Error Resilient Space Partitioning

Orr Dunkelman
Computer Science Department,
University of Haifa, Israel

Zeev Geyzel
Mobileye, an Intel company,
Jerusalem, Israel

Chaya Keller
Department of Computer Science,
Ariel University, Israel

Nathan Keller
Department of Mathematics,
Bar Ilan University, Ramat Gan, Israel

Eyal Ronen
School of Computer Science,
Tel Aviv University, Israel

Adi Shamir
Department of Computer Science,
Weizmann Institute of Science, Rehovot, Israel

Ran J. Tessler
Department of Mathematics,
Weizmann Institute of Science, Rehovot, Israel

Abstract

In this paper we consider a new type of space partitioning which bridges the gap between continuous and discrete spaces in an *error resilient* way. It is motivated by the problem of rounding noisy measurements from some continuous space such as $\mathbb{R}^d$ to a discrete subset of representative values, in which each tile in the partition is defined as the preimage of one of the output points. Standard rounding schemes seem to be inherently discontinuous across tile boundaries, but in this paper we show how to make it perfectly consistent (with *error resilience* ϵ) by guaranteeing that any pair of consecutive measurements X_1 and X_2 whose L_2 distance is bounded by ϵ will be rounded to the same nearby representative point in the discrete output space. We achieve this resilience by allowing a few bits of information about the first measurement X_1 to be unidirectionally communicated to and used by the rounding process of the second measurement X_2. Minimizing this revealed information can be particularly important in privacy-sensitive applications such as COVID-19 contact tracing, in which we want to find out all the cases in which two persons were at roughly the same place at roughly the same time, by comparing cryptographically hashed versions of their itineraries in an error resilient way.

The main problem we study in this paper is characterizing the achievable tradeoffs between the amount of information provided and the error resilience for various dimensions. We analyze the problem by considering the possible colored tilings of the space with k available colors, and use the color of the tile in which X_1 resides as the side information. We obtain our upper and lower bounds with a variety of techniques including isoperimetric inequalities, the Brunn-Minkowski theorem, sphere packing bounds, Sperner's lemma, and Čech cohomology. In particular, we show that when $X_i \in \mathbb{R}^d$, communicating $\log_2(d+1)$ bits of information is both sufficient and necessary (in the worst case) to achieve positive resilience, and when d=3 we obtain a tight upper and lower asymptotic bound of $(0.561\ldots)k^{1/3}$ on the achievable error resilience when we provide $\log_2(k)$ bits of information about X_1's color.

2012 ACM Subject Classification Theory of computation → Randomness, geometry and discrete structures; Theory of computation → Computational geometry; Theory of computation → Error-correcting codes

Keywords and phrases space partition, high-dimensional rounding, error resilience, sphere packing, Sperner's lemma, Brunn-Minkowski theorem, Čech cohomology

Digital Object Identifier 10.4230/LIPIcs.ICALP.2021.4

Category Invited Talk

Related Version *Full Version*: `https://arxiv.org/abs/2008.03675`

Acknowledgements We thank Stephen D. Miller for inspiring discussions on sphere packing.

1 Introduction

Studying various types of space partitioning of a continuous space such as $\mathbb{R}^d$ is a central topic in computational geometry (see, e.g., [11, Chapters 6,12] and the references therein), and each type of partition has different properties and applications within computer science, electrical engineering, and applied mathematics. For example, in *error-correcting codes* (which are extensively used in data communication) we try to squeeze the largest possible number of equal sized disjoint balls into the input space, while in *vector quantization* [16] (which is used extensively in data compression) we try to completely cover the input space with a small number of tiles whose volumes are as similar as possible. In this paper we investigate a new variant which can be viewed as "continuous error correction over the reals". Our main motivation is the problem of rounding noisy analog measurements in $\mathbb{R}^d$, in order to digitally process or store them (see, e.g., [2, 10]). This rounding process seems to be inherently discontinuous across tile boundaries, and this problem is compounded by the fact that for large d almost all input vectors are near boundaries. One natural solution to this discontinuity problem is to try to minimize the fraction of pairs X_1, X_2 with $distance(X_1, X_2) < \epsilon$ which are rounded differently by considering *foam tilings* that minimize the total surface area of unit volume tiles (such a tiling is called "foam" since it emerges in physical collections of soap bubbles). In a beautiful FOCS paper [22] (which was highlighted at CACM [23]), Kindler et al. introduced a clever new construction of such tiles which they called *spherical cubes*. What makes these tiles special is that they have the $O(\sqrt{d})$ surface area of a ball and yet they can tile the whole $\mathbb{R}^d$ space without gaps, which solved an open problem posed by Lord Kelvin in 1887.

In this paper we consider the more ambitious goal of achieving *error resilience* which completely eliminates all the discontinuities in the rounding process rather than reducing their fraction. We call such a rounding scheme *consistent rounding*[1], and make it possible by thinking about X_1 and X_2 as two consecutive noisy measurements of the same X. When the first measurement X_1 is rounded, we allow it to produce a few bits of side information about how it was rounded, and to provide them as an auxiliary input to the process that decides how to round X_2. Note that both X_1 and X_2 are assumed to be real valued vectors which require an infinite number of bits to fully specify them.

To demonstrate the basic idea, consider the one dimensional case in which X_1 and X_2 are real values which have to be consistently rounded to the same nearby integer whenever they are close enough. X_1 is always rounded to the nearest integer, and it produces a single bit of side information which is whether it was rounded to an even or an odd integer P. When X_2 is measured, it is rounded to the nearest integer which has the same parity as P. To demonstrate this process, consider the problematic inputs $X_1 = 0.4999$ and $X_2 = 0.5001$: X_1 is rounded to 0, and X_2 is also rounded to 0 since it is the closest even integer. In fact, X_2 could be anywhere between -1 and 1 and still be consistently rounded to 0, and thus the rounding scheme is resilient to additive errors of up to 0.5. In fact (see Sec. 4.1), this is the highest possible error resilience of any one dimensional consistent rounding scheme; other natural schemes (such as providing one bit of side information about whether X_1 was rounded up or down) provide lower resilience.

The way we think about the problem is to consider a colored tiling of the real line with two colors: All the values in $[-0.5, 0.5)$, $[1.5, 2.5)$, etc. are colored by 1, and all the values in $[0.5, 1.5)$, $[2.5, 3.5)$, etc. are colored by 2. The side information provided about X_1 is the

[1] We note that in statistics, the term "consistent rounding" is used to denote a rounding that is consistent with some external constraints; see [26, p. 237].

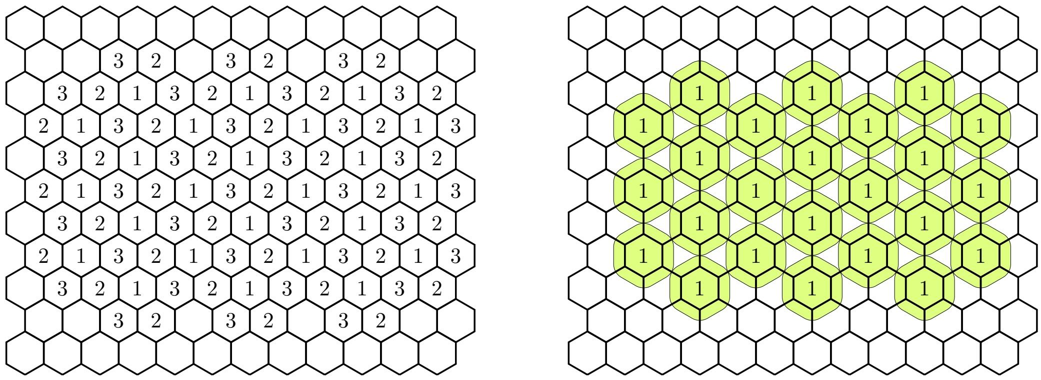

Figure 1 A 3-colored hexagonal tiling of the plane, and a maximal non-intersecting inflation of the tiles colored 1.

color of the tile in which it is located, and the way we process X_2 is to round it to the center of the closest tile which has the same color as that of X_1. The essential property of our partition is that the minimum distance between any two tiles with the same color is 1, and thus we can "inflate" all the tiles of any particular color to their outer parallel body in order to include any erroneously measured value X_2 up to a distance of 0.5 away from the original tile, and still get nonoverlapping tiles which make it possible to uniquely associate such points with original tiles.

To make this perspective clearer, consider the two dimensional plane. In the obvious checkerboard tiling by unit squares, we need at least 4 colors (and thus 2 bits of side information) to color the tiles if we do not allow equi-colored tiles to touch. We can reduce the number of colors to 3 (and thus, provide only $\log_2(3) = 1.58$ bits of side information) by considering the hexagonal partition of the plane depicted in the left part of Figure 1. Given a two dimensional point X_1, we always round it to the center of the hexagon in which it is located, and given X_2 we round it to the center of the nearest hexagon which has X_1's color. To determine the error resilience of this scheme, we inflate all the hexagonal tiles of a particular color by the same amount until they touch each other, as depicted in the right part of Figure 1. As it turns out, this natural scheme is not optimal since the inflated hexagons' corners touch prematurely, leaving large gaps between them. A 3-colored tiling with a higher error resilience will be described in Section 5.1.1, and an asymptotically optimal tiling for a large number of colors will be described in Section 5.2.1.

For inputs $X \in \mathbb{R}^d$, we can provide one bit of side information about each one of its d entries separately, but for large d this is very inefficient. In our colored tiling formulation, it suffices to reveal the color of X_1 in order to consistently round X_2, and thus if we can tile the space with k colors, we need only $\log_2(k)$ bits to specify this color. This naturally leads to the question of what is the minimum number of colors needed to tile $\mathbb{R}^d$ by bounded sized tiles so that no two tiles of the same color will touch (even at a corner). Surprisingly, we could not find any reference to this natural question. As we show in Section 3, there can be no such colored tiling with d colors, and as we show in Section 5.3, $d+1$ colors are sufficient. Consequently, $\log_2(d+1)$ bits of side information about X_1 are necessary and sufficient (in the worst case) to obtain a consistent rounding scheme with positive error resilience. To prove the negative result, we use techniques borrowed from algebraic topology (namely, either a generalization of Sperner's lemma or the Čech cohomology and other cohomology theories), and to prove the positive result we provide an explicit construction of such a colored tiling.

Table 1 Summary of our lower and upper bounds on the error resilience, for different d and k.

Scenario	Lower Bound (LB) on ER	Upper Bound (UB) on ER	Techniques	Source
3 colors in $\mathbb{R}^2$	0.354	0.413	Brunn-Minkowski ineq. (UB), Brick wall tiling (LB)	Sec. 4.1 (UB), Sec. 5.1.1 (LB)
4 colors in $\mathbb{R}^2$	0.5	0.564	Brunn-Minkowski ineq. (UB), Brick wall tiling (LB)	Sec. 4.1 (UB), Sec. 5.1.1 (LB)
k colors in $\mathbb{R}^2$	$0.537\sqrt{k} - O(1)$	$0.537\sqrt{k}$	Circle packing (UB), HCR tiling (LB)	Sec. 4.2 (UB), Sec. 5.2.1 (LB)
4 colors in $\mathbb{R}^3$	0.25	0.365	Brunn-Minkowski ineq. (UB), 3-dim Brick wall (LB)	Sec. 4.1 (UB), Sec. 5.1.2 (LB)
k colors in $\mathbb{R}^3$	$(0.561 - o(1))k^{1/3}$	$0.561k^{1/3}$	Sphere packing (UB), CPB tiling (LB)	Sec. 4.2 (UB), Sec. 5.2.2 (LB)
k colors in $\mathbb{R}^8$	$(0.707 - o(1))k^{1/8}$	$0.707k^{1/8}$	Sphere packing (UB), CPB tiling (LB)	Sec. 4.2 (UB), Sec. 5.2.2 (LB)
k colors in $\mathbb{R}^{24}$	$(1 - o(1))k^{1/24}$	$k^{1/24}$	Sphere packing (UB), CPB tiling (LB)	Sec. 4.2 (UB), Sec. 5.2.2 (LB)
$d+1$ colors in $\mathbb{R}^d$	$\Omega(1/d)$	$O(\log d/\sqrt{d})$	Brunn-Minkowski ineq. (UB), Dimension reducing tiling (LB)	Sec. 4.1 (UB), Sec. 5.3 (LB)

ER – error resilience, LB – lower bound, UB – upper bound,
HCR – honeycomb of rectangles, CPB – close packing of boxes

In addition to minimizing the amount of side information, we study the question of maximizing the error resilience for a given d and k. In the negative direction, in Section 4 we obtain several upper bounds on the achievable error resilience, using different techniques from geometry and analysis (including isoperimetry, the Brunn-Minkowski inequality and results on the sphere packing problem). In the positive direction, we construct in Section 5 a variety of tiling schemes. In particular, while for $d = 2$ and $k = 3$ the hexagonal tiling scheme described above is resilient to additive errors of up to 0.31, we present a tiling with resilience of 0.354, and show that no 3-color tiling can have resilience higher than 0.413. We also show that the maximal resilience achieved by a $(d+1)$-coloring of $\mathbb{R}^d$ is between $\Omega(1/d)$ and $O(\log d/\sqrt{d})$, and use the recent breakthrough results on sphere packing [6, 18, 25] to obtain tight asymptotic lower and upper bounds on the resilience in dimensions $2, 3, 8$, and 24. Our bounds are summarized in Table 1.

Applications. The problem of "continuous error correction over the reals" has numerous applications. For example, in biometric identification, multiple measurements of the same fingerprint are similar but not identical. It would be very helpful if all these slight variants could be represented by the same rounded point P, and we can achieve this by storing a small amount of side information in the biometric database during the initial registration of a new employee.

Another example is the problem of developing a contact tracing app for the COVID-19 pandemic, where we want to record all the cases in which two smart phones were at roughly the same place at roughly the same time. We can do this by measuring in each phone the GPS location, the local time, and perhaps other parameters such as the ambient noise level (in order to rule out the case of people living in different apartments which are separated by a common wall). When someone tests positive for COVID-19, the health authority wants to reveal a list of his measurements, but in order to keep the patient's privacy, it wants to cryptographically hash each measurement before publishing it. Since the measurements are likely to be slightly different for the infected and exposed persons, the health authority can publish the small amount of side information along with the consistently rounded and then hashed measurements.

Finally, the problem may also be relevant to the construction of quantum error correction codes, since the state of a quantum computer is a complex-valued vector in a Hilbert space with exponentially many dimensions which can be perturbed by external noise. Note that the logarithmic number of side information bits we need for error correction can be stored and processed classically. However, describing such potential applications is beyond the scope of this paper.

Related work. A line of study that seems related to our work is *fuzzy* constructions that were widely studied in the cryptographic literature, such as the *fuzzy commitment scheme* of Juels and Wattenberg [20]). Dodis et al. [12] introduced the notions of *fuzzy extractors* and *secure sketches*, which enable two parties to secretly reach a consensus value from multiple noisy measurements of some high entropy source (a recent survey of such techniques can be found in [14]). However, such schemes concentrate on the aspects of cryptographic security (which we do not consider), and produce sketches whose size depends on the number of possible inputs (which is meaningless for real valued inputs). In this sense our consistent rounding scheme can be viewed as an exceptionally efficient reconciliation process, since it can produce for each million entry vector of arbitrarily large real numbers a 20 bit "sketch" (in the form of its color side information), and process this information with trivial point location algorithms.

Open problems. While we fully solved the question of minimizing the amount of side information required for error resilience, several questions remain open regarding the maximal resilience rate that can be achieved for a given amount of side information. In particular, for dimensions $2, 3, 8, 24$ we determined the exact asymptotic resilience when $\log_2 k$ bits of information are allowed, using a connection to the *densest sphere packing* problem. When only very few bits of information are allowed, the situation is much less clear. For example, we do not even know whether the brick wall constructions we present in Section 5.1.1 have the highest error resilience among rounding schemes in $\mathbb{R}^2$ with $\log_2 3$ and $\log_2 4$ bits of side information. It will be interesting to obtain new upper bounds via different techniques or new lower bound constructions.

2 Our Setting

In this section we present the basic setting that will be assumed throughout the paper.

Colored tiling. We study tilings of $\mathbb{R}^d$, where each tile is *connected*, *closed* and *bounded*, and the tiles intersect only in their boundaries. In some of the results we make additional assumptions on the tiles or drop some of the assumptions; such changes are stated explicitly. Each tile is colored in one of k colors.

Error resilience and inflation. In order to compute the error resilience of a given tiling (with respect to the L_2 distance), we consider all tiles of the same color and inflate them (i.e., replace the tile T by the set $T' = \{y : \exists x \in T, |x - y| < r\}$ for some $r > 0$) until they touch each other. Clearly, the error resilience is the maximal r for which such a non-intersecting inflation is possible. We note that in convex geometry, such an inflation T' is called *the outer parallel body of radius* r of T (see [15, p. 943]). The minimal distance between two same-colored points in different tiles is denoted by t, and so, the error resilience is $t/2$.

Breaking ties. A fine point about consistent rounding schemes is how to break ties, and here we deal differently with X_1 and X_2. We want to be able to deal with any X_1, and thus we think about the tiles as being closed sets which include their boundaries. Therefore, points X_1 which are on the boundary between tiles can have more than one possible color. We allow such ties to be broken arbitrarily in the sense that X_1 can be rounded to the center of any one of the tiles that it belong to. However, when we think about X_2 we allow it to be at a distance of strictly less than some bound, and thus the inflated tiles (that contain all the possible X_2's we are interested in) are open sets which have no intersections. Consequently, each X_2 can belong to at most one inflated tile, and is rounded to the center of that tile with no possible ties.

Normalization. The d-dimensional volume of a figure $T \subset \mathbb{R}^d$ is denoted by $\lambda(T)$. We normalize the tiling by assuming that *the volume of each tile is bounded by* 1 (like in the 1-dimensional case presented in the introduction, where all tiles are segments of length 1). We make the natural assumption that any inflated tile T' satisfies $\lambda(\bar{T}') = \lambda(T')$, where $\bar{T}'$ is the topological closure of T'. Normalization with respect to other natural metrics, as well as alternative distance metrics, are discussed in the full version of the paper.

3 The Minimal Number of Colors Required for Error Resilience

In this section we prove that the minimal number of colors required for achieving any positive error resilience in a tiling of $\mathbb{R}^d$ is $d+1$. We provide two proofs, under different additional natural assumptions on the tiles. The first assumes that all tiles are uniformly bounded and relies on a generalization of Sperner's lemma. The second proof assumes that the tiles and their non-empty intersections are contractible (while not having to be uniformly bounded) and uses a more advanced algebraic-topologic argument. The lower bound $d+1$ on the number of required colors is tight; a matching construction for any d is presented in Section 5.3.

3.1 Lower bound for bounded tiles, using Sperner's Lemma

The main result of this subsection is the following.

▶ **Proposition 1.** *For any $m > 0$, the following holds. Let $T_1, T_2, \ldots$ be a colored tiling of $\mathbb{R}^d$ in $d+1$ colors, in which each tile is contained in a box with side length m. If the error resilience of the tiling is $\delta > 0$, then there exist tiles in all $d+1$ colors that intersect at a point.*
Consequently, any tiling with positive error resilience uses at least $d+1$ colors.

We use a generalization of the classical Sperner's lemma [24], called *Bapat's connector-free lemma*. As Bapat's lemma was originally proved only in $\mathbb{R}^2$ and in a discrete setting, we first provide a proof of a continuous version in $\mathbb{R}^d$, and then derive Proposition 1 from it.

3.1.1 Bapat's connector-free lemma – continuous version

Let Δ be a d-simplex in a Euclidean space, i.e., the convex hull of $d+1$ points $x_0, \ldots, x_d$ which do not lie in a d-space. The i'th face of Δ is the span of $\{x_j\}_{j \neq i}$. A *connector* in a d-simplex is a connected set which intersects all its $(d-1)$-dimensional faces.

Bapat's connector-free lemma asserts the following:

▶ **Theorem 2.** *Let $C_0, \ldots, C_d$ be a cover of a d-simplex Δ by closed sets such that the minimal distance between two connected components of the same set is $\delta > 0$. Suppose that the interiors of the sets are disjoint and that no C_i contains a connector. Then $\bigcap_{i=0}^{d} C_i \neq \emptyset$.*

In order to prove the theorem we will reduce it to an analogous discrete claim. The reduction is simple, but requires some more terminology.

A *triangulation* T of a simplex $\Delta \subset \mathbb{R}^d$ is a cover of it by simplices whose interiors are disjoint, such that the intersection of any set of simplices is either empty or the convex hull of some vertices. Note that the vertices of Δ are, in particular, vertices of the triangulation, and that the faces of Δ are endowed by an induced triangulation. The *diameter* of T is the supremum of distances between vertices which share an edge. The 1-*skeleton* of T is the graph formed by the vertices and the edges. A *discrete connector* is a connected subset of the 1-skeleton of T which contains vertices from each facet of Δ. We can now state the discrete version of Theorem 2 (which is the actual statement proved by Bapat, for $d = 2$).

▶ **Theorem 3** (Bapat). *Suppose that the vertices of T are partitioned into disjoint sets $A_0, A_1, \ldots, A_d$ such that no A_i contains a discrete connector. Then there exists a simplex in T whose $d+1$ vertices belong to different sets A_i.*

Proof of Theorem 2, assuming Theorem 3. Assume towards contradiction that there exist sets $C_0, \ldots, C_d$ which cover Δ, such that no C_i contains a connector, but $\bigcap_{i=0}^{d} C_i = \emptyset$. Define the function $f : C_0 \times C_1 \cdots \times C_d \to \mathbb{R}_+$ by setting $f(p_0, \ldots, p_d)$ to be the diameter of the convex hull of $(p_0, \ldots, p_d)$, which is the maximal distance between two p_i's. The domain of the function f compact (here we use the assumption that the minimal distance between two connected components of the same C_i is at least δ) and its range is $\mathbb{R}_+$ (since we assumed $\bigcap_{i=0}^{d} C_i = \emptyset$). Hence, f attains a minimum $\epsilon > 0$.

Let $\eta = \min(\delta, \epsilon)/3$. Let $\tilde{C}_i$ be the η-thickening of C_i in Δ, i.e., $\tilde{C}_i = \{p \in \Delta : \operatorname{dist}(p, C_i) < \eta\}$. Then $\tilde{C}_i$ is an open cover of Δ. By the choice of η, neither $\tilde{C}_i$ contains a connector,[2] and $\bigcap_{i \in \{0,\ldots,d\}} \tilde{C}_i = \emptyset$.

Let T be a triangulation of Δ whose diameter is less than η and all whose vertices lie in the interiors of the sets C_i. (Clearly, such a triangulation exists.) Define A_i as the subset of vertices which lie in $\operatorname{int}(C_i)$. These sets are well defined since the interiors of the different C_i's are disjoint. Since the triangulation is of diameter less than η, each edge between two vertices which belong to the same A_i lies in $\tilde{C}_i$. Indeed, its endpoints are in C_i and any point on the edge is of distance less than η to any endpoint, hence it is in $\tilde{C}_i$. Thus, since $\tilde{C}_i$ contains no connector, A_i contains no discrete connector.

We can now apply Theorem 3 to deduce that there exists a simplex $t = \{v_0, \ldots, v_d\} \in T$ such that $\forall i : v_i \in A_i$. But since $v_i \in C_i$ for all i, this implies $f(v_0, \ldots, v_d) < \eta < \epsilon$, a contradiction to the definition of ϵ. This completes the proof. ◀

[2] To be precise, this relies on the slightly stronger assumption that each connected component of C_i is at least δ-far from one of the facets of Δ. While this extra assumption can be avoided in the proof, it clearly holds in our setting so we make it for simplicity.

To prove Theorem 3, we use the classical Sperner's lemma [24]. To present it, a few more definitions are due.

A *$(d+1)$-labelling* of a triangulation T is of the simplex $\Delta = \text{conv}(e_0, \dots, e_d)$ is a function $\ell : V(T) \to \{0, 1, \dots, d\}$, that is, an assignment of one of $d+1$ colors to each vertex of the triangulation. A $(d+1)$-labelling ℓ is called *proper* if $\ell(e_i) = i$, and for each $v \in T$ that belongs to a lower-dimensional face $\text{conv}(e_{i_1}, \dots, e_{i_r})$, we have $\ell(v) \in \{i_1, \dots, i_r\}$.

▶ **Theorem 4** (Sperner's lemma). *For any triangulation T of Δ, any proper labelling of T contains a simplex all whose vertices have different labels.*

Proof of Theorem 3. We define, using the sets $A_0, \dots, A_d$, a proper labelling ℓ of T. For any j and any $v \in A_j$, $\ell(v)$ is defined as the minimal $i \in \{0, \dots, d\}$ such that the connected component of v in A_j does not intersect the i'th face of Δ. Note that ℓ is well-defined, since each v belongs to a single A_j and no A_j contains a connector. Clearly, $\ell(v) \neq i$, whenever v belongs to the i'th face of Δ. Furthermore, this implies that if $v \in \text{conv}(e_{i_1}, \dots, e_{i_r})$, then $\ell(v) \in \{i_1, \dots, i_r\}$ (as all other colors are forbidden). Hence, ℓ is proper.

By Sperner's lemma, applied to the labelling ℓ, there exists a simplex $\{v_0, \dots, v_d\} \in T$ all whose vertices have different labels. Assume w.l.o.g. that $\ell(v_i) = i$. We want to show that each v_i belongs to a different A_j, which will complete the proof. Assume towards contradiction $v_i, v_k \in A_j$ for $i \neq k$. On the one hand, $\ell(v_i) = i \neq k = \ell(v_k)$. On the other hand, $v_i,\ v_k$ belong to the same connected component in A_j, hence, by the definition of ℓ they must map to the same value. A contradiction, and Theorem 3 follows. ◀

3.1.2 Proof of Proposition 1

We are now ready to prove Proposition 1.

Let $T_1, T_2, \dots$ be a tiling of $\mathbb{R}^d$ that satisfies the assumptions of the proposition. Consider the restriction of the tiling to a large simplex Δ (say, of side length $100m$).

For $i = 0, \dots, d$, denote by $C_i \subset \Delta$ the union of all tiles colored i, restricted to Δ. Clearly, C_i is a closed set and the distance between any two connected components of C_i is at least 2δ. (Indeed, each tile is connected, and as the error resilience of the tiling is δ, the distance betwen two same-colored tiles is at least 2δ).

We claim that no C_i contains a connector. Indeed, a connector cannot include points from different tiles. A single tile is included in a box with side length m, and thus, cannot touch all facets of a simplex with side length $100m$. Hence, there is no single-colored connector.

Therefore, we can apply Theorem 2 to deduce that $\bigcap_{i=0}^{d} C_i \neq \emptyset$, which is exactly the assertion of Proposition 1.

3.2 Lower bound for contractible tiles, using Čech cohomology

Recall that a set in $\mathbb{R}^d$ is called *contractible* if it can be continuously shrunk to a point within the set. (The formal definition is that the identity is homotopic to a constant map.)

Informally, in this section we prove that if the tiles and their non-empty intersections are finite unions of contractible sets (that do not have to be uniformly bounded), then at least $d+1$ colors are required for error resilience. We note that the proof uses a somewhat heavier algebraic-topologic machinery, and so, a reader might prefer to skip it in first reading.
Due to the possibility of pathologies, the formal statement is a bit more cumbersome:

▶ **Proposition 5.** *Let $T_1, T_2, \dots$ be a colored tiling of $\mathbb{R}^d$ with positive error resilience, in which the tiles and all their non-empty intersections are disjoint unions of finitely many closed contractible sets. Assume that the tiling is locally finite (meaning that the number* *of tiles that intersect any bounded ball $B(0,r)$ is finite) and that all T_i's are bounded (not necessarily uniformly). In addition, assume that each T_i has an open neighborhood U_i such that for any I,*

$$\bigcap_{i\in I} U_i \neq \emptyset \Leftrightarrow \bigcap_{i\in I} T_i \neq \emptyset,$$

and the U_i's and their non-empty intersections are disjoint unions of finitely many contractibles. Then the number of colors is at least $d+1$.

A similar method proves an analogous statement for colored tilings of the sphere $\mathbb{S}^d$ (i.e., the unit sphere in $\mathbb{R}^{d+1}$):

▶ **Proposition 6.** *Let $T_1, T_2, \ldots, T_N$ be a colored tiling of $\mathbb{S}^d$ with positive error resilience, in which the tiles and all their non-empty intersections are disjoint unions of finitely many closed contractible sets. Assume that each T_i has an open neighborhood U_i such that for any set of indices I,*

$$\bigcap_{i\in I} U_i \neq \emptyset \Leftrightarrow \bigcap_{i\in I} T_i \neq \emptyset,$$

and the U_i's and their non-empty intersections are disjoint unions of finitely many contractibles. Then the number of colors is at least $d+1$.

▶ Remark 7. We stress that for most natural tilings the additional assumption on the existence of the neighborhoods U_i follows from the existence of T_i's with the corresponding properties. However, there are topological pathologies in which this is not the case.

The proof of Propositions 5 and 6 uses the notion of Čech cohomology and classical results regarding its properties. For the ease of reading, we begin with an intuitive explanation of the proof ideas, and then present the formal proof.

Intuitive proof. The d'th (singular) cohomology group is a topological invariant of a manifold which roughly counts "non trivial holes" of dimension d. A classical result asserts that the d'th cohomology group of a d-dimensional compact oriented manifold like $\mathbb{S}^d$ is $\mathbb{R}$. (This corresponds to the intuitive understanding that $\mathbb{S}^d$ has one d-dimensional hole.) The de-Rham cohomology and the Čech cohomology are analytic and algebro-geometric/combinatorial invariants, that in many cases agree with their topological cousin. In particular, the d'th de-Rham and Čech cohomologies of $\mathbb{S}^d$ are equal to $\mathbb{R}$ as well.

The d'th Čech cohomology with respect to an open cover of the manifold depends on properties of intersections of $d+1$ sets in that cover. In general, it depends on the sets which form the cover, however, it is known that if these sets and their non-empty intersections are finite disjoint unions of contractibles, then the cohomology groups remain the same, independently of the cover. In particular, if the d'th Čech cohomology with respect to such a cover is non trivial, then there must be $d+1$ sets with a non-empty intersection.

Hence, for our cover $U_1, U_2, \ldots$, we know that its d'th Čech cohomology is $\mathbb{R}$. This readily completes the proof of the proposition for $\mathbb{S}^d$, as this implies that there must be a point that belongs to at least $d+1$ of the U_i's. The proof in $\mathbb{R}^d$ works in essentially the same way, with cohomology groups replaced by cohomology groups with compact support.

Formal proof. For the proof we recall the notion of Čech cohomology with values in the constant sheaf $\underline{\mathbb{R}}$, and describe the slightly less standard concept of Čech cohomology with compact support.

Definitions. Let S be either $\mathbb{R}^d$ or a compact manifold such as $\mathbb{S}^d$. Let $\mathcal{U} = \{U_1, U_2, \ldots\}$ be an open cover of S. If S is compact, we assume the collection to be finite. If S is $\mathbb{R}^d$, we assume it to be locally finite and assume in addition that each U_i is bounded.

- A $q-$*simplex* $\sigma = (U_{i_0}, \ldots, U_{i_q})$ of $\mathcal{U}$ is an ordered collection of $q+1$ different sets chosen from $\mathcal{U}$, such that
$$\bigcap_{k=0}^{q} U_{i_k} \neq \emptyset.$$
- For a $q-$simplex $\sigma = (U_{i_k})_{k\in\{0,\ldots,q\}}$, the j'th *partial boundary* is the $(q-1)$-simplex
$$\partial_j \sigma := (U_{i_k})_{k\in\{0,\ldots,q\}\setminus\{j\}},$$
obtained by removing the j'th set from σ.
- A $q-$*cochain* of $\mathcal{U}$ is a function which associates to any $q-$simplex a real number. The $q-$cochains form a vector space denoted by $C^q(\mathcal{U}, \underline{\mathbb{R}})$, with operations
$$(\lambda f + \mu g)(\sigma) = \lambda f(\sigma) + \mu g(\sigma), \quad \text{where} \quad \lambda, \mu \in \mathbb{R},\ f, g \in C^q(\mathcal{U}, \underline{\mathbb{R}}),\ \sigma \ \text{ is a } q-\text{simplex}.$$
Similarly, we define $C^q_c(\mathcal{U}, \underline{\mathbb{R}})$, as the vector space of $q-$cochains *with compact support*, meaning those cochains which assign 0 to all $q-$simplices, except for finitely many.
- There is a *differential map* $\delta_q : C^q(\mathcal{U}, \underline{\mathbb{R}}) \to C^{q+1}(\mathcal{U}, \underline{\mathbb{R}})$ whose application to $f \in C^q(\mathcal{U}, \underline{\mathbb{R}})$ is the $(q+1)-$cochain $\delta_q(f)$ whose value at a $(q+1)-$simplex σ is
$$(\delta_q f)(\sigma) = \sum_{j=0}^{q+1} (-1)^j f(\partial_j \sigma).$$
The restriction of δ_q to $C^q_c(\mathcal{U}, \underline{\mathbb{R}})$ maps it to $C^{q+1}_c(\mathcal{U}, \underline{\mathbb{R}})$.
- It can be easily seen that $\delta_{q+1} \circ \delta_q = 0$.
- The *q'th Čech cohomology group (with compact support)* of S with respect to the cover $\mathcal{U}$ and values in $\underline{\mathbb{R}}$ is
$$\check{H}^q(\mathcal{U}, \underline{\mathbb{R}}) := \mathrm{Ker}(\delta_q)/\mathrm{Image}(\delta_{q-1}),$$
$$\check{H}^q_c(\mathcal{U}, \underline{\mathbb{R}}) := \mathrm{Ker}(\delta_q|_{C^q_c(\mathcal{U}, \underline{\mathbb{R}})})/\mathrm{Image}(\delta_{q-1}|_{C^{q-1}_c(\mathcal{U}, \underline{\mathbb{R}})}).$$
- A cover (by open sets) is *good* if all its sets as well as their multiple intersections are either empty or contractible. It is *almost good* if all non empty intersections are unions of finitely many disjoint contractible components.

Classical results we use. The first result we use is the following:

▶ **Theorem 8.** *If S is a compact smooth orientable manifold (such as $\mathbb{S}^d$), and $\mathcal{U}$ is a good or an almost good finite cover, then*
$$\check{H}^i(\mathcal{U}, \underline{\mathbb{R}}) \simeq H^i_{dR}(S),$$
where the right hand side is the standard de-Rham cohomology group.
Similarly, if $S = \mathbb{R}^d$ and $\mathcal{U}$ is a locally finite good or almost good cover whose sets are bounded, then
$$\check{H}^i_c(\mathcal{U}, \underline{\mathbb{R}}) \simeq H^i_{dR,c}(S),$$
where the right hand side is the i'th de-Rham cohomology group with compact support.

For further reading about de-Rham cohomology, with or without compact support, we refer the reader to [4, Sec. 1]. For further reading about the Čech cohomology, we refer to [4, Sec. 8]. In particular, Theorem 8, for the compact case and good covers is Theorem 8.9 there. The passage to almost good covers is straightforward: In the paragraph which precedes the proof, it is explained that the obstructions to the isomorphism between Čech and de-Rham cohomologies are given by products of the i'th de-Rham cohomology groups, for $i \geq 1$, of the different intersections $\bigcap_{k=0}^{q} U_{i_k}$. Since those intersections are disjoint unions of contractibles, their higher cohomology groups vanish, hence there is no obstruction to the isomorphism.

Regarding the case $S = \mathbb{R}^d$, the proof in [4, Sec. 8] requires a few small changes: In the statement of Proposition 8.5 there, one needs to replace the de-Rham complex of the manifold with the de-Rham complex with compact support, and the direct product with direct sum. The maps r, δ which appear there will still be well defined by our local finiteness assumption on the cover, and the assumption that U_i's are bounded. The proof requires no change. Then, the double complex in the definition of Proposition 8.8 should also be defined using direct sum rather than direct product, but again there is no change in the proof. Given these changes in definitions, the proof of Theorem 8.9 (also for the almost good case) is unchanged.

The second standard result, which is a consequence of Poincaré duality, is the following:

▶ **Theorem 9.** *For a compact smooth oriented manifold S of dimension d (such as $\mathbb{S}^d$),*

$$H^d_{dR}(S) \simeq \mathbb{R}.$$

Similarly, for $S = \mathbb{R}^d$, we have $H^d_{dR,c}(\mathbb{R}^d) \simeq \mathbb{R}$.

See, for example, [4, Sec. 7] for the compact case, and [4, Sec. 4] for $\mathbb{R}^d$.

Theorems 8 and 9 yield:

▶ **Corollary 10.** *If S is a compact smooth orientable manifold (such as $\mathbb{S}^d$), and $\mathcal{U}$ is a good or an almost good finite cover, then*

$$\check{H}^d(\mathcal{U}, \underline{\mathbb{R}}) = \mathbb{R}.$$

Similarly, if $S = \mathbb{R}^d$ and $\mathcal{U}$ is a locally finite good or almost good cover whose sets are bounded, then $\check{H}^d_c(\mathcal{U}, \underline{\mathbb{R}}) = \mathbb{R}$.

Proof of Propositions 5 and 6. We show that there must exist $d+1$ T_i's whose intersection is non-empty. This clearly implies that for achieving any positive error resilience, at least $d+1$ colors are needed.

Assume on the contrary that any $(d+1)$-intersection of the T_i's is empty. Let U_i be as in the statement of the Propositions. Then by definition, they form an almost good cover. All intersections of at least $d+1$ U_i's are empty by our assumptions. Therefore, there are no $d-$simplices, and so $C^d(\mathcal{U}, \underline{\mathbb{R}}) = 0$. Thus, in the compact case, $\check{H}^d(\mathcal{U}, \underline{\mathbb{R}}) = 0$. But on the other hand, by Corollary 10,

$$\check{H}^d(\mathcal{U}, \underline{\mathbb{R}}) \simeq \mathbb{R},$$

a contradiction. For $\mathbb{R}^d$ the same argument works, with $\check{H}^d_c(\mathcal{U}, \underline{\mathbb{R}})$ in place of $\check{H}^d(\mathcal{U}, \underline{\mathbb{R}})$. ◀

4 Upper Bounds on the Error Resilience

In this section we consider tilings of $\mathbb{R}^d$ by tiles $T_1, T_2, \ldots$ of volume at most 1. Each point in $\mathbb{R}^d$ is colored in one of $k \geq d+1$ colors, and our goal is to maximize the minimal distance t between two points of the same color that belong to different tiles. (The maximum is taken over all possible tilings that satisfy the mild regularity conditions stated in Section 2 and over all possible colorings.) Clearly, the error resilience of a rounding scheme based on such a colored tiling is $t/2$.

We present two upper bounds on t, using the Brunn-Minkowski inequality and results on sphere packing. Another bound, using the Minkowski-Steiner formula, is presented in the full version of the paper (see [13]).

The basic idea behind our upper bound proofs is as follows. Assume we have a colored tiling of $\mathbb{R}^d$, with minimal distance t. Pick a single color – say, black – and consider all black tiles inside a large cube S. We obtain a new collection of tiles $T'_1, T'_2, \ldots, T'_m$ that covers part of S. The assumption that the minimal distance between two same-colored points in different tiles is t implies that if we inflate each black tile T'_i into its open parallel outer body of radius $t/2$,

$$T''_i = \{x : \exists y \in T'_i, |x - y| < t/2\}, \tag{1}$$

then the inflations T''_i are pairwise disjoint. Hence, the sum of their volumes essentially cannot exceed the volume of the large cube, and this allows bounding t from above.

4.1 An upper bound using the Brunn-Minkowski inequality

The inflations T''_i can be represented in terms of the *Minkowski sum* of sets in $\mathbb{R}^d$.

▶ **Definition 11.** *For $A, B \subset \mathbb{R}^d$, the Minkowski sum of A, B is $A+B = \{a+b : a \in A, b \in B\}$.*

In terms of this definition, we have

$$T''_i = T'_i + B(0, t/2), \tag{2}$$

where $B(0, t/2)$ is an open ball of radius $t/2$ around the origin. This allows us to lower bound the volume of each T''_i, using the classical Brunn-Minkowski (BM) inequality (see, e.g., [5]). Recall the inequality asserts the following.

▶ **Theorem 12** (Brunn-Minkowski). *Let A, B be compact sets in $\mathbb{R}^d$. Then*

$$\lambda(A+B)^{1/d} \geq \lambda(A)^{1/d} + \lambda(B)^{1/d},$$

where $\lambda(X)$ is the volume of X (formally, the d-dimensional Lebesgue measure of X).

▶ **Proposition 13.** *Let $T_1, T_2, \ldots$ be a k-colored tiling of $\mathbb{R}^d$, with tiles of volume ≤ 1 and minimal distance t. Then*

$$t \leq \left(\frac{2\Gamma(d/2+1)^{1/d}}{\sqrt{\pi}} \right) \cdot (k^{1/d} - 1),$$

where $\Gamma(\cdot)$ is the Gamma function.

Proof. Consider a cube S such that $\lambda(S) = n^d$ (for some "large" n). By the pigeonhole principle, there exists a color (say, black) that covers at least n^d/k of the volume of S. Look at the black tiles whose intersection with S is non-empty, and denote their intersections with S by $T'_1, T'_2, \ldots, T'_m$. Hence, we have m "black" subsets of S, each of volume at most 1, whose total volume is at least n^d/k.

For each T'_i, define $T''_i = T'_i + B(0, t/2)$. By assumption, the regions T''_i are disjoint. Furthermore, they are included in $S + B(0, t/2)$ whose volume is less than $(n+t)^d$. Hence,

$$\sum_i \lambda(T''_i) \leq (n+t)^d. \tag{3}$$

By the Brunn-Minkowski inequality, we have

$$\forall i : \lambda(T''_i)^{1/d} \geq \lambda(T'_i)^{1/d} + (b_{t/2})^{1/d},$$

where $b_{t/2}$ is the volume of the d-dimensional ball $B(0, t/2)$. Thus,

$$\forall i : \lambda(T''_i) \geq \sum_{j=0}^{d} \binom{d}{j} \lambda(T'_i)^{j/d} (b_{t/2})^{1-j/d}.$$

Summing over i and using (3), we get

$$(n+t)^d \geq \sum_{i=1}^{m} \sum_{j=0}^{d} \binom{d}{j} \lambda(T'_i)^{j/d} (b_{t/2})^{1-j/d}. \tag{4}$$

As $0 \leq \lambda(T'_i) \leq 1$, for any $0 \leq j \leq d$ we have $\sum_i \lambda(T'_i)^{j/d} \geq \sum_i \lambda(T'_i) \geq n^d/k$, and hence we obtain

$$(n+t)^d \geq \frac{n^d}{k} \cdot \left(1 + b_{t/2}^{1/d}\right)^d.$$

This implies

$$(1 + \frac{t}{n})k^{1/d} - 1 \geq b_{t/2}^{1/d} = \frac{\pi^{1/2}}{\Gamma(d/2+1)^{1/d}} \cdot \frac{t}{2}.$$

Letting $n \to \infty$ and rearranging, we obtain

$$t \leq \left(\frac{2\Gamma(d/2+1)^{1/d}}{\sqrt{\pi}}\right) \cdot (k^{1/d} - 1),$$

as asserted. ◀

Asymptotic upper bound. For a large number $k \gg d$ of colors, Proposition 13 gives the upper bound

$$t \leq (\sqrt{\frac{2}{\pi e}} + o_d(1))\sqrt{d}k^{1/d},$$

as follows from (5). This bound is not far from being tight. Indeed, its dependence on k is correct, as it can be easily matched by a periodic cubic tiling, in which each tile is a cube with side length 1 and the basic unit is a large cube with side length $k^{1/d}$ that contains each color in exactly one tile (in the same order). Moreover, even regarding the "coefficient" of $k^{1/d}$, the optimal asymptotic upper bounds for $d = 3, 8, 24$ which we obtain below via the sphere packing problem, improve over this bound by only a small factor.

Upper bound for $k = d + 1$ colors. To estimate the upper bound we obtain in this case, note that

$$(d+1)^{1/d} - 1 = (1 + o_d(1))\frac{\ln(d)}{d}, \qquad \text{and} \qquad \Gamma(d/2+1)^{1/d} = \left(\frac{1}{\sqrt{2e}} + o_d(1)\right)\sqrt{d}. \quad (5)$$

Therefore, the bound we obtain in this case is

$$t \leq \left(\sqrt{\frac{2}{\pi e}} + o_d(1)\right)\frac{\ln d}{\sqrt{d}},$$

which implies that the error resilience decreases to zero as d tends to infinity. For comparison, the lower bound we obtain in Section 5.3 is $t \geq \Omega(1/d)$.

Upper bounds for small values of d, k. For $d = 3, k = 4$, the bound is

$$t \leq \frac{2\Gamma(2.5)^{1/3}}{\sqrt{\pi}} \cdot (4^{1/3} - 1) \approx 0.729.$$

For $d = 2$ and $k = 3, 4$, the upper bounds we obtain are $t \leq 0.826$ and $t \leq 1.128$, respectively. For comparison, the best constructions we have in these settings satisfy $t = 0.5$, $t = 1/\sqrt{2}$, and $t = 1$, respectively.

Discussion. The upper bound given by Proposition 13 is loose in two ways. One source of loss is the application of the Brunn-Minkowski inequality. Here, the inequality is tight if the tiles are *balls*, and the farther they are from balls, the larger is the loss. Another source of loss is the space left between the inflations, that is not taken into account in the proof.

Interestingly, there is a dichotomy between these two sources of loss. As follows from the sphere packing problem, when the tiles are balls (and so, there is no loss in the BM inequality), the space between the inflations (and so, the loss of the second type) is relatively large. The space between the inflations can be made smaller if the tiles are taken to be polytopes with a few vertices. However, this comes at the expense of increased loss in the BM inequality, as is demonstrated in the full version of the paper.

Optimality of our 1-dimensional rounding scheme. The argument described above gives an easy proof of the optimality of the 1-dimensional rounding scheme presented in the introduction. Indeed, consider a 2-colored tiling of the line and look at the segment $I = [-n, n]$ for some large n. By the pigeonhole principle, we may assume that black tiles cover at least half of I. By the 1-dimensional Brunn-Minkowski inequality, for each black tile $T_i' \subset I$ and the corresponding inflation $T_i'' = T_i' + (-t/2, t/2)$, we have $\lambda(T_i'') \geq \lambda(T_i') + t$. As $\forall i : \lambda(T_i') \leq 1$, there are at least n tiles. Since the T_i'''s are pairwise disjoint and included in $[-n-1, n+1]$, we obtain

$$2n + 2 \geq \sum_i \lambda(T_i'') \geq \sum_i \lambda(T_i') + \sum_i t \geq n + nt,$$

and thus, $t \leq (n+2)/n$. By letting n tend to infinity, we obtain $t \leq 1$, implying that the error resilience of any two-colored tiling of the line is at most $1/2$.

4.2 An upper bound using the Sphere Packing problem

Our second upper bound uses reduction to the classical *sphere packing* problem, which asks for the maximal possible *density* of a set of non-intersecting congruent spheres in $\mathbb{R}^d$.

▶ **Definition 14.** *The density of a sphere packing (i.e., collection of pairwise disjoint congruent spheres) $P = \cup P_i$ in $\mathbb{R}^d$ is*

$$\limsup_{r\to\infty} \frac{\lambda\left(B(0,r)\cap\bigcup P\right)}{\lambda(B(0,r))}.$$

Intuitively, this measures the fraction of the volume of a large ball covered by the packing.

▶ **Notation 15.** *Denote the maximal density of a sphere packing in $\mathbb{R}^d$ by δ_d, and the volume of the unit ball $B(0,1)\subset\mathbb{R}^d$ by $v_d = \pi^{d/2}/\Gamma(d/2+1)$.*

▶ **Proposition 16.** *Let $T_1, T_2, \ldots$ be a tiling of $\mathbb{R}^d$ in k colors, with tiles of volume ≤ 1 and minimal distance t. Then*

$$t \leq \left(2(\delta_d/v_d)^{1/d}\right)\cdot k^{1/d} = \left(\frac{2\Gamma(d/2+1)^{1/d}\cdot\delta_d^{1/d}}{\sqrt{\pi}}\right)\cdot k^{1/d}.$$

Note that the asymptotic upper bound of Proposition 16 is stronger than the asymptotic upper bound that follows from Proposition 13 by the constant factor $(\delta_d)^{1/d}$. For small values of k, the upper bound given by Proposition 13 is stronger.

Proof. Let T be a k-colored tiling of $\mathbb{R}^d$ that satisfies the assumptions of the proposition, and consider the sequence of balls $\{B(0,n)\}_{n=1,2,3,\ldots}$. By the pigeonhole principle, there exists a color (say, black) such that for each n_ℓ in an infinite subsequence $\{n_\ell\}_{\ell=1,2,\ldots}$, the intersection of the black tiles with the ball $B(0,n_\ell)$ has volume of at least

$$\lambda(B(0,n_\ell))/k = \frac{n_\ell^d\cdot v_d}{k}.$$

As the volume of each tile is at most 1, we know that for each n_ℓ, the number of black tiles that intersect $B(0,n_\ell)$ is at least $n_\ell^d v_d/k$.
Pick some value n_ℓ, denote the intersections of black tiles with $B(0,n_\ell)$ by $T'_1, T'_2, \ldots$, and take one point x_i from each tile T'_i. As the minimal distance between two black points in different tiles is t, balls of radius $t/2$ around the points x_i are pairwise disjoint. Hence, their total volume is at least

$$\frac{n_\ell^d\cdot v_d}{k}\cdot(t/2)^d\cdot v_d.$$

On the other hand, each such ball is contained in the ball $B(0,(n_\ell+t))$ (since its radius is $t/2$, and it contains a point in $B(0,n_\ell)$). This implies that for any $\epsilon>0$ and for a sufficiently large $\ell=\ell(\epsilon)$, the total volume of these balls must be smaller than $(1+\epsilon)\delta_d\cdot\lambda(B(0,n_\ell+t))$, as otherwise, the infinite collection of the balls $B(x_i,t/2)$ (where for each ball $B(0,n_\ell)$ we select x_i's in the way described above, respecting the x_i's selected for smaller values of n_ℓ) would be a sphere packing of $\mathbb{R}^d$ whose density is larger than δ_d. Therefore, for a sufficiently large n_ℓ, we have

$$\frac{n_\ell^d\cdot v_d}{k}\cdot(t/2)^d\cdot v_d \leq (1+\epsilon)(1+\frac{t}{n_\ell})^d\delta_d\cdot n_\ell^d v_d,$$

and letting $\epsilon\to 0$ and $n_\ell\to\infty$, we obtain $t\leq 2(\delta_d/v_d)^{1/d}\cdot k^{1/d}$, as asserted. ◀

Discussion. In the two last decades, there has been a tremendous progress in the research of the sphere packing problem. In 2005, Hales ([18], see also [17]) solved the problem for $d = 3$, proving a 17'th century conjecture of Kepler. Three years ago, in a beautiful short paper, Viazovska [25] solved the problem for $d = 8$, and shortly after, Cohn, Kumar, Miller, Radchenko, and Viazovska [6] used Viazovska's method along with other tools to solve the problem for $d = 24$. For other dimensions, the problem is still open. We can use the results of [6, 18, 25], along with the value of δ_2 that was obtained already by Lagrange, to obtain tight upper bounds on t in dimensions $2, 3, 8$, and 24.

- For $d = 2$, Lagrange (1773) showed that $\delta_2 = \frac{\pi}{2\sqrt{3}} \approx 0.907$. Hence, we obtain the bound $t \leq \frac{2^{1/2}}{3^{1/4}} \cdot k^{1/2} \approx 1.074k^{1/2}$.
- For $d = 3$, Hales [18, 17] showed that $\delta_3 = \frac{\pi}{3\sqrt{2}} \approx 0.740$. Hence, we obtain the bound $t \leq 2^{1/6} \cdot k^{1/3} \approx 1.122k^{1/3}$.
- For $d = 8$, Viazovska [25] showed that $\delta_8 = \frac{\pi^4}{2^4 4!} \approx 0.254$. Hence, we obtain the bound $t \leq \sqrt{2}k^{1/8} \approx 1.414k^{1/8}$.
- For $d = 24$, Cohn et al. [6] showed that $\delta_{24} = \frac{\pi^{12}}{12!} \approx 0.0019$. Hence, we obtain the bound $t \leq 2k^{1/24}$.

As we show in Section 5, all these bounds are asymptotically tight.

Using the same method, we can leverage any upper bound for the sphere packing problem (namely, upper bound on δ_d) into an upper bound on the error resilience of a rounding scheme in the corresponding dimension. The best currently known bound on δ_d for a large d is by Cohn and Zhao [7], who obtained a constant-factor improvement over the classical Kabatiansky-Levenshtein [21] bound $\delta_d \leq 2^{-(0.5990+o(1))d}$. A list of conjectured bounds for $d \leq 10$ can be found in [8].

5 Lower Bounds on the Error Resilience

In this section (like in Section 4), we consider tilings of $\mathbb{R}^d$ by tiles $T_1, T_2, \ldots$ of volume at most 1. Each point in $\mathbb{R}^d$ is colored in one of $k \geq d + 1$ colors, and our goal is to maximize the minimal distance t between two points of the same color that belong to different tiles.

We obtain lower bounds on t for various values of d, k, by constructing explicit tilings. First, we present *brick-wall* tilings, which provide lower bounds for small values of d, k. Then we present tilings based on *close sphere packing*, which show the tightness of our asymptotic bounds for $d = 2, 3, 8, 24$. Finally, we inductively construct a *dimension-reducing* tiling which shows that for any dimension d, positive error resilience can be obtained with $d + 1$ colors.

5.1 Brick-wall tilings

We begin with a tiling of the plane, and then use it to construct a tiling of $\mathbb{R}^3$.

5.1.1 2-dimensional brick wall

In the 2-dimensional brick wall tiling with k colors, demonstrated in Figure 2, each tile is a rectangle with side lengths $\sqrt{(k-2)/2}$ and $\sqrt{2/(k-2)}$ (and so, the area of each tile is 1). The tiling is periodic, where the basic unit is two rows of adjacent rectangles, colored in a round robin fashion. For an even k, the second row is placed exactly below the first row, and the sequence of colors is shifted by $k/2$. For an odd k, the second row is indented by half a brick (making the tiling look like a brick wall), and the sequence is shifted by $(k + 1)/2$.

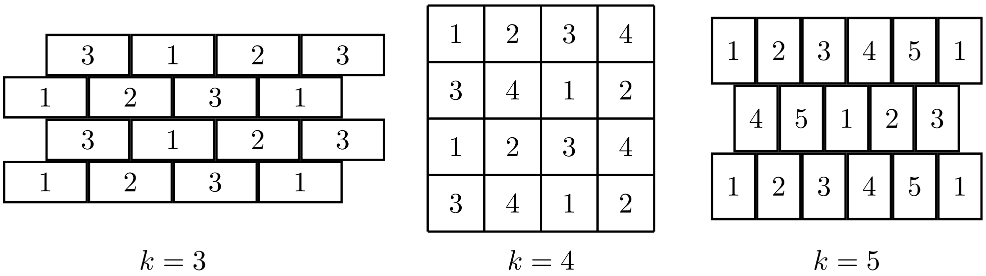

Figure 2 The 2-dimensional brick wall tiling for $k = 3, 4, 5$ colors. The ratio between the width and the height of each tile is $2 : k - 2$.

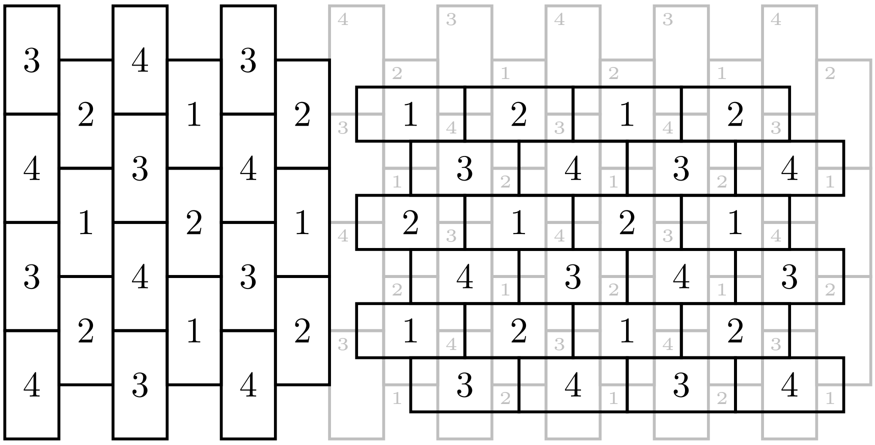

Figure 3 The 3-dimensional brick wall tiling.

It is easy to see that the minimal distance between two same-colored points in different tiles is $\sqrt{(k-2)/2}$. (This distance is attained both in the vertical and in the horizontal directions. Having the same minimal distance in both directions is the optimization that dictates the side lengths of the bricks.) In particular, we obtain the lower bounds $t \geq 1/\sqrt{2}$ for 3 colors, $t \geq 1$ for 4 colors, and $t \geq \sqrt{3/2}$ for 5 colors.

5.1.2 3-dimensional brick wall

The 3-dimensional brick wall (3BW) tiling, demonstrated in Figure 3, is a periodic tiling of $\mathbb{R}^3$, colored in 4 colors. In order to present the tiling, we need an auxiliary notation.

Notation. Consider the slab $D = \mathbb{R} \times \mathbb{R} \times [z_1, z_2] \subset \mathbb{R}^3$. We say that a tiling $T_1, T_2, \ldots$ of D is a *fattened plane tiling* if there exists a tiling $T'_1, T'_2, \ldots$ of the plane such that $\forall i : T_i = T'_i \times [z_1, z_2]$.

The structure of 3BW. The 3BW tiling is periodic, where the basic unit consists of two brick wall layers, placed one on top of the other in the way presented in Figure 3. Each brick wall layer is a fattening of a brick wall tiling of the plane. The underlying plane tiling is a

15	16	13	14	15	16	13	14
1	2	3	4	1	2	3	4
5	6	7	8	5	6	7	8
9	10	11	12	9	10	11	12
13	14	15	16	13	14	15	16
3	4	1	2	3	4	1	2
7	8	5	6	7	8	5	6
11	12	9	10	11	12	9	10
15	16	13	14	15	16	13	14
1	2	3	4	1	2	3	4

Figure 4 The honeycomb of rectangles tiling for $d = 2$ and $k = 16$ colors. The boundaries of the basic "large rectangles" are depicted in bold. The placement of the tiles in each single color corresponds to the honeycomb lattice.

periodic tiling, in which the basic unit consists of four columns of adjacent rectangles, where the even columns are indented by half a brick, making the tiling look like a brick wall. In the lower layer, in odd columns, the colors 1,2 are used alternately, and in even columns, the colors 3,4 are used alternately. Furthermore, the colors in the third and fourth columns are shifted by one, see Figure 3. In the upper layer columns are replaced by rows. Note that once the layers are placed, the coloring of one layer fully determines the coloring of the other.

Analysis given in the full version of the paper shows that by taking $a = 1$ and fixing the height of each layer to be $1/2$, we obtain $t = 1/2$, and thus, error resilience of 0.25.

5.2 Tilings based on close sphere packing

We present the tiling in the case of $\mathbb{R}^2$, where it is easier to describe and analyze, and then we generalize it to higher dimensions.

5.2.1 Honeycomb of rectangles

In the honeycomb of rectangles tiling of the plane with $k = m^2$ colors, each tile is a rectangle with side lengths a and $1/a$, where

$$a = \left(\frac{m^2 - 2m + 1}{\frac{3}{4}m^2 - m}\right)^{1/4} \geq \left(\frac{4}{3} - \frac{8}{3m}\right)^{1/4}. \tag{6}$$

The tiling is periodic, where the basic unit is composed as follows. First, we construct a basic "large rectangle", which is an m-by-m square block of tiles, using all the $k = m^2$ colors (in arbitrary order). Then, the basic unit of the tiling is two "fat rows" of adjacent large rectangles, where the second row is indented by half a large rectangle. The coloring of each large rectangle is the same. The tiling, for $k = 16$, is demonstrated in Figure 4. Note that the tiles colored in some single color form the shape of a honeycomb lattice (including the centers of the hexagons). This is why we call the tiling "honeycomb of rectangles".

It is easy to see that the minimal horizontal and diagonal distances between two same-colored tiles are

$$(m-1)a \qquad \text{and} \qquad \sqrt{\left(\frac{m-1}{a}\right)^2 + \left(\left(\frac{m}{2}-1\right)a\right)^2},$$

respectively. By choosing a such that the two distances are equal, we obtain (6), and the asymptotic lower bound

$$t \geq \left(\frac{4}{3} - \frac{8}{3m}\right)^{1/4} \cdot (m-1) \geq (4/3)^{1/4} \cdot \sqrt{k} - O(1) \approx 1.074\sqrt{k} - O(1),$$

that matches the upper bound obtained above up to an additive $O(1)$ term.

5.2.2 Close packing of boxes

Motivation. This construction, a k-colored tiling of $\mathbb{R}^3$ where $k \gg 3$, is a natural generalization to $\mathbb{R}^3$ of the "honeycomb of rectangles" tiling presented in Section 5.2.1. The idea behind the construction is to choose an optimal sphere packing in $\mathbb{R}^3$, and construct a fattened plane tiling in which the tiles in each color are placed at the centers of the spheres of the packing. (Recall that as the number of colors is large, the size of each tile is negligible with respect to the size of its inflation, and hence, we can treat the tiles as single points.)

We use the classical HCP lattice (one of the most common *closed packings*, see [9]), which corresponds to a periodic sphere packing in which the basic unit is two hexagonal layers of spheres, where in the top layer, each sphere is placed on top in the hollow between three spheres in the bottom layer. The coordinates of the centers of these spheres are:

$$(r,r,r), (3r,r,r), (5r,r,r), \ldots, (2r, r+\sqrt{3}r, r), (4r, r+\sqrt{3}r, r), (6r, r+\sqrt{3}r, r), \ldots$$

for the bottom layer, and

$$(2r, r+\frac{\sqrt{3}}{3}r, r+\frac{2\sqrt{6}}{3}r), (4r, r+\frac{\sqrt{3}}{3}r, r+\frac{2\sqrt{6}}{3}r), \ldots, (r, r+\frac{4\sqrt{3}}{3}r, r+\frac{2\sqrt{6}}{3}r), (3r, r+\frac{4\sqrt{3}}{3}r, r+\frac{2\sqrt{6}}{3}r), \ldots$$

for the top layer.

The structure of the tiling. Assume that the number of colors is $k = m^3$. Each tile is a box with side lengths (a, b, c) to be determined below, and the basic unit is a "large box", which is an $m \times m \times m$ cubic block of tiles, using all the $k = m^3$ colors (in arbitrary order). Then, the basic unit of the tiling is a two-layer fattened plane tiling, in which each layer is a fattened copy of the "honeycomb of rectangles" tiling. The upper layer is shifted by $\frac{m}{2}a$ in the x-coordinate and by $\frac{\sqrt{3}m}{6}a$ in the y-coordinate, so that the corners of the large boxes lie in the coordinates of the sphere centers described above (for $r = \frac{m}{2}a$). A quick calculation shows that in order to make this possible, the proportion $(a : b : c)$ should be *approximately* $(2 : \sqrt{3} : 2\sqrt{6}/3)$ (where we neglect the size of each tile with respect to the size of the inflation, which can be absorbed in an $1 - o(1)$ multiplicative factor in the final value of t). The volume of each tile is clearly abc. In order to make the volumes of all tiles equal to 1, we need

$$a \cdot \frac{\sqrt{3}}{2}a \cdot \frac{\sqrt{6}}{3}a = 1,$$

and thus, $a = (6/\sqrt{18})^{1/3} = 2^{1/6} \approx 1.122$. Hence, the side lengths of each tile are

$$(2^{1/6}, 3^{1/2}/2^{5/6}, 2^{2/3}/3^{1/2}) \approx (1.122, 0.972, 0.916),$$

and the minimal distance between two same-colored points in different tiles is

$$(m-1)a = (2^{1/6} - o(1))k^{1/d},$$

which matches the upper bound proved in Section 4.1.

Generalization to higher dimensions. A similar tiling can be constructed to match any *lattice sphere packing*, assuming the number of colors k is sufficiently large with respect to d. Hence, any dense lattice sphere packing can be translated into a lower bound on the asymptotic error resilience of rounding schemes in the corresponding dimension. In particular, as the E_8 *lattice* and the *Leech lattice* which attain the maximal possible density of sphere packings in dimension 8 and 24 (respectively) are lattice packings, they can be translated to box tilings showing that the asymptotic upper bounds on the error resilience in $\mathbb{R}^8$ and $\mathbb{R}^{24}$ proved in Section 4.1 are tight.

5.3 The dimension reducing tiling

We exemplify the tiling in $\mathbb{R}^3$, but it will be apparent how to generalize it to higher dimensions.

Informal description of the tiling. Informally, the tiling is constructed as follows.

Step 1: We begin with dividing $\mathbb{R}^3$ into basic cubes with side length a (to be determined below), depicted in Figure 5(a). Then, we make the "interior" of each basic cube into a tile and give all these tiles the color 1. In order to keep a minimal distance of t between two points colored 1 in different tiles, we must leave a neighborhood of width $t/2$ in each side of each facet of the basic cube. Hence, we are left with "fattened" walls of total width t. The part of such a wall included in a basic cube is shown in Figure 5(b).

Step 2: We make the "interior" of each wall (i.e., fattened facet) into a tile and give all these tiles the color 2, as is demonstrated in Figure 5(b). (Note that each tile contains points from two adjacent basic cubes.) In order to keep a minimal distance of t between two points colored 2 in different tiles, we must leave a neighborhood of $t/\sqrt{2}$ near each edge (i.e., intersection of facets), as is shown in Figure 5(c). Hence, we are left with a "fattened skeleton".

Step 3: We make the "interior" of each edge of the skeleton into a tile and give all these tiles the color 3, as is shown in Figure 5(d). (Note that each tile contains points from four adjacent basic cubes.) In order to keep a minimal distance of t between two points colored 3 in different tiles, we must leave an additional neighborhood of $t/\sqrt{2}$ near each vertex (i.e., intersection of edges), as is demonstrated in Figure 5(d). Hence, we are left with neighborhoods of the corners (a.k.a. vertices).

Step 4: We make the neighborhoods of the corners into tiles and give them the color 4. (Note that each tile contains points from 8 adjacent basic cubes.) We have to make sure that the distance between each two such "fattened corners" is at least t, and this requirement dictates the choice of t.

The analysis of the tiling (including its formal definition and explanation of the choice of parameters) is presented in the full version of the paper. It is clear that the resulting 4-colored tiling can be generalized into a $(d+1)$-colored tiling of $\mathbb{R}^d$.

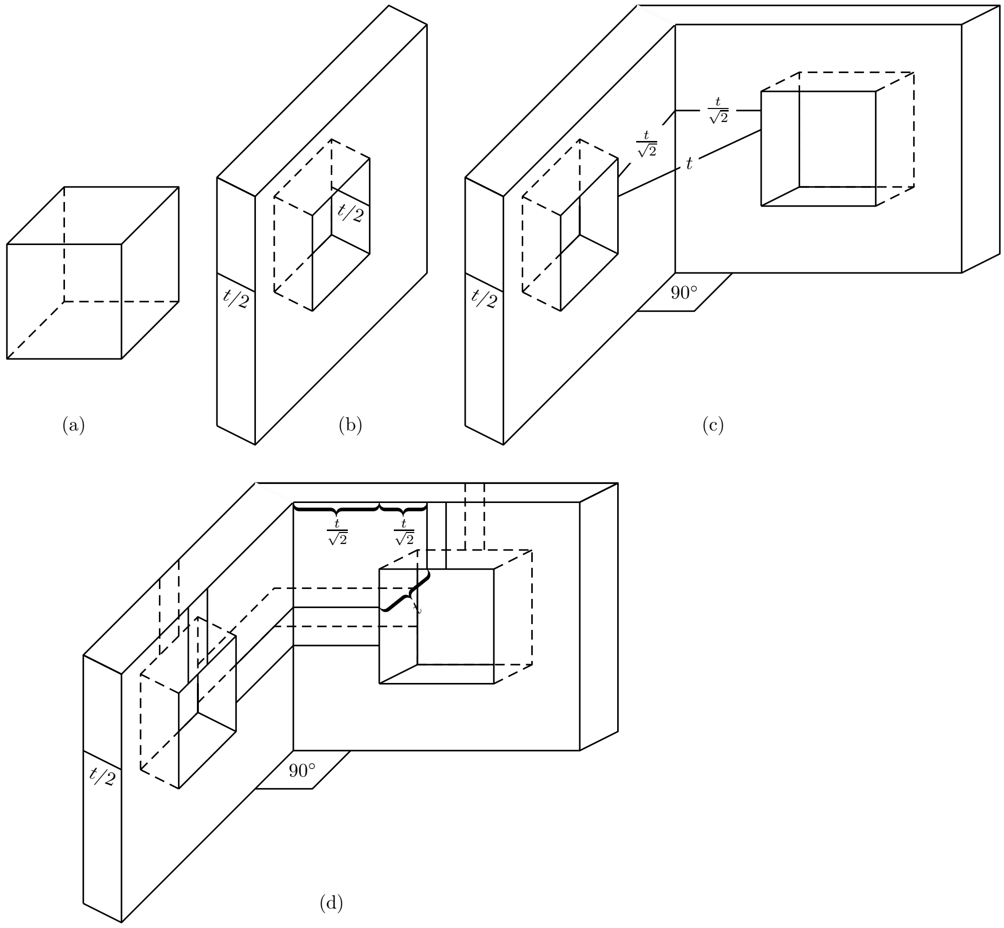

Figure 5 The dimension reducing tiling. Part (a) shows the basic cubes we start with. Part (b) shows a fattened wall of width $t/2$, and its interior that gets the color 2. Part (c) shows where the minimal distance between two tiles colored 2 is attained. Part (d) presents (in full lines) the interiors of the fattened edges that get the color 3 and shows where the minimal distance between two such tiles is attained.

References

1 Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum Key Exchange – A New Hope. In *25th USENIX Security Symposium, USENIX Security 16, Austin, TX, USA, August 10-12, 2016*, pages 327–343, 2016. URL: `https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/alkim`.

2 Michel L. Balinksi. How should data be rounded? In *Lecture Notes – Monograph Series, Vol. 28: Distributions with Fixed Marginals and Related Topics*, pages 33–44. Institute of Mathematical Statistics, 1996.

3 R. B. Bapat. Sperner's lemma with multiple labels. In *Modeling, Computation and Optimization, Chapter 16*, pages 257–262. World Scientific, 2009.

4 Raoul Bott and Loring W. Tu. *Differential Forms in Algebraic Topology.* Springer, 1982.

5 Herbert Busemann. *Convex surfaces.* Interscience publishers, 1958.

6 Henry Cohn, Abhinav Kumar, Stephen D. Miller, Danylo Radchenko, and Maryna Viazovska. The sphere packing problem in dimension 24. *Ann. of Math.*, 185:1017–1033, 2017.

7 Henry Cohn and Yufei Zhao. Sphere packing bounds via spherical codes. *Duke Mathematical Journal*, 163(10):1965–2002, 2014.

8 John H. Conway and Neil J. A. Sloane. What are all the best sphere packings in low dimensions? *Discrete Comput. Geom.*, 13:383–403, 1995.

9 John H. Conway and Neil J. A. Sloane. *Sphere packing, lattices and groups (3rd edition)*. Springer, 2013.

10 Lawrence Cox and Lawrence Ernst. Controlled rounding. *INFOR: Information Systems and Operational Research*, 20(4):423–432, 1982.

11 Mark de Berg, Otfried Cheong, Marc van Kreveld, and Mark Overmars. *Computational Geometry (3rd edition)*. Springer, 2008.

12 Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam D. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM J. Comput.*, 38(1):97–139, 2008. doi:10.1137/060651380.

13 Orr Dunkelman, Zeev Geyzel, Chaya Keller, Nathan Keller, Eyal Ronen, Adi Shamir, and Ran J. Tessler. Error resilient space partitioning (full version). *CoRR*, 2020. arXiv:2008.03675.

14 Benjamin Fuller, Leonid Reyzin, and Adam D. Smith. When are fuzzy extractors possible? *IEEE Trans. Inf. Theory*, 66(8):5282–5298, 2020. doi:10.1109/TIT.2020.2984751.

15 Jacob E. Goodman, Joseph O'Rourke, and Csaba D. Tóth. *Computational Geometry (3rd edition)*. CRC Press, 2017.

16 Robert M. Gray. Vector Quantization. *IEEE ASSP*, 1(2):4–29, 1984.

17 Thomas Hales, Mark Adams, and et al. A formal proof of the kepler conjecture. *Forum of Mathematics, Pi*, 5:e2, 2017.

18 Thomas C. Hales. A proof of the Kepler conjecture. *Ann. of Math.*, 162:1065–1185, 2005.

19 Piotr Indyk, Rajeev Motwani, Prabhakar Raghavan, and Santosh S. Vempala. Locality-preserving hashing in multidimensional spaces. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, El Paso, Texas, USA, May 4-6, 1997*, pages 618–625, 1997. doi:10.1145/258533.258656.

20 Ari Juels and Martin Wattenberg. A fuzzy commitment scheme. In *CCS '99, Proceedings of the 6th ACM Conference on Computer and Communications Security, Singapore, November 1-4, 1999*, pages 28–36, 1999. doi:10.1145/319709.319714.

21 G. A. Kabatyanskii and V. I. Levenshtein. Bounds for packings on a sphere and in space (Russian). *Problemy Peredaci Informacii*, 14:3–25, 1978. English translation in *Prob. Inform. Transmission* **14** (1978), pp. 1—17.

22 Guy Kindler, Ryan O'Donnell, Anup Rao, and Avi Wigderson. Spherical cubes and rounding in high dimensions. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA*, pages 189–198, 2008. doi:10.1109/FOCS.2008.50.

23 Guy Kindler, Anup Rao, Ryan O'Donnell, and Avi Wigderson. Spherical cubes: optimal foams from computational hardness amplification. *Commun. ACM*, 55(10):90–97, 2012. doi:10.1145/2347736.2347757.

24 E. Sperner. Neuer beweis für die invarianz der dimensionszahl und des gebietes (German). *Abh. Math. Seminar Univ. Hamburg*, 6:265–272, 1928.

25 Maryna S. Viazovska. The sphere packing problem in dimension 8. *Ann. of Math.*, 185:991–1015, 2017.

26 Leon Willenborg and Ton de Waal. *Elements of Statistical Disclosure Control*, volume 155 of *Lecture Notes in Statistics*. Springer, 2001.

27 Günter M. Ziegler. *Lectures on Polytopes*. Graduate Texts in Mathematics. Springer, 1995.