

Experimental Study of Energy-Minimizing Point Configurations on Spheres

Brandon Ballinger, Grigoriy Blekherman, Henry Cohn, Noah Giansiracusa, Elizabeth Kelly, and Achill Schürmann

CONTENTS

- 1. Introduction
- 2. Methodology
- 3. Experimental Phenomena
- 4. Conjectured Universal Optima
- 5. Balanced Irreducible Harmonic Optima
- 6. Challenges
- 7. Appendix
- Acknowledgments
- References

In this paper we report on massive computer experiments aimed at finding spherical point configurations that minimize potential energy. We present experimental evidence for two new universal optima (consisting of 40 points in 10 dimensions and 64 points in 14 dimensions), as well as evidence that there are no others with at most 64 points. We also describe several other new polytopes, and we present new geometrical descriptions of some of the known universal optima.

[T]he problem of finding the configurations of stable equilibrium for a number of equal particles acting on each other according to some law of force... is of great interest in connexion with the relation between the properties of an element and its atomic weight. Unfortunately the equations which determine the stability of such a collection of particles increase so rapidly in complexity with the number of particles that a general mathematical investigation is scarcely possible.

—J. J. Thomson, 1897

1. INTRODUCTION

What is the best way to distribute N points over the unit sphere S^{n-1} in \mathbb{R}^n ? Of course the answer depends on the notion of “best.” One particularly interesting case is energy minimization. Given a continuous decreasing function $f: (0, 4] \rightarrow \mathbb{R}$, define the f -potential energy of a finite subset $\mathcal{C} \subset S^{n-1}$ to be

$$E_f(\mathcal{C}) = \frac{1}{2} \sum_{\substack{x, y \in \mathcal{C} \\ x \neq y}} f(|x - y|^2).$$

(The domain of f is only $(0, 4]$ because $|x - y|^2 \leq 4$ when $|x|^2 = |y|^2 = 1$. The factor of $\frac{1}{2}$ is chosen for compatibility with the physics literature, while the use of squared distance is incompatible but more convenient.)

2000 AMS Subject Classification: 52B11, 52C17, 52-04

Keywords: Energy minimization, polytopes, universal optimality

How can one choose $\mathcal{C} \subset S^{n-1}$ with $|\mathcal{C}| = N$ so as to minimize $E_f(\mathcal{C})$?

In this paper we report on lengthy computer searches for configurations with low energy. What distinguishes our approach from most earlier work on this topic (see for example [Altschuler and Pérez-Garrido 05, Altschuler and Pérez-Garrido 06, Altschuler and Pérez-Garrido 07, Altschuler et al. 97, Aste and Weaire 08, Bausch et al. 03, Bowick et al. 02, Bowick et al. 06, Cohn 56, Damelin and Maymeskul 05, Dragnev et al. 02, Edmundson 92, Einert et al. 05, Erber and Hockney 91, Föppl 12, Glasser and Every 92, Hardin and Saff 04, Hardin and Saff 05, Hovinga 04, Katanforoush and Shahshahani 03, Kuijlaars and Saff 98, Kottwitz 91, Livshits and Lozovik 99, Martínez-Finkelshtein et al. 04, Melnyk et al. 77, Morris et al. 96, Pérez-Garrido et al. 97a, Pérez-Garrido et al. 97b, Pérez-Garrido and Moore 99, Rakhmanov et al. 94, Rakhmanov et al. 95, Saff and Kuijlaars 97, Sloane 00, Thomson 97, Whyte 52, Willie 86]) is that we attempt to treat many different potential functions on as even a footing as possible. Much of the mathematical structure of this problem becomes apparent only when one varies the potential function f . Specifically, we find that many optimal configurations vary in surprisingly simple low-dimensional families as f varies.

The most striking possibility is that the family is a single point; in other words, the optimum is independent of f . Cohn and Kumar [Cohn and Kumar 07] defined a configuration to be *universally optimal* if it minimizes E_f for all completely monotonic f (i.e., f is infinitely differentiable and $(-1)^k f^{(k)}(x) \geq 0$ for all $k \geq 0$ and $x \in (0, 4)$, as is the case for inverse power laws). They were able to prove universal optimality only for certain very special arrangements. One of our primary goals in this paper is to investigate how common universal optimality is. Was the limited list of examples in [Cohn and Kumar 07] an artifact of the proof techniques or a sign that these configurations are genuinely rare?

Every universally optimal configuration is an optimal spherical code, in the sense that it maximizes the minimal distance between the points. (Consider an inverse power law $f(r) = 1/r^s$. If there were a configuration with a larger minimal distance, then its f -potential energy would be lower when s is sufficiently large.) However, universal optimality is a far stronger condition than optimality as a spherical code. There are optimal spherical codes of each size in each dimension, but they are rarely universally optimal. In three dimensions, the only universal optima are a single point, two antipodal points,

an equilateral triangle on the equator, or the vertices of a regular tetrahedron, octahedron, or icosahedron.

The universal optimality of these configurations was proved in [Cohn and Kumar 07], building on previous work by Yudin, Kolushov, and Andreev [Yudin 93, Kolushov and Yudin 94, Kolushov and Yudin 97, Andreev 96, Andreev 97], and the completeness of this list follows from a classification theorem due to Leech [Leech 57]. See [Cohn and Kumar 07] for more details.

In higher dimensions much less is known. Cohn and Kumar's main theorem provides a general criterion from which they deduced the universal optimality of a number of previously studied configurations. Specifically, they proved that every spherical $(2m-1)$ -design in which only m distances occur between distinct points is universally optimal. Recall that a spherical d -design in S^{n-1} is a finite subset \mathcal{C} of S^{n-1} such that every polynomial on \mathbb{R}^n of total degree d has the same average over \mathcal{C} as over the entire sphere. This criterion holds for every known universal optimum except one case, namely the regular 600-cell in \mathbb{R}^4 (i.e., the H_4 root system), for which Cohn and Kumar proved universal optimality by a special argument.

A list of all known universal optima is given in Table 1. Here n is the dimension of the Euclidean space, N is the number of points, and t is the greatest inner product between distinct points in the configuration (i.e., the cosine of the minimal angle). For detailed descriptions of these configurations, see [Cohn and Kumar 07, Section 1].

Each known universal optimum is uniquely determined by the parameters listed in Table 1, except for the configurations listed on the last line. For that case, when $q = p^\ell$ with p an odd prime, there are at least $\lfloor (\ell-1)/2 \rfloor$ distinct universal optima (see [Cameron et al. 78] and [Kantor 86]). Classifying these optima is equivalent to classifying generalized quadrangles with parameters (q, q^2) , which is a difficult problem in combinatorics. In the other cases from Table 1, when uniqueness holds, we use the notation $U_{N,n}$ for the unique N -point universal optimum in \mathbb{R}^n .

Each of the configurations in Table 1 had been studied before it appeared in [Cohn and Kumar 07], and was already known to be an optimal spherical code. In fact, when $N \geq 2n+1$ and $n > 4$, the codes on this list are exactly those that have been proved optimal. Cohn and Kumar were unable to determine whether Table 1 is the complete list of universally optimal codes, except when $n \leq 3$.

All that is known in general is that any new universal optimum must have $N \geq 2n+1$ [Cohn and Kumar 07,

n	N	t	Description
2	N	$\cos(2\pi/N)$	N -gon
n	$N \leq n + 1$	$-1/(N - 1)$	simplex
n	$2n$	0	cross polytope
3	12	$1/\sqrt{5}$	icosahedron
4	120	$(1 + \sqrt{5})/4$	regular 600-cell
5	16	$1/5$	hemicube/Clebsch graph
6	27	$1/4$	Schläfli graph/isotropic subspaces
7	56	$1/3$	equiangular lines
8	240	$1/2$	E_8 root system
21	112	$1/9$	isotropic subspaces
21	162	$1/7$	(162, 56, 10, 24) strongly regular graph
22	100	$1/11$	Higman–Sims graph
22	275	$1/6$	McLaughlin graph
22	891	$1/4$	isotropic subspaces
23	552	$1/5$	equiangular lines
23	4600	$1/3$	kissing configuration of the following
24	196560	$1/2$	Leech lattice minimal vectors
$q^{\frac{q^3+1}{q+1}}$	$(q+1)(q^3+1)$	$1/q^2$	isotropic subspaces (q is a prime power)

TABLE 1. The known universal optima.

Proposition 1.4]. It does not seem plausible that the current list is complete, but it is far from obvious where to find any others.

Each known universal optimum is a beautiful mathematical object, connected to various important exceptional structures (such as special lattices or groups). Our long-term hope is to develop automated tools that will help uncover more such objects. In this paper we do not discover any configurations as fundamental as those in Table 1, but perhaps our work is a first step in that direction.

Table 1 shows several noteworthy features. When $n \leq 4$, the codes listed are the vertices of regular polytopes, specifically those with simplicial facets. When $5 \leq n \leq 8$, the table also includes certain semiregular polytopes (their facets are simplices and cross polytopes, with two cross polytopes and one simplex arranged around each $(n - 3)$ -dimensional face). The corresponding spherical codes are all affine cross sections of the minimal vectors of the E_8 root lattice. Remarkably, no universal optima are known for $9 \leq n \leq 20$, except for the simplices and cross polytopes, which exist in every dimension. This gap is troubling—why should these dimensions be disfavored? For $21 \leq n \leq 24$ nontrivial universal optima are known; they are all affine cross sections of the minimal vectors of the Leech lattice (and are no longer the vertices of

semiregular polytopes). Finally, in high dimensions a single infinite sequence of nontrivial universal optima is known.

It is not clear how to interpret this list. For example, is the dimension gap real, or merely an artifact of humanity’s limited imagination? One of our conclusions in this paper is that Table 1 is very likely incomplete but appears to be closer to complete than one might expect.

1.1 Experimental Results

One outcome of our computer searches is two candidates for universal optima, listed in Table 2 and described in more detail in Section 4. These configurations were located through massive computer searches: for each of many pairs (n, N) , we repeatedly picked N random points on S^{n-1} and performed gradient descent to minimize potential energy. We focused on the potential function $f(r) = 1/r^{n/2-1}$, because $x \mapsto 1/|x|^{n-2}$ is the unique nonconstant radial harmonic function on $\mathbb{R}^n \setminus \{0\}$, up to scalar multiplication (recall that distance is squared in the definition of E_f). When $n = 3$, the f -potential energy for this function f is the Coulomb potential energy from electrostatics, and this special case has been extensively studied by mathematicians and other scientists. In higher dimensions, this potential function has frequently

n	N	t	References
10	40	1/6	[Sloane 00, Hovinga 04]
14	64	1/7	[Nordstrom and Robinson 67, de Caen and van Dam 99, Ericson and Zinoviev 01]

TABLE 2. New conjectured universal optima.

been studied as a natural generalization of electrostatics; we call it the *harmonic potential function*.

Because there are typically numerous local minima for harmonic energy, we repeated this optimization procedure many times with the hope of finding the global minimum. For low numbers of points in low dimensions, the apparent global minimum occurs fairly frequently. Figure 1 shows data from three dimensions. In higher dimensions, there are usually more local minima and the true optimum can occur very infrequently.

For each conjectured optimum for harmonic energy, we attempted to determine whether it could be universally optimal. We first determined whether it is in equilibrium under all possible force laws (i.e., “balanced” in the terminology of Leech [Leech 57]). That holds if and only if for each point x in the configuration and each distance d , the sum of all points in the code at distance d from x is a scalar multiple of x . If this criterion fails, then there is some inverse power law under which the code is not even in equilibrium, let alone globally minimal, so it cannot possibly be universally optimal. Most of the time, the code with the lowest harmonic potential energy is not balanced. When it is balanced, we compared several potential functions to see whether we could disprove universal optimality. By [Widder 41, Theorem 9b, p. 154], it suffices to look at the potential functions $f(r) = (4 - r)^k$ with $k \in \{0, 1, 2, \dots\}$ (on each compact subinterval of $(0, 4]$, every completely monotonic function can be approximated arbitrarily closely by positive linear combinations of these potential functions). Because these functions do not blow up at $r = 0$, numerical calculations with them often converge more slowly than they do for inverse power laws (nearby points can ex-

perience only a bounded force pushing them apart), so they are not a convenient choice for initial experimentation. However, they play a fundamental role in detecting universal optima.

To date, our search has led us to 58 balanced configurations with at most 64 points (and at least $2n + 1$ in dimension n) that appear to minimize harmonic energy and were not already known to be universally optimal. In all but two cases, we were able to disprove universal optimality, but the remaining two cases (those listed in Table 2) are unresolved. We conjecture that they are in fact universally optimal.

Figure 2 presents a graphical overview of our data. The triangle of white circles on the upper left represents the simplices, and the diagonal line of white circles represents the cross polytopes. Between them, one can see that the pattern is fairly regular, but as one moves right from the cross polytopes all structure rapidly vanishes. There is little hope of finding a simple method to predict where balanced harmonic optima can be found, let alone universal optima. It also does not seem likely that general universal optima can be characterized by any variant of Cohn and Kumar’s criterion.

Besides the isotropic subspace universal optima from Table 1 and the other universal optima with the same parameters, we can conjecture only one infinite family of balanced harmonic optima with more than $2n$ points

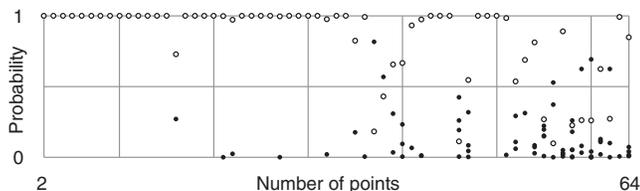


FIGURE 1. Probabilities of local minima for harmonic energy in \mathbb{R}^3 (based on 1000 trials). White circles denote the conjectured harmonic optima.

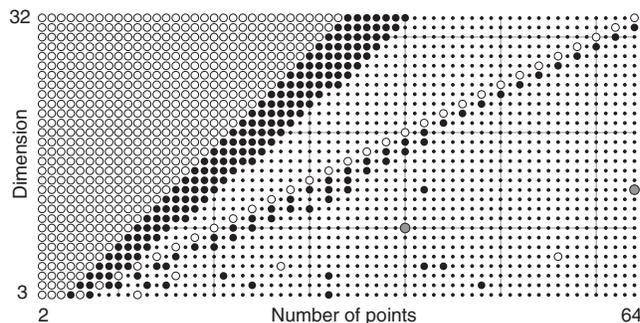


FIGURE 2. Status of conjectured harmonic optima with up to 64 points in at most 32 dimensions: white circle denotes universal optimum, large gray circle denotes conjectured universal optimum, black circle denotes balanced configuration that is not universally optimal, tiny black circle denotes unbalanced configuration.

n	N	t
3	32	$\sqrt{75 + 30\sqrt{5}}/15$
4	10	$1/6$ or $(\sqrt{5} - 1)/4$
4	13	$(\cos(4\pi/13) + \cos(6\pi/13))/2$
4	15	$1/\sqrt{8}$
4	24	$1/2$
4	48	$1/\sqrt{2}$
5	21	$1/\sqrt{10}$
5	32	$1/\sqrt{5}$
$n \geq 6$	$2n + 2$	$1/n$
6	42	$2/5$
6	44	$1/\sqrt{6}$
6	126	$\sqrt{3/8}$
7	78	$3/7$
7	148	$\sqrt{2/7}$
8	72	$5/14$
9	96	$1/3$
14	42	$1/10$
16	256	$1/4$

TABLE 3. Conjectured harmonic optima that are balanced, irreducible, and not universally optimal (see Section 5 for descriptions).

in \mathbb{R}^n , namely the diplo-simplices with $2n + 2$ points in \mathbb{R}^n for $n \geq 6$ (see Section 3.4). Certainly no others are apparent in Figure 2, but the isotropic subspace optima from Table 1 are sufficiently large and exotic that it would be foolish to conjecture that there are no other infinite families.

Table 3 lists the cases in which we found a balanced harmonic optimum but were able to disprove universal optimality, with one systematic exception: we omit configurations that are reducible, in the sense of being unions of orthogonal, lower-dimensional configurations (this terminology is borrowed from the theory of root systems). Reducible configurations are in no sense less interesting or fruitful than irreducible ones. For example, cross polytopes can be reduced all the way to one-dimensional pieces. However, including reducible configurations would substantially lengthen Table 3 without adding much more geometrical content.

n	N	t
7	182	$1/\sqrt{3}$
15	128	$1/5$

TABLE 4. Unresolved conjectured harmonic optima.

Table 4 lists two more unresolved cases. They both appear to be harmonic optima and are balanced, and we have not been able to prove or disprove universal optimality. Unlike the two cases in Table 2, we do not conjecture that they are universally optimal, because each is closely analogous to a case in which universal optimality fails (182 points in \mathbb{R}^7 is analogous to 126 points in \mathbb{R}^6 , and 128 points in \mathbb{R}^{15} is analogous to 256 points in \mathbb{R}^{16}). On the other hand, each is also analogous to a configuration we know or believe is universally optimal (240 points in \mathbb{R}^8 and 64 points in \mathbb{R}^{14} , respectively). We have not been able to disprove universal optimality in the cases in Table 4, but they are sufficiently large that our failure provides little evidence in favor of universal optimality.

Note that the data presented in Tables 3 and 4 may not specify the configurations uniquely. For example, for 48 points in \mathbb{R}^4 there is a positive-dimensional family of configurations with maximal inner product $1/\sqrt{2}$ (which is not the best possible value, according to Sloane’s tables [Sloane 00]). See Section 5 for explicit constructions of the conjectured harmonic optima.

It is worth observing that several famous configurations do not appear in Tables 3 and 4. Most notably, the cubes in \mathbb{R}^n with $n \geq 3$, the dodecahedron, the 120-cell, and the D_5 , E_6 , and E_7 root systems are suboptimal for harmonic energy. Many of these configurations have more than 64 points, but we have included in the tables all configurations we have analyzed, regardless of size.

In each case listed in Tables 3 and 4, our computer programs returned floating-point approximations to the coordinates of the points in the code, but we have been able to recognize the underlying structure exactly. That is possible largely because these codes are highly symmetric, and once one has uncovered the symmetries the remaining structure is greatly constrained. By contrast, for most numbers of points in most dimensions, we cannot even recognize the minimal harmonic energy as an exact algebraic number (although it must be algebraic, because it is definable in the first-order theory of the real numbers).

1.2 New Universal Optima

Both codes listed in Table 2 have been studied before. The first code was discovered by Conway, Sloane, and Smith [Sloane 00] as a conjecture for an optimal spherical code (and discovered independently by Hovinga [Hovinga 04]). The second can be derived from the Nordstrom–Robinson binary code [Nordstrom and Robinson 67] or as a spectral embedding of an association scheme discovered by de Caen and van Dam [de Caen and van Dam 99]

(take $t = 1$ in [de Caen and van Dam 99, Theorem 2 and Proposition 7(i)] and then project the standard orthonormal basis into a common eigenspace of the operators in the Bose–Mesner algebra of the association scheme). We describe both codes in greater detail in Section 4.

Neither code satisfies the condition from [Cohn and Kumar 07] for universal optimality: both are spherical 3-designs (but not 4-designs), with four distances between distinct points in the 40-point code and three in the 64-point code. That leaves open the possibility of an ad hoc proof, similar to the one Cohn and Kumar gave for the regular 600-cell, but the techniques from [Cohn and Kumar 07] do not apply.

To test universal optimality, we have carried out 1000 random trials with the potential function $f(r) = (4 - r)^k$ for each k from 1 to 25. We have also carried out 1000 trials using Hardin and Sloane’s program Gosset [Hardin and Sloane 03] to construct good spherical codes (to take care of the case when k is large). Of course these experimental tests fall far short of a rigorous proof, but the codes certainly appear to be universally optimal.

We believe that they are the only possible new universal optima consisting of at most 64 points, because we have searched the space of such codes fairly thoroughly. By [Cohn and Kumar 07, Proposition 1.4], any new universal optimum in \mathbb{R}^n must contain at least $2n + 1$ points. There are 812 such cases with at most 64 points in dimension at least 4. In each case, we have completed at least 1000 random trials (and usually more). There is no guarantee that we have found the global optimum in any of these cases, because it could have a tiny basin of attraction. However, a simple calculation shows that it is 99.99% likely that in every case we have found every local minimum that occurs at least 2% of the time. We have probably not always found the true optimum, but we believe that we have found every universal optimum within the range we have searched.

We have made our tables of conjectured harmonic optima for up to 64 points in up to 32 dimensions available via the world wide web.¹ They list the best energies we have found and the coordinates of the configurations that achieve them. We would be grateful for any improvements, and we intend to keep the tables up to date, with credit for any contributions received from others.

In addition to carrying out our own searches for universal optima, we have examined Sloane’s tables [Sloane 00] of the best spherical codes known with at most 130 points in \mathbb{R}^4 and \mathbb{R}^5 , and we have verified that they con-

tain no new universal optima. We strongly suspect that there are no undiscovered universal optima of any size in \mathbb{R}^4 or \mathbb{R}^5 , based on Sloane’s calculations as well as our searches, but it would be difficult to give definitive experimental evidence for such an assertion (we see no convincing arguments for why huge universal optima should not exist).

In general, our searches among larger codes have been far less exhaustive than those up to 64 points: we have at least briefly examined well over four thousand different pairs (n, N) , but generally not in sufficient depth to make a compelling case that we have found the global minimum. (Every time we found a balanced harmonic optimum, with the exception of 128 points in \mathbb{R}^{15} and 256 points in \mathbb{R}^{16} , we completed at least 1000 trials to test whether it was really optimal. However, we have not completed nearly as many trials in most other cases, and in any case 1000 trials is not enough when one is studying large configurations.) Nevertheless, our strong impression is that universal optima are rare, and certainly that there are few small universal optima with large basins of attraction.

2. METHODOLOGY

2.1 Techniques

As discussed in the introduction, to minimize potential energy we apply gradient descent, starting from many random initial configurations. That is an unsophisticated approach, because gradient descent is known to perform more slowly in many situations than competing methods such as the conjugate gradient algorithm. However, it has performed adequately in our computations. Furthermore, gradient descent has particularly intuitive dynamics. Imagine particles immersed in a medium with enough viscosity that they never build up momentum. When a force acts on them according to the potential function, the configuration undergoes gradient descent. By contrast, for most other optimization methods the motion of the particles is more obscure, so for example it is more difficult to interpret information such as sizes of basins of attraction.

Once we have approximate coordinates, we can use the multivariate analogue of Newton’s method to compute them to high precision (by searching for a zero of the gradient vector). Usually we do not need to do this, because the results of gradient descent are accurate enough for our purposes, but it is a useful tool to have available.

Obtaining coordinates is simply the beginning of our analysis. Because the coordinates encode not only the

¹Available at <http://aimath.org/data/paper/BBCGKS2006/>.

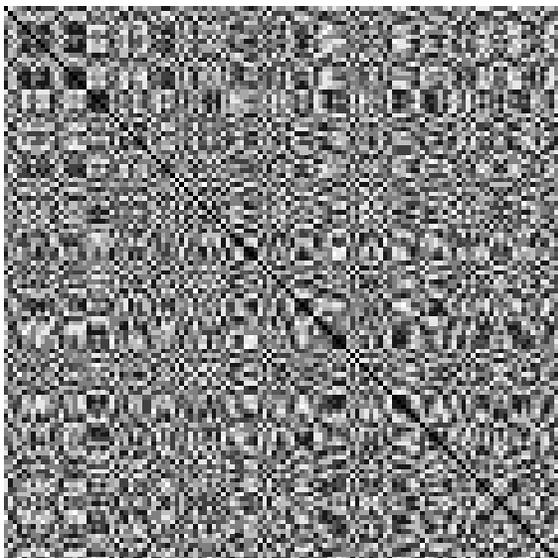


FIGURE 3. The Gram matrix for a regular 600-cell (black denotes 1, white denotes -1 , and gray interpolates between them), with the points ordered as returned by our gradient descent software.

relative positions of the points but also an arbitrary orthogonal transformation of the configuration, interpreting the data can be subtle. A first step is to compute the Gram matrix. In other words, given points x_1, \dots, x_N , compute the $N \times N$ matrix G whose entries are given by $G_{i,j} = \langle x_i, x_j \rangle$. The Gram matrix is invariant under orthogonal transformations, so it encodes almost precisely the information we care about. Its only drawback is that it depends on the arbitrary choice of how the points are ordered. That may sound like a mild problem, but there are many permutations of the points and it is far from clear how to choose one that best exhibits the configuration's underlying structure: compare Figure 3 with Figure 4.

With luck, one can recognize the entries of the Gram matrix as exact algebraic numbers: more frequently than one might expect, they are rational or quadratic irrationals. Once one specifies the entire Gram matrix, the configuration is completely determined, up to orthogonal transformations. Furthermore, one can easily prove that the configuration exists (keep in mind that it may not be obvious that there actually is such an arrangement of points, because it was arrived at via inexact calculations). To do so, one need only check that the Gram matrix is symmetric, it is positive semidefinite, and its rank is at most n . Every such matrix is the Gram matrix of a set of N points in \mathbb{R}^n , and if the diagonal entries are all 1 then the points lie on S^{n-1} .

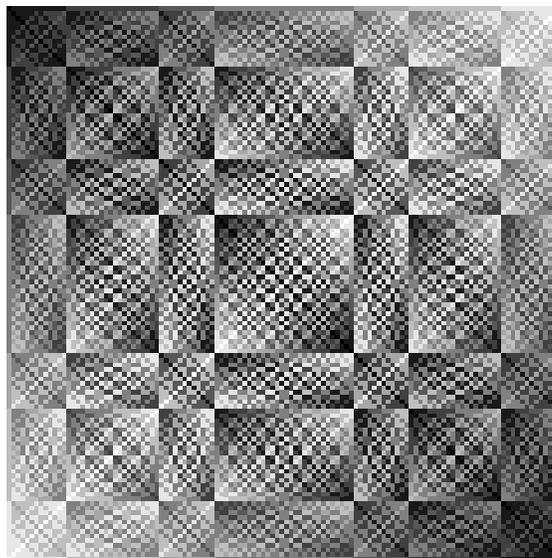


FIGURE 4. The Gram matrix for a regular 600-cell, with the points ordered so as to display structure.

Unfortunately, the exact Gram matrix entries are not always apparent from the numerical data. There is also a deeper reason why simply recognizing the Gram matrix is unsatisfying: it provides only the most “bare bones” description of the configuration. Many properties, such as symmetry or connections to other mathematical structures, are far from apparent given only the Gram matrix, as one can see from Figures 3 and 4.

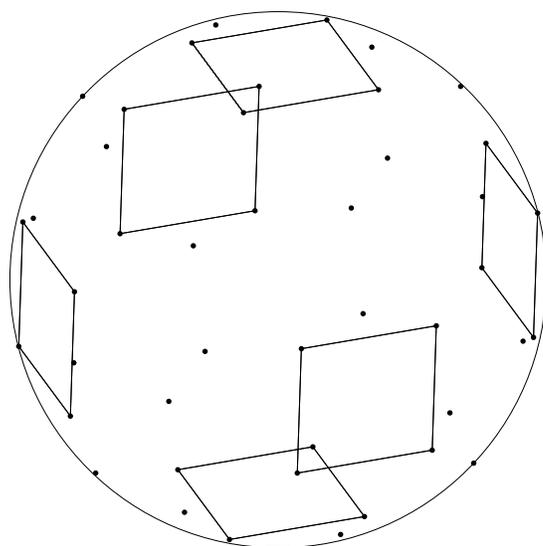


FIGURE 5. An orthogonal projection of the conjectured harmonic optimum with 44 points in \mathbb{R}^3 onto a random plane. Line segments connect points at the minimal distance.

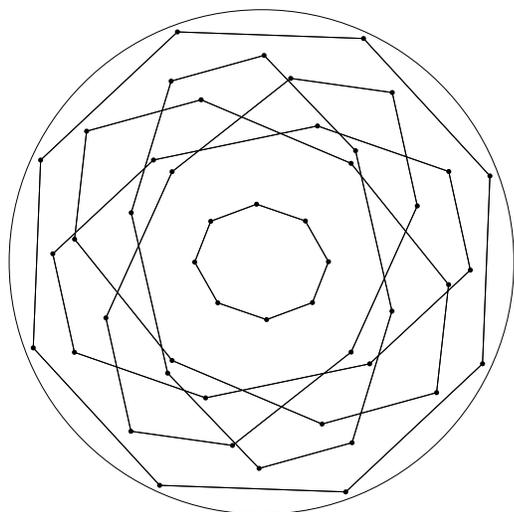


FIGURE 6. An orthogonal projection of the conjectured harmonic optimum with 48 points in \mathbb{R}^4 onto a random plane. Line segments connect points at the minimal distance.

Choosing the right method to visualize the data can make the underlying patterns clearer. For example, projections onto low-dimensional subspaces are often illuminating. Determining the most revealing projection can itself be difficult, but sometimes even a random projection sheds light on the structure. For example, Figures 5 and 6 are projections of the harmonic optima with 44 points in \mathbb{R}^3 and 48 points in \mathbb{R}^4 , respectively, onto random planes. The circular outline is the boundary of the projection of the sphere, and the line segments pair up points separated by the minimal distance. Figure 5 shows a disassembled cube (in a manner described later in this section), while Figure 6 is made up of octagons (see Section 3.5 for a description).

The next step in the analysis is the computation of the automorphism group. In general that is a difficult task, but we can make use of the wonderful software Nauty written by McKay [McKay 81]. Nauty can compute the automorphism group of a graph as a permutation group on the set of vertices; more generally, it can compute the automorphism group of a vertex-labeled graph. We make use of it as follows. Define a combinatorial automorphism of a configuration to be a permutation of the points that preserves inner products (equivalently, distances). If one forms an edge-labeled graph by placing an edge between each pair of points, labeled by their inner product, then the combinatorial automorphism group is the automorphism group of this labeled graph. Nauty is not directly capable of computing such a group, but it is straightforward to reduce the problem to that of computing the

automorphism group of a related vertex-labeled graph. Thus, one can use Nauty to compute the combinatorial automorphism group.

Fortunately, combinatorial automorphisms are the same as geometric symmetries, provided the configuration spans \mathbb{R}^n . Specifically, every combinatorial automorphism is induced by a unique orthogonal transformation of \mathbb{R}^n . (When the points do not span \mathbb{R}^n , the orthogonal transformations are not unique, because there are nontrivial orthogonal transformations that fix the subspace spanned by the configuration.) Thus, Nauty provides an efficient method for computing the symmetry group.

Unfortunately, it is difficult to be certain that one has computed the correct group. Two inner products that appear equal numerically may differ by a tiny amount, in which case the computed symmetry group may be too large. However, that is rarely a problem even with single-precision floating-point arithmetic, and it is difficult to imagine a fake symmetry that appears real to one hundred decimal places.

Once the symmetry group has been obtained, many further questions naturally present themselves. Can one recognize the symmetry group as a familiar group? How does its representation on \mathbb{R}^n break up into irreducibles? What are the orbits of its action on the configuration?

Analyzing the symmetries of the configuration frequently determines much of the structure, but usually not all of it. For example, consider the simplest nontrivial case, namely five points on S^2 . There are two natural ways to arrange them: with two antipodal points and three points forming an equilateral triangle on the orthogonal plane between them, or as a pyramid with four points forming a square in the hemisphere opposite a single point (and equidistant from it). In the first case everything is determined by the symmetries, but in the second there is one free parameter, namely how far the square is from the point opposite it. As one varies the potential function, the energy-minimizing value of this parameter will vary. (We conjecture that for every completely monotonic potential function, one of the configurations described in this paragraph globally minimizes the energy, but we cannot prove it.)

We define the *parameter count* of a configuration to be the dimension of the space of nearby configurations that can be obtained from it by applying an arbitrary radial force law between all pairs of particles. For example, balanced configurations are those with zero parameters, and the family with a square opposite a point has one parameter.

To compute the parameter count for an N -point configuration, start by viewing it as an element of $(S^{n-1})^N$ (by ordering the points). Within the tangent space of this manifold, for each radial force law there is a tangent vector. To form a basis for all these force vectors, look at all distances d that occur in the configuration, and for each of them consider the tangent vector that pushes each pair of points at distance d in opposite directions but has no other effects. All force vectors are linear combinations of these, and the dimension of the space they span is the parameter count for the configuration. (One must be careful to use sufficiently high-precision arithmetic, as when computing the symmetry group.)

This information is useful because in a sense it shows how much humanly understandable structure we can expect to find. For example, in the five-point configuration with a square opposite a point, the distance between them will typically be some complicated number depending on the potential function. In principle one can describe it exactly, but in practice it is most pleasant to treat it as a black box and describe all the other distances in the configuration in terms of it. The parameter count tells how many independent parameters one should expect to arrive at. When the count is zero or one, it is reasonable to search for an elegant description, whereas when the count is twenty, it is likely that the configuration is unavoidably complex.

Figure 7 shows the parameter counts of the conjectured harmonic optima in \mathbb{R}^3 with at most 64 points, compared with the dimension of the full space of all configurations of their size. The counts vary wildly but are often quite a bit smaller than one might expect. Two striking examples are 61 points with 111 parameters, for which there is likely no humanly understandable description, and 44 points with one parameter. The 44-point configuration consists of the vertices of a cube and centers of its edges together with the 24-point orbit (under the cube’s symmetry group) of a point on a diagonal of a face, all projected onto a common sphere. The optimal choice of the point on the diagonal appears complicated.

One subtlety in searching for local minima is that any given potential function will usually not detect all possible families of local minima that could occur for other potential functions. For example, for five points in \mathbb{R}^3 , the family with a square opposite a point does not contain a local minimum for harmonic energy. One can attain a local minimum compared to the other members of the family, but it will be a saddle point in the space of all configurations. Nevertheless, the family does contain local minima for some other completely mono-

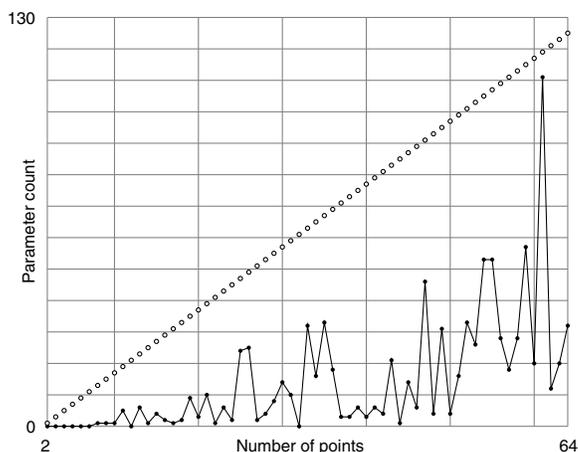


FIGURE 7. Parameter counts for conjectured harmonic optima in \mathbb{R}^3 . Horizontal or vertical lines occur at multiples of ten, and white circles denote the dimension of the configuration space.

tonic potential functions (such as $f(r) = 1/r^s$ with s large).

2.2 Example

For a concrete example, consider Table 5, which shows the results of 10^8 random trials for 27 points in \mathbb{R}^6 (all decimal numbers in tables have been rounded). These parameters were chosen because, as shown in [Cohn and Kumar 07], there is a unique 27-point universal optimum in \mathbb{R}^6 , with harmonic energy 111; it is called the Schläfli configuration. The column labeled “frequency” tells how many times each local minimum occurred. As one can see, the universal optimum occurred more than 99.97% of the time, but we found a total of four others.

Strictly speaking, we have not proved that the local minima listed in Table 5 (other than the Schläfli configuration) even exist. They surely do, because we have computed them to five hundred decimal places and checked that they are local minima by numerically diagonalizing the Hessian matrix of the energy function on the space of configurations. However, we used high-precision floating-point arithmetic, so this calculation does not constitute a rigorous proof, although it leaves no reasonable doubt. It is not at all clear whether there are additional local minima. We have not found any, but the fact that one of the local minima occurs only once in every ten million trials suggests that there might be others with even smaller basins of attraction.

The local minimum with energy 112.736... stands out in two respects besides its extreme rarity: it has many symmetries and it depends on few parameters. That sug-

Harmonic Energy	Frequency	Parameters	Maximal Cosine	Symmetries
111.0000000000	99971504	0	0.2500000000	51840
112.6145815185	653	9	0.4306480635	120
112.6420995468	22993	18	0.3789599707	24
112.7360209988	10	2	0.4015602076	1920
112.8896851626	4840	13	0.4041651631	48

TABLE 5. Local minima for 27 points in \mathbb{R}^6 (with frequencies out of 10^8 random trials).

gests that it should have a simple description, and in fact it does, as a modification of the universal optimum. Only two inner products occur between distinct points in the Schläfli configuration, namely $-\frac{1}{2}$ and $\frac{1}{4}$. In particular, it is not antipodal, so one can define a new code by replacing a single point x with its antipode $-x$. The remaining 26 points can be divided into two clusters according to their distances from $-x$. Immediately after replacing x with $-x$ the code will no longer be a local minimum, but if one allows it to equilibrate, a minimum is achieved. (That is not obvious: the code could equilibrate to a saddle point, because it is starting from an unusual position.) All that changes is the distances of the two clusters from $-x$, while the relative positions within the clusters remain unchanged (aside from rescaling). These two distances are the two parameters of the code. The symmetries of the new code are exactly those of the universal optimum that fix x , so the size of the symmetry group is reduced by a factor of 27.

The Schläfli configuration in \mathbb{R}^6 corresponds to the 27 lines on a smooth cubic surface: there is a natural correspondence between points in the configuration and lines on a cubic surface such that the inner products of $-\frac{1}{2}$ occur between points corresponding to intersecting lines. (This dates back to Schoutte [Schoutte 10]. See also the introduction to [Cohn and Kumar 07] for a brief summary of the correspondence.) One way to view the other local minima in Table 5 is as competitors to this classical configuration. It would be intriguing if they also had interpretations or consequences in algebraic geometry, but we do not know of any.

3. EXPERIMENTAL PHENOMENA

3.1 Analysis of Gram Matrices

For an example of how one might analyze a Gram matrix, consider the case of sixteen points in \mathbb{R}^5 . This case also has a universal optimum, which is the smallest known one that is not a regular polytope (although it is semiregular). It is the five-dimensional hemicube, which consists of half the vertices of the cube. More precisely, it contains the points $(\pm 1, \pm 1, \pm 1, \pm 1, \pm 1)/\sqrt{5}$ with an even

number of minus signs. One can recover the full cube by including the antipode of each point, so the symmetries of the five-dimensional hemicube consist of half of those of the five-dimensional cube (namely, those that preserve the hemicube, rather than exchanging it with its complementary hemicube).

It is essentially an accident of five dimensions that the hemicube is universally optimal. Universal optimality also holds in lower dimensions, but only because the hemicubes turn out to be familiar codes (two antipodal points in two dimensions, a tetrahedron in three dimensions, and a cross polytope in four dimensions). In six dimensions the hemicube appears to be an optimal spherical code, but it does not minimize harmonic energy and is therefore not universally optimal. In seven dimensions, and presumably all higher dimensions, the hemicube is not even an optimal code.

The five-dimensional hemicube has the same structure as the Clebsch graph (see Figure 8). The sixteen points correspond to the vertices of the graph; two distinct points have inner product $-\frac{3}{5}$ if they are connected by an edge in the graph and $\frac{1}{5}$ otherwise. This determines the Gram matrix and hence the full configuration.

For the harmonic potential energy, the hemicube appears to be the only local minimum with sixteen points in \mathbb{R}^5 , but we do not know how to prove that. To construct another local minimum, one can attempt constructions

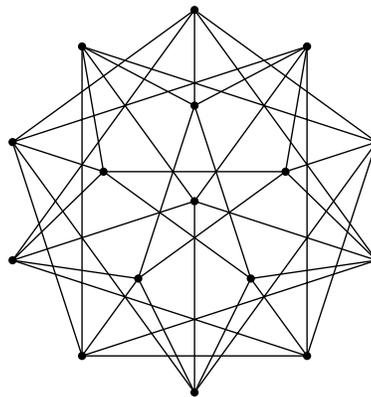


FIGURE 8. The Clebsch graph.

1	a	a	a	a	a	a	b									
a	1	e	e	a^2	a^2	a^2	d	c	c	c	c	d	c	d	c	
a	e	1	e	a^2	a^2	a^2	c	d	c	d	c	c	c	c	d	
a	e	e	1	a^2	a^2	a^2	c	c	d	c	d	c	d	c	c	
a	a^2	a^2	a^2	1	e	e	d	c	c	c	d	c	c	c	d	
a	a^2	a^2	a^2	e	1	e	c	d	c	c	d	d	c	c	c	
a	a^2	a^2	a^2	e	e	1	c	c	d	d	c	c	c	d	c	
b	d	c	c	d	c	c	1	f	f	f	g	g	g	f	g	g
b	c	d	c	c	d	c	f	1	f	g	f	g	g	g	f	g
b	c	c	d	c	c	d	f	f	1	g	g	f	g	g	g	f
b	c	d	c	c	c	d	f	g	g	1	f	f	f	g	g	g
b	c	c	d	d	c	c	g	f	g	f	1	f	g	f	f	g
b	d	c	c	c	d	c	g	g	f	f	f	1	g	g	f	f
b	c	c	d	c	d	c	f	g	g	f	g	g	1	f	f	f
b	d	c	c	c	c	d	g	f	g	g	f	f	f	1	f	f
b	c	d	c	d	c	c	g	g	f	g	g	f	f	f	1	f

TABLE 6. Gram matrix for 16 points in \mathbb{R}^5 ; here $c = ab + (1/2)\sqrt{(1-a^2)(1-b^2)/2}$, $d = ab - \sqrt{(1-a^2)(1-b^2)/2}$, $e = (3a^2 - 1)/2$, $f = (3b^2 - 1)/2$, and $g = (3b^2 + 1)/4$.

such as moving a point to its antipode, as in Section 2.2, but they yield saddle points. However, for other potential functions one sometimes finds other local minima (we have found up to two other nontrivial local minima). To illustrate the techniques from the previous section, we will analyze one of them here. It will turn out to have a fairly simple conceptual description; our goal here is to explain how to arrive at it, starting from the Gram matrix.

The specific example we will analyze arises as a local minimum for the potential function $r \mapsto (4 - r)^{12}$. It is specified by Table 6 with

$$a \approx -0.499890010934$$

and

$$b \approx 0.201039702365$$

(the lines in the table are just for visual clarity).

The first step is to recognize the structure in the Gram matrix. Table 6 highlights this structure, but of course it takes effort to bring the Gram matrix into such a simple form (by recognizing algebraic relations between the Gram matrix entries and reordering the points so as to emphasize patterns). The final form of the Gram matrix exhibits the configuration as belonging to a family specified by two parameters a and b with absolute value less than 1. As described in the table's caption, all the other inner products are simple algebraic functions of a and b . To check that this Gram matrix corresponds to an actual code in S^4 , it suffices to verify that its eigenvalues are 0 (11 times), $1 + 6a^2 + 9b^2$, and $(15 - (6a^2 + 9b^2))/4$ (4 times): there are only five nonzero eigenvalues and they are clearly positive.

Table 6 provides a complete description of the configuration, but it is unilluminating. To describe the code using elegant coordinates, one must have a more conceptual understanding of it. A first step in that direction is the observation that the first point in Table 6 has inner product a or b with every other point. In other words, the remaining 15 points lie on two parallel four-dimensional hyperplanes, equidistant from the first point. A natural guess is that as a and b vary, the structures within these hyperplanes are simply rescaled as the corresponding cross sections of the sphere change in size, and some calculation verifies that this guess is correct.

To understand these two structures and how they relate to each other, set $a = b = 0$, so that they form a 15-point configuration in \mathbb{R}^4 . Its Gram matrix is of course obtained by removing the first row and column of Table 6 and setting $a = b = 0$, $e = f = -1/2$, $g = 1/4$, $c = \sqrt{2}/4$, and $d = -\sqrt{2}/2$. The two substructures consist of the first six points and the last nine, among the fifteen remaining points.

Understanding the 16-point codes in \mathbb{R}^5 therefore simply comes down to understanding this single 15-point code in \mathbb{R}^4 . (It is also the 15-point code from Table 3. Incidentally, Sloane's tables [Sloane 00] show that it is not an optimal spherical code.) The key to understanding it is choosing the right coordinates. The first six points form two orthogonal triangles, and they are the simplest part of this configuration, so it is natural to start with them.

Suppose the points v_1, v_2, v_3 and v_4, v_5, v_6 form two orthogonal equilateral triangles in a four-dimensional vector space. The most natural coordinates to choose for the

1	c	c	d	b	b	$-2a$	a	a	$-2a$	a	a
c	1	c	b	d	b	a	$-2a$	a	a	$-2a$	a
c	c	1	b	b	d	a	a	$-2a$	a	a	$-2a$
d	b	b	1	c	c	$-2a$	a	a	$-2a$	a	a
b	d	b	c	1	c	a	$-2a$	a	a	$-2a$	a
b	b	d	c	c	1	a	a	$-2a$	a	a	$-2a$
$-2a$	a	a	$-2a$	a	a	1	c	c	d	b	b
a	$-2a$	a	a	$-2a$	a	c	1	c	b	d	b
a	a	$-2a$	a	a	$-2a$	c	c	1	b	b	d
$-2a$	a	a	$-2a$	a	a	d	b	b	1	c	c
a	$-2a$	a	a	$-2a$	a	b	d	b	c	1	c
a	a	$-2a$	a	a	$-2a$	b	b	d	c	c	1

TABLE 7. Gram matrix for 12 points in \mathbb{R}^4 ; here $0 < a < \frac{1}{2}$, $b = a - 1$, $c = -3a + 1$, and $d = 4a - 1$.

vector space are the inner products with these six points. Of course the sum of the three inner products with any triangle must vanish (because $v_1 + v_2 + v_3 = v_4 + v_5 + v_6 = 0$), so there are only four independent coordinates, but we prefer not to break the symmetry by discarding two coordinates.

The other nine points in the configuration are determined by their inner products with v_1, \dots, v_6 . Each of them will have inner product d with one point in each triangle and c with the remaining two points. As pointed out above we must have $d + 2c = 0$, and in fact $d = -\sqrt{2}/2$ and $c = \sqrt{2}/4$ because the points are all unit vectors. Note that one can read off all this information from the c and d entries in Table 6.

There is an important conceptual point in the last part of this analysis. Instead of focusing on the internal structure among the last nine points, it is most fruitful to study how they relate to the previously understood subconfiguration of six points. However, once one has a complete description, it is important to examine the internal structure as well.

The pattern of connections among the last nine points in Table 6 is described by the Paley graph on nine vertices, which is the unique strongly regular graph with parameters $(9, 4, 1, 2)$. (The Paley graph is isomorphic to its own complement, so the edges could correspond to inner product either f or g .) Strongly regular graphs, and more generally association schemes, frequently occur as substructures of minimal-energy configurations. It is remarkable to see such highly ordered structures spontaneously occurring via energy minimization.

3.2 Other Small Examples

To illustrate some of the other phenomena that can occur, in this subsection we will analyze the case of 12 points in \mathbb{R}^4 . We have observed two families of local

minima, both of which are slightly more subtle than the previous examples.

For $0 < a < \frac{1}{2}$, set $b = a - 1$, $c = -3a + 1$, and $d = 4a - 1$, and consider the Gram matrix shown in Table 7 (its nonzero eigenvalues are $12a$ and $6 - 12a$, each with multiplicity 2). Unlike the example in Section 3.1, the symmetry group acts transitively on the points, so there are no distinguished points to play a special role in the analysis. Nevertheless, one can analyze it as follows.

Let $v_1, v_2, v_3 \in S^1$ be the vertices of an equilateral triangle in \mathbb{R}^2 , and let v_4 and v_5 be unit vectors that are orthogonal to each other and to each of v_1, v_2 , and v_3 . For $0 < \alpha < 1$, consider the twelve points $\alpha v_i \pm \sqrt{1 - \alpha^2} v_4$ and $-\alpha v_i \pm \sqrt{1 - \alpha^2} v_5$ with $1 \leq i \leq 3$. If one sets $a = \alpha^2/2$ then they have Table 7 as a Gram matrix.

The Gram matrix shown in Table 8 is quite different. There, $0 < a < \frac{1}{3}$, $b = 1 - 12a^2$, $c = 6a^2 - 1$, and $d = 18a^2 - 1$. The nonzero eigenvalues are $\frac{4}{3} + 24a^2$ (with multiplicity 3) and $8 - 72a^2$, which are positive because $a < \frac{1}{3}$.

In this Gram matrix the first four points form a distinguished tetrahedron, and the remaining eight points form two identical tetrahedra. They lie in hyperplanes parallel to and equidistant from the (equatorial) hyperplane containing the distinguished tetrahedron. If one sets $a = \frac{1}{3}$, then all three tetrahedra lie in the same hyperplane, with $b = -\frac{1}{3}$, $c = -\frac{1}{3}$, $d = 1$, and $-3a = -1$. In particular, one can see that the two parallel tetrahedra are in dual position to the distinguished tetrahedron. As the parameter a varies, all that changes is the distance between the parallel hyperplanes. (As a tends to zero some points coincide. One could also use a between 0 and $-\frac{1}{3}$, but that corresponds to using parallel tetrahedra oriented the same way, instead of dually, which generally yields higher potential energy.)

1	$-\frac{1}{3}$	$-\frac{1}{3}$	$-\frac{1}{3}$	$-3a$	a	a	a	$-3a$	a	a	a
$-\frac{1}{3}$	1	$-\frac{1}{3}$	$-\frac{1}{3}$	a	$-3a$	a	a	a	$-3a$	a	a
$-\frac{1}{3}$	$-\frac{1}{3}$	1	$-\frac{1}{3}$	a	a	$-3a$	a	a	a	$-3a$	a
$-\frac{1}{3}$	$-\frac{1}{3}$	$-\frac{1}{3}$	1	a	a	a	$-3a$	a	a	a	$-3a$
$-3a$	a	a	a	1	b	b	b	d	c	c	c
a	$-3a$	a	a	b	1	b	b	c	d	c	c
a	a	$-3a$	a	b	b	1	b	c	c	d	c
a	a	a	$-3a$	b	b	b	1	c	c	c	d
$-3a$	a	a	a	d	c	c	c	1	b	b	b
a	$-3a$	a	a	c	d	c	c	b	1	b	b
a	a	$-3a$	a	c	c	d	c	b	b	1	b
a	a	a	$-3a$	c	c	c	d	b	b	b	1

TABLE 8. Gram matrix for 12 points in \mathbb{R}^4 ; here $0 < a < \frac{1}{3}$, $b = 1 - 12a^2$, $c = 6a^2 - 1$, and $d = 18a^2 - 1$.

This sort of layered structure occurs surprisingly often. One striking example is 74 points in \mathbb{R}^5 . The best such spherical code known consists of a regular 24-cell on the equatorial hyperplane together with two dual 24-cells on parallel hyperplanes as well as the north and south poles. If one chooses the two parallel hyperplanes to have inner products $\pm\sqrt{\sqrt{5}-2}$ with the poles, then the cosine of the minimal angle is exactly $(\sqrt{5}-1)/2$. That agrees numerically with Sloane's tables [Sloane 00] of the best codes known, but of course there is no proof that it is optimal.

There is almost certainly no universally optimal 12-point configuration in \mathbb{R}^4 . Aside from some trivial examples for degenerate potential functions, the two cases we have analyzed in this subsection are the only two types of local minima we have observed. For $f(r) = (4-r)^k$ with $k \in \{1, 2\}$ they both achieve the same minimal energy (along with a positive-dimensional family of other configurations). For $3 \leq k \leq 9$ the first family appears to achieve the global minimum, while for $k \geq 10$ the second appears to. As k tends to infinity the energy minimization problem turns into the problem of finding the optimal spherical code. That problem appears to be solved by taking $a = \frac{1}{4}$ in the second family, so that the minimal angle has cosine $\frac{1}{4}$, which agrees with Sloane's tables [Sloane 00].

We conjecture that one or the other of these two families minimizes each completely monotonic potential function. This conjecture is somewhat difficult to test, but we are not aware of any counterexamples.

The examples we have analyzed so far illustrate three basic principles:

1. Small or medium-sized local minima tend to occur in low-dimensional families as one varies the potential function. The dimension is not usually as low as in these examples, but it is typically far lower than

the dimension of the space of all configurations (see Figure 7).

2. These families frequently contain surprisingly symmetrical substructures (such as regular polytopes or configurations described by strongly regular graphs or other association schemes).
3. The same substructures and construction methods occur in many different families.

3.3 $2n + 1$ Points in \mathbb{R}^n

Optimal spherical codes are known for up to $2n$ points in \mathbb{R}^n (see [Böröczky 04, Theorem 6.2.1]), but not for $2n+1$ points, except in \mathbb{R}^2 and \mathbb{R}^3 . Here we present a natural conjecture for all dimensions.

These codes consist of a single point we call the north pole together with two n -point simplices on hyperplanes parallel to the equator; the simplices are in dual position relative to each other. Each point in the simplex closer to the north pole will have inner product α with the north pole, and the inner product between any two points in the further simplex will be α . The number α can be chosen so that each point in either one of the simplices has inner product α with each point in the other simplex except the point furthest from it. To achieve that, α must be the unique root between 0 and $1/n$ of the cubic equation

$$(n^3 - 4n^2 + 4n)x^3 - n^2x^2 - nx + 1 = 0.$$

As $n \rightarrow \infty$, $\alpha = 1/n - \sqrt{2}/n^{3/2} + O(1/n^2)$.

Let $\mathcal{C}_n \subset S^{n-1}$ be this spherical code, with α chosen as above. The cosine of the minimal angle in \mathcal{C}_n is α .

Conjecture 3.1. *For each $n \geq 2$, the code \mathcal{C}_n is an optimal spherical code. Furthermore, every optimal $(2n+1)$ -point code in S^{n-1} is isometric to \mathcal{C}_n .*

On philosophical grounds it seems reasonable to expect to be able to prove this conjecture: most of the difficulty in packing problems comes from the idiosyncrasies of particular spaces and dimensions, so when a phenomenon occurs systematically one expects a conceptual reason for it. However, we have made no serious progress toward a proof.

One can also construct \mathcal{C}_n as follows. Imagine adding one point to a regular cross polytope by placing it in the center of a facet. The vertices of that facet form a simplex equidistant from the new point, as do the vertices of the opposite facet. The structure is identical to the code \mathcal{C}_n , except for the distances from the new point, and the proper distances can be obtained by allowing the code to equilibrate with respect to increasingly steep potential functions.

It appears that for $n > 2$ these codes do not minimize harmonic energy, so they are not universally optimal. When $n = 4$, something remarkable occurs with the (conjectured) minimum for harmonic energy. That configuration consists of a regular pentagon together with two pairs of antipodal points that are orthogonal to each other and the pentagon. If one uses gradient descent to minimize harmonic energy, it seems to converge with probability 1 to this configuration, but the convergence is very slow, much slower than for any other harmonic energy minimum we have found. The reason is that this configuration is a degenerate minimum for the harmonic energy, in the sense that the Hessian matrix has more zero eigenvalues than one would expect.

Each of the nine points has three degrees of freedom, so the Hessian matrix has twenty-seven eigenvalues. Specifically, they are 0 (ten times), 4, $7/4$ (twice), $9/2$ (four times), 9 (twice), $25/8 \pm \sqrt{209}/8$ (twice each), and $31/8 \pm \sqrt{161}/8$ (twice each). Six of the zero eigenvalues are unsurprising, because they come from the problem's invariance under the six-dimensional Lie group $O(4)$, but the remaining four are surprising indeed.

The corresponding eigenvectors are infinitesimal displacements of the nine points that produce only a fourth-order change in energy, rather than the expected second-order change. To construct them, do not move the antipodal pairs of points at all, and move the pentagon points orthogonally to the plane of the pentagon. Each must be displaced by $(1 - \sqrt{5})/2$ times the sum of the displacements of its two neighbors. This yields a four-dimensional space of displacements, which are the surprising eigenvectors.

This example is noteworthy because it shows that harmonic energy is not always a Morse function on the space

of all configurations. One might hope to apply Morse theory to understand the relationship between critical points for energy and the topology of the configuration space, but the existence of degenerate critical points could substantially complicate this approach.

3.4 $2n + 2$ Points in \mathbb{R}^n

After seeing a conjecture for the optimal $(2n + 1)$ -point code in S^{n-1} , it is natural to wonder about $2n + 2$ points. A first guess is the union of a simplex and its dual simplex (in other words, the antipodal simplex), which was named the diplo-simplex by Conway and Sloane [Conway and Sloane 91]. One can prove using the linear programming bounds for real projective space that this code is the unique optimal antipodal spherical code of its size and dimension (see [Conway and Sloane 99, Chapter 9]), but for $n > 2$ it is not even locally optimal as a spherical code (see Section 7) and we do not have a conjecture for the true answer.

For the problem of minimizing harmonic energy, the diplo-simplex is suboptimal for $3 \leq n \leq 5$ but appears optimal for all other n .

One particularly elegant case is $n = 4$. The midpoints of the edges of a regular simplex form a 10-point code in S^3 with maximal inner product $\frac{1}{6}$, and Bachoc and Vallentin [Bachoc and Vallentin 09] have proved that it is the unique optimal spherical code. It is also the kissing configuration of the five-dimensional hemicube (the universally optimal 16-point configuration in \mathbb{R}^5). In other words, it consists of the ten nearest neighbors of any point in that code. This code appears to minimize harmonic energy, but it is not the unique minimum: two orthogonal regular pentagons have the same harmonic energy.

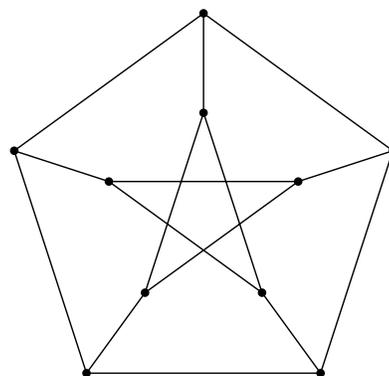


FIGURE 9. The Petersen graph.

As pointed out in the introduction of [Cohn and Kumar 07], this code is not universally optimal, but it nevertheless seems to be an exceedingly interesting configuration. Only the inner products $-\frac{2}{3}$ and $\frac{1}{6}$ occur (besides 1, of course). If one forms a graph whose vertices are the points in the code and whose edges correspond to pairs of points with inner product $-\frac{2}{3}$, then the result is the famous Petersen graph (Figure 9).

Like the nontrivial universal optima in dimensions 5 through 8, this code consists of the vertices of a semiregular polytope that has simplices and cross polytopes as facets, with a simplex and two cross polytopes meeting at each face of codimension 3. Its kissing configuration is also semiregular, with square and triangular facets, but it is a suboptimal code (specifically, a triangular prism).

3.5 48 Points in \mathbb{R}^4

One of the most beautiful configurations we have found is a 48-point code in \mathbb{R}^4 . The points form six octagons that map to the vertices of a regular octahedron under the Hopf map from S^3 to S^2 . Recall that if we identify \mathbb{R}^4 with \mathbb{C}^2 using the inner product $\langle x, y \rangle = \operatorname{Re} \bar{x}^t y$ on \mathbb{C}^2 , then the Hopf map sends (z, w) to $z/w \in \mathbb{C} \cup \{\infty\}$, which we can identify with S^2 via stereographic projection to a unit sphere centered at the origin. The fibers of the Hopf map are the circles given by intersecting S^3 with the complex lines in \mathbb{C}^2 .

Sloane, Hardin, and Cara [Sloane et al. 03] found a spherical 7-design of this form, consisting of two dual 24-cells, and it has the same minimal angle as our code (which is the minimal angle in an octagon), but it is a different code. In \mathbb{C}^2 , the Sloane–Hardin–Cara code is the union of the orbits under multiplication by eighth roots of unity of the points $(1, 0)$, $(0, 1)$, $(\pm 1, 1)/\sqrt{2}$, and $(\pm i, 1)/\sqrt{2}$. Our code is the union of the orbits of $(1, 0)$, $(0, 1)$, $(\pm \zeta, \zeta)/\sqrt{2}$, and $(\pm i\zeta^2, \zeta^2)/\sqrt{2}$, where $\zeta = e^{\pi i/12}$. Each octagon has been rotated by a multiple of $\pi/12$ radians. Because a regular octagon is invariant under rotation by $\pi/4$ radians, there are only three distinct rotations by multiples of $\pi/12$. Each such rotation occurs for the octagons lying over two antipodal vertices of the octahedron in the base space S^2 of the Hopf fibration.

It is already remarkable that performing these rotations yields a balanced configuration with lower harmonic energy than the union of the 24-cell and its dual, but the structure of the code’s convex hull is especially noteworthy. The facets can be computed using the program Polymake [Gawrilow and Joswig 00]. The facets of the dual 24-cell configuration are 288 irregular tetrahedra, all equivalent under the action of the symmetry group

(and each possessing eight symmetries). By contrast, our code has 128 facets forming two orbits under the symmetry group: one orbit of 96 irregular tetrahedra and one of 32 irregular octahedra. The irregular octahedra are obtained from regular ones by rotating one of the facets, which are equilateral triangles, by an angle of $\pi/12$. We will use the term “twisted facets” to denote the rotated facet and its opposite facet (by symmetry, either one could be viewed as rotated relative to the other).

The octahedra in our configuration meet other octahedra along their twisted facets and tetrahedra along their other facets. Grouping the octahedra according to adjacency therefore yields twisted chains of octahedra. Each chain consists of eight octahedra, and they span the 3-sphere along great circles. The total twist amounts to $8\pi/12 = 2\pi/3$, from which it follows that the chains close with facets aligned correctly. The 32 octahedra form four such chains, and the corresponding great circles are fibers in the same Hopf fibration as the vertices of the configuration. These Hopf fibers map to the vertices of a regular tetrahedron in S^2 . It is inscribed in the cube dual to the octahedron formed by the images of the vertices of the code.

Another way to view the facets of this polytope, or any spherical polytope, is as holes in the spherical code. More precisely, the (outer) facet normals of any full-dimensional polytope inscribed in a sphere are the holes in the spherical code (i.e., the points on the sphere that are local maxima for distance from the code). The normals of the octahedral facets are the deep holes in this code (i.e., the points at which the distance is globally maximized). Notice that these points are defined using the intrinsic geometry of the sphere, rather than relying on its embedding in Euclidean space.

The octahedral facets of our code can be thought of as more important than the tetrahedral facets. The octahedra appear to us to have prettier, clearer structure, and once they have been placed, the entire code is determined (the tetrahedra simply fill the gaps). This idea is not mathematically precise, but it is a common theme in many of our calculations: when we examine the facet structure of a balanced code, we often find a small number of important facets and a large number of less meaningful ones.

3.6 Hopf Structure

As in the previous example, many notable codes in S^3 , S^7 , or S^{15} can be understood using the complex, quaternionic, or octonionic Hopf maps (see for example

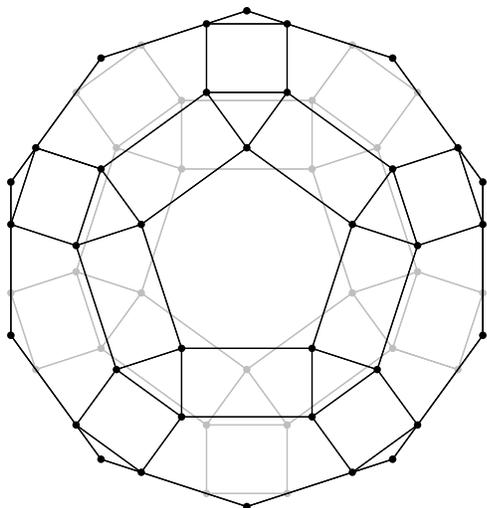


FIGURE 10. The 60-point polytope in \mathbb{R}^3 over which the regular 120-cell fibers.

[Dixon 95] and [Altschuler and Pérez-Garrido 07]). In this subsection, we describe this phenomenon for the regular 120-cell and 600-cell in S^3 . The Hopf structure on the 600-cell is mathematical folklore, but we have not been able to locate it in the published literature, while the case of the 120-cell is more subtle and may not have been previously examined.

The H_4 reflection group (which is the symmetry group of both polytopes) contains elements of order 10 that act on \mathbb{R}^4 with no eigenvalues other than primitive tenth roots of unity. If one chooses such an element, then \mathbb{R}^4 has the structure of a two-dimensional complex vector space such that this element acts via multiplication by a primitive tenth root of unity. The orbits are regular 10-gons lying in Hopf fibers. In the case of the regular 600-cell, this partitions the 120 vertices into 12 regular 10-gons lying in Hopf fibers over the vertices of a regular icosahedron in S^2 . For the regular 120-cell (with 600 vertices), the corresponding polyhedron in S^2 has 60 vertices, but it is far from obvious what it is. We know of no way to determine it without calculation, but computing with coordinates reveals that it is a distorted rhombicosidodecahedron, with the square facets replaced by golden rectangles. Specifically, its facets are 12 regular pentagons, 20 equilateral triangles, and 30 golden rectangles. The golden rectangles meet pentagons along their long edges and triangles along their short edges. Figure 10 shows the orthogonal projection into the plane containing a pentagonal face (gray vertices and edges are on the far side of the polyhedron).

3.7 Facet Structure of Universal Optima

The known low-dimensional universal optima (through dimension 8) are all regular or semiregular polytopes, whose facets are well known to be regular simplices or cross polytopes. However, there seems to have been little investigation of the facets of the higher-dimensional universal optima from Table 1. In this subsection we will look at the smallest higher-dimensional cases: $U_{100,22}$, $U_{112,21}$, $U_{162,21}$, $U_{275,22}$, $U_{552,23}$, and $U_{891,22}$ (recall that $U_{N,n}$ denotes the N -point code in \mathbb{R}^n from Table 1, when it is unique). Each of the first four is a two-distance set given by a spectral embedding of a strongly regular graph. (Recall that a spectral embedding is obtained by orthogonally projecting the standard orthonormal basis into an eigenspace of the adjacency matrix of the graph.) The last two have three distances between distinct points.

These codes have enormous numbers of facets (more than seventy-five trillion for $U_{552,23}$), so it is not feasible to find the facets using general-purpose methods. Instead, one must make full use of the large symmetry groups of these configurations. With Dutour Sikirić's package Polyhedral [Dutour Sikirić 06] for the program GAP [GAP Group 06], this can be done for these configurations. We have used it to compute complete lists of orbits of facets under the action of the symmetry group. (The results are rigorous, because we use exact coordinates for the codes. In particular, when necessary we use the columns of the Gram matrices to embed scalar multiples of these codes isometrically into high-dimensional spaces using only rational coordinates.) Of course, the results of this computation then require analysis by hand to reveal their structure.

For an introductory example, it is useful to review the case of the five-dimensional hemicube (see Section 3.1). It has ten obvious facets contained in the ten facets of the cube. Each is a four-dimensional hemicube, i.e., a regular cross polytope. The remaining facets are regular simplices (one opposite each point in the hemicube).

One can view the five-dimensional hemicube as an antiprism formed by two four-dimensional cross polytopes in parallel hyperplanes. The cross polytopes are arranged so that each one's vertices point toward deep holes of the other. (The deep holes of a cross polytope are the vertices of the dual cube, and in four dimensions the vertices of that cube consist of two cross polytopes. The fact that the deep holes of a four-dimensional cross polytope contain another such cross polytope is crucial for this construction to make sense.) Of course, the distance between the parallel hyperplanes is chosen so as to max-

imize the minimal distance. What is remarkable about this antiprism is that it is far more symmetrical than one might expect: normally the two starting facets of an antiprism play a very different role from the facets formed by taking the convex hull, but in this case extra symmetries occur. The simplest case of such extra symmetries is the construction of a cross polytope as an antiprism made from two regular simplices in dual position.

The three universal optima $U_{100,22}$, $U_{112,21}$, and $U_{162,21}$ are each given by an unusually symmetric antiprism construction analogous to that of the hemicube. In each case, the largest facets (i.e., those containing the most vertices) contain half the vertices. These facets are themselves spectral embeddings of strongly regular graphs (the Hoffman–Singleton graph, the Gewirtz graph, and the unique $(81, 20, 1, 6)$ strongly regular graph). Within the universal optima, the largest facets occur in pairs in parallel hyperplanes, and the vertices of each facet in a pair point toward holes in the other. These holes belong to a single orbit under the symmetry group of the facet, and that orbit is the disjoint union of several copies of the vertices of the facet: two copies for the Hoffman–Singleton and Gewirtz cases and four in the third case. These holes are the deepest holes in the Hoffman–Singleton case; in the other two cases, they are not quite the deepest holes (there are not enough deep holes for the construction to work using them).

Brouwer and Haemers [Brouwer and Haemers 92, Brouwer and Haemers 93] discovered the underlying combinatorics of these constructions (i.e., that the strongly regular graphs corresponding to the universal optima can be naturally partitioned into two identical graphs). However, the geometric interpretation as antiprisms appears to be new.

The universal optima $U_{100,22}$, $U_{112,21}$, and $U_{162,21}$ are antiprisms, but that cannot possibly be true for $U_{275,22}$, because 275 is odd. Instead, the McLaughlin configuration $U_{275,22}$ is analogous to the Schläfli configuration $U_{27,6}$. Both are two-distance sets. In the Schläfli configuration, the neighbors of each point form a five-dimensional hemicube and the nonneighbors form a five-dimensional cross polytope. Both the hemicube and the cross polytope are unusually symmetric antiprisms, and their vertices point toward each other's deep holes. (The deep holes of the hemicube form a cross polytope, and those of the cross polytope form a cube consisting of two hemicubes.) The McLaughlin configuration is completely analogous: the neighbors of each point form $U_{162,21}$ and the nonneighbors form $U_{112,21}$. They point toward each other's deep holes; this is possible because the deep holes

of $U_{112,21}$ consist of four copies of $U_{162,21}$, and its deep holes consist of two copies of $U_{112,21}$. Furthermore, the deep holes in these two universal optima are of exactly the same depth (i.e., distance to the nearest point in the code), as is also the case for the five-dimensional cross polytope and hemicube used to form the Schläfli configuration.

The Schläfli and McLaughlin configurations both have the property that their deep holes are the antipodes of their vertices. Thus, it is natural to form antiprisms from two parallel copies of them, with vertices pointed at each other's deep holes. That yields antipodal configurations of 54 points in \mathbb{R}^7 and 550 points in \mathbb{R}^{23} . If one also includes the two points orthogonal to the parallel hyperplanes containing the original two copies, then this construction gives the universal optima $U_{56,7}$ and $U_{552,23}$.

Each high-dimensional universal optimum has many types of facets of different sizes. For example, the facets of the Higman–Sims configuration $U_{100,22}$ form 123 orbits under the action of the symmetry group (see Table 9). The largest facets, which come from the Hoffman–Singleton graph as described above, are by far the most important, but each type of facet appears to be of interest. They are often more subtle than one might expect. For example, it is natural to guess that the facets with 42 vertices would be regular cross polytopes, based on the number of vertices, but they are not. Instead, when rescaled to the unit sphere they have the following structure:

The facets with 42 vertices are two-distance sets on the unit sphere in \mathbb{R}^{21} , with inner products $1/29$ and $-13/29$. If we define a graph on the vertices by letting edges correspond to pairs with inner product $-13/29$, then this graph is the bipartite incidence graph for points and lines in the projective plane $\mathbb{P}^2(\mathbb{F}_4)$. To embed this graph in \mathbb{R}^{21} , represent the 21 points in $\mathbb{P}^2(\mathbb{F}_4)$ as the permutations of (a, b, \dots, b) , where $a^2 + 20b^2 = 1$ and $2ab + 19b^2 = 1/29$. Specifically, take $a = 0.9977\dots$ and $b = 0.0151\dots$ (these are fourth-degree algebraic numbers). Choose c and d so that $5c^2 + 16d^2 = 1$ and $8cd + c^2 + 12d^2 = 1/29$ (specifically, take $c = -0.4362\dots$ and $d = 0.0550\dots$). Then embed the 21 lines into \mathbb{R}^{21} as permutations of $(c, c, c, c, c, d, \dots, d)$, where the five c entries correspond to the points contained in the line. This embedding gives the inner products of $1/29$ and $-13/29$, as desired (and in fact those are the only inner products for which a construction of this form is possible).

As shown in Table 9, there are 92 different types of simplicial facets in the Higman–Sims configuration. One orbit consists of regular simplices: for each point in the

Vertices	Number of Orbits
22	92
23	13
24	6
25	3
27	3
28	1
30	1
31	1
36	1
42	1
50	1

TABLE 9. Number of orbits of facets of different sizes in the Higman–Sims configuration $U_{100,22}$.

configuration, the 22 points at the furthest distance from it form a regular simplex. All the other simplices are irregular. Nine orbits consist of simplices with no symmetries whatsoever, and the remaining ones have some symmetries but not the full symmetric group.

The universal optima $U_{552,23}$ and $U_{891,22}$ have more elaborate facet structures, but we have completely classified their facets (which form 116 and 422 orbits, respectively). The facets corresponding to their deep holes form single orbits, consisting of $U_{100,22}$ in the first case and $U_{162,21}$ in the second. These results can all be understood in terms of the standard embeddings of these configurations into the Leech lattice, as follows:

Let v be any vector with norm 6 in the Leech lattice Λ_{24} . Among the 196,560 minimal vectors in Λ_{24} (those with norm 4), there are 552 minimal vectors w satisfying $|w - v|^2 = 4$, and they form a copy of the 552-point universal optimum. This shows that $U_{552,23}$ is a facet of $U_{196560,24}$. Taking kissing configurations shows that $U_{275,22}$ is a facet of $U_{4600,23}$ and that $U_{162,21}$ is a facet of $U_{891,22}$. We conjecture that each of these facets corresponds to a deep hole in the code, and that all of the deep holes arise in this way, but we have not proved this conjecture beyond $U_{891,22}$ (the $U_{196560,24}$ case has since been proved in [Dutour Sikirić et al. 09]). The $U_{100,22}$ facets of $U_{552,23}$ can also be seen in this picture: given two vectors $v_1, v_2 \in \Lambda_{24}$ with $|v_1|^2 = |v_2|^2 = 6$ and $|v_1 - v_2|^2 = 4$, the corresponding $U_{552,23}$ facets of $U_{196560,24}$ intersect in a $U_{100,22}$ facet of $U_{552,23}$, which corresponds to a deep hole.

3.8 96 Points in \mathbb{R}^9

Another intriguing code that arose in our computer searches is a 96-point code in \mathbb{R}^9 (see Table 3). This code was known previously: it is mentioned but not described in [Conway and Sloane 99, Table 9.2], which refers to a

paper in preparation that never appeared, and it is described in [Ericson and Zinoviev 01, Appendix D]. Here we describe it in detail, with a different approach from that in [Ericson and Zinoviev 01].

The code is not universally optimal, but it is balanced and it appears to be an optimal spherical code. What makes it noteworthy is that the cosine of its minimal angle is $\frac{1}{3}$. Any such code corresponds to an arrangement of unit balls in \mathbb{R}^{10} that are all tangent to two fixed tangent balls, where the interiors of the balls are not allowed to overlap (this condition forces the cosine of the minimal angle between the sphere centers to be at most $\frac{1}{3}$ when the angle is centered at the midpoint between the fixed balls). The largest such arrangement most likely consists of 96 balls.

To construct the code, consider three orthogonal tetrahedra in \mathbb{R}^9 . Call the points in the first v_1, v_2, v_3, v_4 , in the second v_5, v_6, v_7, v_8 , and in the third $v_9, v_{10}, v_{11}, v_{12}$. Within each of these tetrahedra, all inner products between distinct points are $-\frac{1}{3}$, and between tetrahedra they are all 0. Call these tetrahedra the *basic tetrahedra*.

The points $\pm v_1, \dots, \pm v_{12}$ will all be in the code, and we will identify 72 more points in it. Each of the additional points will have inner product $\pm\frac{1}{3}$ with each of v_1, \dots, v_{12} , and we will determine them via those inner products. Because $v_1 + \dots + v_4 = v_5 + \dots + v_8 = v_9 + \dots + v_{12} = 0$, the inner products with the elements of each basic tetrahedron must sum to zero. In particular, two must be $\frac{1}{3}$ and the other two $-\frac{1}{3}$. That restricts us to $\binom{4}{2} = 6$ patterns of inner products with each basic tetrahedron, so there are $6^3 = 216$ points satisfying all the constraints so far. We must cut that number down by a factor of 3.

The final constraint comes from considering the inner products between the new points. A simple calculation shows that one can reconstruct a point x from its inner products with v_1, \dots, v_{12} via

$$x = \frac{3}{4} \sum_{i=1}^{12} \langle x, v_i \rangle v_i,$$

and inner products are computed via

$$\langle x, y \rangle = \frac{3}{4} \sum_{i=1}^{12} \langle x, v_i \rangle \langle y, v_i \rangle.$$

In other words, if x and y have identical inner products with one of the basic tetrahedra, that contributes $\frac{1}{3}$ to their own inner product. If they have opposite inner products with one of the basic tetrahedra, that contributes $-\frac{1}{3}$. Otherwise the contribution is 0.

Vertices	Automorphisms	Orbit Size
9	16	27648
9	48	13824
9	48	4608
9	96	18432
9	1440	4608
12	1024	864
12	31104	512
16	10321920	18

TABLE 10. Facets of the convex hull of the configuration of 96 points in \mathbb{R}^9 , modulo the action of the symmetry group.

We wish to avoid the situation in which x and y have identical inner products with two basic tetrahedra, or opposite inner products with both, and neither identical nor opposite inner products with the third. In that case, $\langle x, y \rangle = \pm \frac{2}{3}$.

To rule out this situation, we assign elements of $\mathbb{Z}/3\mathbb{Z}$ to quadruples by

$$\pm \left(\frac{1}{3}, \frac{1}{3}, -\frac{1}{3}, -\frac{1}{3} \right) \mapsto 0,$$

$$\pm \left(\frac{1}{3}, -\frac{1}{3}, \frac{1}{3}, -\frac{1}{3} \right) \mapsto 1,$$

and

$$\pm \left(\frac{1}{3}, -\frac{1}{3}, -\frac{1}{3}, \frac{1}{3} \right) \mapsto -1.$$

Consider the 72 points with inner products $\pm \frac{1}{3}$ with each of v_1, \dots, v_{12} such that exactly two inner products with each basic tetrahedron are $\frac{1}{3}$ and furthermore the elements of $\mathbb{Z}/3\mathbb{Z}$ coming from the inner products with the basic tetrahedra sum to 0. Given any two such points, if they have identical or opposite inner products with two basic tetrahedra, then the same must be true with the third. Thus, we have constructed $24 + 72 = 96$ points in \mathbb{R}^9 such that all the inner products between them are $\pm 1, \pm \frac{1}{3}$, or 0.

The facets of this code form eight orbits under the action of its symmetry group; they are listed in Table 10. The most interesting facets are those with 16 vertices, which form regular cross polytopes. These facets and the two orbits with 12 vertices all correspond to deep holes.

3.9 Distribution of Energy Levels

Typically, there are many local minima for harmonic energy. One intriguing question is how the energies of the local minima are distributed. For example, Table 11

Energy	Frequency
5395.000000	186418
5398.650556	4393
5398.687876	2356
5400.842726	18
5400.880057	149
5400.890460	47
5400.894513	26
5400.928674	25
5400.936106	41
5400.940237	28
5400.940550	7
5400.943094	38
5402.029556	7
5402.088248	3
5402.093726	10
5402.116636	1
5402.152619	1
5402.213231	2
5402.366164	1
5402.922701	1
5403.091064	111
5403.115123	1
5403.129076	108
5403.271100	66
5403.319898	157
5403.326719	84
5403.347209	24
5403.455701	7
5403.462898	8
5403.488923	4

TABLE 11. Thirty lowest harmonic energies observed for local minima with 120 points on S^3 ($2 \cdot 10^5$ trials, 5223 different energy levels observed).

shows the thirty lowest energies obtained in $2 \cdot 10^5$ trials with 120 points in \mathbb{R}^4 , together with how often they occurred. The regular 600-cell is the unique universal optimum (with energy 5395), but we found 5223 different energy levels. This table is probably not a complete list of the lowest thirty energies (five of them occurred only once, so it is likely there are more to be found), but we suspect that we have found the true lowest ten.

The most remarkable aspect of Table 11 is the three gaps in it. There are huge gaps from 5395 to 5398.65, from 5398.69 to 5400.84, and from 5400.95 to 5402.02. Each gap is far larger than the typical spacing between energy levels. Perhaps one of these gaps contains some rare local minima, but they appear to be real gaps.

What could cause such gaps? We do not have a complete theory, but we believe the gaps reflect bottlenecks in the process of constructing the code by gradient descent. Figure 11 is a graph of energy as a function of time for gradient descent, starting at a random configu-

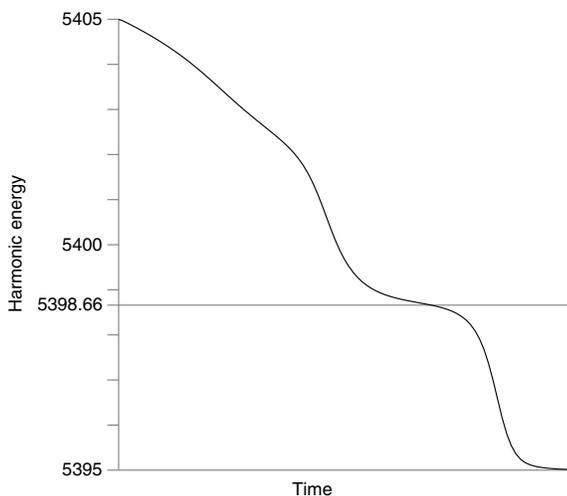


FIGURE 11. Energy as a function of time under gradient descent.

ration of 120 points in S^3 . (The figure represents a single run of the optimization procedure, so it should be viewed as a case study, not a statistical argument.) The graph begins when the energy has just reached 5405 and ends once convergence to the universal optimum is apparent. Of course the convergence is monotonic, but its speed varies dramatically. The rate of decrease is slowest near energy 5398.66, which is indicated by a horizontal line. We do not believe it could be a coincidence that that is very nearly the energy of the two lowest-energy local minima in Table 11: the slowdown probably occurs because of a bottleneck. More precisely, in order to achieve the ground state the system must develop considerable large-scale order and symmetry. Probably short-range order develops first and then slowly extends to long-range order. During this process there may be bottlenecks in which different parts of the system must come into alignment with each other. The local minima correspond to the rare cases in which the system gets stuck in the middle of a bottleneck, but even when it does not get stuck it still slows down.

It is not completely clear why the gaps are separated by several energy levels that are surprisingly close to each other. The most likely explanation is that there are several slightly different ways to get stuck during essentially the same bottleneck, but we have no conceptual understanding of what these ways are. It would be very interesting to have a detailed theory of this sort of symmetry breaking.

4. CONJECTURED UNIVERSAL OPTIMA

4.1 40 Points in \mathbb{R}^{10}

The 40-point code from Table 2 consists of 40 points on the unit sphere in \mathbb{R}^{10} . The only inner products that occur are 1 , $-\frac{1}{2}$, $-\frac{1}{3}$, 0 , and $\frac{1}{6}$; each point has these inner products with 1, 8, 3, 4, and 24 points, respectively. Grouping pairs of points according to their inner product yields a 4-class association scheme, which Bannai, Bannai, and Bannai [Bannai et al. 08] have recently shown is uniquely determined by its intersection numbers.

The 40 points form 10 regular tetrahedra, and more specifically 5 orthogonal pairs of tetrahedra. That accounts for all the inner products of 1 , $-\frac{1}{3}$, and 0 . Each point has inner product $-\frac{1}{2}$ with one point in each of the other 9 tetrahedra, except for the tetrahedron orthogonal to the one containing it. All remaining inner products are $\frac{1}{6}$. The configuration is chiral (i.e., not equivalent to any reflection of itself under the action of $SO(10)$) and has a symmetry group of order 1920. Specifically, the symmetry group is the semidirect product of the symmetric group S_5 with the subgroup of $(\mathbb{Z}/2\mathbb{Z})^5$ consisting of all vectors that sum to zero, where S_5 acts by permuting the five coordinates.

We conjecture that this code is the unique 40-point code in \mathbb{R}^{10} with maximal inner product $\frac{1}{6}$, but that appears difficult to prove. For example, it is an even stronger assertion than optimality as a spherical code. We are unaware of any occurrence of this code in the published literature, but it appears in Sloane's online tables [Sloane 00] with the annotation that it was found by Smith and "beautified" by Conway and Sloane (in the sense of recognizing it using elegant coordinates). It also appears in Hovinga's online report [Hovinga 04].

Conway, Sloane, and Smith construct the code as an explicit list of 40 vectors with entries $-1/\sqrt{6}$, 0 , or $1/\sqrt{6}$. Here we explain the combinatorics underlying this construction. That may well be how Conway and Sloane beautified the code, but [Sloane 00] presents no details about the construction beyond the list of vectors.

Consider the 10-dimensional vector space V spanned by the orthonormal basis vectors $v_{\{i,j\}}$ for all two-element subsets $\{i,j\} \subset \mathbb{Z}/5\mathbb{Z}$ with $i \neq j$. Say a vector in V has type i if for every basis vector v_S , its coefficient vanishes if and only if $i \in S$. Each vector in the code will have type i for some $i \in \mathbb{Z}/5\mathbb{Z}$, and the six nonzero coefficients will equal $\pm 1/\sqrt{6}$. Given such a vector, define a graph on the vertex set $(\mathbb{Z}/5\mathbb{Z}) \setminus \{i\}$ by connecting j to k if the coefficient of $v_{\{j,k\}}$ is $-1/\sqrt{6}$. If the graph has e edges and vertex j has degree d_j , then the vector is in the code if

and only if $d_{i+1} \equiv d_{i+2} \equiv e \pmod{2}$ and $d_{i-1} \equiv d_{i-2} \not\equiv e \pmod{2}$ (where i is the type of the vector).

The facet structure of this code seems surprisingly unilluminating compared to the others we have analyzed. There are 24 orbits of facets: 21 orbits of irregular simplices (with symmetry groups ranging from 2 to 768 in size), one orbit of 11-vertex facets, and two orbits of 12-vertex facets. The most symmetrical facets are those in one of the two orbits of 12-vertex facets. They have 3072 symmetries and are given by the orthogonal union of three identical irregular tetrahedra in \mathbb{R}^3 (with eight symmetries).

Our confidence in this code’s universal optimality is based on detailed numerical experiments. One reason a configuration could be universally optimal is that it has no competitors (i.e., except in degenerate cases there are no other local minima). That is not true for the 40-point code, but there are remarkably few competitors. In particular, it appears to have only one “serious” competitor. We have found five other families of local minima, but four of them are rare and never seem to come close to beating our conjectured universal optimum. The best experimental evidence we can imagine for universal optimality would be to describe explicitly each competing family that has been observed and prove that it never contains the global energy minimum. That might be possible for this code, but we have not completed it. The four rare families are sufficiently complicated that we have not analyzed them explicitly (under the circumstances it did not seem worth the effort). However, we have a complete description of the serious competitor.

That family depends on a parameter α that must satisfy $0 < \alpha^2 \leq 1/27$. The configuration always contains a fixed 16-point subset with the following structure. If we call the 16 points $w_{i,j}$ with $i, j \in \{1, 2, 3, 4\}$, then

$$\langle w_{i,j}, w_{k,\ell} \rangle = \begin{cases} 1 & \text{if } (i, j) = (k, \ell), \\ -\frac{1}{3} & \text{if } i = k \text{ or } j = \ell \text{ but not both, and} \\ \frac{1}{9} & \text{otherwise.} \end{cases}$$

In other words, if we arrange the points in a 4×4 grid, then the rows and columns are regular tetrahedra and all other inner products are $\frac{1}{9}$. To construct such a configuration, take the tensor product in $\mathbb{R}^3 \otimes_{\mathbb{R}} \mathbb{R}^3 = \mathbb{R}^9$ of two regular tetrahedra in \mathbb{R}^3 .

To describe the remaining 24 points, we will specify their inner products with the first 16 points. That will determine their projections into the 9-dimensional subspace containing the 16 points, so the only additional information needed to pin them down will be whether they

are above or below that hyperplane (relative to some orientation).

Each of the 24 points will have inner product -3α with four of the first 16 points and α with each of the others. The only constraint is that it must have inner product -3α with exactly four points, one in each of the eight tetrahedra (i.e., one in each row and column of the 4×4 grid). The 4×4 grid exhibits a one-to-one correspondence with permutations of four elements, so there are $4! = 24$ ways to satisfy these constraints. The points corresponding to even permutations will be placed above the 9-dimensional hyperplane, and those corresponding to odd permutations will be placed below it. This construction yields a 40-point code in \mathbb{R}^{10} provided that $0 < \alpha^2 \leq 1/27$. (When $\alpha = 0$ some points coincide, and when $\alpha^2 > 1/27$ the inner products cannot be achieved by unit vectors.)

We have not proved that the codes in this family never improve on the conjectured universal optimum, but we are confident that it is true. The best spherical code in the family occurs when $\alpha = (\sqrt{109} - 1)/54 = 0.1748\dots$; in this special case, α is also the maximal inner product, which is quite a bit larger than the maximal inner product $\frac{1}{6}$ in the conjectured universal optimum. That implies that when k is sufficiently large, the conjectured optimum is better for the potential function $f(r) = (4 - r)^k$ (because the energy is dominated asymptotically by the contribution from the minimal distance). In principle one could verify the finitely many remaining values of k by a finite computation. We have done enough exploration to convince ourselves that it is true, but we have not found a rigorous proof.

4.2 64 Points in \mathbb{R}^{14}

The simplest construction we are aware of for the 64-point configuration in \mathbb{R}^{14} uses the Nordstrom–Robinson binary code [Nordstrom and Robinson 67, Griess 08]. Shortening that code twice yields a binary code of length 14, size 64, and minimal distance 6, which is known to be unique [MacWilliams and Sloane 77, pp. 74–75]. One can view it as a subset of the cube $\{-1, 1\}^{14}$ instead of $\{0, 1\}^{14}$. Then after rescaling by a factor of $1/\sqrt{14}$ to yield unit vectors, this code is the 64-point configuration in \mathbb{R}^{14} that we conjecture is universally optimal. The same process with less shortening yields the codes from Tables 3 and 4 with 128 points in \mathbb{R}^{15} and 256 points in \mathbb{R}^{16} ; the 64-point and 128-point codes have previously appeared in [Ericson and Zinoviev 01, Appendix D] via the same approach, as conjectures for optimal spherical codes.

Inner Product	Condition
1	$(x_1, x_2) = (y_1, y_2)$
$-\frac{1}{7}$	$x_1 = y_1$ but $x_2 \neq y_2$
$-\frac{3}{7}$	$x_1 \neq y_1$ and $x_2 + y_2 \in \{(x_1 + y_1)^3, x_1 y_1 (x_1 + y_1)\}$
$\frac{1}{7}$	otherwise

TABLE 12. The inner products for the 64-point code.

This construction makes some of the facet structure of the code clear. There are 28 facets with 32 vertices that come from the facets of the cube containing the code. In fact, the code is obtained from exactly the same antiprism construction as described in Section 3.7 (the vertices of these facets point toward deep holes of the opposite facets). There are also 66 other orbits of facets under the action of the symmetry group, but those orbits seem to be less interesting.

An alternative construction of the code is by describing its Gram matrix explicitly. As mentioned above, this construction amounts to forming an association scheme by taking $t = 1$ in [de Caen and van Dam 99, Theorem 2 and Proposition 7(i)], and then performing a spectral embedding. Bannai, Bannai, and Bannai [Bannai et al. 08] have shown that this association scheme is uniquely determined by its intersection numbers.

More concretely, the points correspond to elements of \mathbb{F}_8^2 , where \mathbb{F}_8 is the finite field of order 8, and the inner products are determined by Table 12. The Gram matrix has 14 eigenvalues equal to $32/7$ and the others equal to 0, so it is indeed the Gram matrix of a 14-dimensional configuration.

Unfortunately, this 64-point code has many competitors. We have found over two hundred local minima for harmonic energy and expect that the total number is much larger. That makes it difficult to imagine an iron-clad experimental argument for universal optimality. We suspect that the code is universally optimal for two reasons: we have failed to find any counterexample, and during the process no competitor came close enough to worry us. (By contrast, in many cases in which one can disprove universal optimality, one finds worryingly close competitors before tweaking the construction to complete the disproof.) However, we realize that the evidence is far from conclusive.

5. BALANCED IRREDUCIBLE HARMONIC OPTIMA

In this section we briefly describe each of the configurations in Tables 3 and 4.

32 points in \mathbb{R}^3 : The union of a regular icosahedron and its dual dodecahedron.

10 points in \mathbb{R}^4 : Both configurations are described in Section 3.4.

13 points in \mathbb{R}^4 : In \mathbb{C}^2 with the inner product $\langle x, y \rangle = \operatorname{Re} \bar{x}^t y$, the points are $(\zeta/\sqrt{2}, \zeta^5/\sqrt{2})$, where ζ runs over all 13th roots of unity. This code was discovered by Sloane, Hardin, and Cara [Sloane et al. 03].

For a less compact description, view \mathbb{R}^4 as the orthogonal direct sum of two planes, and let R be the operation of rotating the first plane by $2\pi/13$ and the second by five times that angle. The unit sphere in \mathbb{R}^4 contains the direct product of the circles of radius $1/\sqrt{2}$ in the two planes, and the 13-point code is the orbit of a point in this direct product under the group generated by R .

The factor of 5 is special because $5^2 \equiv -1 \pmod{13}$. In particular, R^8 rotates the second plane by $2\pi/13$ and the first by 8 times that amount, which is the same as 5 times it in the opposite direction. In other words, the two planes play the same role, if one ignores their orientations. Only the square roots of -1 modulo 13 (or, trivially, the square roots of 1) have that property.

For all the points in this configuration, both complex coordinates have absolute value $1/\sqrt{2}$. Therefore the configuration is contained in a flat two-dimensional torus sitting inside S^3 (namely, the product of the circles of radius $1/\sqrt{2}$ in the two complex coordinate axes). Figure 12 shows the complex phases of the 13 points. For each N we can ask whether there is an N -point harmonic optimum in S^3 that is contained in such a torus. For $1 \leq N \leq 10$ it seems that there is. We conjecture that $N = 13$ is the only larger value of N for which this happens.

For other potential functions, similar phenomena can occur in more cases. For example, for the logarithmic potential function $f(r) = -\log r$, Jaron

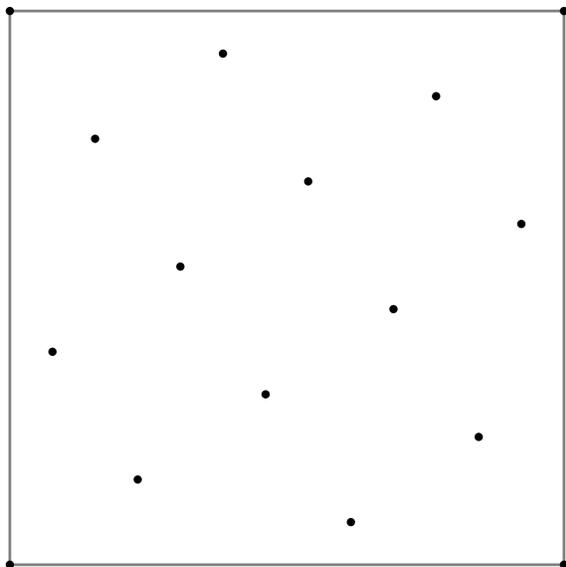


FIGURE 12. The 13-point harmonic optimum in S^3 , drawn on a torus by plotting (ϕ, ψ) for the point $(e^{i\phi}/\sqrt{2}, e^{i\psi}/\sqrt{2})$.

Lanier has conjectured in a private communication that the minimal-energy configuration of 11 points on S^3 also lies on a flat torus. In \mathbb{C}^2 , the points are $(\alpha\zeta, \sqrt{1-\alpha^2}\zeta^4)$, where ζ runs over all 11th roots of unity and α is the unique root between 0 and 1 of

$$5\alpha^8 - 36\alpha^6 + 51\alpha^4 - 4\alpha^2 - 7 = 0.$$

15 points in \mathbb{R}^4 : Described in Section 3.1.

24 points in \mathbb{R}^4 : The regular 24-cell (equivalently, the D_4 root system).

48 points in \mathbb{R}^4 : Described in Section 3.5.

21 points in \mathbb{R}^5 : The edge midpoints and face centers of a regular simplex (rescaled to lie on the same sphere).

32 points in \mathbb{R}^5 : Start with a regular simplex v_1, \dots, v_6 in \mathbb{R}^5 . The 32 points will be these six points and their negatives, along with 20 points determined as follows by their inner products with v_1, \dots, v_6 . Each will have inner product $\pm 1/\sqrt{5}$ with each of v_1, \dots, v_6 , with three plus signs and three minus signs. There are $\binom{6}{3} = 20$ ways to choose these signs.

$2n + 2$ points in \mathbb{R}^n (for $n \geq 6$): Described in Section 3.4.

42 points in \mathbb{R}^6 : The edge midpoints of a regular simplex and their antipodes.

44 points in \mathbb{R}^6 : This code contains plus or minus an orthonormal basis of \mathbb{R}^6 together with the 32 vectors whose coordinates with respect to that basis are $\pm 1/\sqrt{6}$ and where the number of minus signs is even. In other words, it consists of a cross polytope and a hemicube within the cube dual to the cross polytope. This code was previously conjectured to be an optimal spherical code (see [Ericson and Zinoviev 01, Table D.6]).

126 points in \mathbb{R}^6 : The union of the minimal vectors of the E_6 and E_6^* lattices (rescaled to lie on the same sphere). Equivalently, one can project the E_7 root system orthogonally to a minimal vector in E_7^* , followed by rescaling as in the first construction.

78 points in \mathbb{R}^7 : Like the 44 points in \mathbb{R}^6 , this code consists of a cross polytope and a hemicube within the cube dual to the cross polytope.

148 points in \mathbb{R}^7 : The points are all the permutations of

$$\frac{(\pm 1, \pm 1, 0, 0, 0, 0, 0)}{\sqrt{2}}$$

and the hemicube consisting of the points

$$\frac{(\pm 1, \pm 1, \pm 1, \pm 1, \pm 1, \pm 1, \pm 1)}{\sqrt{7}}$$

that have an even number of minus signs. This is a seven-dimensional analogue of a construction of the E_8 root system, but it is less symmetric, because the two types of points form distinct orbits under the action of the symmetry group. This code was previously conjectured to be an optimal spherical code (see [Ericson and Zinoviev 01, Table D.7]).

182 points in \mathbb{R}^7 : The union of the minimal vectors of the E_7 and E_7^* lattices (rescaled to lie on the same sphere). Equivalently, one can project the E_8 root system orthogonally to any root; 126 roots are unchanged, 112 project to 56 nonzero points, and 2 project to the origin. Rescaling the nonzero projections to lie on the unit sphere yields the 182-point configuration.

72 points in \mathbb{R}^8 : The edge midpoints of a regular simplex and their antipodes. This code was previously conjectured to be an optimal spherical code (see [Ericson and Zinoviev 01, Table D.8]).

96 points in \mathbb{R}^9 : Described in Section 3.8.

42 points in \mathbb{R}^{14} : This code consists of seven disjoint five-dimensional regular simplices, whose vertices have inner products $-\frac{1}{5}$ with each other. Each point has inner product $-\frac{1}{2}$ with a unique point in each simplex other than the one containing it, and all other inner products between points in different simplices are $\frac{1}{10}$. Grouping pairs of points according to their inner product yields a three-class association scheme that can be derived from the Hoffman–Singleton graph (see [van Dam 99, Section 5.1]). Let G be the Hoffman–Singleton graph, and let H be its second subconstituent. In other words, take any vertex v in G , and let H be the vertices not equal to or adjacent to v . The vertices in H correspond to points in the 14-dimensional configuration, with inner product $-\frac{1}{2}$ between adjacent vertices, $-\frac{1}{5}$ between nonadjacent vertices with no common neighbor, and $\frac{1}{10}$ between nonadjacent vertices with one common neighbor. This code was previously conjectured to be an optimal spherical code (see [Ericson and Zinoviev 01, Table D.14]).

128 points in \mathbb{R}^{15} : Described in Section 4.2.

256 points in \mathbb{R}^{16} : Described in Section 4.2.

6. CHALLENGES

We conclude with a list of computational and theoretical challenges:

1. How often do 196,560 randomly chosen points on S^{23} converge to the Leech lattice minimal vectors under gradient descent for harmonic energy? For 240 points on S^7 , one frequently obtains the E_8 root system (855 times out of 1000 trials), and the Higman–Sims configuration of 100 points on S^{21} occurs fairly often (257 times out of 1000 trials); by contrast, the universal optimum with 112 points on S^{20} occurs rarely (once in 1000 trials). This approach could be an intriguing construction of the Leech lattice, but we have no intuition for how likely it is to work.
2. What are the potential energy barriers that separate the local minima in Table 5 (or any other case)? In other words, if one continuously transforms one configuration into another, how low can one make the greatest energy along the path connecting them? The lowest possible point of greatest energy will always be a saddle point for the energy function.

3. How many harmonic local minima are there for 120 points on S^3 (Table 11), or even 64 points on S^{13} ? Is the number small enough that one could conceivably compile a complete list? Is the list for 27 points on S^5 in Table 5 complete?
4. For large numbers of points, what can one say (experimentally, heuristically, or rigorously) about the distribution of energy levels for local minima, or about the gaps in the distribution?

7. APPENDIX: LOCAL NONOPTIMALITY OF DIPLO-SIMPLICES

In this appendix we prove that (above dimension two) diplo-simplices are not even locally optimal as spherical codes. The motivation behind the calculations is the case of $n = 3$, where the diplo-simplex is a cube. One can improve this spherical code by rotating one face relative to the opposite face while moving them slightly closer together. To carry out this approach in higher dimensions, we must understand the faces of the diplo-simplex (which is in general quite different from the hypercube). The diplo-simplex in \mathbb{R}^n is the orthogonal projection of the vertices of the cross polytope in \mathbb{R}^{n+1} onto the hyperplane on which the sum of the coordinates is zero. Its dual polytope is therefore the cross section of the hypercube in \mathbb{R}^{n+1} by that hyperplane. The vertices of the cross section differ depending on whether n is odd or even: $(1, \dots, 1, -1, \dots, -1)$ is a vertex when $n+1$ is even, and $(1, \dots, 1, 0, -1, \dots, -1)$ is when $n+1$ is odd. That leads to the following description of the diplo-simplex, based on identifying its faces using this correspondence.

Let v_1, \dots, v_k and w_1, \dots, w_k be unit vectors forming two orthogonal $(k-1)$ -dimensional regular simplices, and let t be a unit vector orthogonal to both simplices. The $(2k-1)$ -dimensional diplo-simplex consists of the points

$$\pm \left(\alpha v_i + \sqrt{1 - \alpha^2} t \right)$$

and

$$\pm \left(\alpha w_i + \sqrt{1 - \alpha^2} t \right),$$

where $\alpha = \sqrt{(2k-2)/(2k-1)}$. To improve its minimal angle, use the points

$$\begin{aligned} & \alpha v_i + \sqrt{1 - \alpha^2} t, \\ & \alpha w_i + \sqrt{1 - \alpha^2} t, \\ & - \left(\alpha \left(\beta v_i + \sqrt{1 - \beta^2} w_i \right) + \sqrt{1 - \alpha^2} t \right), \\ & - \left(\alpha \left(\sqrt{1 - \beta^2} v_i - \beta w_i \right) + \sqrt{1 - \alpha^2} t \right), \end{aligned}$$

with β a small positive number and α slightly greater than $\sqrt{(2k-2)/(2k-1)}$. When $k=1$, choosing α and β optimally leads to the optimal 8-point code in \mathbb{R}^3 , but when $k=3$, Sloane's tables [Sloane 00] show that this approach is suboptimal.

The situation is slightly more complicated for even-dimensional diplo-simplices. Again, let v_1, \dots, v_k and w, \dots, w_k be unit vectors forming two orthogonal $(k-1)$ -dimensional simplices. Let t and z be two unit vectors orthogonal to both simplices and each other. The $2k$ -dimensional diplo-simplex consists of the points $z, -z,$

$$\pm \left(\alpha v_i + \beta z + \sqrt{1 - \alpha^2 - \beta^2} t \right),$$

and

$$\pm \left(\alpha w_i - \beta z + \sqrt{1 - \alpha^2 - \beta^2} t \right),$$

where $\alpha = \sqrt{(2k+1)(2k-2)}/(2k)$ and $\beta = 1/(2k)$. Note that α vanishes when $k=1$ and our construction below will not work. Indeed the hexagon, which is the diplo-simplex in the plane, is an optimal spherical code. To improve the minimal angle for $k \geq 2$ use the points $z, -z,$

$$\begin{aligned} & \alpha v_i + \beta z + \sqrt{1 - \alpha^2 - \beta^2} t, \\ & \alpha w_i - \beta z + \sqrt{1 - \alpha^2 - \beta^2} t, \\ & - \left(\alpha \left(\gamma v_i + \sqrt{1 - \gamma^2} w_i \right) - \beta z + \sqrt{1 - \alpha^2 - \beta^2} t \right), \\ & - \left(\alpha \left(\sqrt{1 - \gamma^2} v_i - \gamma w_i \right) + \beta z + \sqrt{1 - \alpha^2 - \beta^2} t \right), \end{aligned}$$

with γ a small positive number, α slightly larger than $\sqrt{(2k+1)(2k-2)}/(2k)$, and β slightly less than $1/(2k)$. The numbers α and β should be chosen such that $\alpha^2 + 2\beta^2$ increases from its original value of $(2k-1)/(2k)$.

ACKNOWLEDGMENTS

We thank Eiichi Bannai, Christian Borgs, John Conway, Edwin van Dam, Charles Doran, Noam Elkies, Florian Gaisendrees, Robert Griess, Abhinav Kumar, Jaron Lanier, James Morrow, Frank Stillinger, and Salvatore Torquato for helpful discussions.

Ballinger and Giansiracusa were supported by the University of Washington Mathematics Department's NSF VIGRE grant. Schürmann was supported by the Deutsche Forschungsgemeinschaft (DFG) under grant SCHU 1503/4-1.

REFERENCES

- [Altschuler and Pérez-Garrido 05] E. L. Altschuler and A. Pérez-Garrido. "Global Minimum for Thomson's Problem of Charges on a Sphere." *Phys. Rev. E* 71 (2005), 047703:1-4.
- [Altschuler and Pérez-Garrido 06] E. L. Altschuler and A. Pérez-Garrido. "Defect Free Global Minima in Thomson's Problem of Charges on a Sphere." *Phys. Rev. E* 73 (2006), 036108:1-6.
- [Altschuler and Pérez-Garrido 07] E. L. Altschuler and A. Pérez-Garrido, "Symmetric Four-Dimensional Polytope and Visualization Method in Four, Eight, and Sixteen Dimensions Using Hopf Maps." *Phys. Rev. E* 76 (2007), 016705:1-6.
- [Altschuler et al. 97] E. L. Altschuler, T. J. Williams, E. R. Ratner, R. Tipton, R. Stong, F. Dowla, and F. Wooten. "Possible Global Minimum Lattice Configurations for Thomson's Problem of Charges on a Sphere." *Phys. Rev. Lett.* 78 (1997), 2681-2685.
- [Andreev 96] N. N. Andreev. "An Extremal Property of the Icosahedron." *East J. Approx.* 2 (1996), 459-462.
- [Andreev 97] N. N. Andreev. "Location of Points on a Sphere with Minimal Energy" (Russian). *Tr. Mat. Inst. Steklova* 219 (1997), 27-31; translation in *Proc. Steklov Inst. Math.* 219 (1997), 20-24.
- [Aste and Weaire 08] T. Aste and D. Weaire. *The Pursuit of Perfect Packing*, Second edition. New York: Taylor & Francis, 2008.
- [Bachoc and Vallentin 09] C. Bachoc and F. Vallentin. "Optimality and Uniqueness of the (4, 10, 1/6) Spherical Code." *J. Combin. Theory, Ser. A* 116 (2009), 195-204.
- [Bannai et al. 08] E. Bannai, E. Bannai, and H. Bannai, "Uniqueness of Certain Association Schemes." *European J. Combin.* 29 (2008), 1379-1395.
- [Bausch et al. 03] A. R. Bausch, M. J. Bowick, A. Cacciuto, A. D. Dinsmore, M. F. Hsu, D. R. Nelson, M. G. Nikolaides, A. Travesset, and D. A. Weitz. "Grain Boundary Scars and Spherical Crystallography." *Science* 299 (2003), 1716-1718.
- [Böröczky 04] K. Böröczky, Jr. *Finite Packing and Covering*, Cambridge Tracts in Mathematics 154. Cambridge: Cambridge University Press, 2004.
- [Bowick et al. 02] M. Bowick, A. Cacciuto, D. R. Nelson, and A. Travesset. "Crystalline Order on a Sphere and the Generalized Thomson Problem." *Phys. Rev. Lett.* 89 (2002), 185502:1-4.
- [Bowick et al. 06] M. Bowick, A. Cacciuto, D. R. Nelson, and A. Travesset. "Crystalline Particle Packings on a Sphere with Long Range Power Law Potentials." *Phys. Rev. B* 73 (2006), 024115:1-16.
- [Brouwer and Haemers 92] A. E. Brouwer and W. H. Haemers. "Structure and Uniqueness of the (81, 20, 1, 6) Strongly Regular Graph." *Discrete Math.* 106/107 (1992), 77-82.
- [Brouwer and Haemers 93] A. E. Brouwer and W. H. Haemers. "The Gewirtz Graph: An Exercise in the Theory of Graph Spectra." *European J. Combin.* 14 (1993), 397-407.
- [de Caen and van Dam 99] D. de Caen and E. R. van Dam. "Association Schemes Related to Kasami Codes and Kerdock Sets." *Des. Codes Cryptogr.* 18 (1999), 89-102.

- [Cameron et al. 78] P. J. Cameron, J. M. Goethals, and J. J. Seidel. “Strongly Regular Graphs Having Strongly Regular Subconstituents.” *J. Algebra* 55 (1978), 257–280.
- [Cohn 56] H. Cohn. “Stability Configurations of Electrons on a Sphere.” *Math. Tables Aids Comput.* 10 (1956), 117–120.
- [Cohn and Kumar 07] H. Cohn and A. Kumar. “Universally Optimal Distribution of Points on Spheres.” *J. Amer. Math. Soc.* 20 (2007), 99–148.
- [Cohn et al. 07] H. Cohn, J. H. Conway, N. D. Elkies, and A. Kumar. “The D_4 Root System Is Not Universally Optimal.” *Experimental Mathematics* 16 (2007), 313–320.
- [Conway and Sloane 91] J. Conway and N. J. A. Sloane. “The Cell Structures of Certain Lattices.” In *Miscellanea Mathematica*, pp. 71–107. Berlin: Springer, 1991.
- [Conway and Sloane 99] J. Conway and N. J. A. Sloane. *Sphere Packings, Lattices and Groups*, third edition, Grundlehren der Mathematischen Wissenschaften 290. New York: Springer-Verlag, 1999.
- [van Dam 99] E. R. van Dam. “Three-Class Association Schemes.” *Journal of Algebraic Combinatorics* 10 (1999), 69–107.
- [Damelin and Maymeskul 05] S. B. Damelin and V. Maymeskul. “On Point Energies, Separation Radius and Mesh Norm for s -Extremal Configurations on Compact Sets in \mathbb{R}^n .” *J. Complexity* 21 (2005), 845–863.
- [Dixon 95] G. Dixon. “Octonions: E_8 Lattice to Λ_{16} .” Preprint, 1995, arXiv:hep-th/9501007.
- [Dragnev et al. 02] P. D. Dragnev, D. A. Legg, and D. W. Townsend. “Discrete Logarithmic Energy on the Sphere.” *Pacific J. Math.* 207 (2002), 345–358.
- [Dutour Sikirić 06] M. Dutour Sikirić. “Polyhedral,” a package for GAP. Available online (<http://www.liga.ens.fr/~dutour/Polyhedral/index.html>), 2006.
- [Dutour Sikirić et al. 09] M. Dutour Sikirić, A. Schürmann, and F. Vallentin. “The Contact Polytope of the Leech Lattice.” Preprint, 2009, arXiv:0906.1427.
- [Edmundson 92] J. R. Edmundson. “The Distribution of Point Charges on the Surface of a Sphere.” *Acta Cryst.* A48 (1992), 60–69.
- [Einert et al. 05] T. Einert, P. Lipowsky, J. Schilling, M. J. Bowick, and A. Bausch. “Grain Boundary Scars on Spherical Crystals.” *Langmuir* 21 (2005), 12076–12079.
- [Erber and Hockney 91] T. Erber and G. M. Hockney. “Equilibrium Configurations of N Equal Charges on a Sphere.” *J. Phys. A: Math. Gen.* 24 (1991), L1369–L1377.
- [Ericson and Zinoviev 01] T. Ericson and V. Zinoviev. *Codes on Euclidean Spheres*, North-Holland Mathematical Library 63. Amsterdam: North-Holland Publishing Co., 2001.
- [Föppl 12] L. Föppl. “Stabile Anordnungen von Elektronen im Atom.” *J. Reine Angew. Math.* 141 (1912), 251–302.
- [GAP Group 06] The GAP Group. *GAP: Groups, Algorithms, and Programming*, Version 4.4. Available online (<http://www.gap-system.org/>), 2006.
- [Gawrilow and Joswig 00] E. Gawrilow and M. Joswig. “Polymake: A Framework for Analyzing Convex Polytopes.” In *Polytopes—Combinatorics and Computation (Oberwolfach, 1997)*, pp. 43–73, DMV Sem. 29. Basel: Birkhäuser, 2000.
- [Glasser and Every 92] L. Glasser and A. G. Every. “Energies and Spacings of Point Charges on a Sphere.” *J. Phys. A: Math. Gen.* 25 (1992), 2473–2482.
- [Griess 08] R. Griess. “Few-Cosine Spherical Codes and Barnes–Wall Lattices.” *J. Combin. Theory Ser. A* 115 (2008), 1211–1234.
- [Hardin and Saff 04] D. P. Hardin and E. B. Saff. “Discretizing Manifolds via Minimum Energy Points.” *Notices Amer. Math. Soc.* 51 (2004), 1186–1194.
- [Hardin and Saff 05] D. P. Hardin and E. B. Saff. “Minimal Riesz Energy Point Configurations for Rectifiable d -Dimensional Manifolds.” *Adv. Math.* 193 (2005), 174–204.
- [Hardin and Sloane 03] R. H. Hardin and N. J. A. Sloane. “GOSSET: A General-Purpose Program for Designing Experiments.” Available online (<http://www.research.att.com/~njas/gosset/>), 2003.
- [Hovinga 04] S. H. Hovinga. “Polytopes and Optimal Packing of p Points in n Dimensional Spheres.” Available online (<http://presh.com/hovinga/>), 2004.
- [Kantor 86] W. Kantor. “Some Generalized Quadrangles with Parameters q^2, q .” *Math. Z.* 192 (1986), 45–50.
- [Katanforoush and Shahshahani 03] A. Katanforoush and M. Shahshahani. “Distributing Points on the Sphere I.” *Experiment. Math.* 12 (2003), 199–209.
- [Kolushov and Yudin 94] A. V. Kolushov and V. A. Yudin. “On the Korkin–Zolotarev Construction” (Russian). *Diskret. Mat.* 6 (1994), 155–157; translation in *Discrete Math. Appl.* 4 (1994), 143–146.
- [Kolushov and Yudin 97] A. V. Kolushov and V. A. Yudin. “Extremal Dispositions of Points on the Sphere.” *Anal. Math.* 23 (1997), 25–34.
- [Kottwitz 91] D. A. Kottwitz. “The Densest Packing of Equal Circles on a Sphere.” *Acta Cryst.* A47 (1991), 158–165.
- [Kuijlaars and Saff 98] A. B. J. Kuijlaars and E. B. Saff. “Asymptotics for Minimal Discrete Energy on the Sphere.” *Trans. Amer. Math. Soc.* 350 (1998), 523–538.
- [Leech 57] J. Leech. “Equilibrium of Sets of Particles on a Sphere.” *Math. Gaz.* 41 (1957), 81–90.
- [Livshits and Lozovik 99] A. M. Livshits and Yu. E. Lozovik. “Coulomb Clusters on a Sphere: Topological Classification.” *Chemical Physics Letters* 314 (1999), 577–583.
- [MacWilliams and Sloane 77] F. J. MacWilliams and N. J. A. Sloane. *The Theory of Error-Correcting Codes*, North-Holland Mathematical Library 16. Amsterdam: North-Holland Publishing Company, 1977.

- [Martínez-Finkelshtein et al. 04] A. Martínez-Finkelshtein, V. Maymeskul, E. A. Rakhmanov, and E. B. Saff. “Asymptotics for Minimal Discrete Riesz Energy on Curves in \mathbb{R}^d .” *Canad. J. Math.* 56 (2004), 529–552.
- [McKay 81] B. D. McKay. “Practical Graph Isomorphism” (10th Manitoba Conference on Numerical Mathematics and Computing (Winnipeg, 1980)). *Congressus Numerantium* 30 (1981), 45–87. Software available online (<http://cs.anu.edu.au/~bdm/nauty>), 1981.
- [Melnik et al. 77] T. W. Melnyk, O. Knop, and W. R. Smith. “Extremal Arrangements of Points and Unit Charges on a Sphere: Equilibrium Configurations Revisited.” *Canad. J. Chem.* 55 (1977), 1745–1761.
- [Morris et al. 96] J. R. Morris, D. M. Deaven, and K. M. Ho. “Genetic-Algorithm Energy Minimization for Point Charges on a Sphere.” *Phys. Rev. B* 53 (1996), R1740–R1743.
- [Nordstrom and Robinson 67] A. W. Nordstrom and J. P. Robinson. “An Optimum Nonlinear Code.” *Information and Control* 11 (1967), 613–616.
- [Pérez-Garrido and Moore 99] A. Pérez-Garrido and M. Moore. “Symmetric Patterns of Dislocations in Thomson’s Problem.” *Phys. Rev. B* 60 (1999), 15628–15631.
- [Pérez-Garrido et al. 97a] A. Pérez-Garrido, M. J. W. Dodgson, M. A. Moore, M. Ortuño, and A. Díaz-Sánchez. “Comment on ‘Possible Global Minimum Lattice Configurations for Thomson’s Problem of Charges on a Sphere.’” *Phys. Rev. Lett.* 79 (1997), 1417.
- [Pérez-Garrido et al. 97b] A. Pérez-Garrido, M. J. W. Dodgson, and M. Moore. “Influence of Dislocations in Thomson’s Problem.” *Phys. Rev. B* 56 (1997), 3640–3643.
- [Rakhmanov et al. 94] E. A. Rakhmanov, E. B. Saff, and Y. M. Zhou. “Minimal Discrete Energy on the Sphere.” *Math. Res. Lett.* 1 (1994), 647–662.
- [Rakhmanov et al. 95] E. A. Rakhmanov, E. B. Saff, and Y. M. Zhou. “Electrons on the Sphere.” In *Computational Methods and Function Theory*, edited by R. M. Ali, S. Ruscheweyh, and E. B. Saff, pp. 111–127. River Edge, NJ: World Scientific, 1995.
- [Saff and Kuijlaars 97] E. B. Saff and A. B. J. Kuijlaars. “Distributing Many Points on a Sphere.” *Math. Intelligencer* 19 (1997), 5–11.
- [Schoutte 10] P. H. Schoutte. “On the Relation between the Vertices of a Definite Six-Dimensional Polytope and the Lines of a Cubic Surface.” *Proc. Roy. Acad. Amsterdam* 13 (1910), 375–383.
- [Sloane 00] N. J. A. Sloane, with the collaboration of R. H. Hardin, W. D. Smith, and others. “Tables of Spherical Codes.” Available online (<http://www.research.att.com/~njas/packings/>), 2000.
- [Sloane et al. 03] N. J. A. Sloane, R. H. Hardin, and P. Cara. “Spherical Designs in Four Dimensions.” In *Proceedings of the 2003 IEEE Information Theory Workshop*, pp. 253–258. Los Alamitos: IEEE, 2003.
- [Thomson 97] J. J. Thomson. “Cathode Rays.” *Phil. Mag.* 44 (1897), 293–316.
- [Whyte 52] L. L. Whyte. “Unique Arrangements of Points on a Sphere.” *Amer. Math. Monthly* 59 (1952), 606–611.
- [Widder 41] D. V. Widder. *The Laplace Transform*. Princeton: Princeton University Press, 1941.
- [Willie 86] L. T. Willie. “Searching Potential Energy Surfaces by Simulated Annealing.” *Nature* 324 (1986), 46–48.
- [Yudin 93] V. A. Yudin. “Minimum Potential Energy of a Point System of Charges” (Russian). *Diskret. Mat.* 4 (1992), 115–121; translation in *Discrete Math. Appl.* 3 (1993), 75–81.

Brandon Ballinger, Google, Inc., 1600 Amphitheatre Parkway, Mountain View, CA 94043 (brandonb@google.com)

Grigoriy Blekherman, Microsoft Research, One Microsoft Way, Redmond, WA 98052-6399 (grrigg@gmail.com)

Henry Cohn, Microsoft Research New England, One Memorial Drive, Cambridge, MA 02142 (cohn@microsoft.com)

Noah Giansiracusa, Department of Mathematics, Brown University, Providence, RI 02912 (noahgian@math.brown.edu)

Elizabeth Kelly, Department of Mathematics, University of Washington, Seattle, WA 98195 (thebethkelly@gmail.com)

Achill Schürmann, Institute of Applied Mathematics, TU Delft, Mekelweg 4, 2628 CD Delft, The Netherlands (a.schurmann@tudelft.nl)

Received October 18, 2007; accepted in revised form October 8, 2008.