}N}.$$

To see why this is the case, note that an arbitrary set $G \in \mathcal{G}$ contains between $2^k N \mu(G) - \varepsilon 2^k N$ and $2^k N \mu(G) + \varepsilon 2^k N$ elements of P_0 . Thus it contains between

$$\frac{1}{2}(2^k N \mu(G) - \varepsilon 2^k N - h(2^k N)) \quad \text{and} \quad \frac{1}{2}(2^k N \mu(G) + \varepsilon 2^k N + h(2^k N))$$

red elements of P_0 , and consequently between

$$\frac{1}{2}(2^k N \mu(G) - \varepsilon 2^k N - h(2^k N)) \quad \text{and} \quad \frac{1}{2}(2^k N \mu(G) + \varepsilon 2^k N + 2h(2^k N))$$

elements of P_1 , where the upper bound is increased by $h(2^k N)/2$ since we had to add at most so many blue points in order to make sure that P_1 has the desired cardinality. Repeating this procedure, we obtain a point set P_2 of cardinality $2^{k-2}N$ whose discrepancy with respect to μ is at most

$$\frac{D(P_2, \mathcal{G}, \mu)}{2^{k-2}N} \leq \varepsilon + \frac{h(2^k N)}{2^{k-1}N} + \frac{h(2^{k-1}N)}{2^{k-2}N}.$$

We repeat this procedure over and over again, until we arrive at a point set P_k which has cardinality N , and whose discrepancy with respect to μ satisfies

$$D(P_k, \mathcal{G}, \mu) \leq \underbrace{\varepsilon N}_{=h(N)} + \sum_{j=0}^k \frac{h(2^{k-j}N)}{2^{k-j-1}} \leq Ch(N),$$

where we have used the condition that $h(2N) \leq (2 - \delta)h(N)$. Note that the value of C does not depend on the measure μ . This finishes the proof. \square

2.4. Proof of Theorem 1. Theorem 1 now easily follows from the combinatorial discrepancy estimate in Tusnády's problem (Proposition 2) and the transference principle for general measures (Theorem 3) applied in the case $\mathcal{G} = \mathcal{A}$. The only point that needs checking is whether μ satisfies the approximability condition (9) with respect to the collection \mathcal{A} of axis-parallel boxes with one vertex at the origin. But (9) follows trivially from the prior result (4).

A slightly more direct way to prove (9) would be to observe that the collection \mathcal{A} is a VC class (its VC-dimension is d), see e.g. [13] for definitions. This implies that [30] it is a uniform Glivenko–Cantelli class, i.e.

$$\sup_{\mu} \mathbf{E} \sup_{A \in \mathcal{A}} \left| \frac{1}{N} \sum_{n=1}^N \mathbb{1}_A(x_n) - \mu(A) \right| \rightarrow 0 \text{ as } N \rightarrow \infty,$$

where the expectation is taken over independent random points x_1, \dots, x_N with distribution μ , and the supremum is over all probability measures μ on $[0, 1]^d$. This immediately yields (9).

In conclusion we want to make some remarks on possible algorithmic ways of finding a point set satisfying the conclusion of Theorem 1. Following the proof of the transference principle, two steps are necessary. First one has to find the ε -approximation of μ . The existence of such a point set is guaranteed in the proof as a consequence of (4). However, this is not of much practical use since the point set for (4) cannot be constructed explicitly. However, in Corollary 1 of [2] it is proved that a set of $2^{26}d\varepsilon^{-2}$ random points which are sampled randomly and independently from the distribution μ will have, with positive probability, a discrepancy with respect to μ which is less than a given ε . This result is deduced from large deviation inequalities, and thus for a larger value of the constant (say 2^{35} instead of 2^{26}) the probability that the random point set has discrepancy at most ε with respect to μ will be extremely close to 1. We can think of ε as roughly $1/N$; for the cardinality of the random point set, this would give roughly $2^{35}dN^2$. Note, however, that in each iteration of the coloring procedure the

number of points is halved; accordingly, starting with $2^{35}dN^2$ points leads to a point set of cardinality N after a rather small number of steps. Admittedly, by using random points for the ε -approximation the whole problem is only shifted rather than solved; it is typically rather difficult to draw independent random samples from a general multivariate distribution μ .

The second part of the proof (the coloring procedure) is less of a problem from an algorithmic point of view, since in recent years much work has been done on algorithms for actually finding balanced colorings in combinatorial discrepancy theory. In particular, the recent bound [5] for Tusnády's problem due to Bansal and Garg mentioned in Section 2 is algorithmic in the following sense: they describe an efficient randomized algorithm that, given a set V of N points in $[0, 1]^d$, finds a red-blue coloring of V which has discrepancy at most $c_d(\log N)^d$ with probability arbitrarily close to 1. "Efficient" here means that the running time of the algorithm is bounded by a polynomial in N . In another recent preprint, Levy, Ramadas and Rothvoss [21] describe an efficient deterministic algorithm that achieves the same guarantees as the randomized algorithm from [4] for the Komlós problem. Since the techniques used in [5] are very closely related to those of [4], it seems likely that an efficient deterministic algorithm to find colorings for Tusnády's problem with discrepancy bounded by $c_d(\log N)^d$ can be constructed via the methods of [21]. Unfortunately, the stronger bound of $c_d(\log N)^{d-\frac{1}{2}}$ proved in [27] relies on an existential result of Banaszczyk [3], and no efficient algorithm is currently known that achieves this bound.

ACKNOWLEDGEMENTS

This paper was conceived while the authors were taking a walk in the vicinity of Rodin's *Gates of Hell* sculpture on Stanford University campus during the MCQMC 2016 conference. We want to thank the MCQMC organizers for bringing us together. Based on this episode, we like to call the open problem stated in the first section of this paper the *Gates of Hell Problem*.

REFERENCES

- [1] C. Aistleitner and J. Dick. Low-discrepancy point sets for non-uniform measures. *Acta Arith.*, 163(4):345–369, 2014.
- [2] C. Aistleitner and J. Dick. Functions of bounded variation, signed measures, and a general Koksma-Hlawka inequality. *Acta Arith.*, 167(2):143–171, 2015.
- [3] W. Banaszczyk. On series of signed vectors and their rearrangements. *Random Structures Algorithms*, 40(3):301–316, 2012.
- [4] N. Bansal, D. Dadush, and S. Garg. An algorithm for Komlós conjecture matching Banaszczyk's bound. In I. Dinur, editor, *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*, pages 788–799. IEEE Computer Society, 2016.
- [5] N. Bansal and S. Garg. Algorithmic discrepancy beyond partial coloring. Preprint. Available at <http://arxiv.org/abs/1611.01805>.
- [6] J. Beck. Balanced two-colorings of finite sets in the square. I. *Combinatorica*, 1(4):327–335, 1981.
- [7] Beck, J.: Balanced two-colorings of finite sets in the cube. In: Proceedings of the Oberwolfach Meeting "Kombinatorik" (1986), vol. 73, pages 13–25, 1989.

- [8] J. Beck. Some upper bounds in the theory of irregularities of distribution. *Acta Arith.*, 43(2):115–130, 1984.
- [9] J. Beck and W. W. L. Chen. *Irregularities of distribution*, volume 89 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 1987.
- [10] D. Bilyk. Roth’s orthogonal function method in discrepancy theory and some new connections. In *A panorama of discrepancy theory*, volume 2107 of *Lecture Notes in Math.*, pages 71–158. Springer, Cham, 2014.
- [11] D. Bilyk, M. T. Lacey, and A. Vagharshakyan. On the small ball inequality in all dimensions. *J. Funct. Anal.*, 254(9):2470–2502, 2008.
- [12] M. Cambou, M. Hofert, and C. Lemieux. Quasi-random numbers for copula models. Preprint. Available at <http://arxiv.org/abs/1508.03483>.
- [13] B. Chazelle. *The discrepancy method*. Cambridge University Press, Cambridge, 2000.
- [14] W. W. L. Chen. On irregularities of distribution and approximate evaluation of certain functions. *Quart. J. Math. Oxford Ser. (2)*, 36(142):173–182, 1985.
- [15] J. Dick, C. Irrgeher, G. Leobacher, and F. Pillichshammer. On the optimal order of integration in Hermite spaces with finite smoothness. Preprint. Available at <http://arxiv.org/abs/1608.06061>.
- [16] J. Dick and F. Pillichshammer. *Digital nets and sequences*. Cambridge University Press, Cambridge, 2010.
- [17] M. Götz. Discrepancy and the error in integration. *Monatsh. Math.*, 136(2):99–121, 2002.
- [18] E. Hlawka and R. Mück. A transformation of equidistributed sequences. In *Applications of number theory to numerical analysis (Proc. Sympos., Univ. Montréal, Montréal, Que., 1971)*, pages 371–388. Academic Press, New York, 1972.
- [19] E. Hlawka and R. Mück. Über eine Transformation von gleichverteilten Folgen. II. *Computing (Arch. Elektron. Rechnen)*, 9:127–138, 1972.
- [20] K. G. Larsen. On range searching in the group model and combinatorial discrepancy. *SIAM J. Comput.*, 43(2):673–686, 2014.
- [21] A. Levy, H. Ramadas, and T. Rothvoss. Deterministic discrepancy minimization via the multiplicative weight update method. Preprint. Available at <http://arxiv.org/abs/1611.08752>.
- [22] L. Lovász, J. Spencer, and K. Vesztegombi. Discrepancy of set-systems and matrices. *European J. Combin.*, 7(2):151–160, 1986.
- [23] J. Matoušek. On the discrepancy for boxes and polytopes. *Monatsh. Math.*, 127(4):325–336, 1999.
- [24] J. Matoušek. *Geometric discrepancy*, volume 18 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 2010.
- [25] J. Matoušek, A. Nikolov, and K. Talwar. Factorization norms and hereditary discrepancy. Preprint. Available at <https://arxiv.org/abs/1408.1376>.
- [26] R. B. Nelsen. *An introduction to copulas*. Springer Series in Statistics. Springer, New York, second edition, 2006.
- [27] A. Nikolov. Tighter bounds for the discrepancy of boxes and polytopes. Preprint. Available at <http://arxiv.org/abs/1701.05532>.
- [28] M. Rosenblatt. Remarks on a multivariate transformation. *Ann. Math. Statistics*, 23:470–472, 1952.
- [29] K. F. Roth. On irregularities of distribution. *Mathematika*, 1:73–79, 1954.
- [30] V. N. Vapnik and A. Y. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. In *Measures of complexity*, pages 11–30. Springer, Cham, 2015. Reprint of *Theor. Probability Appl.* 16 (1971), 264–280.

CHRISTOPH AISTLEITNER, INSTITUTE OF ANALYSIS AND NUMBER THEORY, TU GRAZ,
AUSTRIA

E-mail address: `christoph.aistleitner@jku.at`

DMITRIY BILYK, SCHOOL OF MATHEMATICS, UNIVERSITY OF MINNESOTA, USA

E-mail address: `dbilyk@math.umn.edu`

ALEKSANDAR NIKOLOV, DEPARTMENT OF COMPUTER SCIENCE, UNIVERSITY OF TORONTO

E-mail address: `anikolov@cs.toronto.edu`