An Improved Discrepancy Approach to Declustering

Benjamin Doerr Nils Hebbinghaus [1] Sören Werth [1]

Mathematisches Seminar, Bereich II, Christian-Albrechts-Universität zu Kiel,
Christian-Albrechts-Platz 4, 24118 Kiel, Germany.
E-mail: {bed,nhe,swe}@numerik.uni-kiel.de

1 Introduction

The last decade saw dramatic improvements in computer processing speed and storage capacities. Nowadays, the bottleneck in data-intensive applications is disk I/O, the time needed to retrieve typically large amount of data from storage devices. One idea to overcome this obstacle is to spread the data on disks of multi-disk systems so that they can be retrieved in parallel. The data allocation is determined by declustering schemes. Their aim is to allocate the data in such a manner that typical requests find their data evenly distributed on the disks.

The declustering problem is to assign data blocks from a multi-dimensional grid system to one of M storage devices in a balanced manner. More precisely, our grid is $V = [n_1] \times \cdots \times [n_d]$ for some positive integers $n_1, \ldots, n_d$.[2] A query Q requests the data assigned to a sub-grid $[x_1..y_1] \times \cdots \times [x_d..y_d]$ for some integers $1 \leq x_i \leq y_i \leq n_i$. We assume that the time to process such a query is proportional to the maximum number of requested data blocks that are stored in a single device. If we represent the assignment of the data blocks to the devices through a mapping $\chi : V \to [M]$, then the query time of the query above is $\max_{i \in [M]} |\chi^{-1}(i) \cap Q|$, where we identify the query Q with its associated sub-grid. Clearly, no declustering scheme can do better than $|Q|/M$. Hence a natural performance measure is the additive deviation from this lower bound.

---

[1] The authors thanks the DFG-Graduiertenkolleg 357 "Effiziente Algorithmen und Mehrskalenmethoden" for supporting this research.

[2] We use the notations $[n] := \{1, 2, \ldots, n\}$ and $[n..m] := \{n, n+1, \ldots, m\}$ for $n, m \in \mathbb{N}$, $n \leq m$.

Preprint submitted to Elsevier Science *14 February 2004*

This makes the problem a combinatorial discrepancy problem in M colors. Denote by $\mathcal{E}$ the set of all sub-grids in V. Then $\mathcal{H} = (V, \mathcal{E})$ is a hypergraph. For a coloring $\chi : V \to [M]$, the positive discrepancy of $\mathcal{H}$ with respect to χ and the positive discrepancy of $\mathcal{H}$ in M colors are

$$\begin{aligned}
\operatorname{disc}^+(\mathcal{H},\chi) &:= \max_{i\in[M],E\in\mathcal{E}} \left(|\chi^{-1}(i)\cap E| - \tfrac{1}{M}|E|\right),\\
\operatorname{disc}^+(\mathcal{H},M) &:= \min_{\chi:V\to[M]} \operatorname{disc}^+(\mathcal{H},\chi)
\end{aligned}$$

Apart from the fact that we only regard positive deviations, these notions were introduced by Srivastav and the first author in [DS03]. Independently, Anstee, Demetrovics, Katona and Sali [ADKS00] and Sinha, Bhatia and Chen [SBC03] proved a lower bound of $\Omega(\log M)$ for the additive error of any declustering scheme in dimension two. Sinha et al. [SBC03] also give the bound $\Omega(\log^{\frac{d-1}{2}} M)$ for arbitrary $d \geq 3$, but their proof contains a crucial error.

The current best upper bounds in arbitrary dimension for the declustering schemes are proposed by C.-M. Chen and C. Cheng [CC02]. They present two schemes for d–dimensional problems with an additive error $O(\log^{d-1} M)$. The first one works if $M = p^k$ for some $k \in \mathbb{N}$ and p a prime such that $p \geq d$, whereas the second works for arbitrary M, but the error increases with N

Our Results: For the upper bounds, we present an improved scheme that yields an additive error of $O(\log^{d-1} M)$ for a broader range of values of M and independent of the data size. Our requirement on M is that if $M = q_1 \dots q_k$, $q_1 < \dots < q_k$ is the canonical factorization of M into prime powers, $d \leq q_1 + 1$ is needed. Thus, in particular, our schemes work for M being a power of two (such that $M \geq d - 1$), which is very useful from the view-point of application. We also show that the latin hypercube construction used by Chen et al. [CC02] is much better than claimed. Where they show that the latin hypercube coloring extended to the whole grid has an error of at most 2^d times the one of the latin hypercube, we show that both errors are the same.

For the lower bound, we present the first correct proof of the $\Omega(\log^{\frac{d-1}{2}} M)$ bound. Again, a more careful analysis shows that that the positive discrepancy is at least $\frac{1}{2d}$ times the normal discrepancy instead of 3^{-d} as used in [SBC03]. Note that in typical applications with M between 16 and 1024, these 2^d and 3^d factors are at least as important as finding the right exponent of the $\log M$ term.

Since a central result of this paper are discrepancy bounds that are independent of the size of the grid, we usually work with the hypergraph $\mathcal{H}_N^d = ([N]^d, \mathcal{E}_N^d)$, $\mathcal{E}_N^d = \{\prod_{i=1}^d [x_i..y_i] \mid 1 \leq x_i \leq y_i \leq N\}$ for some sufficiently large integer N. We prove the following result.

Theorem 1 *Let M, $d \geq 2$ be positive integers and q_1 the smallest prime power in the canonical factorization of M into prime powers. We have*

(i) $\operatorname{disc}^+(\mathcal{H}_N^d, M) = O(\log^{d-1} M)$ *for* $d \leq q_1 + 1$, *independently of* $N \in \mathbb{N}$.

(ii) $\operatorname{disc}^+(\mathcal{H}_N^d, M) = \Omega(\log^{\frac{d-1}{2}} M)$ *for* $N \geq M$.

The combinatorial discrepancy results are shown via strong results from geometric discrepancy theory. The problem of geometric discrepancy in the unit cube $[0,1[^d$ is to distribute $n \in \mathbb{N}$ points evenly with respect to axis-parallel boxes: In every box R should be approximately $n \operatorname{vol}(R)$ points, where $\operatorname{vol}(R)$ denotes the volume of R. Again, discrepancy quantifies the distance to a perfect distribution. The discrepancy of a point set $\mathcal{P}$ with respect to a box $R \subseteq [0,1[^d$ and the set of all axis-parallel boxes $\mathcal{R}_d$ are defined by

$$\begin{aligned} D(\mathcal{P}, R) &= ||\mathcal{P} \cap R| - n \operatorname{vol}(R)|, \\ D(\mathcal{P}, \mathcal{R}_d) &= \sup_{R \in \mathcal{R}_d} |D(\mathcal{P}, R)|. \end{aligned}$$

2 The lower bound

The general idea in the proofs of the lower bound in Sinha et al. [SBC03] and Anstee et al. [ADKS00] is the same, here described in two dimensions:

Starting with an arbitrary M–coloring of $[M]^2$, there is a monochromatic set $\hat{P}$ with M vertices. Based on this set, an M–point set $\mathcal{P}$ in $[0,1[^2$ is constructed. By discrepancy theory [Sch72], there is a rectangle R such that $D(\mathcal{P}, R) = \Omega(\log M)$. Rounding R to the $[M]^2$ grid, they construct a hyperedge $\hat{R}$ that has almost the volume as R. Additionally $\hat{R}$ contains as many vertices of $\hat{P}$ as R points of $\mathcal{P}$. With the help of $\hat{R}$ and a short calculation the lower bound of the additive error $\Omega(\log M)$ is shown.

The small, but crucial mistake in the proof of Sinha et al. [SBC03] is in the transfer from the geometric discrepancy setting back to the combinatorial one. Recall that the authors started with a color class of exactly M^{d-1} points. They down-scaled it by a factor of M to a set in the unit cube (that, note this fact, is a subset of $\{0, \frac{1}{M}, \frac{2}{M}, \ldots, \frac{M-1}{M}\}^d$). Then their geometric discrepancy argument yields a rectangle of polylogarithmic discrepancy. However, the rectangle $[0, \frac{M-1}{M}]^d$ has a much larger discrepancy: It contains all M^{d-1} points, but has a volume of $(\frac{M-1}{M})^d$ only.

This yields a discrepancy of $M^{d-1}(1 - (\frac{M-1}{M})^d) = \Omega(M^{d-2})$. For dimension $d \geq 3$ this is larger than the upper bound, what also indicates an error in the proof of Sinha et al. [SBC03]. The last argument also shows that rounding an arbitrary box to a box in the grid can cause a roundoff error, which is of magnitude larger than the discrepancy. For this reason, a direct generalization using the lower bound of Roth [Rot64] is not possible. A more careful analysis is needed. In particular, we have to ensure the existence of a *small* box having large discrepancy. Using ideas of Beck [BC87], we show

Theorem 2 *For any n–point set $\mathcal{P}$ in the unit cube $[0,1[^d$, there is an axis-parallel cube Q with side at most $n^{-\frac{(2d-3)d}{(d-1)^2(2d+1)}}$ fully contained in $[0,1[^d$ with*

$$D(\mathcal{P}, Q) = \Omega(\log^{\frac{d-1}{2}} n).$$

Now Theorem 1 (ii) follows from Theorem 2 using the roundoff reduction of Anstee et al. [ADKS00] and Sinha et al. [SBC03].

3 The upper bound

We use geometric discrepancies to construct a declustering scheme for the proof of our upper bound. The notation of Niederreiter [Nie87] is used in the following. For an integer $b \geq 2$, an elementary interval in base b is an interval of the form $E = \prod_{i=1}^{d} \left[a_i b^{-d_i}, (a_i+1)b^{-d_i}\right[$, with integers $d_i \geq 0$ and $0 \leq a_i < b^{d_i}$ for $1 \leq i \leq d$. For integers t, m such that $0 \leq t \leq m$, a (t,m,d)–net in base b is a point set of b^m points in $[0,1[^d$ such that all elementary intervals with volume b^{t-m} contain exactly b^t points. Note that any elementary interval with volume b^{t-m} has discrepancy zero in a (t,m,d)–net. Since any subset of an elementary interval of volume b^{t-m} has discrepancy at most b^t and any box can be packed with elementary intervals in a way that the uncovered part can be covered by $O(\log^{d-1} n)$ elementary intervals of volume b^{t-m}, the following is immediate:

Theorem 3 *A (t,m,d)–net $\mathcal{P}$ has discrepancy $D(\mathcal{P}, \mathcal{R}_d) = O(\log^{d-1} n)$.*

The central argument in our proof of the upper bound is the following result of Niederreiter [Nie87] on the existence of $(0,m,d)$–nets. From the view-point of application it is important that his proof is constructive.

Theorem 4 *Let $b \geq 2$ be an arbitrary base and $b = q_1 q_2 \ldots q_u$ be the canonical factorization of b into prime powers such that $q_1 < \cdots < q_u$. Then for any $m \geq 0$ and $d \leq q_1 + 1$ there exists a $(0,m,d)$–net in base b.*

We construct colorings of $\mathcal{H}_N^d$ from $(0,m,d)$–nets with small discrepancy. We start with colorings for $\mathcal{H}_M^d$ in Lemma 5.

Lemma 5 *Let $\mathcal{P}_{net}$ be a $(0, d-1, d)$–net in base M in $[0,1[^d$. Then there is a M–coloring χ_M of $\mathcal{H}_M^d = ([M]^d, \mathcal{E}_M^d)$ such that all rows of $[M]^d$ contain every color exactly once and* $\operatorname{disc}(\mathcal{H}_M^d, \chi_M) \leq 2D(\mathcal{P}_{net}, \mathcal{R}_d)$.

In Lemma 6, we show that it is sufficient to consider the discrepancy of $\mathcal{H}_M^d$ with respect to these colorings for determining the upper bound of the discrepancy of $\mathcal{H}_N^d$. The Lemma 6 is a remarkable improvement of Theorem 4.2 in [CC02] , where $\operatorname{disc}(\mathcal{H}_N^d, \chi) \leq 2^d \operatorname{disc}(\mathcal{H}_M^d, \chi_M)$ is shown. Note that this reduces the implicit constant in the upper bound by factor of 2^d.

Lemma 6 *Let χ_M be a M–coloring of $\mathcal{H}_M^d$ such that all rows of $[M]^d$ contain every color exactly once and χ a coloring of $\mathcal{H}_N^d$ defined by $\chi(x_1, \ldots, x_d) = \chi_M(y_1, \ldots, y_d)$ such that $x_i \equiv y_i \mod M$ for $i \in [d]$, $x_i \in [N]$, $y_i \in [M]$. Then*

$$\operatorname{disc}(\mathcal{H}_N^d, \chi) = \operatorname{disc}(\mathcal{H}_M^d, \chi_M).$$

The upper bound in Theorem 1 follows from the above.

References

[ADKS00] R. Anstee, J. Demetrovics, G. O. H. Katona, and A. Sali. Low discrepancy allocation of two-dimensional data. In *Foundations of Information and Knowledge Systems, First International Symposium*, volume 1762 of *Lecture Notes in Computer Science*, pages 1–12, 2000.

[BC87] J. Beck and W. L. Chen. *Irregularities of distribution*, volume 89 of *Cambridge Tracts in Mathematics.* Cambridge University Press, Cambridge, 1987.

[CC02] C.-M. Chen and C. Cheng. From discrepancy to declustering: near optimal multidimensional declustering strategies for range queries. In *ACM Symp. on Database Principles*, pages 29–38, Madison, WI, 2002.

[DS03] B. Doerr and A. Srivastav. Multicolour discrepancies. *Combinatorics, Probability and Computing*, 12:365–399, 2003.

[Nie87] H. Niederreiter. Point sets and sequences with small discrepancy. *Monatsh. Math.*, 104:273–337, 1987.

[Rot64] K. F. Roth. Remark concerning integer sequences. *Acta Arithmetica*, 9:257–260, 1964.

[SBC03] R. K. Sinha, R. Bhatia, and C.-M. Chen. Asymptotically optimal declustering schemes for 2-dim range queries. *Theoret. Comput. Sci.*, 296:511–534, 2003.

[Sch72] W. M. Schmidt. On irregularities of distribution VII. *Acta Arith.*, 21:45–50, 1972.