

On Metric Ramsey-Type Phenomena

Yair Bartal Nathan Linial Manor Mendel Assaf Naor

Abstract

The main question studied in this article may be viewed as a non-linear analogue of Dvoretzky's Theorem in Banach space theory or as part of Ramsey Theory in combinatorics. Given a finite metric space on n points, we seek its subspace of largest cardinality which can be embedded with a given distortion in Hilbert space. We provide nearly tight upper and lower bounds on the cardinality of this subspace in terms of n and the desired distortion. Our main theorem states that for any $\epsilon > 0$, every n point metric space contains a subset of size at least $n^{1-\epsilon}$ which is embeddable in Hilbert space with $O\left(\frac{\log(1/\epsilon)}{\epsilon}\right)$ distortion. The bound on the distortion is tight up to the $\log(1/\epsilon)$ factor. We further include a comprehensive study of various other aspects of this problem.

Contents

1	Introduction	1
1.1	Results for arbitrary metric spaces	2
1.2	Results for special classes of metric spaces	4
2	Metric Composition	5
2.1	The basic definitions	5
2.2	Generic upper bounds via metric composition	6
3	Metric Ramsey-type Theorems	12
3.1	Ultrametrics and Hierarchically Well-Separated Trees	12
3.2	An overview of the proof of Theorem 1.3	14
3.3	The weighted metric Ramsey problem and its relation with metric composition . . .	19
3.4	Exploiting metrics with bounded aspect ratio	20
3.5	Passing from an ultrametric to a k -HST	27
3.6	Passing from a k -HST to metric composition	29
3.7	Distortions arbitrarily close to 2	30
4	Dimensionality Based Upper Bounds	33
5	Expanders and Poincaré Inequalities	36
6	Markov Type, Girth and Hypercubes	43
6.1	Graphs with large girth	43
6.2	The discrete cube	48

1 Introduction

The philosophy of modern Ramsey Theory states that large systems necessarily contain large, highly structured sub-systems. The classical Ramsey coloring theorem [49, 29] is a prime example of this principle: Here “large” refers to the cardinality of a set, and “highly structured” means being monochromatic.

Another classical theorem, which can be viewed as a Ramsey-type phenomenon, is Dvoretzky’s theorem on almost spherical sections of convex bodies. This theorem, a cornerstone of modern Banach space theory and convex geometry, states that for all $\epsilon > 0$, every n -dimensional normed space X contains a k -dimensional subspace Y with $d(Y, \ell_2^k) \leq 1 + \epsilon$, where $k \geq c(\epsilon) \log n$. Here $d(\cdot, \cdot)$ is the Banach-Mazur distance, which is defined for two isomorphic normed spaces Z_1, Z_2 as:

$$d(Z_1, Z_2) = \inf\{\|T\| \cdot \|T^{-1}\|; T \in GL(Z_1, Z_2)\}.$$

Dvoretzky’s theorem is indeed a Ramsey-type theorem, in which “large” is interpreted as high-dimensional, and “highly structured” means close to Euclidean space in the Banach-Mazur distance.

Dvoretzky’s theorem was proved in [24], and the estimate $k \geq c(\epsilon) \log n$, which is optimal as a function of n , is due to Milman [45]. The dimension of almost spherical sections of convex bodies has been studied in depth by Figiel, Lindenstrauss and Milman in [27], where it was shown that under some additional geometric assumptions, the logarithmic lower bound for $\dim(Y)$ in Dvoretzky’s theorem can be improved significantly. We refer to the books [46, 48] for good expositions of Dvoretzky’s theorem, and to [47, 44] for an “isomorphic” version of Dvoretzky’s theorem.

The purpose of this paper is to study non-linear versions of Dvoretzky’s theorem, or viewed from the combinatorial perspective, metric Ramsey-type problems. In spite of the similarity of these problems, the results in the metric setting differ markedly from those for the linear setting.

Finite metric spaces and their embeddings in other metric spaces have been intensively investigated in recent years. See for example the surveys [30, 36], and the book [42] for an exposition of some of the results.

Let $f : X \rightarrow Y$ be an embedding of the metric spaces (X, d_X) into (Y, d_Y) . We define the *distortion* of f by

$$\text{dist}(f) = \sup_{\substack{x, y \in X \\ x \neq y}} \frac{d_Y(f(x), f(y))}{d_X(x, y)} \cdot \sup_{\substack{x, y \in X \\ x \neq y}} \frac{d_X(x, y)}{d_Y(f(x), f(y))}.$$

We denote by $c_Y(X)$ the least distortion with which X may be embedded in Y . When $c_Y(X) \leq \alpha$ we say that X α -embeds into Y and denote $X \xrightarrow{\alpha} Y$. When there is a bijection f between two metric spaces X and Y with $\text{dist}(f) \leq \alpha$ we say that X and Y are α -equivalent. For a class of metric spaces \mathcal{M} , $c_{\mathcal{M}}(X)$ is the minimum α such that X α -embeds into some metric space in \mathcal{M} . For $p \geq 1$ we denote $c_{\ell_p}(X)$ by $c_p(X)$. The parameter $c_2(X)$ is known as the *Euclidean distortion* of X . A fundamental result of Bourgain [15] states that $c_2(X) = O(\log n)$ for every n -point metric space (X, d) .

A metric Ramsey-type theorem states that a given metric space contains a large subspace which can be embedded with small distortion in some “well-structured” family of metric spaces (e.g., Euclidean). This can be formulated using the following notion:

Definition 1.1 (Metric Ramsey functions). Let \mathcal{M} be some class of metric spaces. For a metric space X , and $\alpha \geq 1$, $R_{\mathcal{M}}(X; \alpha)$ denotes the largest size of a subspace Y of X such that $c_{\mathcal{M}}(Y) \leq \alpha$.

Denote by $R_{\mathcal{M}}(\alpha, n)$ the largest integer m such that any n -point metric space has a subspace of size m that α -embeds into a member of \mathcal{M} . In other words, it is the infimum over X , $|X| = n$, of $R_{\mathcal{M}}(X; \alpha)$.

It is also useful to have the following conventions: For $\alpha = 1$ we allow omitting α from the notation. When $\mathcal{M} = \{X\}$, we write X instead of \mathcal{M} . Moreover when $\mathcal{M} = \{\ell_p\}$, we use R_p rather than R_{ℓ_p} .

In the most general form, let \mathcal{N} be a class of metric spaces and denote by $R_{\mathcal{M}}(\mathcal{N}; \alpha, n)$ the largest integer m such that any n -point metric space in \mathcal{N} has a subspace of size m that α -embeds into a member of \mathcal{M} . In other words, it is the infimum over $X \in \mathcal{N}$, $|X| = n$, of $R_{\mathcal{M}}(X; \alpha)$.

1.1 Results for arbitrary metric spaces

This paper provides several results concerning metric Ramsey Functions. One of our main objectives is to provide bounds on the *Euclidean Ramsey Function*, $R_2(\alpha, n)$.

The first result on this problem, well-known as a non-linear version of Dvoretzky's theorem, is due to Bourgain, Figiel and Milman [17] :

Theorem 1.2 ([17]). *For any $\alpha > 1$ there exists $C(\alpha) > 0$ such that $R_2(\alpha, n) \geq C(\alpha) \log n$. Furthermore, there exists $\alpha_0 > 1$ such that $R_2(\alpha_0, n) = O(\log n)$.*

While Theorem 1.2 provides a tight characterization of $R_2(\alpha, n) = \Theta(\log n)$ for values of $\alpha \leq \alpha_0$ (close to 1), this bound turns out to be very far from the truth for larger values of α (in fact, a careful analysis of the arguments in [17] gives $\alpha_0 \approx 1.023$, but as we later discuss, this is not the right threshold).

Motivated by problems in the field of Computer Science, more researchers [32, 14, 5] have investigated metric Ramsey problems. A close look (see [5]) at the results of [32, 14] as well as [17] reveals that all of these can be viewed as based on Ramsey-type theorems where the target class is the class of ultrametrics (see Section 3.1 for the definition).

The usefulness of such results for embeddings in ℓ_2 stems from the well-known fact [34] that ultrametrics are isometrically embeddable in ℓ_2 . Thus, denoting the class of ultrametrics by UM, we have that $R_2(\alpha, n) \geq R_{\text{UM}}(\alpha, n)$.

The recent result of Bartal, Bollobás and Mendel [5] shows that for large distortions the metric Ramsey function behaves quite differently than the behavior expressed by Theorem 1.2. Specifically, they prove that $R_2(\alpha, n) \geq R_{\text{UM}}(\alpha, n) \geq \exp((\log n)^{1-O(1/\alpha)})$ (in fact, it was already implicit in [14] that a similar bound holds for a particular α). The main theorem in this paper is:

Theorem 1.3 (Metric Ramsey-type Theorem). *For every $\varepsilon > 0$, any n -point metric space has a subset of size $n^{1-\varepsilon}$ which embeds in Hilbert space with distortion $O\left(\frac{\log(1/\varepsilon)}{\varepsilon}\right)$. Stated in terms of the metric Ramsey function, there exists an absolute constant $C > 0$ such that for every $\alpha > 1$ and every integer n :*

$$R_2(\alpha, n) \geq R_{\text{UM}}(\alpha, n) \geq n^{1-C\frac{\log(2\alpha)}{\alpha}}.$$

We remark that the lower bound above for $R_{\text{UM}}(\alpha, n)$ is meaningful only for large enough α . Small distortions are dealt with in Theorem 1.6 (see also Theorem 3.26).

The fact that the subspaces obtained in this Ramsey-type theorem are ultrametrics is not just an artifact of our proof. More substantially, it is a reflection of new embedding techniques that we introduce. Indeed, most of the previous results on embedding into ℓ_p have used what may be called Fréchet-type embeddings: forming coordinates by taking the distance from a fixed subset

of the points. This is the way an arbitrary finite metric space is embedded in ℓ_∞ (attributed to Fréchet). Bourgain’s embedding [15] and its generalizations [41] also fall in this category of embeddings. However, it is possible to show that Fréchet-type embeddings are not useful in the context of metric Ramsey-type problems. More specifically, we show in [6] that such embeddings cannot achieve bounds similar to those of Theorem 1.3.

Ultrametrics have a useful representation as *hierarchically well-separated trees* (HSTs). A k -HST is an ultrametric where vertices in the rooted tree are labelled by real numbers. The labels decrease by a factor $\geq k$ as you go down the levels away from the root. The distance between two leaves is the label of their lowest common ancestor. These *decomposable* metrics were introduced by Bartal [3]. Subsequently, it was shown (see [3, 4, 28]) that any n -point metric can be $O(\log n)$ -probabilistically embedded¹ in ultrametrics. This theorem has found many unexpected algorithmic applications in recent years, mostly in providing computationally efficient approximate solutions for several NP-hard problems (see the survey [30] for more details).

The basic idea in the proof of Theorem 1.3 is to iteratively find large subspaces that are hierarchically structured, gradually improving the distortion between these subspaces and a hierarchically well-separated tree. These hierarchical structures are naturally modelled via a notion (which is a generalization of the notion of k -HST) we call *metric composition* closure. Given a class of metric spaces \mathcal{M} , a metric space in the class $\text{comp}_k(\mathcal{M})$ is obtained by taking a metric space $M \in \mathcal{M}$ and replacing its points with copies of metric spaces from $\text{comp}_k(\mathcal{M})$ dilated so that there is a factor k gap between distances in M and distances within these copies.

Metric compositions are also used to obtain the following bounds on the metric Ramsey function in its more general form:

Theorem 1.4 (Generic bounds on the metric Ramsey function). *Let \mathcal{C} be a proper class of finite metric spaces that is closed under: (i) Isometry, (ii) Passing to a subspace (iii) Dilation. Then there exists $\delta < 1$ such that $R_{\mathcal{C}}(n) \leq n^\delta$ for infinitely many values of n .*

In particular we can apply Theorem 1.4 to the class $\mathcal{C} = \{X; c_{\mathcal{M}}(X) \leq \alpha\}$ where \mathcal{M} is some class of metric spaces. If there exists a metric space Y with $c_{\mathcal{M}}(Y) > \alpha$, then we get that there exists $\delta < 1$ such that $R_{\mathcal{M}}(\alpha, n) < n^\delta$ for infinitely many n ’s.

In the case of ℓ_2 or ultrametrics much better bounds are possible, showing that the bound in Theorem 1.3 is almost tight. For ultrametrics this is a rather simple fact [5]. For embedding into ℓ_2 this follows from bounds for expander graphs, described later in more detail.

Theorem 1.5 (Near Tightness). *There exist absolute constants $c, C > 0$ such that for every $\alpha > 2$ and every integer n :*

$$R_{\text{UM}}(\alpha, n) \leq R_2(\alpha, n) \leq Cn^{1-\frac{c}{\alpha}}.$$

The behavior of $R_{\text{UM}}(\alpha, n)$ and $R_2(\alpha, n)$ exhibited by the bounds in Theorems 1.2 and 1.3 is very different. Somewhat surprisingly, we discover the following phase transition:

Theorem 1.6 (Phase transition). *For every $\alpha > 1$ there exist constants $c, C, c', C', K > 0$ depending only on α such that $0 < c' < C' < 1$ and for every integer n :*

a) *If $1 < \alpha < 2$ then $c \log n \leq R_{\text{UM}}(\alpha, n) \leq R_2(\alpha, n) \leq 2 \log_2 n + C$.*

b) *If $\alpha > 2$ then $n^{c'} \leq R_{\text{UM}}(\alpha, n) \leq R_2(\alpha, n) \leq K n^{C'}$.*

¹A metric space can be α -probabilistically embedded in a class of metric spaces if it is α -equivalent to a convex combination of metric spaces in the class, via a non-contractive Lipschitz embedding [4].

Using bounds on the dimension with which any n point ultrametric is embeddable with constant distortion in ℓ_p [7] we obtain the following corollary:

Corollary 1.7 (Ramsey-type Theorems with Low Dimension). *There exists $0 < C(\alpha) < 1$ such that for every $p \geq 1$, $\alpha > 2$, and every integer n ,*

$$R_{\ell_p^d}(\alpha, n) \geq n^{C(\alpha)},$$

where $C(\alpha) \geq 1 - \frac{c \log \alpha}{\alpha}$, $d = \lceil \lceil \frac{c'}{(\alpha-2)^2} \rceil C(\alpha) \log n \rceil$, and $c, c' > 0$ are universal constants.

This result is meaningful since, although ℓ_2 isometrically embeds into L_p for every $1 \leq p \leq \infty$, there is no known ℓ_p analogue of the Johnson-Lindenstrauss dimension reduction lemma [31] (in fact, the Johnson-Lindenstrauss lemma is known to fail in ℓ_1 [19, 33]). These bounds are almost best possible.

Theorem 1.8 (The Ramsey Problem for Finite Dimensional Normed Spaces). *There exist absolute constants $C, c > 0$ such that for any $\alpha > 2$, every integer n and every finite dimensional normed space X ,*

$$R_X(\alpha, n) \leq Cn^{1-\frac{c}{\alpha}}(\dim X) \log \alpha.$$

For completeness, we comment that a natural question, in our context, is to bound the size of largest subspace of an arbitrary finite metric space that is isometrically embedded in ℓ_p . In [8] we show that $R_p(n) = 3$ for every $1 < p < \infty$ and $n \geq 3$.

Finally, we note that one important motivation for this work is the applicability of metric embeddings to the theory of algorithms. In many practical situations, one encounters a large body of data, the successful analysis of which depends on the way it is represented. If, for example, the data has a natural metric structure (such as in the case of distances in graphs), a low distortion embedding into some normed space helps us draw on geometric intuition in order to analyze it efficiently. We refer to the papers [4, 26, 37] and the surveys [30, 36] for some of the applications of metric embeddings in Computer Science. More about the relevance of Theorem 1.3 to Computer Science can be found in [9] (see also [5, 10]).

1.2 Results for special classes of metric spaces

We provide nearly tight bounds for concrete families of metric spaces: expander graphs, the discrete cube, and high girth graphs. In all cases the difficulty is in providing upper bounds on the Euclidean Ramsey function.

Let $G = (V, E)$ be a d -regular graph, $d \geq 3$, with absolute multiplicative spectral gap γ (i.e. the second largest eigenvalue, in absolute value, of the adjacency matrix of G is less than γd). For such expander graphs it is known [37, 41] that $c_2(G) = \Omega_{\gamma, d}(\log |V|)$ (here, and in what follows, the notation $a_n = \Omega(b_n)$ means that there exists a constant $c > 0$ such that for all n $|a_n| \geq c|b_n|$). When c is allowed to depend on, say, γ and d we use the notation $\Omega_{\gamma, d}$. In Section 5 we prove the following

Theorem 1.9 (Metric Ramsey problem for expanders). *Let $G = (V, E)$ be a d -regular graph, $d \geq 3$ with absolute multiplicative spectral gap $\gamma < 1$. Then for every $p \in [1, \infty)$, and every $\alpha \geq 1$,*

$$|V|^{1-\frac{C}{\alpha \log_d(1/\gamma)}} \leq R_2(G; \alpha) \leq R_p(G; \alpha) \leq Cd|V|^{1-\frac{c \log_d(1/\gamma)}{p\alpha}},$$

where $C, c > 0$ are absolute constants.

The proof of the upper bound in Theorem 1.9 involves proving certain Poincaré inequalities for power graphs of G .

Let $\Omega_d = \{0, 1\}^d$ be the discrete cube equipped with the Hamming metric. It was proved by Enflo, [25], that $c_2(\Omega_d) = \sqrt{d}$. Both Enflo's argument, and subsequent work of Bourgain, Milman and Wolfson [18], rely on non-linear notions of type. These proofs strongly use the structure of the whole cube, and therefore seem not applicable for subsets of the cube. In Section 6.2 we prove the following strengthening of Enflo's bound:

Theorem 1.10 (Metric Ramsey problem for the discrete cube). *There exist absolute constants C, c such that for every $\alpha > 1$:*

$$2^{d\left(1 - \frac{\log(C\alpha)}{\alpha^2}\right)} \leq R_2(\Omega_d; \alpha) \leq C2^{d\left(1 - \frac{c}{\alpha^2}\right)}.$$

The lower bounds on the Euclidean Ramsey function mentioned above are based on the existence of large subsets of the graphs which are within distortion α from forming an equilateral space. In particular for the discrete cube this corresponds to a code of large relative distance. Essentially, our upper bounds on the Euclidean Ramsey function show that for a fixed size, no other subset achieves significantly better distortions.

In [38] it was proved that if $G = (V, E)$ is a d -regular graph, $d \geq 3$, with girth g , then $c_2(G) \geq c^{\frac{d-2}{d}}\sqrt{g}$. In Section 6.1 we prove the following strengthening of this result:

Theorem 1.11 (Metric Ramsey problem for large girth graphs). *Let $G = (V, E)$ be a d -regular graph, $d \geq 3$, with girth g . Then for every $1 \leq \alpha < \frac{\sqrt{g}}{6}$,*

$$R_2(G; \alpha) \leq C(d-1)^{-\frac{cg}{\alpha^2}}|V|,$$

where $C, c > 0$ are absolute constants.

The proofs of Theorem 1.10 and Theorem 1.11 use the notion of Markov type, due to K. Ball [2]. In addition, we need to understand the algebraic properties of the graphs involved (Krawtchouk polynomials for the discrete cube and Geronimus polynomials in the case of graphs with large girth).

2 Metric Composition

In this section we introduce the notion of metric composition, which plays a basic role in proving both upper and lower bounds on the metric Ramsey problem. Here we introduce this construction and use it to derive some non-trivial upper bounds. The bounds achievable by this method are generally not tight. For the Ramsey problem on ℓ_p , better upper bounds are given in Sections 4 and 5. In Section 3 we use metric composition in the derivation of lower bounds.

2.1 The basic definitions

Definition 2.1 (Metric Composition). Let M be a finite metric space. Suppose that there is a collection of disjoint finite metric spaces N_x associated with the elements x of M . Let $\mathcal{N} = \{N_x\}_{x \in M}$. For $\beta \geq 1/2$, the β -composition of M and \mathcal{N} , denoted by $C = M_\beta[\mathcal{N}]$, is a metric space on the disjoint union $\dot{\cup}_x N_x$. Distances in C are defined as follows. Let $x, y \in M$ and $u \in N_x, v \in N_y$, then:

$$d_C(u, v) = \begin{cases} d_{N_x}(u, v) & x = y \\ \beta d_M(x, y) & x \neq y. \end{cases}$$

where $\gamma = \frac{\max_{z \in M} \text{diam}(N_z)}{\min_{x \neq y \in M} d_M(x, y)}$. It is easily checked that the choice of the factor $\beta\gamma$ guarantees that d_C is indeed a metric. If all the spaces N_x over $x \in M$ are isometric copies of the same space N , we use the simplified notation $C = M_\beta[N]$.

Informally stated, a metric composition is created by first multiplying the distances in M by $\beta\gamma$, and then replacing each point x of M by an isometric copy of N_x .

A related notion is the following:

Definition 2.2 (Composition Closure). Given a class \mathcal{M} of finite metric spaces, we consider $\text{comp}_\beta(\mathcal{M})$, its closure under $\geq \beta$ -compositions. Namely, this is the smallest class \mathcal{C} of metric spaces that contains all spaces in \mathcal{M} , and satisfies the following condition: Let $M \in \mathcal{M}$, and associate with every $x \in M$ a metric space N_x that is isometric to a space in \mathcal{C} . Also, let $\beta' \geq \beta$. Then $M_{\beta'}[\mathcal{N}]$ is also in \mathcal{C} .

2.2 Generic upper bounds via metric composition

We need one more definition:

Definition 2.3. A class \mathcal{C} of finite metric spaces is called a *metric class* if it is closed under isometries. \mathcal{C} is said to be *hereditary*, if $M \in \mathcal{C}$ and $N \subset M$ imply $N \in \mathcal{C}$. The class is said to be *dilation invariant* if $(M, d) \in \mathcal{C}$ implies that $(M, \lambda d) \in \mathcal{C}$ for every $\lambda > 0$.

Let $\mathcal{M}^\alpha = \{X; c_{\mathcal{M}}(X) \leq \alpha\}$, the class of all metric spaces that α -embed into some metric space in \mathcal{M} . Clearly, \mathcal{M}^α , is a hereditary, dilation invariant metric class.

We recall that $R_{\mathcal{C}}(X)$ is the largest cardinality of a subspace of X that is isometric to some metric space in the class \mathcal{C} .

Proposition 2.4. *Let \mathcal{C} be a hereditary, dilation invariant metric class of finite metric spaces. Then, for every finite metric space M and a class $\mathcal{N} = \{N_x\}_{x \in M}$, and every $\beta \geq 1/2$,*

$$R_{\mathcal{C}}(M_\beta[\mathcal{N}]) \leq R_{\mathcal{C}}(M) \cdot \max_{x \in M} R_{\mathcal{C}}(N_x).$$

In particular, for every finite metric space N ,

$$R_{\mathcal{C}}(M_\beta[N]) \leq R_{\mathcal{C}}(M)R_{\mathcal{C}}(N).$$

Proof. Let $m = R_{\mathcal{C}}(M)$ and $k = \max_{x \in M} R_{\mathcal{C}}(N_x)$. Fix any $X \subseteq \dot{\cup}_x N_x$ with $|X| > mk$. For every $z \in M$ let $X_z = X \cap N_z$. Set $Z = \{z \in M; X_z \neq \emptyset\}$. Note that $|X| = \sum_{z \in Z} |X_z|$ so that if $|Z| \leq m$ then there is some $y \in M$ with $|X_y| > k$. In this case, the set X_y consists of more than k elements in X , the metric on which is isometric to a subspace of N_y , and therefore is not in \mathcal{C} . Since \mathcal{C} is hereditary this implies that $X \notin \mathcal{C}$. Otherwise, $|Z| > m$. Fix for each $z \in Z$ some arbitrary point $u_z \in X_z$ and set $Z' = \{u_z; z \in Z\}$. Now, Z' consists of more than m elements in X , the metric on which is a $\beta\gamma$ -dilation of a subspace of M , hence not in \mathcal{C} . Again, the fact that \mathcal{C} is hereditary implies that $X \notin \mathcal{C}$. \square

In what follows let $R_{\mathcal{C}}(\mathcal{A}, n) = R_{\mathcal{C}}(\mathcal{A}; 1, n)$. Recall that $R_{\mathcal{C}}(\mathcal{A}; 1, n) \geq t$ iff for every $X \in \mathcal{A}$ with $|X| = n$, there is a subspace of X with t elements that is isometric to some metric space in the class \mathcal{C} .

Lemma 2.5. *Let \mathcal{C} be a hereditary, dilation invariant metric class of finite metric spaces. Let \mathcal{A} be a class of metric spaces, and let $\delta \in (0, 1)$. If there exist an integer $m > 1$ such that $R_{\mathcal{C}}(\mathcal{A}, m) \leq m^\delta$, then for any $\beta \geq 1/2$, and infinitely many integers n :*

$$R_{\mathcal{C}}(\text{comp}_\beta(\mathcal{A}), n) \leq n^\delta.$$

Proof. Fix some $\beta \geq 1/2$. Let $A \in \mathcal{A}$ such that $|A| = m > 1$ and $R_{\mathcal{C}}(A) \leq m^\delta$. Define inductively a sequence of metric spaces in $\text{comp}_\beta(\mathcal{A})$ by: $A_1 = A$ and $A_{i+1} = A_\beta[A_i]$. Proposition 2.4 implies that $R_{\mathcal{C}}(A_{i+1}) \leq R_{\mathcal{C}}(A_i)R_{\mathcal{C}}(A) \leq R_{\mathcal{C}}(A_i)m^\delta$. It follows that $R_{\mathcal{C}}(A_i) \leq m^{i\delta} = |A_i|^\delta$. \square

Lemma 2.6. *Let \mathcal{C} be a non-empty hereditary, dilation invariant metric class of finite metric spaces. Let \mathcal{A} be a class of finite metric spaces, such that $R_{\mathcal{C}}(\mathcal{A}, m) < m$ for some integer m (i.e., there is some space $A \in \mathcal{A}$ with no isometric copy in \mathcal{C} .) Then there exists $\delta \in (0, 1)$, such that for any $\beta \geq 1/2$, and infinitely many integers n :*

$$R_{\mathcal{C}}(\text{comp}_\beta(\mathcal{A}), n) \leq n^\delta.$$

Proof. Let m be the least cardinality of a space $A \in \mathcal{A}$ of with no isometric copy in \mathcal{C} . Since \mathcal{C} is nonempty and hereditary, $m \geq 2$. Define δ by $m - 1 = m^\delta$. Now apply Lemma 2.5. \square

Lemma 2.6 can be applied to obtain non-trivial bounds on various metric Ramsey functions.

Corollary 2.7. *Let \mathcal{C} be a hereditary, dilation invariant metric class which contains some, but not all finite metric spaces. Then there exists a $\delta \in (0, 1)$, such that $R_{\mathcal{C}}(n) \leq n^\delta$ for infinitely many integers n .*

Proof. We use Lemma 2.6 with $\mathcal{A} = \text{comp}_\beta(\mathcal{A}) =$ the class of all metric spaces. \square

Let \mathcal{M} be a fixed class of metric spaces and $\alpha \geq 1$. The following corollary follows when we apply Corollary 2.7 with $\mathcal{C} = \mathcal{M}^{\overset{\alpha}{\circ}}$ defined above.

Corollary 2.8. *Let \mathcal{M} be a metric class of finite metric spaces and $\alpha \geq 1$. The following assertions are equivalent:*

- a) *There exists an integer n , such that $R_{\mathcal{M}}(\alpha, n) < n$.*
- b) *There exists $\delta \in (0, 1)$, such that $R_{\mathcal{M}}(\alpha, n) \leq n^\delta$ for infinitely many integers n .*

In order to state our next result recall that a normed space X is said to have cotype q if there is a positive constant C such that for every finite sequence $x_1, \dots, x_m \in X$,

$$\left(\mathbb{E} \left\| \sum_{i=1}^m \varepsilon_i x_i \right\|^2 \right)^{1/2} \geq C \left(\sum_{i=1}^m \|x_i\|^q \right)^{1/q},$$

where $\varepsilon_1, \dots, \varepsilon_m$ are i.i.d. ± 1 Bernoulli random variables. It is well known (see [46]) that for $2 \leq q < \infty$, ℓ_q has cotype q (and it doesn't have cotype q' for any $q' < q$).

Corollary 2.9. *Let X be a normed space. Then the following assertions are equivalent:*

- a) *X has finite cotype.*
- b) *For any $\alpha > 1$, there exists $\delta \in (0, 1)$ such that for infinitely many integers n , $R_X(\alpha, n) \leq n^\delta$.*
- c) *There exists $\alpha > 1$ and an integer n such that $R_X(\alpha, n) < n$.*

Proof. To prove the implication **a)** \implies **b)**, fix $\alpha > 1$. Now, since X has finite cotype, there is an integer h such that $d(\ell_\infty^h, Z) > \alpha$ for every h -dimensional subspace Z of X , where $d(\cdot, \cdot)$ is the Banach-Mazur distance. This implies that for some $\epsilon > 0$, an ϵ -net \mathcal{E} in the unit ball of ℓ_∞^h does not α -embed into X . This follows from a standard argument in non-linear Banach space theory. Indeed, a compactness argument would imply that otherwise B_∞^h , the unit ball of ℓ_∞^h , α -embeds into X . By Rademacher's theorem, (see for example [12]) such an embedding must be differentiable in an interior point of B_∞^h . The derivative, T , is a linear mapping which is easily seen to satisfy $\|T\| \cdot \|T^{-1}\| \leq \alpha$, so that $d(\ell_\infty^h, Z) \leq \alpha$ for the subspace $Z = T(\ell_\infty^h)$. Apply Corollary 2.8 with $\mathcal{M} = X$, and $n = |\mathcal{E}|$ to conclude that **b)** holds.

The implication **b)** \implies **c)** is obvious, so we turn to prove that **c)** \implies **a)**. Assume that X does not have finite co-type, and fix some $0 < \epsilon < \alpha - 1$. By the Maurey-Pisier theorem (see [43] or Theorem 14.1 in [23]), it follows that for every n , ℓ_∞^n $(\alpha - \epsilon)$ -embeds into X . Since ℓ_∞^n contains an isometric copy of every n -point metric space, we deduce that for each n , $R_X(\alpha, n) = n$, contrary to our assumption **c)**. \square

We now need the following variation on the theme of metric composition.

Definition 2.10. A family of metric spaces \mathcal{N} is called *nearly closed under composition*, if for every $\lambda > 1$, there exists some $\beta \geq 1/2$ such that $c_{\mathcal{N}}(X) \leq \lambda$ for every $X \in \text{comp}_\beta(\mathcal{N})$. In other words,

$$\text{comp}_\beta(\mathcal{N}) \subseteq \mathcal{N}^{\stackrel{\lambda}{\leftrightarrow}}.$$

We have the following variant of Corollary 2.8:

Lemma 2.11. *Let \mathcal{M} be a metric class of finite metric spaces and let \mathcal{N} be some class of finite metric spaces which is nearly closed under composition. Assume that there is some space in \mathcal{N} which does not α -embed into any space in \mathcal{M} . Then there exists $\delta \in (0, 1)$, such that for every $1 \leq \alpha' < \alpha$, $R_{\mathcal{M}}(\mathcal{N}; \alpha', n) \leq n^\delta$ for infinitely many integers n .*

Proof. Fix some $\alpha' < \alpha$ and let $\lambda = \alpha/\alpha'$. As \mathcal{N} is nearly closed under composition there exists $\beta \geq 1/2$ such that $\text{comp}_\beta(\mathcal{N}) \subseteq \mathcal{N}^{\stackrel{\lambda}{\leftrightarrow}}$. This means that for every $Z \in \text{comp}_\beta(\mathcal{N})$ there exists some $N \in \mathcal{N}$ that is λ -equivalent to Z .

For every integer p let $k(p) = R_{\mathcal{M}}(\mathcal{N}; \alpha', p)$. If $|Z| = |N| = n$, then there is $X \subseteq N$ such that $c_{\mathcal{M}}(X) \leq \alpha'$ and $|X| \geq k(n)$. Let $Y \subseteq Z$ be the set corresponding to X under the λ -equivalence between Z and N . Then, $|Y| = |X| \geq k(n)$ and by composition of maps, $c_{\mathcal{M}}(Y) \leq \lambda\alpha' = \alpha$. That is, every n -set Z in $\text{comp}_\beta(\mathcal{N})$ contains a $k(n)$ subset Y that α -embeds into a space in \mathcal{M} , i.e. $Y \in \mathcal{M}^{\stackrel{\alpha}{\hookrightarrow}}$. In our notation, this means that $k(n) \leq R_{\mathcal{C}}(\text{comp}_\beta(\mathcal{N}), n)$, where $\mathcal{C} = \mathcal{M}^{\stackrel{\alpha}{\hookrightarrow}}$.

The assumption made in the lemma about \mathcal{N} means that $R_{\mathcal{C}}(\mathcal{N}, m) < m$ for some integer m . By Lemma 2.6 there exists $\delta \in (0, 1)$ such that for infinitely many integers n ,

$$R_{\mathcal{M}}(\mathcal{N}; \alpha', n) = k(n) \leq R_{\mathcal{C}}(\text{comp}_\beta(\mathcal{N}), n) < n^\delta,$$

as claimed. \square

Next, we give a several results that demonstrate the applicability of Lemma 2.11.

Proposition 2.12. *Let $(X, \|\cdot\|)$ be a normed space. The class \mathcal{M} of finite subsets of X is nearly closed under composition.*

Proof. Fix some $\lambda > 1$. Let $Z \in \text{comp}_\beta(\mathcal{M})$ for some $\beta > 1/2$ to be determined later. We prove that Z can be λ -embedded in X . The proof is by induction on the number of steps taken in composing Z from spaces in \mathcal{M} . If $Z \in \mathcal{M}$ there is nothing to prove. Otherwise, it is possible to express Z as $Z = M_\beta[\mathcal{N}]$, where $M \in \mathcal{M}$ and $\mathcal{N} = \{N_z\}_{z \in M}$ such that each of the spaces N_z is in $\text{comp}_\beta(\mathcal{M})$ and can be created by a shorter sequence of composition steps. By induction we assume that there exists β for which N_z can be λ -embedded in X . Fix for every $z \in M$, $\phi_z : N_z \rightarrow X$ satisfying:

$$\forall u, v \in N_z \quad d_{N_z}(u, v) \leq \|\phi_z(u) - \phi_z(v)\| \leq \lambda d_{N_z}(u, v),$$

and for all $u \in N_z$, $\|\phi_z(u)\| \leq \lambda \text{diam}(N_z)$ (this can be assumed by an appropriate translation).

Define $\phi : Z \rightarrow X$ as follows: for every $u \in Z$, let $z \in M$ be such that $u \in N_z$, then $\phi(u) = \beta\gamma \cdot z + \phi_z(u)$, where $\gamma = \frac{\max_z \text{diam}(N_z)}{\min_{x \neq y \in M} \|x - y\|}$.

We now bound the distortion of ϕ . Assume $\beta > 2\lambda$. Consider first $u, v \in N_z$ for some $z \in M$.

$$d_Z(u, v) = d_{N_z}(u, v) \leq \|\phi_z(u) - \phi_z(v)\| \leq \lambda d_{N_z}(u, v) = \lambda d_Z(u, v).$$

Now, let $u \in N_x, v \in N_y$, for $x \neq y \in M$,

$$\begin{aligned} \|\phi(u) - \phi(v)\| &\leq \beta\gamma\|x - y\| + \|\phi_x(u) - \phi_y(v)\| \\ &\leq \beta\gamma\|x - y\| + \lambda(\text{diam}(N_x) + \text{diam}(N_y)) \\ &\leq \gamma(\beta + 2\lambda)\|x - y\| = \frac{\beta + 2\lambda}{\beta} d_Z(u, v), \end{aligned}$$

Similarly,

$$\begin{aligned} \|\phi(u) - \phi(v)\| &\geq \beta\gamma\|x - y\| - \|\phi_x(u) - \phi_y(v)\| \\ &\geq \beta\gamma\|x - y\| - \lambda(\text{diam}(N_x) + \text{diam}(N_y)) \\ &\leq \gamma(\beta - 2\lambda)\|x - y\| = \frac{\beta - 2\lambda}{\beta} d_Z(u, v). \end{aligned}$$

Hence if $\beta \geq 2\lambda \frac{\lambda+1}{\lambda-1}$, we have,

$$\text{dist}(\phi) \leq \max \left\{ \lambda, \frac{\beta + 2\lambda}{\beta - 2\lambda} \right\} = \lambda.$$

□

Recall that a normed space X is said to be λ finitely representable in a normed space Y if for any finite dimensional linear subspace $Z \subset X$ and every $\eta > 0$ there is a subspace W of Y such that $d(Z, W) \leq \lambda + \eta$.

Corollary 2.13. *Let X and Y be normed spaces and $\alpha > 1$. The following are equivalent:*

- 1) X is not α finitely representable in Y .
- 2) There are $\eta > 0$ and $\delta \in (0, 1)$ such that $R_Y(X; \alpha + \eta, n) < n^\delta$ for infinitely many integers n .
- 3) There is some $\eta > 0$ and an integer n such that $R_Y(X; \alpha + \eta, n) < n$.

Proof. If X is not α -finitely representable in Y then there is a finite dimensional linear subspace Z of X whose Banach-Mazur distance from any subspace of Y is greater than α . As in the proof of Corollary 2.9, a combination of a compactness argument and a differentiation argument imply that there is a finite subset S of X which does not $(\alpha + 2\eta)$ embed in Y for some $\eta > 0$. Since the subsets of X are nearly closed under composition, by applying Lemma 2.11, we deduce the implication **1)** \implies **2)**.

The implication **2)** \implies **3)** is obvious, so we turn to show **3)** \implies **1)**. Let $A \subset X$ be a finite subset that does not $\alpha + \eta$ embed in Y , and let Z be A 's linear span. Clearly $d(Z, W) > \alpha + \eta$ for any linear subspace W of Y . It follows that X is not α finitely representable in Y . \square

Recall that a graph H is called a *minor* of a graph G if H is obtained from G by a sequence of steps, each of which is either a contraction or a deletion of an edge. We say that a family \mathcal{F} of graphs is minor-closed if it is closed under taking minors. The Wagner conjecture famously proved by Robertson and Seymour [51], states that for any non-trivial minor-closed family of graphs \mathcal{F} , there is a *finite* set of graphs, \mathcal{H} , such that $G \in \mathcal{F}$ if and only if no member of \mathcal{H} is a minor of G . We say then that \mathcal{F} is characterized by the list \mathcal{H} of forbidden minors. For example, planar graphs are precisely the graphs which do not have $K_{3,3}$ or K_5 as minors, and the set of all trees is precisely the set of all connected graphs with no K_3 minor.

There is a graph-theoretic counterpart to composition. Namely, let $G = (V, E)$ be a graph, and suppose that to every vertex $x \in V$ corresponds a graph $H_x = (V_x, E_x)$ with a marked vertex $r_x \in V_x$, where the H_x are disjoint. The corresponding *graph composition*, denoted $G[\{H_x\}_{x \in V}]$, is a graph with vertex set $\dot{\cup}_{x \in V} V_x$, and edge set:

$$E = \{[u, v]; x \in V, [u, v] \in E_x\} \cup \{[r_x, r_y]; [x, y] \in E\}.$$

The composition closure of a family of graphs \mathcal{F} can be defined similarly to Definition 2.2, and family \mathcal{F} is said to be closed under composition if it equals its closure.

Recall that a connected graph G is called bi-connected if it stays connected after we delete any single vertex from G (and erase all the edges incident with it). The maximal bi-connected subgraphs of G are called its *blocks*.

We make the following elementary graph-theoretic observation:

Proposition 2.14. *Let H be a bi-connected graph with ≥ 3 vertices that is a minor of a graph G . Then H is a minor of a block of G .*

Proof. Consider a sequence of steps in which edges in G are being shrunk to form H . If there are two distinct blocks B_1, B_2 in G that are not shrunk to a single vertex, then the resulting graph is not bi-connected. Indeed, there is a cut-vertex a in G that separates B_1 from B_2 , and the vertex into which a is shrunk still separates the shrunk versions of B_1, B_2 . This observation means that in shrinking G to H , only a single block B of G retains more than one vertex. But then H is a minor of B , as claimed. \square

In the graph composition described above, each vertex $r_x \in V_x$ is a cut vertex. Consequently, each block of the composition is either a block of G (The subgraph induced by the vertices $\{r_x; x \in V\}$ is isomorphic with G) or of one of the H_x (that is isomorphic with the subgraph induced on V_x). We conclude:

Proposition 2.15. *Let \mathcal{F} be a minor-closed family of graphs characterized by a list of bi-connected forbidden minors. Then \mathcal{F} is closed under graph composition.*

Let \mathcal{F} again be a family of undirected graphs. A metric space M is said to be *supported* on \mathcal{F} if there exists a graph $G \in \mathcal{F}$ and positive weights on the edges of G such that M is the geodetic, or shortest path metric on a subset of the vertices of the weighted G .

Here is the metric counterpart of Proposition 2.15:

Proposition 2.16. *Let \mathcal{F} be a minor-closed family of graphs characterized by a list of bi-connected forbidden minors. Then the class of metrics supported on \mathcal{F} is nearly closed under composition.*

Proof. Fix some $\lambda > 1$. Let \mathcal{F}' be the class of metrics supported on \mathcal{F} . Let $X \in \text{comp}_\beta(\mathcal{F}')$ for some $\beta > 1/2$ to be determined later. We prove that X can be λ -embedded in \mathcal{F}' . The proof is by induction on the number of steps taken in composing X from spaces in \mathcal{F}' . If $X \in \mathcal{F}'$ there is nothing to prove.

Otherwise, there exists a weighted graph $G = (V, E, w)$ in \mathcal{F} . For simplicity, we identify G with a metric space in \mathcal{F}' , equipped with the geodetic metric defined by its weights. It is possible to express X as $X = G_\beta[\mathcal{H}']$, where $\mathcal{H}' = \{H'_z\}_{z \in V}$ such that each of the metric spaces H'_z is in $\text{comp}_\beta(\mathcal{F}')$. By induction we assume that there exists β for which each H'_z can be λ -embedded in \mathcal{F}' . Therefore there exists a family of disjoint weighted graphs $\{H_z = (V_z, E_z, w_z)\}_{z \in V}$, such for every $z \in V$, there is a non-contractive Lipschitz bijection, $\phi_z : H'_z \rightarrow V_z$, satisfying for any $u, v \in H'_z$, $d_{H'_z}(u, v) \leq d_{H_z}(\phi_z(u), \phi_z(v)) \leq \lambda d_{H'_z}(u, v)$.

Let $Y = G[\{H_z\}_{z \in V}]$ be the graph composition of the above graphs. Define weights w' on the edges of Y as follows: For any $z \in V$, $[u, v] \in E_z$, let $w'([u, v]) = w_z([u, v])$. For $[x, y] \in E$, let $w'([r_x, r_y]) = \beta\gamma w([x, y])$, where $\gamma = \frac{\max_{z \in V} \text{diam}(H'_z)}{\min_{x \neq y \in V} d_G(x, y)}$ (as in the definition of metric composition). For simplicity, we identify Y with the weighted graph defined above as well as the geodetic metric defined by this graph. The proof shows that if β is large enough, then the geodetic metric on the graph composition Y is λ -equivalent (and thus arbitrarily close) to the metric β -composition X . Proposition 2.15 implies that Y belongs to \mathcal{F}' , which proves the claim.

Indeed, define the bijection $\phi : X \rightarrow \dot{\cup}_{u \in V} V_u$ as follows: for $z \in V$, if $u \in H'_z$, then $\phi(u) = \phi_z(u)$. The geodetic path between any two vertices $u', v' \in V_z$ is exactly the same path as in H_z , since the cost of every step outside of V_z exceeds $\text{diam}(H_z)$ (by definition of γ). This implies that

$$d_X(u, v) = d_{H'_z}(u, v) \leq d_{H_z}(\phi_z(u), \phi_z(v)) = d_Y(\phi(u), \phi(v)) \leq \lambda d_{H'_z}(u, v) = \lambda d_X(u, v).$$

Also, the distance in the graph composition between $u' \in V_x$ and $v' \in V_y$ with $x \neq y \in V$, is at most $\beta\gamma d_G(x, y) + 2\lambda \max_z \text{diam}(H_z) \leq \gamma(\beta + 2\lambda)d_G(x, y)$. It follows that for $u \in H'_x$ and $v \in H'_y$,

$$d_X(u, v) = \beta\gamma d_G(x, y) \leq d_Y(\phi(u), \phi(v)) \leq \gamma(\beta + 2\lambda)d_G(x, y) = \left(\frac{\beta + 2\lambda}{\beta}\right) d_X(u, v).$$

Hence if $\beta \geq \frac{2\lambda}{\lambda-1}$, we have, $\text{dist}(\phi) \leq \max\left\{\lambda, \frac{\beta+2\lambda}{\beta}\right\} = \lambda$. □

Recall that a Banach space X is called superreflexive if it admits an equivalent uniformly convex norm. A finite-metric characterization of such spaces had been found by Bourgain [16]. Namely, X is superreflexive if and only if for every $\alpha > 0$ there is an integer h such that the complete binary tree of depth h doesn't α -embed into X . Let TREE denote the set of metrics supported on trees. Since any weighted tree is almost isometric to a subset of a deep enough complete binary tree we conclude using Lemma 2.11.

Corollary 2.17. *Let X be a Banach space. Then the following assertions are equivalent:*

a) X is superreflexive.

b) For any $\alpha > 1$ there exists $\delta < 1$ such that for infinitely many integers n ,

$$R_X(\text{TREE}; \alpha, n) \leq n^\delta.$$

c) For any $\alpha > 1$ there exists an integer n such that

$$R_X(\text{TREE}; \alpha, n) < n.$$

3 Metric Ramsey-type Theorems

In this section we prove Theorem 1.3, i.e. we give an $n^{\Omega(1)}$ lower bound on $R_2(\alpha, n)$ for $\alpha > 2$.

The proof actually establishes a lower bound on $R_{\text{UM}}(\alpha, n)$. The bound on R_2 follows since ultrametrics embed isometrically in ℓ_2 . The lower bound for embedding into ultrametrics utilizes their representation as hierarchically well-separated trees. We begin with some preliminary background on ultrametrics and hierarchically well-separated trees in Section 3.1. We also note that our proof of the lower bound makes substantial use of the notions of metric composition and composition closure which were introduced in Section 2.

We begin with a description of the lemmas on which the proof of the lower bound is based and the way they are put together to prove the main theorem. This is done in Section 3.2. Detailed proofs of the main lemmas appear in Sections 3.3–3.6. Most of the proof is devoted to the case where α is a fixed, large enough constant. In Section 3.7, we extend the proof to apply for every $\alpha > 2$.

3.1 Ultrametrics and Hierarchically Well-Separated Trees

Recall that an *ultrametric* is a metric space (X, d) such that for every $x, y, z \in X$,

$$d(x, z) \leq \max\{d(x, y), d(y, z)\}.$$

A more restricted class of metrics with an inherently hierarchical structure plays a key role in the sequel. Such spaces have already figured prominently in earlier work on embedding into ultrametric spaces [3, 5].

Definition 3.1. [3] For $k \geq 1$, a *k-hierarchically well-separated tree* (*k-HST*) is a metric space whose elements are the leaves of a rooted tree T . To each vertex $u \in T$ there is associated a label $\Delta(u) \geq 0$ such that $\Delta(u) = 0$ iff u is a leaf of T . It is required that if a vertex u is a child of a vertex v then $\Delta(u) \leq \Delta(v)/k$. The distance between two leaves $x, y \in T$ is defined as $\Delta(\text{lca}(x, y))$, where $\text{lca}(x, y)$ is the least common ancestor of x and y in T .

A *k-HST* is said to be *exact* if $\Delta(u) = \Delta(v)/k$ for every two internal vertices where u is a child of v .

First, note that an ultrametric on a finite set and a (finite) 1-HST are identical concepts. Any *k-HST* is also a 1-HST, i.e., an ultrametric. However, when $k > 1$ a *k-HST* is a stronger notion which has a hierarchically clustered structure. More precisely, a *k-HST* with diameter D decomposes into subspaces of diameter at most D/k and any two points at distinct subspaces are at distance exactly D . Recursively, each subspace is itself a *k-HST*. It is this hierarchical decomposition that makes *k-HSTs* useful.

When we discuss *k-HSTs*, we freely use the tree T as in Definition 3.1, *the tree defining the HST*. An internal vertex in T with out-degree 1 is said to be *degenerate*. If u is non-degenerate, then

$\Delta(u)$ is the diameter of the sub-space induced on the subtree rooted by u . Degenerate nodes do not influence the metric on T 's leaves, hence we may assume that all internal nodes are non-degenerate (note that this assumption need not hold for *exact* k -HSTs).

We need some more notation:

Notation 3.2. Let UM to denote the class of ultrametrics, and k -HST denotes the class of k -HSTs. Also let EQ denote the class of equilateral spaces.

The following simple observation is not required for the proof, but may help direct the reader's intuition. More complex connections between these concepts do play an important role in the proof.

Proposition 3.3. *The class of k -HSTs is the k -composition closure of the class of equilateral spaces; i.e., $k\text{-HST} = \text{comp}_k(\text{EQ})$.*

In particular, the class of ultrametrics is the 1-composition closure of the class of equilateral spaces; i.e., $\text{UM} = \text{comp}_1(\text{EQ})$.

We recall the following well known fact (e.g. [34]), that allows us to reduce the Euclidean Ramsey problem to the problem of embedding into ultrametrics:

Proposition 3.4. *Any ultrametric is isometrically embeddable in ℓ_2 . In particular,*

$$R_2(\alpha, n) \geq R_{\text{UM}}(\alpha, n).$$

This proposition can be proved by induction of the structure of the tree defining the ultrametric. It is shown inductively that each rooted subtree embeds isometrically into a sphere with radius proportional to the subtree's diameter, and that any two subtrees rooted at an internal vertex are mapped into orthogonal subspaces.

When considering Lipschitz embeddings, the k -HST representation of an ultrametric comes naturally into play. This is expressed by the following variant on a proposition from [4]:

Lemma 3.5. *For any $k > 1$, any ultrametric is k -equivalent to an exact k -HST.*

Lemma 3.5 is proved via a simple transformation of the tree defining the ultrametric. This is done by coalescing consecutive internal vertices, whose labels differ by a factor which is less than k . The complete proof of Lemma 3.5 appears in Section 3.5

We end this section with a proposition on embeddings into ultrametrics, which is implicit in [3]. Although this proposition is *not* used in the proofs, it is useful for obtaining efficient algorithms out of these theorems.

Lemma 3.6. *Every n -point metric space is n -equivalent to an ultrametric.*

Proof. Let M be an n -point metric space. We inductively construct an n -point HST X with $\text{diam}(X) = \text{diam}(M)$ and a non-contracting n -Lipschitz bijection between M and X .

Define a graph with vertex set M in which $[u, v]$ is an edge if and only if $d_M(u, v) < \frac{\text{diam}(M)}{n}$. Clearly, this graph is disconnected. Let A_1, \dots, A_m be the vertex sets of the connected components. By induction there are HSTs X_1, \dots, X_m with $\text{diam}(X_i) = \text{diam}((A_i, d_M)) < \text{diam}(M)$ and bijections $f_i : A_i \rightarrow X_i$ such that for every $u, v \in A_i$, $d_M(u, v) \leq d_{X_i}(f_i(u), f_i(v)) \leq |A_i|d_M(u, v) < nd_M(u, v)$. Let T_i be the tree defining X_i . We now construct the HST X whose defining labelled tree T is rooted at z . The root's label is $\Delta(z) = \text{diam}(M)$ and it has m children, where the i th child, u_i , is a root of a labelled tree isomorphic to T_i . Since $\Delta(u_i) = \text{diam}(X_i) < \text{diam}(M) = \text{diam}(X) = \Delta(z)$, the resulting tree T indeed defines an HST. Finally, if $u \in A_i$ and $v \in A_j$ for $i \neq j$ then $d_M(u, v) \geq \text{diam}(M)/n$. Since $\text{diam}(X) = \Delta(z) = \text{diam}(M)$, the inductive hypothesis implies the existence of the required bijection. \square

3.2 An overview of the proof of Theorem 1.3

In this section we describe the proof of the following theorem:

Theorem 3.7. *There exists an absolute constant $C > 0$ such that for every $\alpha > 2$,*

$$R_{\text{UM}}(\alpha, n) \geq n^{1-C\frac{\log \alpha}{\alpha}}.$$

By Proposition 3.4, the same bound holds true for $R_2(\alpha, n)$.

We begin with an informal description and motivation. The main lemmas needed for the proof are stated, and it is shown how they imply the theorem. Detailed proofs for most of these lemmas appear in subsequent subsections.

Our goal is to show that for any $\alpha > 2$, every n point metric space X contains a subspace which is α -equivalent to an ultrametric of cardinality $\geq n^{\psi(\alpha)}$, where $\psi(\alpha)$ is independent of n . In much of the proof we pursue an even more illusive goal. We seek large subsets that embed even into k -HSTs (Recall that this is a restricted class of ultrametrics). A conceptual advantage of this is that it directs us towards seeking *hierarchical* substructures within the given metric space. Such structures can be described as the composition closure of some class of metric spaces \mathcal{M} . A metric space in $\text{comp}_\beta(\mathcal{M})$ is composed of a hierarchy of dilated copies of metric spaces from \mathcal{M} , and the proof iteratively finds such large structures. The class \mathcal{M} varies from iteration to iteration, gradually becoming more restricted, and getting closer to the class EQ. When \mathcal{M} is approximately EQ this procedure amounts to finding a k -HST (due to Proposition 3.3). It is therefore worthwhile to consider a special case of the general problem, where $X \in \text{comp}_\beta(\mathcal{M})$, and we seek a subspace of X that is α -equivalent to a k -HST.

It stands to reason that if spaces in \mathcal{M} have large Ramsey numbers, then something similar should hold true also for spaces in $\text{comp}_\beta(\mathcal{M})$. After all, if β is large, then the copies of dilated metric spaces from \mathcal{M} are hierarchically well-separated. This would have reduced the problem of estimating Ramsey numbers for spaces in $\text{comp}_\beta(\mathcal{M})$ to the same problem for the much simpler class \mathcal{M} .

While this argument is not quite true, a slight modification of it does indeed work. For the purpose of this intuitive discussion, it is convenient to think of β being large, in particular with respect to k and α . Consider that X is the β -composition of $M \in \mathcal{M}$ and a set of $|M|$ disjoint metric spaces $\{N_i\}_{i \in M}$, $N_i \in \text{comp}_\beta(\mathcal{M})$. Assume (inductively) that each N_i contains a subspace N'_i that is α -equivalent to a k -HST H_i of size $|N_i|^\psi$. Find a subspace M' of M that is also α -equivalent to a k -HST K and attach the roots of the appropriate H_i 's to the corresponding leaves of K (with an appropriate dilation). This yields a k -HST H , and by the separation property of compositions with large β , we obtain a subspace X' of X which is α -equivalent to H . However, the size of the final subspace $X' = \dot{\cup}_{i \in M'} N'_i$ depends not only on the size of M' , the subspace we find in M , but also on how large the chosen N'_i 's are. Therefore, the correct requirement is that M' satisfies:

$$\sum_{i \in M'} |N_i|^\psi \geq \left(\sum_{i \in M} |N_i| \right)^\psi.$$

This gives rise to the following definition:

Definition 3.8 (Weighted Ramsey Function). Let \mathcal{M}, \mathcal{N} be classes of metric spaces. Denote by $\psi_{\mathcal{M}}(\mathcal{N}, \alpha)$ the largest $0 \leq \psi \leq 1$ such that for every metric space $X \in \mathcal{N}$ and any weight function $w : X \rightarrow \mathbb{R}^+$, there is a subspace Y of X that α -embeds in \mathcal{M} and satisfies:

$$\sum_{x \in Y} w(x)^\psi \geq \left(\sum_{x \in X} w(x) \right)^\psi. \quad (*)$$

When \mathcal{N} is the class of all metric spaces, it is omitted from the notation.

In what follows the notion of a weighted metric space refers to a pair (X, w) , where X is a metric space and $w : X \rightarrow \mathbb{R}^+$ is a weight function.

The following is an immediate consequence of Definition 3.8 (by considering the constant weight function $w(x) \equiv 1$).

Proposition 3.9.

$$R_{\mathcal{M}}(\mathcal{N}; \alpha, n) \geq n^{\psi_{\mathcal{M}}(\mathcal{N}, \alpha)}.$$

In particular,

$$R_{\mathcal{M}}(\alpha, n) \geq n^{\psi_{\mathcal{M}}(\alpha)}.$$

We note that it is possible to show, via the results of Section 2, that in the setting of embedding into composition classes, and in particular in our case of k -HSTs or ultrametrics, the last inequality in Proposition 3.9 holds with equality for infinitely many n 's.

The entire proof is thus dedicated to bounding the weighted Ramsey function when the target metric class is the class of ultrametrics. The proofs in the sequel produce embeddings into k -HSTs and ultrametrics. In this context, the following conventions for $\psi_{\mathcal{M}}(\mathcal{N}, \alpha)$ are useful:

- $\psi_k(\mathcal{N}, \alpha) = \psi_{k\text{-HST}}(\mathcal{N}, \alpha)$. In particular, $\psi_k(\alpha) = \psi_{k\text{-HST}}(\alpha)$.
- $\psi(\mathcal{N}, \alpha) = \psi_1(\mathcal{N}, \alpha) = \psi_{\text{UM}}(\mathcal{N}, \alpha)$. In particular, $\psi(\alpha) = \psi_{\text{UM}}(\alpha)$.

The following strengthening of Theorem 3.7 is the main result proved in this section.

Theorem 3.7'. There exists an absolute constant $C > 0$ such that for every $\alpha > 2$,

$$\psi(\alpha) \geq 1 - C \frac{\log \alpha}{\alpha}.$$

Our goal can now be rephrased as follows: given an arbitrary weighted metric space (X, w) , find a subspace of X , satisfying the weighted Ramsey condition (*) with $\psi(\alpha)$ as in Theorem 3.7', that is α -equivalent to an ultrametric.

Before continuing with the outline of the proof, we state a useful property of the weighted Ramsey function. When working with the regular Ramsey question it is natural to perform a procedure of the following form: first find a subspace which is α_1 -embedded in some ‘‘nice’’ class of metric spaces, then find a smaller subspace of this subspace which is α_2 equivalent to our target class of metric spaces, thus obtaining overall $\alpha_1 \alpha_2$ distortion. If the first subspace has size $n' \geq n^{\psi_1}$ and the second is of size $n'' \geq n'^{\psi_2}$ then $n'' \geq n^{\psi_1 \psi_2}$.

The weighted Ramsey problem has the same super-multiplicativity property:

Lemma 3.10. *Let $\mathcal{M}, \mathcal{N}, \mathcal{P}$ be classes of metric spaces and $\alpha_1, \alpha_2 \geq 1$. Then*

$$\psi_{\mathcal{M}}(\mathcal{P}, \alpha_1 \alpha_2) \geq \psi_{\mathcal{M}}(\mathcal{N}, \alpha_1) \cdot \psi_{\mathcal{N}}(\mathcal{P}, \alpha_2).$$

The interpretation of this lemma (proved in Section 3.3) is as follows: Suppose that we are given a metric space in \mathcal{P} and we seek a subspace that embeds with low distortion in \mathcal{M} , and satisfies condition (*). We can first find a subspace which α_1 -embeds in \mathcal{N} and then a subspace of which that α_2 -embeds in \mathcal{M} . In the course of this procedure we multiply the distortions and the ψ s of the corresponding classes.

The discussion in the paragraph preceding Definition 3.8 on how Ramsey-type properties of class \mathcal{M} carry over to $\text{comp}_\beta(\mathcal{M})$, leads to the following proposition: if for every $X \in \mathcal{M}$ and every $w : X \rightarrow \mathbb{R}^+$ there is a subspace $Y \subset X$, satisfying the weighted Ramsey condition (*) with parameter ψ , which is α -equivalent to a k -HST, then the same holds true for every $M \in \text{comp}_\beta(\mathcal{M})$. In our notation, we have the following lemma (proved in Section 3.3):

Lemma 3.11. *Let \mathcal{M} be a class of metric spaces. Let $k \geq 1$ and $\alpha \geq 1$. Then for any $\beta \geq \alpha k$,*

$$\psi_k(\text{comp}_\beta(\mathcal{M}), \alpha) = \psi_k(\mathcal{M}, \alpha),$$

In particular for $\beta \geq \alpha$,

$$\psi(\text{comp}_\beta(\mathcal{M}), \alpha) = \psi(\mathcal{M}, \alpha).$$

The following simple notion is used extensively in the sequel.

Definition 3.12. The *aspect ratio* of a finite metric space M , is defined as:

$$\Phi(M) = \frac{\text{diam}(M)}{\min_{x \neq y} d_M(x, y)}.$$

When $|M| = 1$ we use the convention $\Phi(M) = 1$. We note that $\Phi(M)$ can be viewed as M 's normalized diameter, or as its Lipschitz distance from an equilateral metric space.

Again, it is helpful to consider the k -HST representation of an ultrametric Y . In particular, notice that in this hierarchical representation, the number of levels is $O(\log_k \Phi(Y))$. In view of this fact, it seems reasonable to expect that when $\Phi(X)$ is small it would be easier to find a large subspace of X that is close to an ultrametric. This is, indeed, shown in Section 3.4.

Definition 3.13. The class of all metric spaces M with aspect ratio $\Phi(M) \leq \Phi$, for some given parameter Φ , is denoted by $\mathcal{N}(\Phi)$. Two more conventions that we use are: For every real $\Phi \geq 1$,

- $\psi(\Phi, \alpha) = \psi(\mathcal{N}(\Phi), \alpha)$. Similarly $\psi_k(\Phi, \alpha) = \psi_k(\mathcal{N}(\Phi), \alpha)$, and in general where \mathcal{M} is a class of metric spaces, $\psi_{\mathcal{M}}(\Phi, \alpha) = \psi_{\mathcal{M}}(\mathcal{N}(\Phi), \alpha)$.
- $\text{comp}_\beta(\Phi) = \text{comp}_\beta(\mathcal{N}(\Phi))$.

The main idea in bounding $\psi(\Phi, \alpha)$ is that the metric space can be decomposed into a small number of subspaces, the number of which can be bounded by a function of Φ , such that we can find among these, subspaces that are far enough from each other and contain enough weight to satisfy the weighted Ramsey condition (*). Such a decomposition of the space yields the recursive construction of a hierarchically well-separated tree, or an ultrametric. This is done in the proof of the following lemma. A more detailed description of the ideas involved in this decomposition and the proof of the lemma can be found in Section 3.4.

Lemma 3.14. *There exists an absolute constant $C' > 0$ such that for every $\alpha > 2$ and $\Phi \geq 1$:*

$$\psi(\Phi, \alpha) \geq 1 - C' \frac{\log \alpha + \log \log(4\Phi)}{\alpha}.$$

Note that for the class of metric spaces with aspect ratio $\Phi \leq \exp(O(\alpha))$, Lemma 3.14 yields the bound stated in Theorem 3.7'.

Combining Lemma 3.14 with Lemma 3.11 gives an immediate consequence on β -composition classes: for $\beta \geq \alpha$,

$$\psi(\text{comp}_\beta(\Phi), \alpha) = \psi(\Phi, \alpha) \geq 1 - C' \frac{\log \alpha + \log \log(4\Phi)}{\alpha}. \quad (1)$$

We now pass to a more detailed description of the proof of Theorem 3.7'. Let X be a metric space and assume that for some specific value of α we can prove the bound in the theorem (e.g., this trivially holds for $\alpha = \Phi(X)$ where we have $\psi(X, \alpha) = 1$).

Let \hat{X} be an arbitrary metric space and let X be a subspace of \hat{X} that is α -equivalent to an ultrametric, satisfying the weighted Ramsey condition $(*)$ with $\psi = \psi(\hat{X}, \alpha)$. We will apply the following ‘‘distortion refinement’’ procedure: find a subspace of X that is $(\alpha/2)$ -equivalent to an ultrametric, satisfying condition $(*)$ with $\psi' \geq (1 - C'' \frac{\log \alpha}{\alpha})\psi$. This implies that $\psi(\hat{X}, \alpha/2) \geq (1 - C'' \frac{\log \alpha}{\alpha})\psi(\hat{X}, \alpha)$. Theorem 3.7' now follows: we start with $\alpha = \Phi(\hat{X})$ and then apply the above distortion refinement procedure iteratively until we reach a distortion below our target. It is easy to verify that this implies the bound stated in the theorem.

The distortion refinement uses the bound in (1) on $\psi(\text{comp}_\beta(\Phi), \alpha'')$, in the particular case $\alpha'' < \alpha/2$ and $\Phi \leq \exp(O(\alpha))$. This is useful due the following claim: if X is α -equivalent to an ultrametric then it contains a subspace X' that is $(1 + 2/\beta)$ -equivalent to a metric space Z in $\text{comp}_\beta(\Phi)$, for $\Phi \leq \exp(O(\alpha))$, and which satisfies condition $(*)$ with $\psi'' \geq (1 - 2 \frac{\log \alpha}{\alpha})\psi$. By (1) we obtain a subspace Z' of Z which is α'' -equivalent to an ultrametric. By appropriately choosing all the parameters, it now follows from Lemma 3.10 that there is a subspace X'' of X' which is $(\alpha/2)$ -equivalent to an ultrametric, and the desired bound on $\psi(\hat{X}, \alpha/2)$ is achieved.

The proof of the above claim is based on two Lemmas relating ultrametrics, k -HSTs and metric compositions. Let X be α -equivalent to an ultrametric Y . The subspace X' is produced via a Ramsey-type result for ultrametrics which states that every ultrametric Y contains a subspace Y' which is α' -equivalent to a k -HST with $k > \alpha'$ (Lemma 3.5 can be viewed as a non-Ramsey result of this type when $k = \alpha'$). Moreover, we can ensure that condition $(*)$ is satisfied for the pair $Y' \subset Y$ with the bound stated below.

Lemma 3.15. *For every $k \geq \alpha > 1$,*

$$\psi_k(\text{UM}, \alpha) \geq 1 - \frac{\log(k/\alpha)}{\log \alpha}.$$

The proof of this lemma involves an argument on general tree structures described in Section 3.5.

Now, by Lemma 3.10 we obtain a subspace X' that is $\alpha'\alpha$ -equivalent to a k -HST for $k > \alpha'$. If k is large enough then the subtrees of the k -HST impose a clustering of X' . That is, each subtree corresponds to a subspace of X' of very small diameter, whereas the α distortion implies that the aspect ratio of the metric reflected by inter-cluster distances is bounded by α . By a recursive application of this procedure we obtain a metric space in $\text{comp}_\beta(\alpha)$, with the exact relation between k, α , and β stated in the lemma below. The details of this construction are given in Section 3.6.

Lemma 3.16. *For any $\alpha, \beta \geq 1$, if a metric space M is α -equivalent to an $\alpha\beta$ -HST then M is $(1 + 2/\beta)$ -equivalent to a metric space in $\text{comp}_\beta(\alpha)$.*

The distortion refinement process described above is formally stated in the following lemma:

Lemma 3.17. *There exists an absolute constant $C''' > 0$ such that for every metric space \hat{X} and any $\alpha > 8$,*

$$\psi\left(\hat{X}, \frac{\alpha}{2}\right) \geq \psi(\hat{X}, \alpha) \left(1 - C''' \frac{\log \alpha}{\alpha}\right).$$

Proof. Fix a weight function $w : \hat{X} \rightarrow \mathbb{R}^+$, let X be a subspace of \hat{X} that is α equivalent to an ultrametric Y , and satisfies the weighted Ramsey condition $(*)$ with $\psi(\hat{X}, \alpha)$. Fix two numbers $\alpha', \beta \geq 1$ which will be determined later, and set $k = \alpha\alpha'\beta$. Lemma 3.15 implies that Y contains a subspace Y' which is α' -equivalent to a k -HST, and Y' satisfies condition $(*)$ with $\psi_k(\text{UM}, \alpha') \geq 1 - \frac{\log(k/\alpha')}{\log \alpha'}$. By mapping X into an ultrametric Y , and then mapping the image of X in Y into a k -HST, we apply Lemma 3.10, obtaining a subspace X' of X that is $\alpha'\alpha$ -equivalent to a k -HST W , which satisfies condition $(*)$ with exponent $\psi_k(\text{UM}, \alpha') \cdot \psi(\hat{X}, \alpha) \geq \left(1 - \frac{\log(k/\alpha')}{\log \alpha'}\right) \psi(\alpha)$. Denote $\Phi = \alpha'\alpha$. We have that X' is Φ -equivalent to a $\Phi\beta$ -HST and therefore by Lemma 3.16, X' is $(1 + 2/\beta)$ equivalent to a metric space Z in $\text{comp}_\beta(\Phi)$. Now, we can use the bound in (1) to find a subspace Z' of Z that is β -equivalent to an ultrametric, and satisfies condition $(*)$ with exponent $\psi(\text{comp}_\beta(\Phi), \beta)$. By mapping X' into $Z \in \text{comp}_\beta(\Phi)$ and finally to an ultrametric, we apply Lemma 3.10 again, obtaining a subspace X'' of \hat{X} that is $\beta(1 + 2/\beta) = \beta + 2$ equivalent to an ultrametric U , satisfying condition $(*)$ with exponent

$$\left(1 - C' \frac{\log \beta + \log \log(4\Phi)}{\beta}\right) \left(1 - \frac{\log(\alpha\beta)}{\log \alpha'}\right) \psi(\hat{X}, \alpha).$$

Finally, if we choose $\beta = \alpha/2 - 2$ and let $\Phi = 2^{2\alpha}$ (which determines α'), we get that

$$\psi\left(\hat{X}, \frac{\alpha}{2}\right) \geq \left(1 - 9C' \frac{\log \alpha}{\alpha}\right) \left(1 - 2 \frac{\log \alpha}{\alpha}\right) \psi(\hat{X}, \alpha) \geq \left(1 - C'' \frac{\log \alpha}{\alpha}\right) \psi(\hat{X}, \alpha),$$

for an appropriate choice of C'' □

Theorem 3.7' is a straightforward consequence of Lemma 3.17:

Proof of Theorem 3.7'. By an appropriate choice of C we may assume that $\alpha > 8$. Let X be a metric space and set $\Phi = \Phi(X)$. Recall that $\psi(X, \Phi) = 1$. Let $m = \lfloor \log \alpha \rfloor$ and $M = \lceil \log \Phi \rceil$. Lemma 3.17 implies that $\psi(X, \alpha/2) \geq \psi(X, \alpha) - C'' \frac{\log \alpha}{\alpha}$, and so by an iterative application of this lemma we get

$$\psi(X, \alpha) \geq \psi(X, 2^m) \geq \psi(X, 2^M) - C'' \sum_{i=m+1}^M \frac{i}{2^i} \geq 1 - C'' \sum_{i=m+1}^{\infty} \frac{i}{2^i} = 1 - C'' \frac{m+2}{2^m} \geq 1 - 6C'' \frac{\log \alpha}{\alpha}.$$

□

This completes the overview of the proof of Theorem 3.7'. Sections 3.3–3.6 contain the proofs of the Lemmas described above.

Additionally, in Section 3.7 we describe in detail how to achieve Ramsey-type theorems for arbitrary values of $\alpha > 2$. The main ideas that make this possible are first, replacing Lemma 3.14 with another lemma that can handle distortions $2 + \epsilon$ and second, providing a more delicate application of our Lemmas, using the fact that we can find k -HSTs with large k ($\approx 1/\epsilon$) rather than just ultrametries, to ensure that accumulated losses in the distortion are small.

We end with a discussion on the algorithmic aspects of the metric Ramsey problem. Given a metric space X on n points, it is natural to ask whether we can find in polynomial time a subspace Y of X with n^ψ points which is α -equivalent to an ultrametric, for ψ as in Theorem 3.7. It is easily checked that the proofs of our lemmas yield polynomial time algorithms to find the corresponding subspaces. Thus, the only obstacle in achieving a polynomial time algorithm, is the fact, that

the proof of Theorem 3.7' involves $O(\log \Phi)$ iterations of an application of Lemma 3.17. We seek, however, a polynomial dependence only on n . This is remedied as follows: It is easily seen that using Lemma 3.6 we can start from $\psi(X, |X|) = 1$ rather than $\psi(X, \Phi(X)) = 1$. Thus we replace the bound of Φ with n , and end up with at most $O(\log n)$ iterations of Lemma 3.17. This implies a polynomial time algorithm to solve the metric Ramsey problem.

3.3 The weighted metric Ramsey problem and its relation with metric composition

In this section we prove Lemma 3.11 and Lemma 3.10. We begin with Lemma 3.10, which allows us to move between different classes of metric spaces while working with the weighted Ramsey problem.

Lemma 3.10. Let $\mathcal{M}, \mathcal{N}, \mathcal{P}$ be classes of metric spaces and $\alpha_1, \alpha_2 \geq 1$. Then

$$\psi_{\mathcal{M}}(\mathcal{P}, \alpha_1 \alpha_2) \geq \psi_{\mathcal{M}}(\mathcal{N}, \alpha_1) \cdot \psi_{\mathcal{N}}(\mathcal{P}, \alpha_2).$$

Proof. Let $\psi_1 = \psi_{\mathcal{M}}(\mathcal{N}, \alpha_1)$ and $\psi_2 = \psi_{\mathcal{N}}(\mathcal{P}, \alpha_2)$. Take $P \in \mathcal{P}$ and a weight function $w : P \rightarrow \mathbb{R}^+$. There is a subspace P' of P and an α_2 -embedding $f : P' \rightarrow N$, where $N \in \mathcal{N}$, and

$$\sum_{x \in P'} w(x)^{\psi_2} \geq \left(\sum_{x \in P} w(x) \right)^{\psi_2},$$

Similarly, for every weight function $w' : N \rightarrow \mathbb{R}^+$ there exists a subspace N' of N and an α_1 -embedding $g : N' \rightarrow M$, where $M \in \mathcal{M}$, and

$$\sum_{y \in N'} w'(y)^{\psi_1} \geq \left(\sum_{y \in N} w'(y) \right)^{\psi_1}.$$

By letting $P'' = f^{-1}(N')$, and for $y \in N$, $w'(y) = w(f^{-1}(y))^{\psi_2}$, we get that

$$\sum_{x \in P''} w(x)^{\psi_1 \psi_2} \geq \left(\sum_{x \in P'} w(x)^{\psi_2} \right)^{\psi_1} \geq \left(\sum_{x \in P} w(x) \right)^{\psi_1 \psi_2}.$$

Define $h : P'' \rightarrow M$ by $h(x) = g(f(x))$, then h is an $\alpha_1 \alpha_2$ -embedding. \square

Lemma 3.11 shows that the weighted Ramsey function stays unchanged as we pass from a class \mathcal{M} of metric spaces to its composition closure.

Lemma 3.11. Let \mathcal{M} be a class of metric spaces. Let $k \geq 1$ and $\alpha \geq 1$. Then for any $\beta \geq \alpha k$,

$$\psi_k(\text{comp}_{\beta}(\mathcal{M}), \alpha) = \psi_k(\mathcal{M}, \alpha).$$

Proof. Since $\mathcal{M} \subseteq \text{comp}_{\beta}(\mathcal{M})$, clearly $\psi_k(\text{comp}_{\beta}(\mathcal{M}), \alpha) \leq \psi_k(\mathcal{M}, \alpha)$. In what follows we prove the reverse inequality.

Let $\psi = \psi_k(\mathcal{M}, \alpha)$. Let $X \in \text{comp}_{\beta}(\mathcal{M})$. We prove that for any $w : X \rightarrow \mathbb{R}^+$ there exists a subspace Y of X and a k -HST H such that Y is α -equivalent to H via a non-contractive α -Lipschitz embedding, and:

$$\sum_{x \in Y} w(x)^{\psi} \geq \left(\sum_{x \in X} w(x) \right)^{\psi}.$$

The proof is by structural induction on the metric composition. If $X \in \mathcal{M}$ then this holds by definition of ψ . Otherwise, let $M \in \mathcal{M}$ and $\mathcal{N} = \{N_z\}_{z \in M}$ be such that $X = M_\beta[\mathcal{N}]$.

By induction, for each $z \in M$, there is a subspace Y_z of N_z that is α -equivalent to a k -HST H_z , defined by the tree T_z , via a non-contractive α -Lipschitz embedding, and

$$\sum_{u \in Y_z} w(u)^\psi \geq \left(\sum_{u \in N_z} w(u) \right)^\psi.$$

For a point $z \in M$ let $w'(z) = \sum_{u \in N_z} w(u)$. There exists a subspace Y_M of M that is α -equivalent to a k -HST H_M , defined by T_M , via a non-contractive α -Lipschitz embedding, and

$$\sum_{z \in Y_M} w'(z)^\psi \geq \left(\sum_{z \in M} w'(z) \right)^\psi = \left(\sum_{x \in X} w(x) \right)^\psi.$$

Let $Y = \cup_{z \in Y_M} Y_z$. It follows that

$$\sum_{u \in Y} w(x)^\psi = \sum_{z \in Y_M} \sum_{u \in Y_z} w(u)^\psi \geq \sum_{z \in Y_M} \left(\sum_{u \in N_z} w(u) \right)^\psi \geq \left(\sum_{x \in X} w(x) \right)^\psi.$$

We now construct a k -HST H that is defined by a tree T , as follows. Start with a tree T' that is isomorphic to T_M and has labels $\Delta(u) = \beta\gamma \cdot \Delta_{T_M}(u)$ (where $\gamma = \frac{\max_{z \in M} \text{diam}(N_z)}{\min_{x \neq y \in M} d_M(x, y)}$, as in Definition 2.1). At each leaf of the tree corresponding to a point $z \in M$, create a labelled subtree rooted at z that is isomorphic to T_z with labels as in T_z . Denote the resulting tree by T . Since we have a non-contractive α -embedding of Y_z in H_z , it follows that $\Delta(z) = \text{diam}(H_z) \leq \alpha \text{diam}(Y_z) \leq \alpha \text{diam}(N_z)$. Let p be a parent of z in T_M . Since we have a non-contractive α -embedding of Y_M into H_M , it follows that $\Delta_{T_M}(p) \geq d_M(x, y)$ for some $x, y \in Y_M$. Therefore $\Delta(p) \geq \beta\gamma \cdot \min\{d_M(x, y); x \neq y \in M\}$. Consequently, $\Delta(p)/\Delta(z) \geq \beta/\alpha \geq k$. Since H_M and H_z are k -HSTs, it follows that T also defines a k -HST.

It is left to show that Y is α -equivalent to H . Recall that for each $z \in M$ there is a non-contractive Lipschitz bijection $f_z : Y_z \rightarrow H_z$ that satisfies for every $u, v \in Y_z$, $d_{Y_z}(u, v) \leq d_{H_z}(f_z(u), f_z(v)) \leq \alpha d_{Y_z}(u, v)$. Define $f : Y \rightarrow H$ as follows. If $z \in M$ and $u \in N_z$, set $f(u) = f_z(u)$. Then for $u, v \in M$ such that $u, v \in N_z$ we have

$$d_Y(u, v) = d_{Y_z}(u, v) \leq d_{H_z}(f_z(u), f_z(v)) = d_H(f(u), f(v)) \leq \alpha d_{Y_z}(u, v) = \alpha d_Y(u, v).$$

Additionally, we have a non-contractive Lipschitz bijection $f_M : Y_M \rightarrow H_M$ that satisfies for every $x, y \in Y_M$, $d_{Y_M}(x, y) \leq d_{H_M}(f_M(x), f_M(y)) \leq \alpha d_{Y_M}(x, y)$. Hence for $x \neq y \in M$, and $u \in N_x, y \in N_y$,

$$d_Y(u, v) = \beta\gamma d_{Y_M}(x, y) \leq \beta\gamma d_{H_M}(f_M(x), f_M(y)) = d_H(f(u), f(v)) \leq \alpha\beta\gamma d_{Y_M}(x, y) = \alpha d_Y(u, v).$$

□

3.4 Exploiting metrics with bounded aspect ratio

In this section we prove Lemma 3.14 (Section 3.2). That is, we give lower bounds on $\psi = \psi(\Phi, \alpha)$ which depend on the aspect ratio of the metric space, Φ .

The proof of the lemma starts by obtaining lower bounds for a restricted class of weight functions w . These bounds are then extended to general weights. The class of “nice” weight functions is

itself divided into two classes. In one class we have a lower bound on the minimal weight relative to the total overall weight, and the other is the constant weight function. This is formally defined as follows:

Definition 3.18. Fix some $q \geq 1$. A sequence $x = \{x_i\}_{i=1}^{\infty}$ of non-negative real numbers will be called q -decomposable if there exists $\omega > 0$ such that:

$$\{i \in \mathbb{N}; x_i > 0\} = \left\{ i \in \mathbb{N}; x_i \geq \frac{1}{q} \sum_{j=1}^{\infty} x_j \right\} \cup \{i \in \mathbb{N}; x_i = \omega\}.$$

We will prove the following lemma:

Lemma 3.19. Let $q \geq 2$, and $t \geq 8$ be an integer. Let (X, d) be an n -point metric space and let $w : X \rightarrow \mathbb{R}^+$, be a weight function such that $\{w(x)\}_{x \in X}$ is q -decomposable. Then there exists a subspace $Y \subseteq X$ that is $4t$ -equivalent to an ultrametric and satisfies:

$$\sum_{x \in Y} w(x)^\psi \geq \left(\sum_{x \in X} w(x) \right)^\psi,$$

where $\psi = [t \log(4q\Phi(X))]^{-2/t}$.

The proof of Lemma 3.19 uses a decomposition of the metric space X into a small number of subspaces. This type of strategy has been used in several earlier papers in combinatorics and theoretical computer science, but the argument closest in spirit to ours is in [5]. The idea is to consider two diametrical points, split the space into shells according to the distance from one of these two points, and discard the points in one of the shells. Intuitively, we would like to discard a shell with small weight. The exact choice is somewhat more sophisticated, tailored to ensure the weighted Ramsey condition (*). The other shells form two subsets of the space that are substantially separated. By an appropriate choice of the parameters, we can guarantee that the union of the inner layers has diameter smaller than a constant factor of the diameter of the whole space, and hence a smaller aspect ratio. The role of q -decomposable weights is as follows: This argument works fairly well for uniform weights, and a slight modification of it yields bounds as a function of q (in addition to Φ) when in the weighted case only a few points carry each at least $\frac{1}{q}$ of the total weight. Here the argument splits according to the diameter of the set of "heavy" points. If the diameter is small, the previous argument is started from a point that resides far away from the heavy points. This guarantees that none of the "heavyweights" get eliminated in the above-described process. If their diameter is proportional to that of the whole space, it is possible to argue similarly to the uniform-weight case, except that we now obtain better bounds, since we can make estimates in terms of q (rather than the cardinality of the space n).

The extension of Lemma 3.19 to arbitrary weight functions requires a lemma on numerical sequences. This lemma allows us to reduce the case of general sequences of weights to q -decomposable ones.

Lemma 3.20. Fix $q \geq 16$ and let $x = \{x_i\}_{i=1}^{\infty}$ be a sequence of non-negative real numbers. Denote $p = 1 - \frac{\log_2 \log_2 q}{\log_2 q}$. There exists a sequence $y = \{y_i\}_{i=1}^{\infty}$ such that $y_i \leq x_i$ for all $i \geq 1$, $\sum_{i \geq 1} y_i^p \geq \left(\sum_{i \geq 1} x_i \right)^p$, and the sequence $\{y_i^p\}_{i=1}^{\infty}$ is q -decomposable.

Together these lemmas imply our main lemma:

Lemma 3.14. For every $\alpha > 2$ and every $\Phi \geq 1$:

$$\psi(\Phi, \alpha) \geq 1 - C \frac{\log \alpha + \log \log 4\Phi}{\alpha},$$

where C is a universal constant.

Proof. Clearly we may assume that $\alpha \geq 32$. Let X be a metric space with $\Phi(X) \leq \Phi$, and $w : X \rightarrow \mathbb{R}^+$ a weight function. Set $t = \lfloor \alpha/4 \rfloor$. By applying Lemma 3.20 to the sequence $\{w(x)\}_{x \in X}$ with $q = 2^t$, we obtain a weight function w' such that $w'(x) \leq w(x)$ for all $x \in X$, the sequence $\{w'(x)^p\}_{x \in X}$ is q -decomposable, and

$$\sum_{x \in X} w'(x)^p \geq \left(\sum_{x \in X} w(x) \right)^p,$$

where $p = 1 - \frac{\log_2 t}{t}$.

Let $\beta = [t \log(4q\Phi(X))]^{-2/t}$ and apply Lemma 3.19 to the space X and weights $w'' = w'^p$. We obtain a subspace Y which is $4t$ -equivalent to an ultrametric, such that

$$\sum_{x \in Y} w(x)^{p\beta} \geq \sum_{x \in Y} w'(x)^{p\beta} \geq \left(\sum_{x \in X} w'(x)^p \right)^\beta \geq \left(\sum_{x \in X} w(x) \right)^{p\beta}.$$

Therefore,

$$\psi(\Phi, \alpha) \geq p\beta \geq \left(1 - \frac{\log_2 t}{t}\right) \left(1 - \frac{4 \log t}{t} - \frac{2 \log \log(4\Phi)}{t}\right) \geq 1 - C \frac{\log \alpha + \log \log(4\Phi)}{\alpha},$$

for an appropriate choice of C . □

We now pass to the proof of Lemma 3.20. Let $x = \{x_i\}_{i=1}^\infty$ be a sequence of non-negative real numbers which isn't identically zero. Let $p \geq 0$. Recall that the (p, ∞) norm of x is defined by $\|x\|_{p, \infty} = \sup_{i \geq 1} i^{1/p} x_i^*$, where $\{x_i^*\}_{i=1}^\infty$ is the non-increasing rearrangement of the sequence $(|x_i|)_{i=1}^\infty$. We will require the following numerical fact:

Lemma 3.21. For every $x \in \ell_1$ as above and every $0 < p < 1$:

$$\|x\|_{p, \infty} \geq \left(\frac{1-p}{2-p}\right)^{1/p} \cdot \frac{\|x\|_1^{1/p}}{\|x\|_\infty^{(1-p)/p}}.$$

Proof. We may assume without loss of generality that $\|x\|_1 = 1$ and $\|x\|_\infty = x_1 \geq x_2 \geq \dots \geq 0$. Obviously $\|x\|_{p, \infty} \geq x_1$ so that if $x_1 \geq [(1-p)/(2-p)]^{1/p} x_1^{-(1-p)/p}$ we are done. Assume therefore that the reverse inequality holds, i.e., $x_1 < \frac{1-p}{2-p}$. Set $\alpha = \|x\|_{p, \infty}^p$ and denote $j = \lceil \alpha/x_1^p \rceil + 1$. Note that for every $i \geq 1$, $x_i \leq (\alpha/i)^{1/p}$. Therefore,

$$\sum_{i=1}^{j-1} x_i \leq (j-1)x_1 \leq \left\lceil \frac{\alpha}{x_1^p} \right\rceil x_1 \leq x_1 \left(\frac{\alpha}{x_1^p} + 1 \right) = \alpha x_1^{1-p} + x_1,$$

and

$$\begin{aligned} \sum_{i=j}^{\infty} x_i &\leq \sum_{i=j}^{\infty} \frac{\alpha^{1/p}}{i^{1/p}} \leq \alpha^{1/p} \int_{j-1}^{\infty} z^{-1/p} dz \\ &\leq \alpha^{1/p} \frac{p}{1-p} \cdot \left[\frac{\alpha}{x_1^p} \right]^{-\frac{1-p}{p}} \leq \alpha^{1/p} \frac{p}{1-p} \left(\frac{x_1^p}{\alpha} \right)^{\frac{1-p}{p}} = \frac{p}{1-p} \alpha x_1^{1-p}. \end{aligned}$$

By summing both inequalities and using the bound on x_1 we get

$$\frac{1}{1-p} \alpha x_1^{1-p} + \frac{1-p}{2-p} \geq 1,$$

which simplifies to give the required result. \square

Proof of Lemma 3.20. We may assume that x is a non-increasing sequence of non-negative real numbers and that $\|x\|_1 = 1$.

We will prove below that there exist indexes $0 \leq l \leq b$ such that $x_l^p \geq \frac{2}{q}$ and:

$$S = \sum_{i=1}^l x_i^p + (b-l)x_b^p \geq 1. \quad (2)$$

If $b = l$ assume that l is the minimal index for which (2) holds. It follows that in this case $S = \sum_{i=1}^l x_i^p < 1 + x_l^p \leq 2$. Similarly if $b > l$, fix l and assume that b is the minimal index for which (2) holds. It follows that $S = \sum_{i=1}^l x_i^p + (b-l)x_b^p \leq \sum_{i=1}^l x_i^p + (b-1-l)x_{b-1}^p + x_b^p < 1 + x_b^p \leq 2$.

Define the sequence $\{y_i\}_{i=1}^{\infty}$ so that $y_i = x_i$ for $i \leq l$, $y_i = x_b$ for $l < i \leq b$ and $y_i = 0$ for $i > b$. It follows that $y_i \leq x_i$ for all $i \geq 1$ and $\sum_{i \geq 1} y_i^p = S \geq 1 = (\sum_{i \geq 1} x_i)^p$. Since for $j \leq l$ $y_j^p = x_j^p \geq \frac{2}{q} \geq \frac{S}{q} = \frac{1}{q} \sum_{i \geq 1} y_i^p$, for $l < j \leq b$, $y_j^p = x_b^p$ and for $j > b$, $y_j^p = 0$, we get that $\{y_i^p\}_{i=1}^{\infty}$ is q -decomposable.

It remains to prove (2). Let $l \geq 0$ be the largest integer for which $x_l^p \geq \frac{2}{q}$. If $\sum_{i=1}^l x_i^p \geq 1$ we are done. Otherwise, consider the sequence $z = (x_{l+1}, x_{l+2}, \dots)$. By the choice of l , for $i > l$, $x_i < (2/q)^{1/p}$, and thus $\|z\|_{\infty} \leq (2/q)^{1/p}$. Moreover, $\frac{1-p}{(2-p)\|z\|_{\infty}^{1-p}} \geq \frac{\log_2 \log_2 q}{2} \cdot \left(\frac{\log_2 q}{2} \right)^{(1-p)/p} \geq 1$, so by applying Lemma 3.21 to z we get that $\|z\|_{p,\infty}^p / \|z\|_1 \geq 1$, i.e. there is an integer $b > l$ such that:

$$(b-l)x_b^p \geq \|z\|_1 = 1 - \sum_{i=1}^l x_i \geq 1 - \sum_{i=1}^l x_i^p.$$

\square

We are now in position to prove the main technical lemma.

Proof of Lemma 3.19. For simplicity denote $\beta(\Phi) = [t \log_2(4q\Phi)]^{-2/t}$. We will prove by induction on n that any n point weighted metric space (M, d, w) contains a subspace $N \subset M$ such that:

$$\sum_{x \in N} w(x)^{\beta(\Phi(M))} \geq \left(\sum_{x \in M} w(x) \right)^{\beta(\Phi(M))},$$

and for which there is a non-contractive, $4t$ -Lipschitz embedding of N into an ultrametric H with $\text{diam}(H) = \text{diam}(M)$.

In what follows for every $S \subset M$ we denote $w(S) = \sum_{x \in S} w(x)$.

Let (M, d, w) be a weighted n -point metric space such that w is q -decomposable. Without loss of generality we may assume that $w(M) = 1$ and $\min_{x \neq y \in M} d(x, y) = 1$. Denote $\Phi = \Phi(M)$. The latter assumption implies that $\Phi = \text{diam}(M)$. In what follows we denote for $r > 0$ and $x \in M$, $B(x, r) = \{y \in M; d(y, x) < r\}$. The proof proceeds by proving the following claim:

Claim 3.22. *There exists $i \in \{1, \dots, t\}$ and $x_0 \in M$ such that if we set $A = \{x_0\} \cup B\left(x_0, \frac{(i-1)\Phi}{4t}\right)$ and $B = M \setminus B\left(x_0, \frac{i\Phi}{4t}\right)$ then:*

$$\max \left\{ \frac{w(A)^{\beta(\Phi/2)}}{[\max_{y \in A} w(y)]^{\beta(\Phi/2) - \beta(\Phi)}}, w(A)^{(\log_2 q)^{-1/(t-1)}} \right\} + w(B) \geq 1. \quad (3)$$

Before proving Claim 3.22 we will show how it implies the required result. Let i, A, B be as in Claim 3.22. Note that $A \neq \emptyset$ and $\text{diam}(A) < \frac{\Phi}{2} < \text{diam}(M)$. In particular it follows that, $|A|, |B| < n$ so that by the induction hypothesis there are subspaces $A' \subset A$ and $B' \subset B$ such that

$$\sum_{x \in A'} w(x)^{\beta(\Phi(A))} \geq w(A)^{\beta(\Phi(A))} \quad \text{and} \quad \sum_{x \in B'} w(x)^{\beta(\Phi(B))} \geq w(B)^{\beta(\Phi(B))} \geq w(B),$$

HSTs X and Y with $\text{diam}(X) = \text{diam}(A)$, $\text{diam}(Y) = \text{diam}(B)$, and non-contractive embeddings $f : A' \rightarrow X$, $g : B' \rightarrow Y$ which are $4t$ -Lipschitz. Let T be the tree defining X and u be its root. Let S be the tree defining Y and v be its root. Define a tree R as follows: its root is r and the only two subtrees emerging from it are isomorphic to T and S . Label the root of R by setting $\Delta(r) = \text{diam}(M)$, and leave the labels of T and S unchanged. Note that

$$\begin{aligned} \Delta(r) &= \text{diam}(M) \\ &\geq \max\{\text{diam}(A), \text{diam}(B)\} = \max\{\text{diam}(X), \text{diam}(Y)\} = \max\{\Delta(u), \Delta(v)\}, \end{aligned}$$

so that with these definitions the leaves of R , $X \cup Y$, form a HST with $\text{diam}(X \cup Y) = \Phi = \text{diam}(M)$. Define $h : A' \cup B' \rightarrow X \cup Y$ by $h|_{A'} = f$ and $h|_{B'} = g$. If $a \in A'$ and $b \in B'$ then $d(h(a), h(b)) = \Phi \geq d(a, b)$. Hence h is non-contracting. Additionally:

$$d(a, b) \geq d(a, x_0) - d(b, x_0) \geq \frac{\Phi i}{4t} - \frac{\Phi(i-1)}{4t} = \frac{\Phi}{4t} = \frac{d_R(h(a), h(b))}{4t},$$

so that h is $4t$ -Lipschitz.

Observe that since $\beta(\Phi) \leq \beta(\Phi(A)) \leq (\log_2 q)^{-1/(t-1)}$ and $w(x) \leq 1$ (point-wise),

$$\sum_{x \in A'} w(x)^{\beta(\Phi)} \geq \sum_{x \in A'} w(x)^{\beta(\Phi(A))} \geq w(A)^{\beta(\Phi(A))} \geq w(A)^{(\log_2 q)^{-1/(t-1)}}.$$

Moreover, since $\Phi(A) \leq \Phi/2$:

$$\begin{aligned} \sum_{x \in A'} w(x)^{\beta(\Phi)} &\geq \frac{1}{[\max_{y \in A} w(y)]^{\beta(\Phi(A)) - \beta(\Phi)}} \sum_{x \in A'} w(x)^{\beta(\Phi(A))} \\ &\geq [\max_{y \in A} w(y)]^{\beta(\Phi)} \left[\frac{w(A)}{\max_{y \in A} w(y)} \right]^{\beta(\Phi(A))} \\ &\geq [\max_{y \in A} w(y)]^{\beta(\Phi)} \left[\frac{w(A)}{\max_{y \in A} w(y)} \right]^{\beta(\Phi/2)}. \end{aligned}$$

We deduce that:

$$\sum_{x \in A'} w(x)^{\beta(\Phi)} \geq \max \left\{ \frac{w(A)^{\beta(\Phi/2)}}{[\max_{y \in A} w(y)]^{\beta(\Phi/2) - \beta(\Phi)}}, w(A)^{(\log_2 q)^{-1/(t-1)}} \right\},$$

so that by (3),

$$\sum_{x \in A' \cup B'} w(x)^{\beta(\Phi)} \geq \sum_{x \in A'} w(x)^{\beta(\Phi)} + w(B) \geq 1,$$

as required.

Proof of Claim 3.22: The fact that w is q -decomposable implies that we can split $M = N_1 \dot{\cup} N_2$, so that $w(x) \geq \frac{1}{q}$ for every $x \in N_1$ and there is $\omega > 0$ such that $w(x) = \omega$ for every $x \in N_2$. We distinguish between two cases:

Case 1: $\text{diam}_M(N_1) > \frac{\Phi}{2}$. In this case there are $x_0, x'_0 \in N_1$ such that $d(x_0, x'_0) > \frac{\Phi}{2}$. It follows in particular that $B(x_0, \Phi/4) \cap B(x'_0, \Phi/4) = \emptyset$ so that by interchanging the roles of x_0 and x'_0 , if necessary, we may assume that $w(B(x_0, \Phi/4)) \leq \frac{w(M)}{2} = \frac{1}{2}$. Since $x_0 \in N_1$, $w(x_0) \geq \frac{1}{q}$. This implies that there exist $i \in \{1, \dots, t\}$ such that

$$w \left(\{x_0\} \cup B \left(x_0, \frac{(i-1)\Phi}{4t} \right) \right)^{(\log_2 q)^{-1/(t-1)}} \geq w \left(B \left(x_0, \frac{i\Phi}{4t} \right) \right),$$

since otherwise:

$$\begin{aligned} \frac{1}{2} &\geq w \left(B \left(x_0, \frac{\Phi}{4} \right) \right) \\ &> w \left(\{x_0\} \cup B \left(x_0, \frac{(t-1)\Phi}{4t} \right) \right)^{(\log_2 q)^{-1/(t-1)}} \\ &> \dots > w \left(B \left(x_0, \frac{\Phi}{4t} \right) \right)^{(\log_2 q)^{-1}} \\ &\geq w(x_0)^{(\log_2 q)^{-1}} \geq \frac{1}{q^{(\log_2 q)^{-1}}} = \frac{1}{2}, \end{aligned}$$

which is a contradiction. Fixing such an index i and defining A, B as in the statement of Claim 3.22 we get that:

$$\begin{aligned} w(A)^{(\log_2 q)^{-1/(t-1)}} + w(B) &= \\ &= w \left(\{x_0\} \cup B \left(x_0, \frac{(i-1)\Phi}{4t} \right) \right)^{(\log_2 q)^{-1/(t-1)}} + \left[1 - w \left(B \left(x_0, \frac{i\Phi}{4t} \right) \right) \right] \geq 1, \end{aligned}$$

which proves (3).

Case 2: $\text{diam}_M(N_1) \leq \frac{\Phi}{2}$. In this case take $x_0 \in M$ such that $d(x_0, N_1) = \max_{x \in M} d(x, N_1)$. We claim that this implies that $N_1 \cap B(x_0, \Phi/4) = \emptyset$. Indeed, otherwise it will follow that $d(x_0, N_1) < \Phi/4$ so that by the choice of x_0 , for every $x, y \in M$,

$$d(x, y) \leq d(x, N_1) + d(y, N_1) + \text{diam}(N_1) < 2d(x_0, N_1) + \frac{\Phi}{2} < \Phi,$$

which is a contradiction.

Set $m = |N_2|$ and denote for $i \in \{0, \dots, t\}$:

$$\epsilon_i = \frac{|(\{x_0\} \cup B(x_0, \frac{i\Phi}{4t})) \cap N_2|}{m}.$$

Note that since $x_0 \in N_2$, $m^{-1} = \epsilon_0 \leq \epsilon_t \leq 1$. We claim that this implies that there is some $i \in \{1, \dots, t\}$ such that:

$$\epsilon_{i-1}^{\beta(\Phi/2)} m^{\beta(\Phi/2) - \beta(\Phi)} \geq \epsilon_i \quad (4)$$

Indeed, if we set $a = \log_2(2q\Phi) \geq 1$, then if there is no such i we have for every $i \in \{1, \dots, t\}$:

$$\epsilon_{i-1} < \left(\frac{\epsilon_i}{m^{\frac{1}{(ta)^{2/t}} - \frac{1}{[t(a+1)]^{2/t}}}} \right)^{(ta)^{2/t}}.$$

Denote $b = m^{\frac{1}{(ta)^{2/t}} - \frac{1}{[t(a+1)]^{2/t}}}$ and $c = (ta)^{2/t}$. The above inequality then becomes $\epsilon_{i-1} < (\epsilon_i/b)^c$. Iterating this t times we get:

$$\frac{1}{m} = \epsilon_0 < \frac{\epsilon_t^{c^t}}{b^{c+c^2+\dots+c^t}} = \frac{\epsilon_t^{c^t}}{b^{\frac{c}{c-1}(c^t-1)}} \leq \frac{1}{b^{c^t-1}}.$$

So that:

$$m^{(t^2a^2-1)} \left[\frac{1}{(ta)^{2/t}} - \frac{1}{[t(a+1)]^{2/t}} \right] < m,$$

but an application of the mean value theorem gives a contradiction, since:

$$(t^2a^2 - 1) \left[\frac{1}{(ta)^{2/t}} - \frac{1}{[t(a+1)]^{2/t}} \right] \geq \frac{t^2a^2}{2} \frac{2}{t^{1+2/t}(a+1)^{1+2/t}} \geq t^{1-2/t} \left(\frac{a}{a+1} \right)^2 \geq 8^{3/4} \cdot \frac{1}{4} \geq 1.$$

Choose an index $i \in \{1, \dots, t\}$ satisfying (4) and let A, B be as in the statement of Claim 3.22 for this particular i . Observe that:

$$B = M \setminus B\left(x_0, \frac{i\Phi}{4t}\right) \supset M \setminus B\left(x_0, \frac{\Phi}{4}\right) \supset N_1,$$

so that:

$$\begin{aligned} \frac{w(A)^{\beta(\Phi/2)}}{[\max_{y \in A} w(y)]^{\beta(\Phi/2) - \beta(\Phi)}} + w(B) &= \frac{(\omega \epsilon_{i-1} m)^{\beta(\Phi/2)}}{\omega^{\beta(\Phi/2) - \beta(\Phi)}} + w(N_1) + (1 - \epsilon_i)m\omega \\ &= \omega^{\beta(\Phi)} (\epsilon_{i-1} m)^{\beta(\Phi/2)} + w(N_1) + (1 - \epsilon_i)m\omega \\ &\geq (m\omega)^{\beta(\Phi)} \epsilon_i + w(N_1) + (1 - \epsilon_i)m\omega \\ &\geq m\omega \epsilon_i + w(N_1) + (1 - \epsilon_i)m\omega = w(M) = 1. \end{aligned}$$

This concludes the proof of Claim 3.22. □

3.5 Passing from an ultrametric to a k -HST

In what follows we show that every ultrametric contains large subsets which are embeddable in a k -HST with distortion $\alpha < k$.

An unweighted version of the following result was proved in [5]. The bound for the weighted Ramsey function is a straightforward modification of the proof in [5]:

Lemma 3.23 ([5]). *For every $k > \alpha > 1$,*

$$\psi_k(\text{UM}, \alpha) \geq \frac{1}{\lceil \log_\alpha k \rceil}.$$

If k is large with respect to α then the bound of [5] provides a good approximation for $\psi_k(\text{UM}, \alpha)$. In fact, this is how Lemma 3.23 is used in Section 3.7.

However, when k is close to α the bound in Lemma 3.23 is not good enough for proving our main theorem. We obtain bounds for this range of parameters in the following lemma (stated in Section 3.2 in slightly weaker form).

Lemma 3.15. For every $k > \alpha > 1$,

$$\psi_k(\text{UM}, \alpha) \geq 1 - \frac{1}{\lceil \log_{k/\alpha} \alpha \rceil}.$$

Before proving Lemma 3.15 we require some lemmas concerning unweighted trees.

Definition 3.24. Let $h > 1$ be an integer and $i \in \{0, \dots, h-1\}$. We say that a rooted tree T is (i, h) -periodically sparse if for every $l \equiv i \pmod{h}$, every vertex at depth l in T is degenerate. T is called h periodically sparse if there exists $i \in \{0, \dots, h-1\}$ for which T is (i, h) -periodically sparse.

In what follows we always use the convention that a subtree T' of a rooted tree T is rooted at the root of T and that the leaves of T' are also leaves of T . We denote by $\text{lvs}(T)$ the leaves of T .

Lemma 3.25. *Fix an integer $h > 1$. Let T be a finite rooted tree. Then for any $w : \text{lvs}(T) \rightarrow \mathbb{R}^+$ there exists a subtree of T , T' , which is h periodically sparse and:*

$$\sum_{v \in \text{lvs}(T')} w(v)^{\frac{h-1}{h}} \geq \left(\sum_{v \in \text{lvs}(T)} w(v) \right)^{\frac{h-1}{h}}.$$

The techniques we use in the proof of Lemma 3.25 are similar to those used in the proof of Lemma 3.23 in [5]. It can also be derived from a result in [11] concerning influences in multi-stage games. These facts are also closely related to an isoperimetric inequality of Loomis and Whitney [39].

Proof. For every $i \in \{0, 1, \dots, h-1\}$ let $f_i(T)$ be the maximum of $\sum_{v \in \text{lvs}(T')} w(v)^{\frac{h-1}{h}}$ over all the (i, h) -periodically sparse subtrees, T' , of T . We will prove by induction on the maximal depth of T that:

$$\prod_{i=0}^{h-1} f_i(T) \geq \left(\sum_{v \in \text{lvs}(T)} w(v) \right)^{h-1},$$

from which it will follow that $\max_{0 \leq i \leq h-1} f_i(T) \geq \left(\sum_{v \in \text{lvs}(T)} w(v) \right)^{\frac{h-1}{h}}$, as required.

For a tree T of depth 0, consisting of a single node v , we have that $f_i(T) = w(v)^{\frac{h-1}{h}}$ and therefore

$$\prod_{i=0}^{h-1} f_i(T) \geq \left(w(v)^{\frac{h-1}{h}} \right)^h = w(v)^{h-1}.$$

Assume that the maximal depth of T is at least 1, let r be the root of T and denote by v_1, \dots, v_l its children. Denote by T_j the subtree of T rooted at v_j . Observe that:

$$f_0(T) \geq \max_{1 \leq j \leq l} f_{h-1}(T_j),$$

and for $i \in \{1, \dots, h-1\}$:

$$f_i(T) = \sum_{j=1}^l f_{i-1}(T_j),$$

By repeated application of Hölder's inequality:

$$\sum_{j=1}^l \prod_{i=0}^{h-2} [f_i(T_j)]^{\frac{1}{h-1}} \leq \left(\prod_{i=0}^{h-2} \sum_{j=1}^l f_i(T_j) \right)^{\frac{1}{h-1}}.$$

Therefore, by the induction hypothesis:

$$\begin{aligned} \prod_{i=0}^{h-1} f_i(T) &\geq \max_{1 \leq j \leq l} f_{h-1}(T_j) \cdot \prod_{i=1}^{h-1} \sum_{j=1}^l f_{i-1}(T_j) \\ &\geq \max_{1 \leq j \leq l} f_{h-1}(T_j) \cdot \left(\sum_{j=1}^l \left(\prod_{i=0}^{h-2} f_i(T_j) \right)^{\frac{1}{h-1}} \right)^{h-1} \\ &\geq \left(\sum_{j=1}^l \left(\prod_{i=0}^{h-1} f_i(T_j) \right)^{\frac{1}{h-1}} \right)^{h-1} \geq \left(\sum_{j=1}^l \sum_{v \in \text{lvs}(T_j)} w(v) \right)^{h-1} = \left(\sum_{v \in \text{lvs}(T)} w(v) \right)^{h-1}. \end{aligned}$$

□

Before proving Lemma 3.15, we prove the following variant of a proposition from [4]:

Lemma 3.5. For any $k > 1$, any ultrametric is k -equivalent to an exact k -HST.

Proof. Let T be a labelled tree rooted at r . Define a new labelled tree T' as follows. Let u be a minimal depth vertex in T that has a child v of for which $\Delta(u) \neq k\Delta(v)$. Let $0 \leq i \in \mathbb{N}$ be defined via $k^i \leq \frac{\Delta(u)}{\Delta(v)} < k^{i+1}$. Relabel v by setting $\Delta'(v) = \frac{\Delta(u)}{k^i} \geq \Delta(v)$, and replace the edge $[u, v]$ by a path of length i whose labels decrease by a factor k at each step. Denote the tree thus obtained by T' . If we start out with an HST X with defining tree T , then the tree T' produced in this procedure defines a new HST. Iterating this construction as long as possible, we arrive at a tree \tilde{T} which defines an exact k -HST. To prove that we have distorted the metric by a factor of at most k observe that by construction, for any $x, y \in X$, $\text{lca}_T(x, y) = \text{lca}_{\tilde{T}}(x, y)$ and that for any $v \in T \cap \tilde{T}$, $\Delta_T(v) \leq \Delta_{\tilde{T}}(v) \leq k\Delta_T(v)$. □

Proof of Lemma 3.15. Let $h = \lceil \log_{k/\alpha} \alpha \rceil$ and let $s = k^{1/h}$. By Lemma 3.5, X is s -equivalent to some exact s -HST Y via a non-contractive Lipschitz embedding. Let T be the tree defining Y . Lemma 3.25 yields a subtree S of T which is (i, h) -periodically sparse for some $i \in \{0, \dots, h-1\}$, such that

$$\sum_{v \in S} w(g^{-1}(v))^{\frac{h-1}{h}} \geq \left(\sum_{x \in X} w(x) \right)^{\frac{h-1}{h}}.$$

By attaching a path of length $h-1-i$ to the root of S we may assume that S is $(h-1, h)$ -periodically sparse. Similarly, by adding appropriate paths to the leaves of S we may assume that there is an integer m such that all the leaves of S are at depth mh . Denote by r the root of S . We change the tree S as follows. For every integer $0 \leq j < m$ delete all the vertices of S whose depth is in the interval $[jh+1, (j+1)h-1]$ and connect every vertex of depth jh directly to all its descendants of depth $(j+1)h$. Denote the tree thus obtained by S' and denote by Y' the metric space induced by S' on Y . It is evident that Y' is an exact s^{h-1} -HST. We claim that Y' is s^{h-1} equivalent to a subspace of X via a non-contractive s^{h-1} Lipschitz embedding. Indeed, fix $u, v \in Y'$ and let w be their least common ancestor in S . If we denote by q the depth of w in S then the key observation is that since S is $(h-1, h)$ -periodically sparse, $q \not\equiv (h-1) \pmod{h}$. We can therefore write $q = i + jh$ for some $i \in \{0, \dots, h-2\}$ and $j \geq 0$. If we denote by w' the least common ancestor of u, v in S' then by the construction, w' is in depth jh in S . Hence $d_Y(u, v) = \frac{\Delta(r)}{s^{i+jh}}$ and $d_{Y'}(u, v) = \frac{\Delta(r)}{s^{jh}}$, so that:

$$d_Y(u, v) \leq d_{Y'}(u, v) \leq s^i d_Y(u, v) \leq s^{h-2} d_Y(u, v).$$

This shows that Y' is s^{h-2} equivalent to Y via a non-contractive s^{h-2} Lipschitz embedding. Since Y is s equivalent to a subspace of X via a non-contractive s Lipschitz bijection we have that Y' is s^{h-1} equivalent to a subspace of X .

Recall that $s^h = k$, and it remains to show that $s^{h-1} \leq \alpha$. Indeed by our choice of h , $h-1 \leq \log_{k/\alpha} \alpha$, or $\frac{1}{h-1} \geq \log_\alpha(k/\alpha)$. Therefore $\alpha^{\frac{h}{h-1}} \geq k$, and so $s^{h-1} = k^{\frac{h-1}{h}} \leq \alpha$. \square

3.6 Passing from a k -HST to metric composition

In this section we prove the if a metric space is close to a k -HST then it is very close to a metric space in the composition closure of a class of metric spaces with low aspect ratio.

Lemma 3.16. For any $\alpha, \beta \geq 1$, if a metric space L is α -equivalent to a $\beta\alpha$ -HST then L is $(1 + 2/\beta)$ -equivalent to a metric space in $\text{comp}_\beta(\alpha)$.

Proof. Let L be a metric space. Let $k = \beta\alpha$. Let X be a k -HST such that there is an α Lipschitz non-contractive bijection $f : L \rightarrow X$. Namely, for every $x, y \in L$, $d_L(x, y) \leq d_X(f(x), f(y)) \leq \alpha d_L(x, y)$.

Let T be the tree defining X . For a vertex $u \in T$, let T_u be the subtree of T rooted at u . Let X_u (a subspace of X) denote the HST defined by T_u and $L_u = f^{-1}(X_u)$. Then $\text{diam}(L_u) \leq \text{diam}(X_u) = \Delta(u)$.

Our goal is to build a metric space $Z \in \text{comp}_\beta(\alpha)$ along with a non-contractive Lipschitz bijection $g : L \rightarrow Z$ which satisfies for every $x, y \in L$, $d_L(x, y) \leq d_Z(g(x), g(y)) \leq \left(1 + \frac{2}{\beta}\right) d_L(x, y)$. We prove this by induction on the size of L . The inductive hypothesis needs to be further strengthened with the requirement that $\text{diam}(Z) \leq \text{diam}(L) = \Delta$.

Let r be the root of T , with $\Delta(r) = \Delta$. Let C denote the set of children of r . By induction, there exists for each child $u \in C$ a metric space $N_u \in \text{comp}_\beta(\alpha)$ and a non-contractive Lipschitz bijection

$g_u : L_u \rightarrow N_u$ which satisfies for every $x, y \in L_u$, $d_{L_u}(x, y) \leq d_{N_u}(g_u(x), g_u(y)) \leq \left(1 + \frac{2}{\beta}\right) d_{L_u}(x, y)$. Also $\text{diam}(N_u) \leq \text{diam}(L_u) = \Delta(u)$.

Define a metric space $M = (C, d_M)$ by setting for every $u \neq v \in C$,

$$d_M(u, v) = \max\{d_L(x, y); x \in L_u, y \in L_v\}.$$

Fix $u \neq v \in C$ and $x \in L_u, y \in L_v$. Since $d_X(f(x), f(y)) = \Delta$, we have that $\Delta/\alpha \leq d_L(x, y) \leq \Delta$. It follows that for every $u \neq v \in C$ and $x \in L_u, y \in L_v$,

$$\frac{\Delta}{\alpha} \leq d_L(x, y) \leq d_M(u, v) \leq \text{diam}(L) = \Delta.$$

Therefore $\Phi(M) \leq \alpha$ and $\text{diam}(M) \leq \Delta$. Also for every u, v, x, y as above,

$$\begin{aligned} d_L(x, y) &\leq d_M(u, v) \leq d_L(x, y) + \text{diam}(L_u) + \text{diam}(L_v) \\ &= d_L(x, y) + \Delta(u) + \Delta(v) \leq d_L(x, y) + 2\frac{\Delta}{k} \\ &\leq d_L(x, y) + 2\frac{\alpha d_L(x, y)}{\beta\alpha} \leq \left(1 + \frac{2}{\beta}\right) d_L(x, y). \end{aligned}$$

Now, we let

$$\gamma = \frac{\max_{u \in C} \text{diam}(N_u)}{\min_{u, v \in C} d_M(u, v)}, \quad \text{and} \quad \beta' = \frac{1}{\gamma} \geq \frac{\Delta/\alpha}{\Delta/k} = \beta.$$

Define $Z \in \text{comp}_\beta(\alpha)$, by letting $Z = M_{\beta'}[\mathcal{N}]$, where $\mathcal{N} = \{N_u\}_{u \in C}$. Also define for every $u \in C$ and $x \in X_u$, $g(x) = g_u(x)$.

Let $u, v \in C$ and $x \in L_u, y \in L_v$. When $u = v$ the bound on the distortion of g follows from our induction hypothesis. For $u \neq v$, $d_Z(g(x), g(y)) = \beta' \gamma d_M(u, v) = d_M(u, v)$, which implies the required bound on the distortion of g , and the requirement $\text{diam}(Z) \leq \Delta$. \square

3.7 Distortions arbitrarily close to 2

Our goal in this section is to prove the following theorem:

Theorem 3.26. *There is an absolute constant $c > 0$ such that for any $k \geq 1$ and $0 < \epsilon < 1$, for any integer n :*

$$R_{k\text{-HST}}(2 + \epsilon, n) \geq n^{\frac{c\epsilon}{\log(2k/\epsilon)}}.$$

In particular,

$$R_{\text{UM}}(2 + \epsilon, n) \geq n^{\frac{c\epsilon}{\log(2/\epsilon)}}.$$

By Proposition 3.4 the same bound holds for $R_2(2 + \epsilon, n)$.

As in the case of large α , we derive Theorem 3.26 from the following stronger claim.

Theorem 3.26'. *There is an absolute constant $c > 0$ such that for any $k \geq 1$ and $0 < \epsilon < 1$:*

$$\psi_k(2 + \epsilon) \geq \frac{c\epsilon}{\log(2k/\epsilon)}.$$

The proof of Theorem 3.26' uses most of the techniques developed for the case of large α . The basic idea is to first apply Theorem 3.7' to obtain some constant α' for which there is a constant bound on $\psi(\alpha')$, e.g. $1/2$. So, our goal is to find a further subspace for which we can improve the distortion from α' to $2 + \epsilon$. Again, we would like to exploit metric spaces with low aspect ratio Φ . For large α we could do this with Φ bounded with respect to α . Since now we started with some constant α' we can expect Φ to be constant as well. However, the bound of Lemma 3.14 does not apply for small values of α . Thus, our first step is to obtain meaningful lower bounds on $\psi_k(\Phi, 2 + \epsilon)$ for every $\epsilon > 0$. This is done by giving a lower bound on $\psi_{\text{EQ}}(\Phi, 2 + \epsilon)$, that is by finding a large equilateral subspace, which is a special case of a k -HST. We can now apply Lemma 3.11 to get lower bounds on $\psi_k(\text{comp}_\beta(\Phi), 2 + \epsilon)$. For large α we were able to extend such bounds by finding a subspace close to a k -HST, and therefore very close to a metric space in $\text{comp}_\beta(\Phi)$ via Lemma 3.16. In the present case, "very close" means distortion $\approx 1 + \epsilon$, which implies that k and β must be in the range of $1/\epsilon$. This is achieved by initially applying Lemma 3.23 to get a bound on $\psi_k(\alpha')$.

We begin with a proof of the bound on $\psi_k(\Phi, 2 + \epsilon)$, which is based on bounds on embedding into an equilateral space. We start with the following result:

Lemma 3.27. *Let $\alpha > 2$, $s \geq 2$ be real numbers and $t \geq 1$ be an integer. Let M be an n point metric space. Then at least one of the following two conditions holds:*

1. *M contains a subspace N of size at least s that is α -equivalent to an equilateral space.*
2. *M contains a subspace N of size at least n/s^t , such that $\text{diam}(N) < (\alpha/2)^{-t} \text{diam}(M)$.*

Proof. By induction on t . Suppose that M has no subspace of size s that is α equivalent to an equilateral space. For $t = 1$ let $N_0 = M$. For $t > 1$ we get by the induction hypothesis there a subspace $N_{t-1} \subseteq M$ which contains at least n/s^{t-1} points and $\text{diam}(N_{t-1}) \leq (\alpha/2)^{-t+1} \text{diam}(M)$.

Let $\{c_1, \dots, c_t\}$ be a maximal subset of N_{t-1} such that $d(c_i, c_j) \geq \text{diam}(N_{t-1})/\alpha$ for $i \neq j$. Since $\{c_1, \dots, c_t\}$ is α equivalent to an equilateral space, our assumption implies that $t \leq s$. Let $C_i = N_{t-1} \cap B(c_i, \text{diam}(N_{t-1})/\alpha)$. By the maximality of t , $\cup_{i=1}^t C_i = N_{t-1}$, and so if we set N_t to be the largest C_i , we have that its cardinality is at least $|N_{t-1}|/t \geq n/s^t$. Now:

$$\text{diam}(N_t) \leq \text{diam}(B(c_i, \text{diam}(N_{t-1})/\alpha)) < \frac{2}{\alpha} \text{diam}(N_{t-1}) \leq \left(\frac{2}{\alpha}\right)^t \text{diam}(M).$$

□

This implies a bound on the cardinality of a subspace that is α -equivalent to an equilateral space.

Corollary 3.28. *Fix $\alpha > 2$ and an integer $n \geq 4$. Let M be a metric space of size n . Then,*

$$R_{\text{EQ}}(M; \alpha, n) \geq \left(\frac{n}{2}\right)^{\lceil \log_{\alpha/2} \Phi(M) \rceil^{-1}} \geq n^{\frac{1}{2} \lceil \log_{\alpha/2} \Phi(M) \rceil^{-1}}.$$

Proof. Apply Lemma 3.27 with $t = \lceil \log_{\alpha/2} \Phi(M) \rceil$ and $s = (n/2)^{1/t}$. We obtain a subspace N of M . All we have to do is verify that with these parameters the second condition in Lemma 3.27 cannot hold. Indeed, otherwise $|N| \geq n/s^t = 2$ so that $\text{diam}(N) \geq \min_{x \neq y} d_M(x, y)$, and it follows that $(\alpha/2)^t < \Phi$, which contradicts the choice of t . □

We show next that Corollary 3.28 implies bounds for the weighted Ramsey problem, and so we can bound $\psi_{\text{EQ}}(\Phi, \alpha)$ for any $\alpha > 2$. Since an equilateral is in particular a k -HST, we get a bound on $\psi_k(\Phi, \alpha)$. To obtain this we need to extend the bound in Corollary 3.28 to hold for the weighted Ramsey problem. To achieve this we make use of another lemma from [5], which is similar in flavor to Lemma 3.20:

Lemma 3.29 ([5]). *Let $x = \{x_i\}_{i=1}^{\infty}$ be a sequence of non-negative real numbers. Then there exists a sequence $y = \{y_i\}_{i=1}^{\infty}$ such that $y_i \leq x_i$ for all $i \geq 1$ and:*

$$\sum_{i \geq 1} y_i^{1/2} \geq \left(\sum_{i \geq 1} x_i \right)^{1/2}.$$

Moreover, one of the following two cases holds true:

1. For all $i > 2$, $y_i = 0$.
2. There exists $\omega > 0$ such that for all $i \geq 1$ either $y_i = \omega$ or $y_i = 0$.

Corollary 3.30. *For any $k \geq 1, \alpha > 2$ and $\Phi > 1$,*

$$\psi_k(\Phi, \alpha) \geq \psi_{\text{EQ}}(\Phi, \alpha) \geq \frac{1}{4} \left[\log_{\alpha/2} \Phi \right]^{-1}.$$

Proof. Let M be an n -point metric space with aspect ratio $\Phi(M) \leq \Phi$. Let $w : M \rightarrow \mathbb{R}^+$ be a weight function normalized so that $\sum_{x \in M} w(x) = 1$. Apply Lemma 3.29 to the sequence $\{w(x)\}_{x \in M}$ to obtain a sequence $\{w'(x)\}_{x \in M}$ such that for all $x \in M$, $w'(x) \leq w(x)$. In addition, either (i) There are $u, v \in M$ such that $w(u)^{1/2} + w(v)^{1/2} \geq w'(u)^{1/2} + w'(v)^{1/2} \geq 1$, or there is a subset $N \subseteq M$ such that for all $x \in N$, $w(x) \geq w'(x) = \omega > 0$, and $|N|\omega^{1/2} \geq 1$. In the first case the subset $\{u, v\}$ is isometric to an equilateral space and we are done. In the second case, if $|N| \leq 4$ then $\omega \geq 1/16$. Hence, we can choose two points $u', v' \in N$ such that $w(u')^{1/4} + w(v')^{1/4} \geq 2\omega^{1/4} \geq 1$. Again, $\{u', v'\}$ is isometric to an equilateral space. Otherwise, $|N| > 4$ and by Corollary 3.28 there is a subspace $N' \subseteq N$ which is α -equivalent to an equilateral space and $|N'| \geq |N|^{\frac{1}{2}} \left[\log_{\alpha/2} \Phi \right]^{-1}$. Hence:

$$\begin{aligned} \sum_{x \in N'} w(x)^{\frac{1}{4}} \left[\log_{\alpha/2} \Phi \right]^{-1} &\geq |N|^{\frac{1}{2}} \left[\log_{\alpha/2} \Phi \right]^{-1} \omega^{\frac{1}{4}} \left[\log_{\alpha/2} \Phi \right]^{-1} \\ &= (|N|\omega^{1/2})^{\frac{1}{2}} \left[\log_{\alpha/2} \Phi \right]^{-1} \geq 1. \end{aligned}$$

□

Proof of Theorem 3.26'. Theorem 3.7' implies in particular that there is a constant θ for which $\psi(\theta/2) \geq 1/2$. In other words, given a metric space X , there exists a subspace X' of X which is $(\theta/2)$ -equivalent to an ultrametric Y and satisfies the weighted Ramsey condition (*) with $\psi = 1/2$.

Let $\beta = 8k/\epsilon$ and $k' = \theta\beta$. It follows from Lemma 3.23 that Y contains a subspace Y' which is 2-equivalent to a k' -HST and satisfies condition (*) with $\psi_{k'}(\text{UM}, 2) \geq \lceil \log k' \rceil^{-1}$. By mapping X into an ultrametric Y and its image in Y into a k' -HST, we can apply Lemma 3.10 to obtain a subspace X'' of X that is $(\theta/2) \cdot 2 = \theta$ -equivalent to a k' -HST, and satisfies condition (*) with

$$\psi_{k'}(\theta) \geq \psi_{k'}(\text{UM}, 2) \cdot \psi(\theta/2) \geq \frac{1}{2 \lceil \log k' \rceil}.$$

Now, X'' is θ -equivalent to a $\theta\beta$ -HST and so Lemma 3.16 implies that it is $(1 + 2/\beta)$ -equivalent to a metric space Z in $\text{comp}_\beta(\theta)$. Therefore

$$\psi_{\text{comp}_\beta(\theta)}(1 + 2/\beta) \geq \psi_{k'}(\theta) \geq \frac{1}{2\lceil \log k' \rceil}.$$

Additionally, using Lemma 3.11 and the bound of Lemma 3.30, we have that there is a constant c' such that

$$\psi_k\left(\text{comp}_\beta(\theta), 2 + \frac{\epsilon}{4}\right) = \psi_k\left(\theta, 2 + \frac{\epsilon}{4}\right) \geq \frac{c'\epsilon}{\log \theta}.$$

It follows that Z contains a subspace Z' which is $(2 + \epsilon/4)$ -equivalent to a k -HST.

By mapping X into $Z \in \text{comp}_\beta(\theta)$, and then its image in Z into a k -HST, we can apply Lemma 3.10 to obtain a subspace of X that is $(2 + \epsilon/4)(1 + 2/\beta) \leq (2 + \epsilon)$ -equivalent to a k -HST and which satisfies the weighted Ramsey condition $(*)$ with

$$\psi_k(2 + \epsilon) \geq \psi_k\left(\text{comp}_\beta(\theta), 2 + \frac{\epsilon}{4}\right) \cdot \psi_{\text{comp}_\beta(\theta)}\left(1 + \frac{2}{\beta}\right) \geq \frac{c'\epsilon}{2\log \theta \lceil \log(8\theta k/\epsilon) \rceil},$$

which implies the theorem by an appropriate choice of c . □

4 Dimensionality Based Upper Bounds

In this section we prove some upper bounds on the Ramsey function of low dimensional spaces. In particular, these imply bounds on the Euclidean Ramsey function $R_2(\alpha, n)$. In addition, these bounds show that the lower bounds for low dimensional ℓ_p spaces from Corollary 1.7 in the introduction are nearly tight. Our upper bounds on $R_2(\alpha, n)$ for $\alpha < 2$ improve the results of [17] by showing that for any $\alpha < 2$, $R_2(\alpha, n) \leq 2\log_2 n + C(\alpha)$. The bounds obtained on $R_2(\alpha, n)$ for $2 < \alpha \leq \log n / \log \log n$, are also possibly tight.

The proof technique we employ here originates from a counting argument by Bourgain [15] and later variants (see [42]). A different argument, based on geometric considerations, uses expander graphs. Expander graphs, in fact yield the best upper bound we have on $R_p(\alpha, n)$ for $\alpha \geq 2$ and all $p \geq 1$. This is shown in Section 5.

In this section we prove the following bounds:

Theorem 4.1. *Let X be an h -dimensional normed space and n be an integer. Then*

- For any $1 < \alpha < 2$, $R_X(\alpha, n) \leq 2\log_2 n + 2h\log_2\left(\frac{C}{2-\alpha}\right)$.
- For any $\alpha \geq 2$, $R_X(\alpha, n) \leq Cn^{1-c/\alpha}h\log \alpha$

where $c, C > 0$ are some absolute constants.

Using the Johnson-Lindenstrauss dimension reduction Lemma [31] we derive the following bounds for the Euclidean Ramsey function $R_2(\alpha, n)$.

Corollary 4.2. *There are absolute constants $c, C > 0$ such that for every integer n ,*

- For any $0 < \epsilon \leq 1$, $R_2(2 - \epsilon, n) \leq 2\log_2 n + C\frac{\log^2(2/\epsilon)}{\epsilon^2}$.
- For any $2 \leq \alpha \leq \frac{\log n}{\log \log n}$, $R_2(\alpha, n) \leq Cn^{1-c/\alpha}$.

The counting argument presented below is based on the existence of dense graphs for which all metrics defined on subgraphs are very far from each other.

Let A be a set of vertices in the graph $G = (V, E)$. We denote by E_A the set of edges in G with both vertices in A , and the cardinality of E_A by e_A . The *density* of A is $\frac{e_A}{\binom{|A|}{2}}$.

We first explain the relevance of large girth and high density to our problem. Let $G = (V, E)$ be a large graph of girth g with no large sparse subgraphs. With every $H \subseteq E$ we associate a metric on V defined by $\rho_H(u, v) = \min\{g - 1, d_H(u, v)\}$, where d_H is the shortest path metric in the subgraph of G with edge set H . Below we show that among these ρ_H are metrics that cannot be embedded with small distortion in any low-dimensional normed space.

Lemma 4.3. *If there exists a graph $G = (V, E)$ of size n with girth at least g , in which every set of $\geq s$ vertices has density at least q , then for every h -dimensional normed space X and every real $1 \leq \alpha < g - 1$,*

$$R_X(\alpha, n) \leq \max \left\{ s, \frac{2}{q} \left[h \log_2 \left(\frac{14\alpha g}{g - \alpha - 1} \right) + \log_2 \left(\frac{n}{s} \right) \right] \right\}$$

Proof. To prove the theorem, we may certainly assume that $R_X(\alpha, n) \geq s$. Let $k = R_X(\alpha, n)$, namely, every n point metric space contains a subset of size k that α -embeds in X . In particular, for every $H \subseteq E$ there is a set of k vertices $A_H \subseteq V$ such that (A_H, ρ_H) α -embeds into X . Therefore there is a certain set A of k vertices, that is suitable for many sets $H \subseteq E$. Namely, there is a class \mathcal{H} of at least $2^{|E|} / \binom{n}{k}$ subsets $H \subseteq E$ for which $A = A_H$, and therefore (A, ρ_H) α -embeds into X . Consider $H_1, H_2 \in \mathcal{H}$ equivalent if $H_1 \cap E_A = H_2 \cap E_A$. There are at most $2^{|E| - e_A}$ members in \mathcal{H} that are equivalent to a given set H . Consequently, there are at least $2^{e_A} / \binom{n}{k}$ subsets $H \subseteq E$ which are mutually inequivalent and for which (A, ρ_H) α -embeds into X . Let $f_H : A \rightarrow X$ be such an embedding, i.e., for every $u, v \in A$:

$$\frac{1}{\alpha} \rho_H(u, v) \leq \|f_H(u) - f_H(v)\|_X \leq \rho_H(u, v).$$

Since ρ_H takes values in $\{0, 1, \dots, g - 1\}$, by applying an appropriate translation we may assume that $f_H(A) \subseteq B_X(0, g)$. We now "round" the images $f_H(A)$ to the points of a δ -net in $B_X(0, g)$, where δ will be determined soon. Let \mathcal{N} be a δ -net of $B_X(0, g)$, and define $\phi_H(v)$ to be the closest point in \mathcal{N} to $f_H(v)$. We claim that if $H_1, H_2 \subseteq E$ are inequivalent, i.e., $H_1 \cap E_A \neq H_2 \cap E_A$, then $\phi_{H_1} \neq \phi_{H_2}$. Indeed, we may assume that there are $u, v \in A$ such that $(u, v) \in H_1 \setminus H_2$. Since the girth of G is at least g , this implies that $\rho_{H_2}(u, v) = g - 1$, whereas $\rho_{H_1}(u, v) = 1$. Now, if $\phi_{H_1}(u) = \phi_{H_2}(u)$ and $\phi_{H_1}(v) = \phi_{H_2}(v)$ then:

$$\begin{aligned} \frac{g - 1}{\alpha} &= \frac{\rho_{H_2}(u, v)}{\alpha} \\ &\leq \|f_{H_2}(u) - f_{H_2}(v)\|_X \leq 2\delta + \|\phi_{H_2}(u) - \phi_{H_2}(v)\|_X \\ &= 2\delta + \|\phi_{H_1}(u) - \phi_{H_1}(v)\|_X \leq 4\delta + \|f_{H_1}(u) - f_{H_1}(v)\|_X \\ &\leq 4\delta + \rho_{H_1}(u, v) = 4\delta + 1. \end{aligned}$$

We select $\delta = \frac{g - \alpha - 1}{5\alpha}$ so this becomes a contradiction. It follows that each of the aforementioned $2^{e_A} / \binom{n}{k}$ inequivalent sets $H \in \mathcal{H}$ gives rise to a distinct function $\phi_H : A \rightarrow \mathcal{N}$.

By standard volume estimates, $|\mathcal{N}| \leq \left(\frac{2g}{\delta}\right)^h$. Hence there are at most $|\mathcal{N}|^k \leq \left(\frac{2g}{\delta}\right)^{kh}$ distinct functions from A to \mathcal{N} . Consequently,

$$\left(\frac{2g}{\delta}\right)^{kh} \geq \frac{2^{e_A}}{\binom{n}{k}} \geq \frac{2^q \binom{k}{2}}{\binom{n}{k}}.$$

By estimating $\binom{n}{k} \leq \left(\frac{ne}{k}\right)^k \leq \left(\frac{ne}{s}\right)^k$ we have that

$$h \log_2 \left(\frac{2g}{\delta} \right) \geq \frac{(k+1)q}{2} - \log_2 \left(\frac{en}{s} \right),$$

which yields the claimed bound on k . \square

Such graphs do exist as we now show:

Lemma 4.4. *For every integer $g \geq 4$, there exist graphs $G = (V, E)$ of arbitrarily large order n and girth at least g in which every set $A \subseteq V$ of cardinality at least $n^{1-\frac{1}{8g}}$ has density at least $n^{-1+\frac{1}{2g}}$.*

Proof. This is a standard construction from random graph theory. Let $N \geq C^g$ be an arbitrarily large integer, where C is an appropriately chosen constant. Let $\eta = \frac{1}{4g}$. Pick a random graph in $G(N, p)$ where $p = 2 \cdot N^{-1+2\eta}$. We claim that (i) With probability $\geq \frac{1}{2}$ this graph has fewer than $\frac{N}{2}$ cycles of length $< g$, and (ii) With almost certainty, every set of cardinality $\geq \left(\frac{N}{2}\right)^{1-\frac{\eta}{2}}$ has density $\geq \left(\frac{N}{2}\right)^{-1+2\eta}$. The theorem now follows by taking a graph with these two properties and removing $\frac{N}{2}$ vertices, including at least one vertex from each cycle of length $< g$. The resulting graph has $\frac{N}{2}$ vertices, it has no short cycles, and satisfies the density condition.

The expected number of cycles of length $< g$ is

$$\sum_{i=3}^{g-1} \frac{1}{2^i} p^i N(N-1) \dots (N-i+1) \leq \frac{1}{6} \sum_{i=3}^{g-1} (pN)^i \leq (pN)^{g-1} = \left(2 \cdot N^{\frac{1}{2g}}\right)^{g-1} \leq \frac{N}{4}.$$

In the last inequalities we use the facts that $pN \geq 2$, $N \geq C^g$, and $C \geq 4$. It follows that with probability $\geq \frac{1}{2}$, there are no more than $N/2$ cycles shorter than g .

The expected density in every set of vertices is, of course, p . To estimate the deviation, we use the Chernoff bound:

$$\Pr[e_A \leq \frac{1}{2} \binom{|A|}{2} p] \leq e^{-\binom{|A|}{2} \frac{p}{8}}.$$

Thus, the probability that there exists a set of cardinality $\geq k$ and density $\leq p/2$ does not exceed $2^N \exp(-\binom{k}{2} p/8)$. For $k = \left(\frac{N}{2}\right)^{1-\frac{\eta}{2}}$, $p = 2 \cdot N^{-1+2\eta}$, and assuming $N \geq C^g$, this easily seen to be $o(1)$. The claim follows. \square

Theorem 4.1 now follows easily:

Proof of Theorem 4.1. The claim for $\alpha \geq 2$ is obtained combining Lemma 4.3 and Lemma 4.4 with $g = \lceil \alpha + 2 \rceil$. As $1/q \leq s \leq n^{1-\frac{1}{8\alpha}}$, we obtain that for an appropriate choice of constant C , $R_X(\alpha, n) \leq C n^{1-\frac{1}{8\alpha}} (h \log \alpha + \log(n^{\frac{1}{8\alpha}}))$. We choose $c = 1/16$, so that $n^{1-c/\alpha} = n^{1-\frac{1}{16\alpha}} \geq n^{1-\frac{1}{8\alpha}} \log(n^{\frac{1}{8\alpha}})$. The claim for $\alpha \geq 2$ now follows.

For the case $\alpha < 2$ we use, instead of Lemma 4.4, the (trivial) analogous statement for the complete graph K_n . I.e., we apply Lemma 4.3 with $g = 3$, $s = 2$, and $q = 1$. \square

We are now ready to prove the promised upper bounds on $R_2(\alpha, n)$.

Proof of Corollary 4.2. The result follows using the Johnson-Lindenstrauss dimension reduction lemma [31] for ℓ_2 . Let $\alpha \geq 1$, and let $k = R_2(\alpha, n)$, i.e., every n -point metric space M contains a k -point subspace that α -embeds into ℓ_2 . By [31], for any $0 < \delta \leq 1$, this subspace $\alpha(1+\delta)$ -embeds

into ℓ_2^h with $h \leq \frac{C \log k}{\delta^2}$. Hence, $R_2(\alpha, n) \leq R_{\ell_2^h}(\alpha(1 + \delta), n)$. Now apply Theorem 4.1. The claim for $\alpha \geq 2$ follows by taking $\delta = 1$. For $\alpha = 2 - \epsilon$ we set $\delta = \epsilon/4$. Then

$$k = R_2(2 - \epsilon, n) \leq R_{\ell_2^h} \left(2 - \frac{\epsilon}{2}, n \right) \leq 2 \log_2 n + \frac{C \log k \log \left(\frac{2}{\epsilon} \right)}{\epsilon^2},$$

which implies the bound in the proposition. \square

Another interesting consequence of Theorem 4.1 are upper bounds for metrics defined by planar graphs. This may be interesting in view of the fact that the target metrics in our lower bounds in Section 3 are ultrametrics (and thus planar).

Theorem 4.5. *Let \mathcal{F} be a family of graphs, none of which contains a fixed minor H on r vertices. Then for every integer n and every $\alpha \geq 1$:*

$$R_{\mathcal{F}}(\alpha, n) \leq Cr^3 n^{1-c/\alpha} \log^2 n \log \alpha,$$

where $c, C > 0$ are universal constants.

Proof. It is implicit in [50] that the Euclidean embedding that Rao constructed is also a good low dimensional embedding into ℓ_∞ . More precisely, if $F \in \mathcal{F}$ is a graph on n points then it embeds with distortion C into ℓ_∞^h , with $h \leq Cr^3 \log^2 n$. Thus $R_{\mathcal{F}}(\alpha, n) \leq R_{\ell_\infty^h}(C\alpha, n)$. The result now follows from Theorem 4.1. \square

5 Expanders and Poincaré Inequalities

In this section we prove lower bounds for the metric Ramsey function in the case of expanders. The proof is based on generalizations of Poincaré inequalities used by Matoušek to prove lower bounds on the Euclidean distortion of expander graphs. To obtain these inequalities we pass to a power of the graph and delete vertices with small degree. The argument shows that large subsets of expanders contain large sub-subsets which satisfy an appropriate Poincaré inequality (see Lemma 5.4 below). First, we recall some basic concepts on graphs.

Let $G = (V, E)$ be a d -regular graph, and let A be its adjacency matrix, i.e. $A_{uv} = 1$ if $[u, v] \in E$ and $A_{uv} = 0$ otherwise. Let $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ be the eigenvalues of A . It is easy to observe that $\lambda_1 = d$. Also, $\text{trace}(A) = 0$, so $\lambda_n < 0$. We occasionally write $\lambda_i(G)$ to specify that G is the graph under consideration. We define G 's multiplicative spectral gap as:

$$\gamma(G) = \frac{\lambda_2(G)}{d}.$$

We also define the absolute multiplicative spectral gap of G as:

$$\gamma_+(G) = \frac{\max_{i \geq 2} |\lambda_i(G)|}{d} = \frac{\max\{\lambda_2(G), -\lambda_n(G)\}}{d}.$$

In what follows we will use the following standard estimate: $\gamma_+(G) \geq 1/d$. To verify it observe that $nd = \text{trace}(A^2) = \sum_{i=1}^n \lambda_i(G)^2 \leq d^2 + (n-1)[d\gamma_+(G)]^2$ and use the fact that $1 \leq d \leq n-1$. We remark that this elementary bound is weaker than the Alon-Boppana bound [1], but it is sufficient for our purposes, and hold for all d (while the Alon-Boppana bound only holds for small enough d).

The main statement of this section is:

Theorem 5.1. *Let $G = (V, E)$ be a d -regular graph, $d \geq 3$. Let $\gamma = \gamma_+(G)$. Then for every $p, \alpha \geq 1$:*

$$R_p(G; \alpha) \leq Cd|V|^{1-c\frac{\log_d(1/\gamma)}{p\alpha}}.$$

Where C, c are absolute constants.

Given $S, T \subseteq V$, we denote by $E(S, T)$ the set of directed edges between vertices in S and T , i.e.

$$E(S, T) = \{(u, v) \in S \times T; [u, v] \in E\}.$$

We also denote by $E(S)$ the set of edges in the subgraph induced by G on S , i.e.

$$E(S) = \{\{u, v\}; u, v \in S, [u, v] \in E\}.$$

With this notation, $|E(S)| = \frac{|E(S, S)|}{2}$.

The ‘‘Expander Mixing Lemma’’ [1] states that:

Lemma 5.2 (Expander Mixing Lemma). *Let $G = (V, E)$ be a d -regular graph (which may have loops and/or parallel edges). Then for every $S, T \subseteq V$,*

$$\left| |E(S, T)| - \frac{d|S||T|}{|V|} \right| \leq \gamma_+(G)d\sqrt{|S||T|}.$$

In particular,

$$\left| \frac{2|E(S)|}{|S|} - \frac{|S|}{|V|}d \right| \leq \gamma_+(G)d.$$

Lemma 5.3. *Let $G = (V, E)$ be a d -regular graph, $d \geq 3$. Let $\gamma = \gamma_+(G)$. Then for any $B \subset V$ satisfying $|B| \geq 8\gamma|V|$, there exists $C \subset B$ such that $|C| \geq |B|/3$, and for any $u \in C$,*

$$d\frac{|B|}{8|V|} \leq \deg_C(u) \leq d\frac{4|B|}{|V|}.$$

Proof. Denote $k = |B|$. By the expander mixing lemma,

$$|E(B)| \leq \frac{dk^2}{2n} \left(1 + \frac{1}{8} \right) \leq \frac{dk^2}{n}.$$

Set $B' = \{v \in B; \deg_B(v) \leq (4dk)/n\}$. Since the graph induced by G on B contains $k - |B'|$ vertices of degree greater than $(4dk)/n$, it follows that:

$$\frac{dk^2}{n} \geq |E(B)| \geq \frac{k - |B'|}{2} \cdot \frac{4dk}{n},$$

so that $|B'| \geq k/2$. Again, by the expander mixing lemma,

$$\frac{2|E(B')|}{k} \geq \frac{2|E(B')|}{2|B'|} \geq \frac{dk}{4n}.$$

We now apply an iterative procedure which produces a sequence $B' = B_0 \supset B_1 \supset B_2 \supset \dots$ as follows: if $\min_{v \in B_i} \deg_{B_i}(v) \leq \frac{dk}{8n}$ then B_{i+1} is obtained from B_i by throwing away a vertex $u \in B_i$

with $\deg_{B_i}(u) = \min_{v \in B_i} \deg_{B_i}(v)$. Otherwise $B_{i+1} = B_i$. This procedure eventually ends, and we are left with a subset $C \subseteq B'$. Since at each step we delete at most $\frac{dk}{8n}$ edges from B' , we have that:

$$|E(C)| \geq |E(B')| - \frac{dk^2}{8n} \geq \frac{dk^2}{4n} - \frac{dk^2}{8n} = \frac{dk^2}{8n}.$$

Note that the graph induced on C by G has minimal degree at least $\frac{dk}{8n}$. To estimate $|C|$ we apply the expander mixing lemma to get that:

$$\frac{2|E(C)|}{|C|} \leq d \left(\frac{|C|}{n} + \gamma \right).$$

Thus

$$\frac{d|C|^2}{2n} + \frac{\gamma d|C|}{2} \geq |E(C)| \geq \frac{dk^2}{8n}.$$

Since $k \geq 8\gamma n$, $\frac{\gamma d|C|}{2} \leq \frac{dk^2}{16n}$. Hence:

$$\frac{d|C|^2}{2n} \geq \frac{dk^2}{16n},$$

so that $|C| \geq \frac{k}{3}$. □

The following is the Poincaré inequality used in the proof of Theorem 5.1.

Lemma 5.4. *Let $G = (V, E)$ be a d -regular graph, $d \geq 3$. Let $\gamma = \gamma_+(G)$. Then for any $B \subset V$ satisfying $|B| \geq 8\gamma|V|$, there exists $C \subset B$ such that $|C| \geq |B|/3$ and the following holds true: For any $p \geq 1$ and for every $f : C \rightarrow \ell_p$:*

$$\sum_{u,v \in C} \|f(u) - f(v)\|_p^p \leq \frac{(32p)^p |V|}{d} \sum_{[u,v] \in E(C)} \|f(u) - f(v)\|_p^p.$$

The proof of Lemma 5.4 proceeds by first proving a slightly stringer version of it for $p = 2$ and then extrapolating to the general case via the following lemma based on an extrapolation argument which was used in [41]. Its proof of it is delayed to the end of the section.

Lemma 5.5 (Extrapolation lemma for Poincaré inequalities). *Let $G = (V, E)$ be a graph with maximal degree at most Δ . Fix $p \geq 1$ and let $A > 0$ be a constant such that for every $f : V \rightarrow \mathbb{R}$,*

$$\sum_{u,v \in V} |f(u) - f(v)|^p \leq (Ap)^p \frac{|V|}{\Delta} \sum_{[u,v] \in E} |f(u) - f(v)|^p. \quad (5)$$

Then for every $0 < q \leq p$ and for every $f : V \rightarrow \ell_q$,

$$\sum_{u,v \in V} \|f(u) - f(v)\|_q^q \leq (Ap)^p \frac{|V|}{\Delta} \sum_{[u,v] \in E} \|f(u) - f(v)\|_q^q.$$

Additionally, for every $p < q < \infty$ and every $f : V \rightarrow \ell_q$:

$$\sum_{u,v \in V} \|f(u) - f(v)\|_q^q \leq (4Aq)^q \frac{|V|}{\Delta} \sum_{[u,v] \in E} \|f(u) - f(v)\|_q^q.$$

Proof of Lemma 5.4. Denote $n = |V|$ and $k = |B|$. By Lemma 5.3, there exists $C \subset B$ with $|C| \geq k/3$, such that the induced subgraph of G on C , has minimal degree at least $kd/8n$ and maximal degree at most $\Delta = 4kd/n$.

We first prove that the following inequality hold true for every $f : V \rightarrow \ell_2$:

$$\sum_{u,v \in C} \|f(u) - f(v)\|_2^2 \leq \frac{32n}{d} \sum_{[u,v] \in E(C)} \|f(u) - f(v)\|_2^2 = \frac{32 \cdot 4k}{\Delta} \sum_{[u,v] \in E(C)} \|f(u) - f(v)\|_2^2. \quad (6)$$

By summation we may clearly assume that $f : C \rightarrow \mathbb{R}$. By translation we may assume that $\sum_{v \in C} f(v) = 0$. Extend f to V by letting $f(u) = 0$ for $u \notin C$. Now,

$$\begin{aligned} \sum_{u,v \in C} [f(u) - f(v)]^2 &= 2(|C| + 1) \sum_{v \in C} f(v)^2 - 2 \sum_{u,v \in C} f(u)f(v) \\ &= 2(|C| + 1) \sum_{v \in V} f(v)^2 - \left(\sum_{v \in C} f(v) \right)^2 \\ &= 2(|C| + 1) \sum_{v \in V} f(v)^2 \leq 4k \sum_{v \in V} f(v)^2. \end{aligned}$$

Since $\sum_{v \in V} f(v) = 0$ we can use the spectral gap of H to get that:

$$\begin{aligned} \sum_{[u,v] \in E(C)} [f(u) - f(v)]^2 &= 2 \sum_{v \in C} \deg_{E(C)}(v) f(v)^2 - 2 \sum_{[u,v] \in E(C)} f(u)f(v) \\ &\geq 2 \sum_{v \in V} \frac{dk}{8n} f(v)^2 - 2 \sum_{[u,v] \in E} f(u)f(v) \\ &= \frac{dk}{4n} \sum_{v \in V} f(v)^2 - \langle A^t f, f \rangle \\ &\geq \frac{dk}{4n} \sum_{v \in V} f(v)^2 - \gamma d \sum_{v \in V} f(v)^2 \\ &\geq \left(\frac{k}{4n} - \gamma \right) d \sum_{v \in V} f(v)^2 \geq \frac{k}{8n} d \sum_{v \in V} f(v)^2, \end{aligned}$$

which implies inequality (6). The Poincaré inequalities for $p \geq 1$ now follow immediately from inequality (6) via Lemma 5.5, and by substituting the value of Δ . \square

Proof of Theorem 5.1. The proof proceeds by showing that for every $B \subseteq V$ satisfying $c_p(B) \leq \alpha$,

$$|B| \leq 100d|V|^{1 - \frac{\log_d(1/\gamma)}{2561p\alpha}}.$$

Set $k = |B|$ and $n = |V|$. Define:

$$t = \left\lfloor \frac{\log(8n)}{2560p\alpha \log d + \log\left(\frac{1}{\gamma}\right)} \right\rfloor$$

Note that $t \leq \text{diam}(G)$ since it is well known that $\text{diam}(G) \geq \log_d(n)$. We may also assume that $t \geq 1$, since otherwise, using the fact that $1/\gamma \leq d$, we get that $n < e^{2561p\alpha \log d}$, in which case the required result holds vacuously.

Denote by A the adjacency matrix of G . Let H be the multi-graph with adjacency matrix A^t . In other words, the number of H -edges between two vertices in V is the number of distinct paths of length t joining them (and it is 0 if no such path exists). The multi-graph H is d^t regular and by the spectral theorem, $\gamma_+(H) = \gamma^t$. We may assume that $k \geq 8\gamma^t n$ (otherwise the conclusion of the theorem is trivial).

It follows from Lemma 5.4 that there exists $C \subset B$ such that $|C| \geq k/3$, and for every $f : C \rightarrow \ell_p$:

$$\sum_{u,v \in C} \|f(u) - f(v)\|_p^p \leq \frac{(32p)^p n}{d^t} \sum_{[u,v] \in E_H(C)} \|f(u) - f(v)\|_p^p.$$

Let $f : B \rightarrow \ell_p$ be an embedding such that for all $u, v \in B$, $\frac{d_G(u,v)}{\alpha} \leq \|f(u) - f(v)\|_p \leq d_G(u, v)$. Then:

$$\sum_{[u,v] \in E_H(C)} \|f(u) - f(v)\|_p^p \leq \sum_{[u,v] \in E_H(B)} d_G(u, v)^p \leq |E_H(B)| t^p \leq \frac{d^t k^2 t^p}{n},$$

where the last inequality follows from an application of the expander mixing lemma:

$$|E_H(B)| \leq \frac{d^t k^2}{2n} \left(1 + \frac{1}{8}\right) \leq \frac{d^t k^2}{n}.$$

Let $s = \lfloor \log_d \left(\frac{k}{12}\right) \rfloor$. We may clearly assume that $s > 1$. The number of vertices of distance at most s from a given vertex $v_0 \in G$ is bounded by:

$$1 + d + \dots + d^s \leq 2d^s \leq \frac{k}{6} \leq \frac{|C|}{2}.$$

Hence:

$$\sum_{u,v \in C} \|f(u) - f(v)\|_p^p \geq \frac{1}{\alpha^p} \sum_{u,v \in C} d_G(u, v)^p \geq \frac{|C|^2 s^p}{2\alpha^p} \geq \frac{k^2 s^p}{18\alpha^p} \geq \frac{k^2}{80\alpha^p} \left[\log_d \left(\frac{k}{12}\right) \right]^p. \quad (7)$$

Plugging this into the Poincaré inequality we get that:

$$\frac{k^2}{80\alpha^p} \left[\log_d \left(\frac{k}{12}\right) \right]^p \leq \frac{(32p)^p n}{d^t} \cdot \frac{d^t k^2 t^p}{n},$$

which gives

$$\log_d \left(\frac{k}{12}\right) \leq 2560p\alpha t \implies k \leq 12d^{2560p\alpha t}.$$

Since $t \leq \frac{\log(8n)}{2560p\alpha \log d + \log\left(\frac{1}{\gamma}\right)}$, it follows that:

$$k \leq 12 \exp \left[\frac{\log(8n) \cdot 2560p\alpha \log d}{2560p\alpha \log d + \log\left(\frac{1}{\gamma}\right)} \right] \leq 100n^{1 - \frac{\log(1/\gamma)}{2560p\alpha \log d + \log(1/\gamma)}} \leq 100n^{1 - \frac{\log_d(1/\gamma)}{2561p\alpha}},$$

where we have used once more the estimate $\log(1/\gamma) \leq \log d$. □

It remains to prove Lemma 5.5.

Proof of Lemma 5.5. The case $0 < q \leq p$ is simple. Coordinate-wise summation of (5) shows that for every $f : V \rightarrow \ell_p$:

$$\sum_{u,v \in V} \|f(u) - f(v)\|_p^p \leq (Ap)^p \frac{|V|}{\Delta} \sum_{[u,v] \in E} \|f(u) - f(v)\|_p^p.$$

Since ℓ_2 is isometric to a subspace of L_p , it follows that for every $f : V \rightarrow \ell_2$,

$$\sum_{u,v \in V} \|f(u) - f(v)\|_2^p \leq (Ap)^p \frac{|V|}{\Delta} \sum_{[u,v] \in E} \|f(u) - f(v)\|_2^p.$$

Since $(\mathbb{R}, |x - y|^{q/p})$ is isometric to a subset of ℓ_2 ([52, 22]), the required inequality follows.

We now pass to the case $p < q$. In this case the following standard numerical inequality holds true for every $a, b \in \mathbb{R}$ (see Lemma 4 in [41]):

$$\left| |a|^{q/p} \text{sign}(a) - |b|^{q/p} \text{sign}(b) \right| \leq \frac{q}{p} |a - b| \left(|a|^{\frac{q}{p}-1} + |b|^{\frac{q}{p}-1} \right). \quad (8)$$

It suffices to prove the claims coordinate-wise, i.e. for functions $f : V \rightarrow \mathbb{R}$. Fix some $f : V \rightarrow \mathbb{R}$. By continuity there is some $c \in \mathbb{R}$ such that:

$$\sum_{v \in V} |f(v) + c|^{q/p} \text{sign}(f(v) + c) = 0.$$

Hence, by replacing f with $f + c$ we may assume that:

$$\sum_{v \in V} |f(v)|^{q/p} \text{sign}(f(v)) = 0.$$

Now:

$$\begin{aligned} \sum_{v \in V} |f(v)|^q &= \\ &= \sum_{v \in V} \left| |f(v)|^{q/p} \text{sign}(f(v)) - \frac{1}{|V|} \sum_{u \in V} |f(u)|^{q/p} \text{sign}(f(u)) \right|^p \\ &\leq \frac{1}{|V|} \sum_{u,v \in V} \left| |f(u)|^{q/p} \text{sign}(f(u)) - |f(v)|^{q/p} \text{sign}(f(v)) \right|^p \\ &\leq \frac{(Ap)^p}{\Delta} \sum_{[u,v] \in E} \left| |f(u)|^{q/p} \text{sign}(f(u)) - |f(v)|^{q/p} \text{sign}(f(v)) \right|^p \\ &\leq \frac{(Aq)^p}{\Delta} \sum_{[u,v] \in E} |f(u) - f(v)|^p \left(|f(u)|^{\frac{q}{p}-1} + |f(v)|^{\frac{q}{p}-1} \right)^p, \end{aligned}$$

where in the last two steps we have used (5) and (8), respectively. An application of Hölder's inequality gives that:

$$\begin{aligned} \sum_{[u,v] \in E} |f(u) - f(v)|^p \cdot \left(|f(u)|^{\frac{q}{p}-1} + |f(v)|^{\frac{q}{p}-1} \right)^p \\ \leq \left(\sum_{[u,v] \in E} |f(u) - f(v)|^q \right)^{p/q} \left(\sum_{[u,v] \in E} \left(|f(u)|^{\frac{q}{p}-1} + |f(v)|^{\frac{q}{p}-1} \right)^{\frac{pq}{q-p}} \right)^{1-\frac{p}{q}}. \end{aligned}$$

Using the assumption on the maximal degree we get that:

$$\sum_{[u,v] \in E} \left(|f(u)|^{\frac{q}{p}-1} + |f(v)|^{\frac{q}{p}-1} \right)^{\frac{pq}{q-p}} \leq 2^{\frac{qp}{q-p}-1} \sum_{[u,v] \in E} (|f(u)|^q + |f(v)|^q) \leq 2^{\frac{qp}{q-p}} \Delta \sum_{v \in V} |f(v)|^q.$$

Summarizing, we have shown that:

$$\sum_{v \in V} |f(v)|^q \leq \frac{(2Aq)^p}{\Delta^{\frac{p}{q}}} \left(\sum_{[u,v] \in E} |f(u) - f(v)|^q \right)^{p/q} \left(\sum_{v \in V} |f(v)|^q \right)^{1-\frac{p}{q}}.$$

This inequality simplifies to:

$$\sum_{v \in V} |f(v)|^q \leq \frac{(2Aq)^q}{\Delta} \sum_{[u,v] \in E} |f(u) - f(v)|^q.$$

We conclude by noting that:

$$\sum_{u,v \in V} |f(u) - f(v)|^q \leq 2^q |V| \sum_{v \in V} |f(v)|^q.$$

□

We now show that the interplay between the Euclidean distortion and the cardinality in Theorem 5.1, for $p = 2$, is tight, up to the dependence on d and γ . We require an upper estimate for the diameter of an n -point expander. It is well known that the diameter is $O(\log n)$, but here we will be a little bit more accurate.

We need the following bound on the diameter of expander graphs [20]:

Proposition 5.6. *Let $G = (V, E)$ be a n -vertex, d -regular graph. Denote $\gamma = \gamma_+(G)$. Then the diameter of G is at most $\log_{1/\gamma} n + 1$.*

Proposition 5.7. *Let $G = (V, E)$ be a n -vertex, d -regular graph, $d \geq 3$, and set $\gamma = \gamma_+(G)$. Then, there is an absolute constant $C > 0$ such that for any $\alpha > 1$,*

$$R_2(G; \alpha) \geq R_{\text{EQ}}(G; \alpha) \geq n^{1-\frac{C}{\alpha \log_d(1/\gamma)}}.$$

Proof. Iteratively, extract a point $x \in V$ together with a ball of radius $r = \text{diam}(G)/\alpha$ around x . Each such ball contains at most $d + d(d-1) + \dots + d(d-1)^{\lfloor r \rfloor} \leq 3(d-1)^{r+1}$ points, and thus we can repeat this process at least $n/(3(d-1)^{r+1})$ time, and get the desired set. Its size is at least

$$\frac{n}{3(d-1)^{r+1}} = \frac{n}{3} (d-1)^{-\left(\frac{\text{diam}(G)}{\alpha} + 1\right)} \geq \frac{n}{3(d-1)^2} (d-1)^{-\frac{\log_{1/\gamma} n}{\alpha}} = \frac{1}{3(d-1)^2} n^{1-\frac{1}{\alpha \log(d-1)(1/\gamma)}}.$$

□

6 Markov Type, Girth and Hypercubes

Markov type was defined in [2] and was applied in [38] to obtain lower bounds for the Euclidean distortion of regular graphs with large girth. This concept plays a key role in our analysis of the metric Ramsey problem for the discrete cube and graphs with large girth. Let (X, d) be a metric space. We shall say that $\{M_k\}_{k=0}^\infty$ is a stationary time-reversible Markov chain on X if there are $x_1, \dots, x_n \in X$, an $n \times n$ stochastic matrix A and a stationary distribution $\pi = (\pi_1, \dots, \pi_n)$ of A such that for every i, j , $\pi_i A_{ij} = \pi_j A_{ji}$, $\{M_k\}_{k=0}^\infty$ is a Markov chain with transition matrix A and M_0 is distributed according to π . (X, d) is said to have Markov type $p > 0$ with constant C if for any stationary time-reversible Markov chain on X , and for any time s :

$$\mathbb{E}[d(Z_s, Z_0)^p] \leq C^p s \mathbb{E}[d(Z_1, Z_0)^p].$$

In [2] (see also [38]) it was shown that Hilbert space has Markov type 2 with constant 1. Actually, these references deal with the special case in which A is symmetric and π is the uniform distribution on the states x_1, \dots, x_n , but the proof is easily seen to carry over to stationary time-reversible Markov chains.

6.1 Graphs with large girth

For later applications, it will be convenient to introduce a notion of ‘‘Euclidean distortion at small distances’’ as follows. Let $(X, d_X), (Y, d_Y)$ be metric spaces and $s > 0$. For every injective $f : X \rightarrow Y$ define:

$$\text{dist}^{(s)}(f) = \left(\sup_{0 < d_X(x, y) \leq s} \frac{d_Y(f(x), f(y))}{d_X(x, y)} \right) \cdot \left(\sup_{0 < d_X(x, y) \leq s} \frac{d_X(x, y)}{d_Y(f(x), f(y))} \right),$$

and:

$$c_Y^{(s)}(X) = \inf \left\{ \text{dist}^{(s)}(f); f : X \rightarrow Y \right\}.$$

As before, we write $c_2^{(s)}(X) = c_{\ell_2}^{(s)}(X)$.

Let $G = (V, E)$ be a graph. In what follows we denote by $\delta(G)$ the average degree of G , i.e.

$$\delta(G) = \frac{\sum_{v \in V} \deg(v)}{|V|} = \frac{2|E|}{|V|}.$$

We begin with the following strengthening of a result from [38].

Theorem 6.1. *Let $G = (V, E)$ be a graph with girth g and average degree $\delta = \delta(G)$, then for every integer $1 < s < g/2$, $c_2^{(s)}(G) \geq \frac{\delta-2}{\delta} \sqrt{s}$. In particular,*

$$c_2(G) \geq \frac{\delta-2}{\delta} \sqrt{\left\lfloor \frac{g}{2} \right\rfloor - 1}.$$

Proof. Assume first that G is connected. Consider the reversible Markov chain $\{Z_k\}_{k=0}^\infty$ that corresponds to the canonical random walk on G . Recall that $\pi_v = \deg(v)/(\delta n)$ is a stationary distribution of this Markov chain.

For every $1 < s < g/2$ we have

$$\begin{aligned}
\mathbb{E}[d_G(Z_s, Z_0)] &\geq \mathbb{E}_{v \in V} \left[\frac{\deg(v) - 1}{\deg(v)} (\mathbb{E}[d_G(Z_{s-1}, Z_0) | Z_{s-1} = v] + 1) + \right. \\
&\quad \left. + \frac{1}{\deg(v)} (\mathbb{E}[d_G(Z_{s-1}, Z_0) | Z_{s-1} = v] - 1) \right] \\
&= \mathbb{E}_{v \in V} \left[\frac{\deg(v) - 2}{\deg(v)} + \mathbb{E}[d_G(Z_{s-1}, Z_0) | Z_{s-1} = v] \right] \\
&= \mathbb{E}[d_G(Z_{s-1}, Z_0)] + 1 - \sum_v \pi_v \frac{2}{\deg(v)} \\
&= \mathbb{E}[d_G(Z_{s-1}, Z_0)] + 1 - \sum_v \frac{\deg(v)}{\delta n} \frac{2}{\deg(v)} \\
&= \mathbb{E}[d_G(Z_{s-1}, Z_0)] + \frac{\delta - 2}{\delta}.
\end{aligned}$$

By induction $\mathbb{E}[d_G(Z_s, Z_0)] \geq s \frac{\delta - 2}{\delta}$. Therefore

$$\mathbb{E}[d_G(Z_s, Z_0)^2] \geq [\mathbb{E}d_G(Z_s, Z_0)]^2 \geq s^2 \left(\frac{\delta - 2}{\delta} \right)^2.$$

On the other hand, since Hilbert space has Markov type 2 with constant 1,

$$\mathbb{E}[d_G(Z_s, Z_0)^2] \leq c_2^{(s)}(G)^2 s \mathbb{E}[d_G(Z_1, Z_0)^2] = c_2^{(s)}(G)^2 s.$$

So $c_2^{(s)}(G) \geq \frac{\delta - 2}{\delta} \sqrt{s}$.

If G is disconnected, there is a connected component C of G in which the average degree is at least $\delta = \delta(G)$. The theorem follows by applying the above proof to the connected graph spanned by C . \square

Let $G = (V, E)$ be a d -regular graph, and let A be its adjacency matrix, and $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ be the eigenvalues of A . Recall (see Section 5) that the multiplicative spectral gap of G is $\gamma(G) = \frac{\lambda_2(G)}{d}$ and the absolute multiplicative spectral gap of G is $\gamma_+(G) = \frac{\max\{\lambda_2(G), -\lambda_n(G)\}}{d}$. For $S \subseteq V$, let $E(S) = \{\{u, v\}; u, v \in S, [u, v] \in E\}$. Recall that the Expander Mixing Lemma implies that

$$\frac{2|E(S)|}{|S|} \geq d \left[\frac{|S|}{|V|} - \gamma_+(G) \right].$$

This statement motivates the following useful definition:

Definition 6.2 (Self Mixing Parameter). Let $G = (V, E)$ be a d -regular graph. The self mixing parameter of G is defined as:

$$\mu(G) = \max \left\{ \frac{|S|}{|V|} - \frac{2|E(S)|}{d|S|}; S \subseteq V \right\}.$$

The Expander Mixing Lemma implies that $\mu(G) \leq \gamma_+(G)$. We have in fact the following estimate:

Lemma 6.3. *Let $G = (V, E)$ be an d -regular n -vertex graph, let A be G 's adjacency matrix and let $d = \lambda_1 \geq \dots \geq \lambda_n$ be its eigenvalues. Then:*

$$\mu(G) \leq \frac{-\lambda_n}{d}$$

Proof. Let w_1, \dots, w_n be an orthonormal system of eigenvectors for A with $Aw_i = \lambda_i w_i$ for $i = 1, \dots, n$. Let $\mathbf{1} = \mathbf{1}_V$ be the all-ones vector, then $w_1 = \frac{1}{\sqrt{|V|}}\mathbf{1}$. Let $\mathbf{1}_S$ be the indicator of some subset $S \subseteq V$. Then,

$$\begin{aligned} 2|E(S)| &= \langle A\mathbf{1}_S, \mathbf{1}_S \rangle \\ &= \left\langle A \sum_{i=1}^n \langle \mathbf{1}_S, w_i \rangle w_i, \sum_{i=1}^n \langle \mathbf{1}_S, w_i \rangle w_i \right\rangle \\ &= \sum_{i=1}^n \langle \mathbf{1}_S, w_i \rangle^2 \lambda_i \\ &\geq \langle \mathbf{1}_S, w_1 \rangle^2 \lambda_1 + \lambda_n \sum_{i=1}^n \langle \mathbf{1}_S, w_i \rangle^2 \\ &= \frac{|S|^2}{n} \cdot d + \lambda_n |S|. \end{aligned}$$

□

Lemma 6.4. *Let $G = (V, E)$ be a d -regular graph with girth g and put $\mu = \mu(G)$. Fix $B \subseteq V$, $1 \leq s < g/2$ and denote $\alpha = c_2^{(s)}(B, d_G)$. Assume that $\alpha^2 < s$. Then,*

$$|B| \leq \mu|V| + \frac{2|V|}{d \left(1 - \frac{\alpha}{\sqrt{s}}\right)}.$$

Proof. Set $|B| = k$ and $|V| = n$. By the definition of the self mixing parameter,

$$2|E(B)| \geq \frac{dk^2}{n} - \mu dk.$$

Consider the graph on B induced by G (i.e. the edges are the edges of G which are also in $B \times B$). Its girth is not less than g and its average degree is:

$$\delta = \delta(B) = \frac{2|E(B)|}{k} \geq \frac{dk}{n} - \mu d.$$

Moreover, since G has girth g and $s < g/2$, if $d_B(u, v) \leq s$, for some two vertices $u, v \in B$, then $d_B(u, v) = d_G(u, v)$. Consequently, $c_2^{(s)}(B, d_B) \leq c_2^{(s)}(B, d_G) = \alpha$. An application of Theorem 6.1 yields:

$$\alpha \geq \left(1 - \frac{2}{\delta}\right) \sqrt{s},$$

so that,

$$\frac{dk}{n} - \mu d \leq \delta \leq \frac{2}{1 - \frac{\alpha}{\sqrt{s}}},$$

which gives:

$$k \leq \mu n + \frac{2n}{d \left(1 - \frac{\alpha}{\sqrt{s}}\right)}.$$

□

Let $G = (V, E)$ be a graph and $1 \leq t \leq \text{diam}(G)$ be an integer. We define the t -distance graph of G as $G^{(t)} = (V, E^{(t)})$ where $[u, v] \in E^{(t)}$ if and only if $d_G(u, v) = t$. We collect below some properties of $G^{(t)}$ (part **6** of the lemma below will not be applied in the sequel, and is included here for possible future reference).

Lemma 6.5. *Let $G = (V, E)$ be a d -regular graph, $d \geq 3$, with girth g and let $1 \leq t < g/2$. Then:*

1) $G^{(t)}$ is a $d(d-1)^{t-1}$ regular graph.

2) The girth of $G^{(t)}$ is not less than g/t .

3) For every $u, v \in V$,

$$d_{G^{(t)}}(u, v) < \frac{g}{2t} \implies d_{G^{(t)}}(u, v) = \frac{d_G(u, v)}{t}.$$

4) For every $B \subseteq V$ and $1 \leq s < \frac{g}{2t}$, $c_2^{(s)}(B, d_{G^{(t)}}) \leq c_2(B, d_G)$.

5) If t is even then $\mu(G^{(t)}) \leq 8(d-1)^{-t/4}$.

6) If t is odd then $\gamma(G^{(t)}) \leq 8e^{-(1-\gamma(G))t/8}$.

Proof. Since G has girth g and $t < g/2$, the number of vertices with distance t from a given vertex is the number of leaves of a d -regular tree of depth t , which is $d(d-1)^{t-1}$. This proves **1**). The statements **2**) and **3**) are also simple consequences of the fact that G has girth g . Assertion **4**) follows immediately from assertion **3**).

To prove assertion **5**), note that the adjacency matrix of $G^{(t)}$, $A^{(t)}$, is the t -distance matrix of G , i.e. $A_{uv}^{(t)} = 1$ if $d_G(u, v) = t$ and 0 otherwise. If we denote by A the adjacency matrix of G then there exists a polynomial P_t of degree t such that $A^{(t)} = P_t(A)$ (this is the so-called Geronimus polynomial. The properties of these polynomials that we use here can be found e.g. in [13, 38]). The polynomial P_t has degree t , all its roots are real and reside in the interval $[-2\sqrt{d-1}, 2\sqrt{d-1}]$. An explicit trigonometric expression for P_t is:

$$P_t(2\sqrt{d-1} \cos \theta) = (d-1)^{t/2-1} \frac{(d-1) \sin((t+1)\theta) - \sin((t-1)\theta)}{\sin \theta}.$$

Finally, if t is even, then P_t is an even function and if t is odd, then P_t is an odd function. The spectral theorem shows that $\{P_t(\lambda_i(G))\}$ are the eigenvalues of $A^{(t)}$.

We turn to estimate the smallest eigenvalue of $A^{(t)}$, for t even. This eigenvalue must be negative, but P_t is positive outside the interval $[-2\sqrt{d-1}, 2\sqrt{d-1}]$. In other words, if $P_t(x) < 0$, then $x \in [-2\sqrt{d-1}, 2\sqrt{d-1}]$ and $x = 2\sqrt{d-1} \cos \theta$ for some $\theta \in [-\pi, \pi]$. Therefore,

$$P_t(x) = (d-1)^{t/2-1} \frac{(d-1) \sin((t+1)\theta) - \sin((t-1)\theta)}{\sin \theta}.$$

Using the elementary estimate $|\sin r\alpha| \leq r|\sin \alpha|$ for $\alpha \in [-\pi, \pi]$ and $r \geq 1$, it follows that:

$$|P_t(x)| \leq d(t+1)(d-1)^{t/2-1}.$$

Hence, by Lemma 6.3,

$$\mu(G^{(t)}) \leq \frac{d(t+1)(d-1)^{t/2-1}}{d(d-1)^{t-1}} = (t+1)(d-1)^{-t/2} \leq 8(d-1)^{-t/4},$$

To prove assertion **6**) we distinguish between two cases:

Case one: $\lambda_2(G^{(t)}) = P_t(\lambda_2(G))$. In this case we apply the mean value theorem and find some $a \in (\lambda_2(G), \lambda_1(G))$ such that:

$$\begin{aligned} \log \left[\frac{1}{\gamma(G^{(t)})} \right] &= \log \left[\frac{P_t(d)}{P_t(\lambda_2(G))} \right] \\ &= [d - \lambda_2(G)] \frac{P_t'(a)}{P_t(a)} \\ &= [1 - \gamma(G)] d \sum_{i=1}^t \frac{1}{a - y_i} \end{aligned}$$

where y_i are the roots of P_t . Therefore,

$$\log \left[\frac{1}{\gamma(G^{(t)})} \right] \geq [1 - \gamma(G)] t \frac{d}{d + 2\sqrt{d-1}} \geq \frac{[1 - \gamma(G)] t}{2},$$

as claimed.

Case two: $\lambda_2(G^{(t)}) = P_t(\lambda_i(G))$ for some $i \geq 3$. We claim that λ_i must be in the interval $[-2\sqrt{d-1}, 2\sqrt{d-1}]$. Recall that all the zeros of P_t are in this interval. It is impossible that $\lambda_i < -2\sqrt{d-1}$, since $P_t < 0$ there (t is odd). Also, $\lambda_i > 2\sqrt{d-1}$ is impossible, since P_t is increasing on $[2\sqrt{d-1}, \infty)$ and $\lambda_2(G) \geq \lambda_i(G)$, whereas $P_t(\lambda_i(G)) > P_t(\lambda_2(G))$. Therefore, $\lambda_i(G) \in [-2\sqrt{d-1}, 2\sqrt{d-1}]$, and as in the proof of **5**), we estimate:

$$P_t(\lambda_i(G)) \leq d(t+1)(d-1)^{t/2-1}.$$

Hence:

$$\gamma(G^{(t)}) = \frac{P_t(\lambda_i(G))}{d(d-1)^{t-1}} \leq (t+1)(d-1)^{-t/2} \leq (t+1) \cdot 2^{-t/2},$$

which implies the required result. □

We can now prove an upper bound for the Ramsey problem for graphs with large girth.

Theorem 6.6. *Let $G = (V, E)$ be a d -regular graph, $d \geq 3$ with girth g . Let $1 \leq \alpha < \frac{\sqrt{g}}{6}$. There is an absolute constant $c > 0$ such that*

$$R_2(G; \alpha) \leq 12(d-1)^{-c \frac{g}{\alpha^2}} |V|$$

.

Proof. The proof proceeds by showing that for every $B \subseteq V$ such that $c_2(B, d_G) \leq \alpha$, the following estimate holds:

$$|B| \leq 12(d-1)^{-\frac{g}{64\alpha^2}} |V|.$$

Let t be the unique even defined by $\frac{g}{8\alpha^2} - 2 \leq t < \frac{g}{8\alpha^2}$. Put $s = 4\alpha^2$. Now, since $s < \frac{g}{2t}$, part 4) of Lemma 6.5 implies that:

$$c_2^{(s)}(B, d_{G^{(t)}}) \leq c_2(B, d_G) \leq \alpha.$$

By Lemma 6.5, $\text{girth}(G^{(t)}) \geq g/t$. Also, $s < \frac{g}{2t}$, so we can apply Lemma 6.4 to $G^{(t)}$. Combined with assertion 5) of Lemma 6.5 we deduce:

$$\begin{aligned} |B| &\leq \mu(G^{(t)})|V| + \frac{2|V|}{d(d-1)^{t-1} \left(1 - \frac{\alpha}{\sqrt{4\alpha^2}}\right)} \\ &\leq \left[8(d-1)^{-t/4} + 4(d-1)^{-t}\right] |V| \\ &\leq 12(d-1)^{-t/4} |V| \\ &\leq 12(d-1)^{-\frac{1}{4} \left(\frac{g}{8\alpha^2} - 2\right)} |V| \leq 12(d-1)^{-\frac{g}{64\alpha^2}} |V|. \end{aligned}$$

□

6.2 The discrete cube

The solution for the metric Ramsey problem for the discrete cube is also based on the notion of Markov type. The discrete cube has a small girth, so other ideas are called for. Our analysis utilizes another family of orthogonal polynomial - the Krawtchouk polynomials which appear in many studies related to the discrete cube.

Let $k \leq d$. The degree- k Krawtchouk polynomial for the d -dimensional cube is:

$$K_k^{(d)}(x) = \sum_{j=0}^k (-1)^j \binom{x}{j} \binom{d-x}{k-j}.$$

Again we need an estimate for the smallest value that this polynomial takes.

Lemma 6.7. *Let $1 \leq k \leq \frac{d}{2}$ be even. Then:*

$$K_k^{(d)}(x) \geq - \left(\frac{64k}{d}\right)^{k/2} \binom{d}{k},$$

Proof. It is known (see for example [35]) that all k zeros of $K_k^{(d)}$ are real and belong to the interval:

$$\left[\frac{d}{2} - \sqrt{(k-1)(d-k+2)}, \frac{d}{2} + \sqrt{(k-1)(d-k+2)} \right].$$

Since k is even, $K_k^{(d)}$ is symmetric around $d/2$ (i.e. $K_k^{(d)}(x) = K_k^{(d)}(d-x)$). It is also easily checked that the leading coefficient of $K_k^{(d)}$ is $\frac{(-2)^k}{k!}$. So $K_k^{(d)}(x) \leq 0$ only for x in the above interval. Let z_1, z_2, \dots, z_k be the zeros of $K_k^{(d)}(x)$,

$$K_k^{(d)}(x) = \frac{2^k}{k!} \prod_{i=1}^k (z_i - x),$$

and since x, z_1, \dots, z_k are all in an interval of length $2\sqrt{(k-1)(d-k+2)} \leq 4\sqrt{k(d-k)}$, it follows that for x in the interval above:

$$|K_k^{(d)}(x)| \leq \frac{2^k}{k!} (4\sqrt{k(d-k)})^k \leq \left(\frac{64k}{d}\right)^{k/2} \cdot \binom{d}{k}.$$

To verify this inequality, note that after clearing equal terms, it reduces to $[d(d-k)]^{k/2} \leq d(d-1) \dots (d-k+1)$. This follows by multiplying the inequality $d(d-k) \leq (d-j)(d+j-k+1)$ over $j = 0, \dots, k/2$. \square

Let $\Omega_d = \{0, 1\}^d$ be the graph of the d -dimensional cube. (Two vectors are adjacent if and only if they differ in exactly one coordinate). As before, we consider the t -distance graph on the cube $\Omega_d^{(t)}$. It is well known (e.g. [21]) and easy to show ² that the eigenvalues of the graph $\Omega_d^{(t)}$ are the numbers $K_t^{(d)}(i)$ for $i = 0, \dots, d$ where the i -th eigenvalue appears with multiplicity $\binom{d}{i}$. This graph is $\binom{d}{t}$ -regular and so its largest eigenvalue is $\binom{d}{t}$. Lemmas 6.3 and 6.7 now yield an estimate for the self mixing parameter $\mu(\Omega_d^{(t)})$

Lemma 6.8. *For every even integer $1 \leq t < d/2$,*

$$\mu(\Omega_d^{(t)}) \leq \left(\frac{64t}{d}\right)^{t/2}.$$

To prove the main result of this section, we need an additional estimate.

Lemma 6.9. *Let t, x, d be integers such that $2t \leq x \leq d/2$. Then:*

$$\sum_{j=t/3}^t \binom{x}{j} \binom{d-x}{t-j} \leq 2 \left(\frac{150x}{d}\right)^{t/3} \binom{d}{t}.$$

Proof. Clearly, $\binom{d}{t} = \sum_j \binom{x}{j} \binom{d-x}{t-j}$, so it suffices to consider the range $x \leq \frac{d}{150}$. In other words, we assume $\frac{t}{3} \leq j \leq t \leq \frac{x}{2} \leq \frac{d}{300}$. In this range, the terms decrease geometrically, $\binom{x}{j+1} \binom{d-x}{t-(j+1)} \leq \frac{1}{10} \binom{x}{j} \binom{d-x}{t-j}$. It therefore suffices to show that $\binom{x}{t/3} \binom{d-x}{2t/3} \leq \left(\frac{150x}{d}\right)^{t/3} \binom{d}{t}$. Recall the following elementary and well-known estimates of binomial coefficients: For every $1 \leq k \leq n$,

$$\left(\frac{n}{k}\right)^k \leq \binom{n}{k} \leq \left(\frac{en}{k}\right)^k.$$

We plug this into the expression and simplify, to conclude that the inequality holds. \square

Theorem 6.10. *There are absolute constants $C, c, c' > 0$ such that for every integer d and for every $\alpha \geq 1$,*

$$2^{\left(1 - \frac{\log(c'\alpha)}{\alpha^2}\right)d} \leq R_2(\Omega_d; \alpha) \leq C 2^{\left(1 - \frac{c}{\alpha^2}\right)d}.$$

²To show this, recall the 2^d Walsh functions $\{W_S \mid S \in \Omega_d\}$ that are defined via $W_S(T) = (-1)^{\langle S, T \rangle}$. It is not hard to see that they form a complete set of eigenfunctions for $\Omega_d^{(t)}$ and the eigenvalue corresponding to W_S is $K_t^{(d)}(n - 2|S|)$.

Proof. We start with the lower bound. An easy fact from coding theory, called the Gilbert-Varshamov bound, [40], states that there exists a subset $B \subseteq \Omega_d$ such that every $u \neq v \in B$, are at distance $\geq \frac{d}{\alpha^2}$, and:

$$|B| \geq \frac{2^d}{\sum_{m \leq d/\alpha^2} \binom{d}{m}} \geq 2^{\left(1 - \frac{\log(c'\alpha)}{\alpha^2}\right)d},$$

where the last inequality follows from standard estimates on binomial coefficients. Note that for every $u, v \in \Omega_d$, $\|u - v\|_2 = \sqrt{\rho(u, v)}$, where ρ stands for the Hamming distance. But for every distinct $u, v \in B$, $\frac{d}{\alpha^2} \leq \rho(u, v) \leq d$, so $\frac{\sqrt{d}}{\alpha} \leq \frac{\rho(u, v)}{\|u - v\|_2} \leq \sqrt{d}$. Consequently, $c_2(B) \leq \alpha$.

To motivate the proof of the upper bound, let us sketch a proof based on Markov Type for (a weakening of) the classical fact [25] that $c_2(\Omega_d) \geq a\sqrt{d}$ for some absolute $a > 0$ (In fact, $c_2(\Omega_d) = \sqrt{d}$). The random walk on Ω_d almost surely drifts with constant speed from its point of origin for time $> a'd$ for some absolute $a' > 0$. This is true because a constant fraction of the coordinates stay unchanged for this duration. On the other hand, the fact that Hilbert space has Markov-type 2 implies that the corresponding walk on an image of Ω_d will typically drift only $O(\sqrt{d})$ away from its origin. This discrepancy implies a metric distortion $\geq c\sqrt{d}$, as claimed. The spirit of the proof we present is similar. We only have a subset $B \subseteq \Omega_d$, so we consider (a dense connected component of) the graph $\Omega_d^{(t)}$. The main technical effort is in estimating the typical rate of drift from the walk's origin.

We wish to show, then, that if $B \subseteq \Omega_d$, satisfies $c_2(B) \leq \alpha$, then $|B| \leq C2^{\left(1 - \frac{c}{\alpha^2}\right)d}$. Let $n = 2^d$ and $k = |B|$. We seek an upper bound on k . As in the proof of Theorem 6.6, we investigate the random walk on the distance t graph of the graph in question, namely $\Omega_d^{(t)}$. We define t as the even integer nearest to $\frac{d}{K\alpha^2}$, where K is a suitably large absolute constant to be specified later. It can be verified that $2^{\left(1 - \frac{c}{\alpha^2}\right)d} \geq 2n \left(\frac{32t}{d}\right)^{t/2}$ for this choice of t . Therefore, we may assume that:

$$\frac{k}{n} \geq 2 \left(\frac{64t}{d}\right)^{t/2},$$

or else the required upper bound on k already holds.

Denote by $E_t(B)$ the number of unordered pairs of points of distance t in B . In terms of the graph $\Omega_d^{(t)}$ this is: $E_t(B) = |E_{\Omega_d^{(t)}}(B)|$. By Lemma 6.3:

$$2E_t(B) \geq \frac{\binom{d}{t}k^2}{n} - \mu(\Omega_d^{(t)}) \binom{d}{t}k \geq \frac{\binom{d}{t}k^2}{n} - \left(\frac{64t}{d}\right)^{t/2} \binom{d}{t}k.$$

So that:

$$\delta = \delta_{\Omega_d^{(t)}}(B) = \frac{2E_t(B)}{k} \geq \binom{d}{t} \left[\frac{k}{n} - \left(\frac{64t}{d}\right)^{t/2} \right] \geq \binom{d}{t} \cdot \frac{k}{2n}$$

There is a connected component C of the subgraph of $\Omega_d^{(t)}$ spanned by B that has average degree $\delta' \geq \delta$, i.e.,

$$\delta' = \delta_{\Omega_d^{(t)}}(C) \geq \delta \geq \binom{d}{t} \cdot \frac{k}{2n}.$$

Let $\{Z_r\}_{r=0}^\infty$ be the random walk on $\Omega_d^{(t)}$ restricted to C . We start the walk at the stationary distribution, viz.,

$$P(Z_0 = v) = \frac{\deg_C^t(v)}{\delta'|C|},$$

where $\deg_C^t(v)$ is the degree of vertex v in $\Omega_d^{(t)}$ restricted to C (i.e. the number of elements of C with Hamming distance t to v).

Suppose that our random walk starts from $S \in C$ and reaches, after some time, a vertex T with $x = \rho(S, T)$. Say that we next step from T to W . We seek an upper bound on the probability that $\rho(S, W) \leq x + \frac{t}{3}$. The total number of neighbors W of T in $\Omega_d^{(t)}$ for which this holds is

$$A(x) = \sum_{j \geq \lceil t/3 \rceil} \binom{x}{j} \binom{d-x}{t-j}.$$

By Lemma 6.9, $A(x) \leq 2 \left(\frac{150x}{d}\right)^{t/3} \binom{d}{t}$ when $2t \leq x \leq \frac{d}{2}$.

Now for every possible walk, $\rho(Z_r, Z_0) \leq rt$ holds for every integer $r > 1$. For times $2 \leq r \leq \frac{d}{2t}$ we are able to show that the walk tends to drift at least $a''t$ per step away from its origin for some absolute $a'' > 0$.

$$\begin{aligned} & \mathbb{E}[\rho(Z_{r+1}, Z_0)] \\ & \geq \mathbb{E} \left[\frac{\deg_C^t(Z_r) - A(\rho(Z_r, Z_0))}{\deg_C^t(Z_r)} \left(\rho(Z_r, Z_0) + \frac{t}{3} \right) + \frac{A(\rho(Z_r, Z_0))}{\deg_C^t(Z_r)} (\rho(Z_r, Z_0) - t) \right] \\ & = \mathbb{E}[\rho(Z_r, Z_0)] + \frac{t}{3} - \frac{4t}{3} \mathbb{E} \left[\frac{A(\rho(Z_r, Z_0))}{\deg_C^t(Z_r)} \right] \\ & \geq \mathbb{E}[\rho(Z_r, Z_0)] + \frac{t}{3} - \frac{8t}{3} \left(\frac{150rt}{d} \right)^{t/3} \binom{d}{t} \mathbb{E} \left[\frac{1}{\deg_C^t(Z_r)} \right] \\ & = \mathbb{E}[\rho(Z_r, Z_0)] + \frac{t}{3} - \frac{8t}{3} \left(\frac{150rt}{d} \right)^{t/3} \binom{d}{t} \sum_{v \in C} \frac{1}{\deg_C^t(v)} \cdot \frac{\deg_C^t(v)}{\delta'|C|} \\ & = \mathbb{E}[\rho(Z_r, Z_0)] + \frac{t}{3} - \frac{8t}{3\delta'} \left(\frac{150rt}{d} \right)^{t/3} \binom{d}{t} \\ & \geq \mathbb{E}[\rho(Z_r, Z_0)] + \frac{t}{3} - \frac{16t}{3} \left(\frac{150rt}{d} \right)^{t/3} \frac{n}{k}. \end{aligned}$$

As in the proof of Theorem 6.6, we now contrast this estimate with the fact that Hilbert space has Markov type 2. Namely, that for every r ,

$$\alpha^2 r t^2 \geq c_2(B)^2 r t^2 \geq c_2(C)^2 r \mathbb{E}[\rho^2(Z_1, Z_0)] \geq \mathbb{E}[\rho^2(Z_r, Z_0)] \geq [\mathbb{E}\rho(Z_r, Z_0)]^2.$$

Consequently,

$$\alpha t \sqrt{r} \geq \mathbb{E}[\rho(Z_r, Z_0)] \geq \frac{rt}{3} - \frac{16nt}{3k} \left(\frac{150t}{d} \right)^{t/3} \sum_{r \geq j \geq 1} j^{t/3} \geq \frac{rt}{3} - \frac{16n}{k} \left(\frac{150t}{d} \right)^{t/3} r^{\frac{t}{3}+1}.$$

We set $r = \lceil 36\alpha^2 \rceil$ and, as stated above, choose t as the even integer nearest to $\frac{d}{5500\alpha^2}$ to conclude the proof of desired result. \square

Remark 6.11. It is known that for every $1 \leq p \leq 2$ the metric space $(\ell_p, \|x - y\|_p^{p/2})$ embeds isometrically into ℓ_2 (see [52]). It follows that ℓ_p has Markov type p with constant 1. We can therefore apply the above arguments and conclude that Theorem 6.6 and Theorem 6.10 remain true when dealing with embeddings into ℓ_p , $1 < p \leq 2$. The only necessary modification is that in the upper bound on the Ramsey function α^2 should be replaced by $\alpha^{p/(p-1)}$.

Acknowledgments. The authors would like to express their gratitude to Guy Kindler, Robi Krauthgamer, Avner Magen and Yuri Rabinovich for some helpful discussions.

References

- [1] N. Alon. Eigenvalues and expanders. *Combinatorica*, 6(2):83–96, 1986.
- [2] K. Ball. Markov chains, Riesz transforms and Lipschitz maps. *Geom. Funct. Anal.*, 2(2):137–172, 1992.
- [3] Y. Bartal. Probabilistic approximation of metric spaces and its algorithmic applications. In *37th Annual Symposium on Foundations of Computer Science*, pages 184–193, 1996.
- [4] Y. Bartal. On approximating arbitrary metrics by tree metrics. In *the 30th Annual ACM Symposium on Theory of Computing*, pages 183–193, 1998.
- [5] Y. Bartal, B. Bollobás, and M. Mendel. Ramsey-type theorems for metric spaces with applications to online problems. In *the 42nd Annual Symposium on Foundations of Computer Science*, pages 396–405, 2001.
- [6] Y. Bartal, N. Linial, M. Mendel, and A. Naor. Limitations to Frechét’s Metric Embedding Method. Preprint, 2003.
- [7] Y. Bartal, N. Linial, M. Mendel, and A. Naor. Low Dimensional Embeddings of Ultrametrics. *European Journal of Combinatorics*, to appear, 2003.
- [8] Y. Bartal, N. Linial, M. Mendel, and A. Naor. On Some Low Distortion Metric Ramsey Problems. *Discrete and Computational Geometry*, to appear, 2003.
- [9] Y. Bartal, N. Linial, M. Mendel, and A. Naor. On Metric Ramsey-type Phenomena: Conference version. In *the 35th Annual ACM Symposium on Theory of Computing*, pages 463–472, 2003.
- [10] Y. Bartal and M. Mendel. Multi-embeddings and path-approximation of metric spaces. In *14th Ann. ACM-SIAM Symposium on Discrete Algorithms*, pages 424–433, 2003.
- [11] M. Ben-Or and N. Linial. *Collective Coin Flipping* (S. Micali ed.), Academic Press, New York. pages 91–115, 1989.
- [12] Yoav Benyamini and Joram Lindenstrauss. *Geometric nonlinear functional analysis. Vol. 1*, volume 48 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 2000.
- [13] N. Biggs. *Algebraic graph theory*. Cambridge University Press, London, 1974. Cambridge Tracts in Mathematics, No. 67.
- [14] A. Blum, H. Karloff, Y. Rabani, and M. Saks. A decomposition theorem for task systems and bounds for randomized server problems. *SIAM J. Comput.*, 30(5):1624–1661 (electronic), 2000.
- [15] J. Bourgain. On Lipschitz embedding of finite metric spaces in Hilbert space. *Israel J. Math.*, 52(1-2):46–52, 1985.

- [16] J. Bourgain. The metrical interpretation of superreflexivity in Banach spaces. *Israel J. Math.*, 56(2):222–230, 1986.
- [17] J. Bourgain, T. Figiel, and V. Milman. On Hilbertian subsets of finite metric spaces. *Israel J. Math.*, 55(2):147–152, 1986.
- [18] J. Bourgain, V. Milman, and H. Wolfson. On type of metric spaces. *Trans. Amer. Math. Soc.*, 294(1):295–317, 1986.
- [19] B. Brinkman and M. Charikar. On the impossibility of dimension reduction in ℓ_1 . In *to appear in Proceedings of the 44th Annual IEEE Conference on Foundations of Computer Science*, pages 514–523. ACM, 2003.
- [20] F. R. K. Chung. Diameters and eigenvalues. *J. Amer. Math. Soc.*, 2(2):187–196, 1989.
- [21] P. Delsarte. An algebraic approach to the association schemes of coding theory. *Philips Res. Rep. Suppl.*, (10):vi+97, 1973.
- [22] M. M. Deza and M. Laurent. *Geometry of cuts and metrics*. Springer-Verlag, Berlin, 1997.
- [23] J. Diestel, H. Jarchow, and A. Tonge. *Absolutely summing operators*. Cambridge University Press, Cambridge, 1995.
- [24] A. Dvoretzky. Some results on convex bodies and Banach spaces. In *Proc. Internat. Sympos. Linear Spaces (Jerusalem, 1960)*, pages 123–160. Jerusalem Academic Press, Jerusalem, 1961.
- [25] P. Enflo. On the nonexistence of uniform homeomorphisms between L_p -spaces. *Ark. Mat.*, 8:103–105 (1969), 1969.
- [26] U. Feige. Approximating the bandwidth via volume respecting embeddings. *J. Comput. System Sci.*, 60(3):510–539, 2000. 30th Annual ACM Symposium on Theory of Computing (Dallas, TX, 1998).
- [27] T. Figiel, J. Lindenstrauss, and V. D. Milman. The dimension of almost spherical sections of convex bodies. *Acta Math.*, 139(1-2):53–94, 1977.
- [28] J. Fakcharoenphol, S. Rao and K. Talwar. A tight bound on approximating arbitrary metrics by tree metrics. In *Proc. of the 35th Ann. ACM Symp. on Theory of Computing* pages 448–455, 2003.
- [29] R. L. Graham, B. L. Rothschild, and J. H. Spencer. *Ramsey theory*. John Wiley & Sons Inc., New York, second edition, 1990.
- [30] P. Indyk. Algorithmic applications of low-distortion geometric embeddings. In *the 42nd Annual Symposium on Foundations of Computer Science*, pages 10–33, 2001.
- [31] W. B. Johnson and J. Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. In *Conference in modern analysis and probability (1982)*, pages 189–206. AMS, 1984.
- [32] H. Karloff, Y. Rabani, and Y. Ravid. Lower bounds for randomized k -server and motion-planning algorithms. *SIAM J. Comput.*, 23(2):293–312, 1994.
- [33] J. R. Lee and A. Naor. Embedding the diamond graph in L_p and dimension reduction in L_1 . *To appear in Geometric and Functional Analysis*, 2003.

- [34] Alex J. Lemin. Isometric embedding of ultrametric (non-Archimedean) spaces in Hilbert space and Lebesgue space. In *p-adic functional analysis (Ioannina, 2000)*, volume 222 of *Lecture Notes in Pure and Appl. Math.*, pages 203–218. Dekker, New York, 2001.
- [35] Vladimir I. Levenshtein. Krawtchouk polynomials and universal bounds for codes and designs in Hamming spaces. *IEEE Trans. Inform. Theory*, 41(5):1303–1321, 1995.
- [36] N. Linial. Finite metric spaces- combinatorics, geometry and algorithms. In *Proceedings of the ICM*, 2002.
- [37] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.
- [38] N. Linial, A. Magen, and A. Naor. Girth and euclidean distortion. *Geometric and Functional Analysis*, 12(2):380–394, 2002.
- [39] L. H. Loomis and H. Whitney. An inequality related to the isoperimetric inequality. *Bull. Amer. Math. Soc.*, 55:961–962, 1949.
- [40] F. J. MacWilliams and N. J. A. Sloane. *The theory of error-correcting codes. I*. North-Holland Publishing Co., Amsterdam, 1977. North-Holland Mathematical Library, Vol. 16.
- [41] J. Matoušek. On embedding expanders into l_p spaces. *Israel J. Math.*, 102:189–197, 1997.
- [42] J. Matoušek. *Lectures on discrete geometry*. Springer-Verlag, New York, 2002.
- [43] B. Maurey and G. Pisier. Séries de variables aléatoires vectorielles indépendantes et propriétés géométriques des espaces de Banach. *Studia Math.*, 58(1):45–90, 1976.
- [44] V. Milman and G. Schechtman. An “isomorphic” version of Dvoretzky’s theorem. II. In *Convex geometric analysis (Berkeley, CA, 1996)*, pages 159–164. Cambridge Univ. Press, Cambridge, 1999.
- [45] V. D. Milman. A new proof of A. Dvoretzky’s theorem on cross-sections of convex bodies. *Funkcional. Anal. i Priložen.*, 5(4):28–37, 1971.
- [46] V. D. Milman and G. Schechtman. *Asymptotic theory of finite-dimensional normed spaces*. Springer-Verlag, Berlin, 1986. With an appendix by M. Gromov.
- [47] V. D. Milman and G. Schechtman. An “isomorphic” version of Dvoretzky’s theorem. *C. R. Acad. Sci. Paris Sér. I Math.*, 321(5):541–544, 1995.
- [48] G. Pisier. *The volume of convex bodies and Banach space geometry*. Cambridge University Press, Cambridge, 1989.
- [49] F. P. Ramsey. On a problem of formal logic. *Proc. London Math. Soc.*, 48:122–160, 1930.
- [50] S. Rao. Small distortion and volume preserving embeddings for planar and Euclidean metrics. In *Proceedings of the Fifteenth Annual Symposium on Computational Geometry (Miami Beach, FL, 1999)*, pages 300–306 (electronic), New York, 1999. ACM.
- [51] Neil Robertson and P. D. Seymour. Graph minors. VIII. A Kuratowski theorem for general surfaces. *J. Combin. Theory Ser. B*, 48(2):255–288, 1990.

[52] J. H. Wells and L. R. Williams. *Embeddings and extensions in analysis*. Springer-Verlag, New York, 1975. *Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 84*.

2000 AMS Mathematics Subject Classification: 52C45, 05C55, 54E40, 05C12, 54E40.

Yair Bartal, Institute of Computer Science, Hebrew University, Jerusalem 91904, Israel.
yair@cs.huji.ac.il

Supported in part by a grant from the Israeli Science Foundation (195/02).

Nathan Linial, Institute of Computer Science, Hebrew University, Jerusalem 91904, Israel.
nati@cs.huji.ac.il

Supported in part by a grant from the Israeli National Science Foundation.

Manor Mendel, Institute of Computer Science, Hebrew University, Jerusalem 91904, Israel.
mendelma@yahoo.com

Supported in part by the Landau Center and by a grant from the Israeli Science Foundation (195/02).

Assaf Naor, Theory Group, Microsoft Research, One Microsoft Way 113/2131, Redmond WA 98052-6399, USA.

anaor@microsoft.com