

Random triangulations of the d -sphere with minimum volume

Agelos Georgakopoulos^{*1}, John Haslegrave², and Joel Larsson Danielsson³

¹University of Warwick, UK

²Lancaster University, UK

³Chalmers University of Technology, Sweden

ABSTRACT

We study a higher-dimensional analogue of the Random Travelling Salesman Problem: let the complete d -dimensional simplicial complex K_n^d on n vertices be equipped with i.i.d. volumes on its facets, uniformly random in $[0, 1]$. What is the minimum volume $M_{n,d}$ of a sub-complex homeomorphic to the d -dimensional sphere \mathbb{S}^d , containing all vertices? We determine the growth rate of $M_{n,2}$, and prove that it is well-concentrated. For $d > 2$ we prove such results to the extent that current knowledge about the number of triangulations of \mathbb{S}^d allows.

We remark that this can be thought of as a model of random geometry in the spirit of Angel & Schramm's UIPT, and provide a generalised framework that interpolates between our model and the uniform random triangulation of \mathbb{S}^d .

1 INTRODUCTION

In this paper we study the following analogue of the *Random Travelling Salesman Problem*, in line with a recent trend of higher-dimensional generalisations of graph-theoretic results. Let the complete d -dimensional simplicial complex K_n^d on n vertices be equipped with i.i.d. costs on its facets, uniform in $[0, 1]$. What is the minimum volume $M_{n,d}$ of a *spanning* sub-complex of K_n^d (that is, containing all vertices) homeomorphic to the d -dimensional sphere \mathbb{S}^d ? (The *volume* of a complex is just the sum of costs of its facets.) We determine the growth rate of $M_{n,2}$, and prove that $M_{n,2}$ is well-concentrated around its mean. For $d > 2$ we prove upper and lower bounds and concentration results, but closing the gaps will require progress on the enumeration of triangulations of \mathbb{S}^d , which is notoriously difficult.

^{*}Supported by EPSRC grants EP/V048821/1 and EP/V009044/1. All authors have been supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant Agreement No. 639046).

1.1 BACKGROUND AND MOTIVATION

In the graph case, $d = 1$, our question asks for the minimum cost of a Hamilton cycle of the complete graph K_n . This problem was studied by Frieze [12], who showed that it has constant cost and is close to the minimum 2-factor with high probability. Subsequently, Wästlund determined the limiting constant precisely [36]. A similar, though simpler, problem asks for the spanning tree of minimum cost in K_n . This is known to have expectation $\zeta(3) + o(1)$ [11]. For the directed case, Karp [20] proved a relationship to the assignment problem, which Aldous [1] used to establish that the expectation tends to $\zeta(2) = \pi^2/6$.

Part of the motivation for studying the case $d \geq 2$ comes from the emerging random geometry. The influential work of Angel & Schramm [3] advanced the idea of considering a random triangulation $R_n^{(2)}$ of the 2-sphere chosen uniformly among all such triangulations with n vertices, as a model of random planar geometry. Angel & Schramm [3] proved that $R_n^{(2)}$ converges locally, in the sense of Benjamini & Schramm [9]. The limit is a random triangulation of \mathbb{R}^2 , called the Uniform Infinite Planar Triangulation. We expect our minimum-volume spanning sphere to converge in the same sense, and lead to a different model of random geometry. Another well-known convergence result about $R_n^{(2)}$, by Le Galle [26] and Miermont [32], is that it has a scaling limit which coincides with the *Brownian map*, a model of random fractal geometry studied extensively due to connections with mathematical physics ([15]).

In fact one can interpolate between our model and the uniform triangulation $R_n^{(2)}$ by considering a Boltzmann distribution with an (inverse) temperature parameter β . The two extremes, for $\beta = 0$ and $\beta = \infty$ of this distribution correspond to $R_n^{(2)}$ and our model respectively. It would be very interesting to understand the phase-transition phenomena as β varies. We provide concrete definitions and questions in Section 8.

There has been much previous work on extending notions of (Hamilton) cycles to hypergraphs in an essentially one-dimensional way; see [22] for a survey of results in this area. The topological approach we consider here has emerged recently as the focus of several papers [13, 17, 21, 23, 29, 30], and fits naturally into Linial's 'higher-dimensional combinatorics' programme, see e.g. [28].

Linial & Meshulam studied the homology of 2-dimensional random simplicial complexes [27], which was subsequently extended to higher dimensions by Meshulam & Wallach [31]. This triggered a lot of research on random simplicial complexes analogous to work on Erdős–Rényi random graphs. In particular, Luria & Tessler [30] determined the threshold for the emergence of a spanning 2-sphere in the Linial–Meshulam model, using a delicate second-moment argument. Moving away from spanning structures, Benjamini, Lubetzky & Peled determined the minimum cost of a disc triangulation with boundary u, v, w [8], which may be thought of as a 2-dimensional analogue of the shortest path between two vertices.

Besides probabilistic results, there has also been an interest in extending extremal results from graphs to higher dimensions topologically, by replacing cycles by triangula-

tions of \mathbb{S}^d , $d \geq 2$. For example, Georgakopolous, Haslegrave, Montgomery & Narayanan [13] asymptotically determined the minimum codegree guaranteeing a spanning 2-sphere in a 3-graph. See e.g. [17, 21, 23, 29] for further results with a similar flavour.

1.2 OUR SETUP

Instead of the complex K_n^d , we will work with the complete $(d+1)$ -uniform hypergraph $K_n^{(d)}$ on n vertices. Let $\mathcal{S}_n^{(d)}$ be the set of sub-hypergraphs of $K_n^{(d)}$ that form a combinatorial d -sphere spanning all n vertices. In other words, $\mathcal{S}_n^{(d)}$ is the set of spanning subgraphs S such that the simplicial complex formed by taking the downward closure of $E(S)$ is homeomorphic to \mathbb{S}^d . We endow the (hyper)edges of $K_n^{(d)}$ with i.i.d. random costs, uniformly distributed in $[0, 1]$. We write W_e for the random cost of the edge $e \in E(K_n^{(d)})$, and W_S for the total *cost*, or *volume*, $\sum_{e \in E(S)} W_e$ of a subgraph S , and $M(\mathcal{S})$ for the minimum volume $\min_{S \in \mathcal{S}} W_S$ of a family \mathcal{S} of subgraphs of $K_n^{(d)}$.

We call $\mathcal{S}_n^{(d)}$ the set of *spanning spheres*, and refer to the spanning sphere $S \in \mathcal{S}_n^{(d)}$ achieving this minimum (which is almost surely unique) as the *minimal spanning sphere*. In this paper, we prove upper and lower bounds on $\mathbb{E}(M(\mathcal{S}_n^{(d)}))$. For $d = 2$ these match up to a constant factor. For dimensions 2 and 3 we also show that this random variable is sharply concentrated around its mean.

Key difficulties in dimensions $d > 2$ are the absence of exact asymptotics for the number of spanning spheres, and the fact that the number of facets in a spanning sphere is not fixed. Enumerating the spheres is trivial in dimension 1, and is a classical result of Tutte [34] in dimension 2.

1.3 OUR RESULTS

Our main results are concentration inequalities as well as upper and lower bounds for $M(\mathcal{S}_n^{(d)})$ for $d = 2, 3$. These together imply sharp concentration: that $M(\mathcal{S}_n^{(d)})/\mu_n^{(d)} \rightarrow 1$ in probability as $n \rightarrow \infty$, where $\mu_n^{(d)}$ denotes the median of $M(\mathcal{S}_n^{(d)})$.

The upper bound holds for $d \geq 4$ as well, but for the lower bounds and concentration inequality we need to make an additional assumption on the number of d -spheres with a given number of facets. Under an even stronger assumption, we prove matching upper and lower bounds as well as concentration for any $d \geq 2$.

Upper bounds are proven using an explicit construction, while lower bounds are proven using a first moment method. For sharp concentration, we use the Talagrand inequality. In dimension 2, however, this turns out to not be sharp enough, and we instead use a recent concentration inequality by the third author [25] which builds on the Talagrand inequality.

THEOREM 1.1 ($d = 2$): *Let $\mu_n^{(2)}$ denote the median of $M(\mathcal{S}_n^{(2)})$.*

(i) *Letting $\alpha := \sqrt{3^3/4^3}e \approx 0.394$, we have*

$$\alpha - o(1) \leq \frac{\mu_n^{(2)}}{\sqrt{n}} \leq \frac{e}{2}\alpha + o(1).$$

(ii) With probability at least $1 - \exp(-n^{0.02})$,

$$|M(\mathcal{S}_n^{(2)}) - \mu_n^{(2)}| \leq n^{0.4}.$$

We remark that the constants of the lower and upper bound of Theorem 1.1 (i) differ by a multiplicative factor of only $e/2 \approx 1.36$.

We also consider a variant of our model where the (independent, uniform) costs are put on the pairs of vertices of K_n^2 ; see Theorem 4.1 for details.

We next consider higher dimensions. For each $d \geq 2$ let β_d be a constant such that at most $2^{O(m)} m^{\beta_d m}$ spheres in $\mathcal{S}_n^{(d)}$ have m facets, for every m ; such constants are known to exist but the optimal values are unknown for $d \geq 3$ [33].

THEOREM 1.2: *The following hold for any $d \geq 2$.*

- (i) *There exists a constant $b = b(d)$ such that $M(\mathcal{S}_n^{(d)}) \leq bn^{1-1/d}$ with high probability.*
- (ii) *If $\beta_d \leq 1$, then $M(\mathcal{S}_n^{(d)}) \geq an^{1-\beta_d}$ for some constant $a = a(d)$, w.h.p.*
- (iii) *If $\beta_d < 1/2$, then there exists some $c > 0$ such that $\mathbb{P}(|M(\mathcal{S}_n^{(d)}) - \mu_n^{(d)}| > t) \leq e^{-ct^2/n}$ for any $t > 0$.*

It is known that $\beta_4 \leq 1$ [33], and so (ii) applies in dimension 4, but for $d \geq 5$ this is not known. The third author has proved that $\beta_3 \leq 8/21$ [24]. Combining this with the lower bound in Proposition 5.4 and the upper bound and concentration in Theorem 1.2 above, we obtain the following.

THEOREM 1.3 ($d = 3$): *There are constants $c, C > 0$ such that for all sufficiently large n ,*

$$cn^{13/21} \leq \mu_n^{(3)} \leq Cn^{2/3},$$

and furthermore

$$|M(\mathcal{S}_n^{(3)}) - \mu_n^{(3)}| = O_{\mathbb{P}}(\sqrt{n}).$$

In particular, $M(\mathcal{S}_n^{(3)})/\mu_n^{(3)} \rightarrow 1$ in probability as $n \rightarrow \infty$.

All lower bounds are proved using a first moment method. For this technique to work, we need good upper bounds on the number of n -vertex spheres with a given number of facets m . (We discuss this in section 3.) For fixed n , while spheres with a large number of facets tend to have larger cost, there can be many more such spheres than spheres with few facets. Thus, depending on how quickly the number of spheres grow with m , the expected number of ‘cheap’ spheres might be dominated by spheres with few or many facets.

A well-known question of Gromov [14] asks whether the number of isomorphism classes of spheres in $\mathcal{S}_n^{(d)}$ with m facets grows at most exponentially fast as a function of m . If so, we would be able to prove much stronger results. Furthermore, if we require our sphere to lie in a suitable subfamily $\mathcal{P} \subseteq \mathcal{S}_n^{(d)}$ for which this holds, then our results

apply to the minimum-cost sphere in this family; for such families, the minimum cost tends to be achieved by a sphere with close to the minimum number of facets. Here it is natural to restrict our attention to *symmetrical* families, that is, those that are unions of isomorphism classes. We also require a specific sphere $S_{n,d}^*$ of a particularly simple form, used in proving our upper bound, to lie in the family. We define $S_{n,d}^*$ as follows: its facets are precisely those $(d+1)$ -sets whose intersection with $\{v_1, \dots, v_{n-d}\}$ lies in $\{\{v_1, v_2\}, \dots, \{v_{n-d-1}, v_{n-d}\}, \{v_{n-d}, v_1\}\}$. We will subsequently show that this is indeed a sphere.

We prove matching upper and lower bounds of order $n^{1-\frac{1}{d}}$, and sharp concentration, for such \mathcal{P} in any dimension d .

THEOREM 1.4: *Assume $\mathcal{P} \subseteq \mathcal{S}_n^{(d)}$ is a symmetrical subfamily including $S_{n,d}^*$ and containing at most K^m isomorphism classes of spheres with m facets for some K and every m . Then there is a $\mu = \Theta(n^{1-\frac{1}{d}})$ such that for some constant $c > 0$ and any $0 < t \leq \sqrt{n}$,*

$$\mathbb{P}(|M(\mathcal{P}) - \mu| > t\sqrt{n}) \leq e^{-ct^2},$$

and the optimal spanning sphere has at most $(1 + \varepsilon)dn$ facets with high probability (for any $\varepsilon > 0$).

In particular, Theorem 1.4 holds when \mathcal{P} is the subfamily of locally constructible spheres, the definition and motivation of which is discussed in Section 7.2. Furthermore, if Gromov's aforementioned question has a positive answer, then Theorem 1.4 holds for $\mathcal{P} = \mathcal{S}_n^{(d)}$.

1.4 OUTLINE OF THE PAPER

In Section 3 we discuss bounds on the number of spheres, which will be relevant to Theorems 1.2 and 1.3. We then proceed to prove the upper and lower bounds constituting Theorem 1.1 (i) in Section 4 and Theorem 1.2 (i) and (ii) in Section 5. We prove Theorem 1.1 (ii) in Section 6 and Theorem 1.2 (iii) in Section 7.1. Theorem 1.3 follows from our other results as detailed above. We discuss locally constructible spheres, and prove Theorem 1.4, in Section 7.2.

2 DEFINITIONS & PRELIMINARIES

2.1 HYPERGRAPHS AND SIMPLICIAL SPHERES

A *hypergraph* $G = (V, H)$ consists of a set V of *vertices*, and a set H of pairwise distinct subsets of V , called (*hyper*)*edges*. If each element of H has the same number of vertices k , we say that G is k -uniform. Let $K_n^{(d)}$ denote the hypergraph with vertex set $[n]$, where each $(d+1)$ -tuple of vertices forms a hyperedge.

A *simplicial d -sphere* is a triangulation of \mathbb{S}^d , i.e. a simplicial complex S homeomorphic to \mathbb{S}^d . We will sometimes drop the term *simplicial* for brevity. We can realise S as a hypergraph, by using the same vertex set, and declaring the set of vertices of each

facet of S to be a hyperedge. All the simplicial d -spheres we will consider in this paper are realised as spanning sub-hypergraphs of $K_n^{(d)}$, i.e. hypergraphs on $[n]$ with set of hyperedges contained in that of $K_n^{(d)}$. We define a *simplicial d -ball* analogously.

2.2 SUMS OF UNIFORM RANDOM VARIABLES

For $p > 0$, we let $U(0, p)$ denote a random variable uniformly distributed in $[0, p]$. We will be using the following straightforward bound on the sum of uniform random variables.

LEMMA 2.1: *Let X_1, \dots, X_m be i.i.d. $U(0, 1)$ random variables. Then*

$$\mathbb{P}\left(\sum_{i=1}^m X_i \leq L\right) \leq L^m/m! \leq (Le/m)^m.$$

PROOF. Let χ_A be the characteristic function of the set $A := \{x \in \mathbb{R}_+^m : \sum_{i=1}^m x_i \leq L\}$. Then

$$\mathbb{P}\left(\sum_{i=1}^m X_i \leq L\right) = \int_{[0,1]^m} \chi_A d\mu \leq \int_{[0,L]^m} \chi_A d\mu = L^m/m!,$$

where μ is the Lebesgue measure on \mathbb{R}^m . ■

3 ON THE NUMBER OF SIMPLICIAL SPHERES

In section 5.2, we give a generic lower bound on the random variable $M(\mathcal{S}_n^{(d)})$, using a first moment argument by finding an $L \in \mathbb{R}$ such that the expected number of d -spheres with cost at most L is $o(1)$. However, for this to work we need a good enough upper bound on the total number $A_{n,m}^{(d)}$ of labelled simplicial d -spheres with a given number of vertices n and facets m . Previous bounds often in fact consider spheres up to vertex permutations (often called combinatorially distinct spheres); we will refer to these as *isomorphism classes* and denote their number by $B_{n,m}^{(d)}$. On the one hand, clearly $A_{n,m}^{(d)} \leq n!B_{n,m}^{(d)}$. But on the other hand, since each sphere has at most $(d+1)!m$ automorphisms – an arbitrary facet can be mapped to any of m facets in any of $(d+1)!$ orientations, and this determines the automorphism – we also have that $n!B_{n,m}^{(d)} \leq (d+1)!mA_{n,m}^{(d)}$. Hence $A_{n,m}^{(d)} = m^{O(1)}n!B_{n,m}^{(d)}$.

For $d = 2$ we have $m = 2n - 4$,¹ and so these numbers are merely functions of n . Tutte [34] determined the number of isomorphism classes, which is growing exponentially with n :

$$B_{n,2n-4}^{(2)} = (1 + o(1)) \frac{1}{16} \sqrt{\frac{3\pi}{2}} n^{-5/2} \left(\frac{4^4}{3^3}\right)^{n+1} = 2^{O(n)}. \quad (1)$$

Thus, for sufficiently large n ,

$$A_{n,2n-4}^{(2)} \leq n!B_{n,2n-4}^{(2)} \leq (4^4 n/3^3 e)^n. \quad (2)$$

¹Indeed, Euler's formula implies that each simplicial 2-sphere with n vertices has exactly $m = 2n - 4$ faces and $3m/2 = 3n - 6$ edges.

A first moment argument using (2) (see Proposition 5.4) provides a lower bound on $M(\mathcal{S}_n^{(2)})$ within a constant factor of the upper bound we prove in section 5.1. That upper bound can be slightly improved (while still not matching the lower bound) as an easy corollary of a theorem on the related threshold problem by Luria & Tessler [30]. For details, see the proof of Theorem 1.1.

In higher dimensions, the number m of facets varies among simplicial spheres with n vertices. There is a trade-off where, for a fixed n , there are many more spheres with large m , but these tend to have higher cost and so are individually much less likely to have cost at most L . If the second factor dominates the former, the minimum-cost sphere typically has few facets. In order to understand this interplay we need bounds on the asymptotic growth rate of $A_{n,m}^{(d)}$ as m becomes large.

Our main interest, in view of Theorem 1.2, is how small we can take β_d such that $A_{n,m}^{(d)} \leq 2^{O(m)} m^{\beta_d m}$. However, previous work in this area has concentrated on the growth of the number of isomorphism classes of simplicial spheres with m facets and any number of vertices. This is unknown for $d \geq 3$. In fact, while it is easy to show that there are $2^{\Omega(m)}$ spheres with m facets, it is a major open question of Gromov [14] whether their number grows at most exponentially fast.

Upper bounds on the number of spheres with n labelled vertices and any number of facets, of roughly $\exp(n^{\lfloor d/2 \rfloor + o(1)})$ have been obtained [19], but these are insufficient for our purposes, as they give no control over the relationship between n and m , and are dominated by triangulations with $n \ll m$. Previous bounds on the number of isomorphism classes of spheres with m facets are of the form $m^{\gamma_d m}$ for some constants γ_d [33]. Since this also gives an upper bound on the number of isomorphism classes with m facets and exactly n vertices, and labelling the vertices adds a multiplicative factor of at most $n! \ll m^{m/d}$, we obtain a bound of the desired form with $\beta_d = \gamma_d + 1/d$. Unfortunately, the best known upper bounds on γ_d are too large for our purposes when $d \geq 3$. For instance, it was shown in [33] that $\gamma_3 \leq 1/3$ and $\gamma_4 \leq 3/4$, and these bounds are low enough to imply that $M(\mathcal{S}_n^{(3)}) \gg 1$ and $M(\mathcal{S}_n^{(4)}) = \Omega(1)$ (as we will show in section 5.2). However, these are neither anywhere near our upper bound $n^{1-1/d}$ in section 5.1, nor large enough for our upper bound on the variance $\text{Var}(M(\mathcal{S}_n^{(d)})) = O(n)$ in section 7.1 to imply sharp concentration. Moreover, for $d \geq 5$ the best known upper bounds have $\gamma_d > 1$, which render our first moment method of section 5.2 ineffective.

In a forthcoming paper, the third author improves on the upper bound in the case $d = 3$. This increases the lower bound on $M(\mathcal{S}_n^{(d)})$ enough for the argument in section 7.1 to imply sharp concentration, yet not quite enough to match the upper bound in section 5.1.

THEOREM 3.1 ([24]): *For sufficiently large n , we have $A_{n,m}^{(3)} \leq m^{\frac{2}{7}(m+n)}$.*

Since in this case $n \leq m/3 + O(1)$, it follows that $\beta_3 \leq \frac{2}{7}(1 + 1/3) = 8/21$.

In section 7.2 we discuss a subfamily of more combinatorially tractable spheres, called the *locally constructible* (LC) spheres. All 2-spheres are LC, but for every $d \geq 3$ there are d -spheres that are not. It has been shown [10] that there are only exponentially many LC-spheres with a given number of facets, which allows us to prove much stronger results for this class.

4 UPPER AND LOWER BOUNDS FOR 2-SPHERES

In this section we provide simple bounds that establish Theorem 1.1(i).

PROOF OF THEOREM 1.1(i). For the lower bound, fix $L \in \mathbb{R}$, and define the random variable X to be the number of $S \in \mathcal{S}_n^{(2)}$ whose cost W_S is at most L .

By Tutte's formula (1) there are $o(\gamma^n n!)$ simplicial 2-spheres on n vertices, where $\gamma := 4^4/3^3$. Each simplicial 2-sphere has $m = 2n - 4$ faces, hence by Lemma 2.1 it has cost at most L with probability at most $(Le/m)^m = n^{O(1)}(Le/2n)^{2n}$. By linearity of expectation we thus have

$$\mathbb{E}(X) \leq \gamma^n n! n^{O(1)} (Le/2n)^{2n} = n^{O(1)} (L^2 e \gamma / 4n)^n. \quad (3)$$

Let $L := (1 - \varepsilon) \sqrt{4n/e\gamma}$ for some $\varepsilon > 0$, so that the second factor is $(1 - \varepsilon)^{2n}$. It follows from the Markov inequality that $M(\mathcal{S}_n^{(2)}) > L$, and hence with high probability

$$M(\mathcal{S}_n^{(2)}) > (1 - \varepsilon) \sqrt{\frac{3^3 n}{4^4 e}}.$$

For the upper bound, we use the following result from [30]: there is a sharp threshold for the appearance of a spanning 2-sphere in the Linial–Meshulam model² at $p_c := \sqrt{e/\gamma n}$. Letting $p := (1 + \varepsilon)p_c$, we conclude that $M(\mathcal{S}_n^{(2)}) \leq mp$, because one way to sample the Linial–Meshulam model is to assign i.i.d. $U(0, 1)$ random costs to the 2-simplices, and delete those of cost greater than p . We can, however, improve this bound by a factor of 2 as follows.

Let $R \subseteq \mathcal{S}_n^{(2)}$ be the random set of 2-spheres with no face having cost more than p . For any $S \in \mathcal{S}_n^{(2)}$, conditionally on $\{S \in R\}$, the conditional distribution of the costs of the faces in S are i.i.d. $U(0, p)$. Thus the conditional distribution $(W_S \mid S \in R)$ is the sum of m i.i.d. random variables with expected value $p/2$, and hence Hoeffding's inequality yields $\mathbb{P}(W_S > (1 + \varepsilon)mp/2 \mid S \in R) \leq e^{-\varepsilon^2 m/2}$.

Thus whenever R is non-empty, an arbitrary $S \in R$ has cost at most $(1 + \varepsilon)mp/2$ w.h.p., and R is non-empty w.h.p. since $p > (1 + \varepsilon)p_c$. Hence $M(\mathcal{S}_n^{(2)})$ is at most $(1 + \varepsilon)mp/2 \leq (1 + \varepsilon)^2 \sqrt{ne/\gamma}$ w.h.p. Since $M(\mathcal{S}_n^{(2)}) \leq m$ with the remaining probability, the result follows. ■

Next, we obtain a result analogous to Theorem 1.1(i) for the variant of our model where we put costs on the 1-cells (edges) instead of 2-cells: we define $N(\mathcal{S}_n^{(2)})$ just like $M(\mathcal{S}_n^{(2)})$ except that we put independent $U(0, 1)$ costs on the pairs of vertices of $K_n^{(d)}$, and define the cost of $S \in \mathcal{S}_n^{(2)}$ to be the sum of the costs of the pairs appearing in the edges (i.e. triples of vertices) of S .

²This is the natural generalisation of the binomial random graph to higher dimensions: given some parameter $p \in [0, 1]$, and a dimension $d \in \mathbb{N}$, we keep each d -simplex with vertices in $[n]$ independently with probability p .

THEOREM 4.1: Let $\nu_n^{(2)}$ denote the median of $N(\mathcal{S}_n^{(2)})$. There is a constant $C \in \mathbb{R}$ such that

$$\frac{3^2}{4^{4/3}e^{2/3}} - o(1) \leq \frac{\nu_n^{(2)}}{n^{2/3}} \leq C.$$

Furthermore, $N(\mathcal{S}_n^{(2)})$ is sharply concentrated around $\nu_n^{(2)}$: for some small fixed $\delta > 0$,

$$\mathbb{P}(|N(\mathcal{S}_n^{(2)}) - \nu_n^{(2)}| \geq (\nu_n^{(2)})^{1-\delta}) \leq \exp(-n^\delta).$$

PROOF. We follow the lines of the proof of Theorem 1.1(i). Each simplicial 2-sphere with n vertices has exactly $3n - 6$ edges, thus again by Lemma 2.1 it has cost at most L with probability at most $n^{O(1)}(Le/3n)^{3n}$. The calculation of (3) now becomes

$$\mathbb{E}(X) \leq n^{O(1)}\gamma^n n!(Le/3n)^{3n} \approx \left(\frac{\gamma n L^3 e^3}{e 3^3 n^3}\right)^n = \left(\frac{\gamma L^3 e^2}{3^3 n^2}\right)^n = \left(\frac{4^4 L^3 e^2}{3^6 n^2}\right)^n.$$

Choosing L so that this equals $(1-\varepsilon)^n$, it follows as above that, with high probability, $N(\mathcal{S}_n^{(2)}) > (1-\varepsilon)\left(\frac{3^2}{4^{4/3}}\right)(n/e)^{2/3}$.

For the upper bound, we use an explicit construction. The proof idea is to start with a 2-sphere on 4 vertices (the boundary of a tetrahedron) and $n - 4$ free vertices. We repeatedly make barycentric subdivisions of its faces until the sphere has n vertices.

A heuristic argument for the growth rate $N(\mathcal{S}_n^{(2)}) \sim n^{2/3}$ is as follows. When there are i vertices in the construction so far, we must choose one of the remaining $n - i$ free vertices to use for the barycentric subdivision. The cost of picking any given vertex is a sum of three i.i.d. $U(0, 1)$ random variables. The expected minimum over $n - i$ such random variables is of order $(n - i)^{-1/3}$, and summing over $i < n$ the total expected cost is of order $n^{2/3}$. We want to bound the upper tail of this sum.

However, in order to make this heuristic rigorous we need to avoid problems of dependencies between the costs of the edges added in the current barycentric subdivision and the previous ones. For this we use two tricks: First, similarly to the proof of Theorem 1.1(ii), we partition the vertex set $[n]$ into four sets $\{V_i\}_{i=1}^4$. For the sake of convenience, we assume that a vertex $j \in [n]$ lies in V_i if and only if $j \equiv i \pmod{4}$. Next, we choose the faces to perform barycentric subdivisions of in a particular way. Starting with a tetrahedron $v_1 v_2 v_3 v_4 := 1234$, make a barycentric subdivision of the face $v_2 v_3 v_4$ using some vertex v_5 , then a barycentric subdivision of the face $v_3 v_4 v_5$ using some vertex v_6 , and so on. This 2-sphere can be encoded by the word $v_1 v_2 \dots v_n$, and the set of edges in its 1-skeleton are $\{v_i v_j : |i - j| \leq 3\}$. (In hypergraph terms, it corresponds to the shadow of a tight path in a 4-uniform hypergraph.) Furthermore, we pick each such v_i from the V_j which has $i \equiv j \pmod{4}$. (That is, $v_5 \in V_1$, $v_6 \in V_2$ and so on.)

For $i \geq 5$, let $Z_{i,u} := X_{v_{i-3}u} + X_{v_{i-2}u} + X_{v_{i-1}u}$ denote the cost of the edges added when using the vertex u for a barycentric subdivision of $v_{i-3}v_{i-2}v_{i-1}$, and $Z_i := \min_{u \in V_i} Z_{i,u}$ (where the indices for V are taken modulo 4). For $i \leq 4$, define $Z_1 = 0$, $Z_2 = X_{12}$, $Z_3 = X_{13} + X_{23}$ and $Z_4 = X_{14} + X_{24} + X_{34}$, so that $\sum_{i=1}^4 Z_i \leq 6$ is the total cost of the edges in the tetrahedron 1234. Note that $v_1 v_2 \dots v_n$ encodes a 2-sphere with cost $\sum_{i=1}^n Z_i \geq N(\mathcal{S}_n^{(2)})$.

CLAIM 1: *The random variables Z_i are independent.*

Proof of claim. To prove the claim, consider an edge $ww' \in \binom{[n]}{2}$. How many times is its cost $X_{ww'}$ explored? Assume without loss of generality that w is added to the growing sphere before w' , and $w = u_i$ for some i . Let $k \in \{0, 1, 2, 3\}$ be such that $w \in V_i, w' \in V_{i+k}$ (where indices are again taken modulo 4). If $k = 0$, the edge ww' is never considered – the final 2-sphere contain no edge with both endpoints in the same set V_i . If $k = 1, 2$ or 3 , the edge ww' is explored only when choosing between candidates for $v_{i+k} := \arg \min_{u \in V_{i+k}} Z_{i+k,u}$. So in either case, ww' is only explored at most once, and hence the Z_i 's are independent. ■

It only remains to show that $\sum_{i=1}^n Z_i \leq Cn^{2/3}$ whp for some constant C . Note that the random variables Z_i are not identically distributed, but that the distributions of Z_i and Z_j only differ slightly if i and j are close. We'll therefore group the terms into at most k blocks with k terms each (picking k as the smallest multiple of 4 such that $k^2 \geq n$), and separately upper bound the sum of the Z_i 's in each block. More precisely, we partition $\{1, 2, \dots, n\}$ into sets $\{A_j\}_{j=0}^k$ by setting $A_j := \{s \in \mathbb{N} : jk \leq n - s < (j+1)k\}$. Note that $|A_j| \leq k$ for all j (in fact $|A_j| = k$ for all but a bounded number of j). Let $Y_j := \sum_{i \in A_j} Z_i$, so that $\sum_{i=1}^n Z_i = \sum_{j=0}^k Y_j$.

For $i \in A_j, j \geq 1$, each random variable Z_i is the minimum of at least $jk/4$ independent random variables $Z_{i,u}$. Applying Claim 2 to Y_j with $k_0 = |A_j| \leq k, n_0 = jk/4$ and $b = 25k$, we get that

$$\mathbb{P}\left(Y_k \geq \underbrace{25k(jk/4)^{-1/3}}_{=: a_j}\right) \leq e^{-k}. \quad (4)$$

Noting that $\sum_{j=1}^k j^{-1/3} < (3/2 + o(1))k^{2/3}$ by an integral comparison and recalling that $k = (1 + o(1))\sqrt{n}$, the sum $\sum_{j=1}^k a_j$ is at most $ctn^{2/3}$ for some constant $c < 60$. Next, for $j = 0$, note that Y_0 is the sum of the Z_i 's in the *last* block. Hence $Y_0 \leq |A_0| = k$, because $Z_i \leq 1$ surely. In other words, the inequality (4) holds for $k = 0$ with $a_0 := k$. Then $\sum_{j=0}^k a_j < 60n^{2/3}$ for n sufficiently large. By a union bound,

$$\mathbb{P}\left(\sum_{i=1}^n Z_i \geq 60n^{2/3}\right) \leq ke^{-k} = e^{-\Omega(\sqrt{n})}.$$

Recall now that $\sum_{i=1}^n Z_i$ was the cost of the sphere found with our explicit construction, which is an upper bound on the optimal cost. So with high probability, $N(\mathcal{S}_n^{(2)}) \leq 60n^{2/3}$.

For concentration, we apply Theorem 7.1. There are at most $K^n n^n$ 2-spheres on n vertices, for some constant K , and each has $3d - 6$ edges. So Theorem 7.1 gives, with $\beta = 1/3$ and $t = 3$, that there exists a small constant $\delta > 0$ such that

$$\mathbb{P}(|N(\mathcal{S}_n^{(2)}) - \nu| \geq \nu^{1-\delta}) \leq \exp(-n^\delta). \quad \blacksquare$$

5 BOUNDS FOR d -SPHERES

5.1 UPPER BOUND

In this section we show that there exists a sufficiently cheap sphere with high probability; in order to fulfil the requirements of Theorem 1.4 we will show this is true even when restricted to spheres isomorphic to $S_{n,d}^*$. Recall that the facets of $S_{n,d}^*$ are the $(d+1)$ -sets whose intersection with $\{v_1, \dots, v_{n-d}\}$ lies in $\{\{v_1, v_2\}, \dots, \{v_{n-d-1}, v_{n-d}\}, \{v_{n-d}, v_1\}\}$. For $d = 2$, this is easy to visualize, being the boundary of a bipyramid. In general, it consists of an $n - d$ cycle and d poles; in fact, we consider the subclass Θ_d of isomorphic spheres where the cycle vertices may be permuted, but the poles remain v_1, \dots, v_d . That is, we take an arbitrary cycle C on vertex set v_{d+1}, \dots, v_n . For each such cycle C , define $f_{xy} = \{v_1, \dots, v_d, x, y\}$ for each $xy \in E(C)$, let $B(C)$ be the closure of $\{f_{xy}\}_{xy \in E(C)}$, and let $S_C := \partial B(C)$. Let Θ_d be the family of all such S_C .

PROPOSITION 5.1: S_C is a locally constructible simplicial d -sphere with $d(n - d)$ facets, given by

$$\{f_{xy} - \{z_i\} : xy \in E(C), 1 \leq i \leq d\}.$$

In particular, taking $C = v_{d+1}v_{d+2} \cdots v_nv_{d+1}$, we have $S_C = S_{n,d}^*$ is an LC d -sphere.

Since it requires the definition of locally constructible, we defer the proof of Proposition 5.1 to section 7.2.

PROPOSITION 5.2: For $d \geq 2$, there exist $c_d > 0$ such that there exists an $S \in \mathcal{S}_n^{(d)}$ with cost at most $c_d n^{1-\frac{1}{d}}$, and $d(n - d)$ facets, with high probability.³

PROOF. We will apply the first moment method to the family Θ^d defined above. Our aim is now to choose a Hamilton cycle C of a complete graph G equipped with suitable edge costs on vertex set $\{v_{d+1}, \dots, v_n\}$ such that the resulting d -sphere S_C has low cost W_{S_C} . The contribution to W_{S_C} of the edge xy is $w_{xy} := \sum_{i=1}^d W_{f_{xy} - \{z_i\}}$, and we assign these w_{xy} as the edge-costs on G . Note that the w_{xy} are i.i.d., and that the total cost of the sphere S_C is $\sum_{xy \in E(C)} w_{xy}$, so this is an instance of the traveling salesman problem on the complete graph with i.i.d. edge costs.

These edge costs are of so-called *pseudo-dimension* d , by which we mean that $\mathbb{P}(w_{xy} \leq t)/t^d \rightarrow c$ as $t \rightarrow 0$ for some constant $c \in (0, \infty)$. (In fact, $\mathbb{P}(w_{xy} \leq t) = t^d/d!$ for $t \leq 1$.) In [35] (an expanded version of [37]), Wästlund studied the traveling salesman problem on K_n with i.i.d. edge costs of pseudodimension $d > 1$. He showed that the minimum cost of a Hamilton cycle is sharply concentrated around $(1 + o(1))c^{-1}\beta_{\text{TSP}}(d)n^{1-1/d}$, where c is as above, and $\beta_{\text{TSP}}(d)$ is some constant only depending on d . (We have K_{n-d} instead of K_n , but the asymptotics do not change.) The proposition follows by picking any $c_d > d!\beta_{\text{TSP}}(d)$. ■

REMARK 5.3: Instead of using the result in [35], a more elementary proof method is to analyze a greedy algorithm for finding a Hamilton cycle, and apply a Chernoff bound similar to Claim 2. The resulting upper bound is only worse by a small constant factor.

³For $d = 2$ section 4 provides a better bound that relies on [30].

5.2 LOWER BOUND

Here we give a first moment lower bound on $M(\mathcal{F})$, similar to that of section 4, for a family \mathcal{F} of subgraphs of $K_n^{(d)}$, given some bound on the size of \mathcal{F} and the sizes of members of \mathcal{F} . We make no other assumptions on the structure of \mathcal{F} .

PROPOSITION 5.4: *Let the hyperedges of $K_n^{(d)}$ be equipped with i.i.d. $U(0,1)$ -costs. Suppose \mathcal{F} is a family of subgraphs of $K_n^{(d)}$ such that*

1. *each $F \in \mathcal{F}$ has at least some m_0 edges, and*
2. *$|\{F \in \mathcal{F} : |E(F)| = m\}| \leq b^m m^{\beta m}$ for some constants $b > 0$ and $0 < \beta \leq 1$,*

and let $c > 0$ be an arbitrary constant. Then, there exists an $r = r(b, c)$ such that $M(\mathcal{F}) \geq r m_0^{1-\beta}$ with probability at least $1 - e^{-cm_0}$.

REMARK 5.5: *Every d -sphere on n vertices has at least $m_0 = dn - (d-1)(d+2)$ facets (see e.g. [18], but note that our d corresponds to $d-1$ there). Part (ii) of Theorem 1.2 is a special case of Proposition 5.4 with $\mathcal{F} = \mathcal{S}_n^{(d)}$ and β such that $A_{n,m} \leq 2^{O(m)} m^{\beta m}$.*

For a set $F \in \mathcal{F}$ of m edges, the probability that their cost W_F is below a given value $L = \Theta(n^\alpha)$ (for some constant $\alpha < 1$) decays superexponentially fast as a function of m . On the other hand, the number of $F \in \mathcal{F}$ with m edges might increase superexponentially fast. If $\alpha + \beta \leq 1$, then the former decay rate outperforms the latter growth rate, so that F 's with few edges dominate the expected number of 'cheap' subgraphs.

PROOF. We will apply the first moment method to the number of 'cheap' $F \in \mathcal{F}$. For some $L = L(n)$ to be determined later, let X_m be the (random) number of graphs $F \in \mathcal{F}$ with precisely m facets and with $W_F \leq L$. Then $\mathbb{E}X_m \leq b^m m^{\beta m} \cdot \mathbb{P}(W_F \leq L)$, where F is an arbitrary subgraph with m edges.

By Lemma 2.1, $\mathbb{P}(W_F \leq L) \leq (Le/m)^m$. Let $L := \frac{\delta}{eb} \cdot m_0^{1-\beta}$, for some $\delta \in (0, 1)$. Since $\beta \leq 1$, we have $m_0^{1-\beta} \leq m^{1-\beta}$, so that $\mathbb{E}X_m \leq (bm^\beta Le/m)^m \leq \delta^m$. By Markov's inequality, $\mathbb{P}(M(\mathcal{F}) \leq L)$ is then at most $\sum_{m \geq m_0} \mathbb{E}X_m \leq \sum_{m \geq m_0} \delta^m < \delta^{m_0}/(1-\delta)$. This is at most e^{-cm_0} if we pick $\delta = \delta(c)$ sufficiently small. \blacksquare

6 CONCENTRATION FOR 2-SPHERES

To prove concentration in the case $d = 2$, i.e. part (ii) of Theorem 1.1, we will use a concentration inequality from [25], for which we need the following definitions.⁴ For any $\mathcal{F} \subseteq \mathcal{S}_n^{(d)}$ and sub-hypergraphs G, P of $K_n^{(d)}$, we say that P is a G -patch if $G \cup P$ contains an $S \in \mathcal{F}$. We define $\rho(G) = \rho_{\mathcal{F}}(G)$ as the minimum number of edges in a G -patch:

$$\rho_{\mathcal{F}}(G) := \min\{e(P) : P \text{ is a } G\text{-patch}\} = \min\{e(S - G) : S \in \mathcal{F}\}.$$

⁴These definitions are given in [25] for general families of sub-hypergraphs, but we present them only for the family $\mathcal{S}_n^{(d)}$ and subsets of it.

We think of G as having $\rho_{\mathcal{F}}(G)$ many ‘holes’, that are ‘patched’ with P . Given a sub-hypergraph G , let the random variable $\text{Patch}_{\mathcal{F}}(G)$ be the minimum cost of a G -patch, i.e. $\text{Patch}_{\mathcal{F}}(G) := \min\{W_P : P \text{ is a } G\text{-patch}\}$.

DEFINITION 6.1: *We say that a sub-hypergraph G of $K_n^{(d)}$ is (λ, ε) -patchable (with respect to the random costs W_e and family \mathcal{F}) if $\mathbb{P}(\text{Patch}_{\mathcal{F}}(G) > \lambda) < \varepsilon$. The family \mathcal{F} is said to be $(r, \lambda, \varepsilon)$ -patchable if every G with $\rho_{\mathcal{F}}(G) \leq r$ is (λ, ε) -patchable.*

We let $\ell(\mathcal{F})$ denote the maximum size of a member of \mathcal{F} , i.e. $\ell(\mathcal{F}) := \max\{e(F) : F \in \mathcal{F}\}$, where $e(F)$ denotes the number of facets in F .

THEOREM 6.2 ([25, Theorem 3.1], special case $q = 1$): *Let \mathcal{F} be a family of sub-hypergraphs of $K_n^{(d)}$, and let Q be the quantile function of $M(\mathcal{F})$, i.e. $\mathbb{P}(M(\mathcal{F}) \leq Q(p)) = p$ for all $p \in (0, 1)$. If \mathcal{F} is $(k, \lambda, \varepsilon)$ -patchable for some $\lambda, \varepsilon > 0$ and $k \geq \sqrt{8 \log(\varepsilon^{-1}) \cdot \ell(\mathcal{F})}$, then with probability at least $1 - 2\varepsilon$,*

$$M(\mathcal{F}) \leq \left(\sqrt{Q(\varepsilon)} + \sqrt{\lambda} \right)^2. \quad (5)$$

In particular, if μ is the median of $M(\mathcal{F})$, and we have $\varepsilon < 1/4$ and $\lambda \leq \mu$, then with probability at least $1 - 3\varepsilon$ both $M(\mathcal{F})$ and μ lie in the interval $[Q(\varepsilon), (\sqrt{Q(\varepsilon)} + \sqrt{\lambda})^2]$, and hence

$$|M(\mathcal{F}) - \mu| \leq 3\sqrt{\lambda\mu}.$$

In order to prove part (ii) of Theorem 1.1 using Theorem 6.2, we need to establish that $\mathcal{S}_n^{(2)}$ satisfies a patchability condition. This is the main effort towards part (ii), which will follow easily from the following lemma.

LEMMA 6.3: *Suppose $k : \mathbb{N} \rightarrow \mathbb{N}$ satisfies $k = \Omega(\sqrt{n})$ and $k = o(n)$. Then for any $r > 0$ there is a constant $A = A(r)$ such that $\mathcal{S}_n^{(2)}$ is $(k, Ak^{\frac{3}{4}}/n^{\frac{1}{12}}, r^k)$ -patchable.*

In particular for $r = 1/2$ and $k = \Theta(n^{0.51})$, there is a $\lambda = o(n^{0.3})$ such that for any subgraph $H \subseteq K_n^{(3)}$ with $\rho(H) \leq k$,

$$\mathbb{P}(\text{Patch}(H) \geq \lambda) \leq 2^{-k}.$$

Before proving Lemma 6.3, let us see how it implies part (ii) of Theorem 1.1:

PROOF OF THEOREM 1.1(ii). Recall that a 2-sphere with n vertices has $2n - 4$ faces, or in other words $\ell(\mathcal{S}_n^{(2)}) = m = 2n - 4$. By Theorem 1.1(i), the median μ of $M(\mathcal{S}_n^{(2)})$ satisfies $\mu = \Theta(\sqrt{n})$.

Let $\varepsilon = \exp(-n^{0.02})/3$, and $k = \sqrt{8m \log(\varepsilon^{-1})} = \Theta(n^{0.51})$. Lemma 6.3, applied with $r = 1/2$, implies that $\mathcal{S}_n^{(2)}$ is $(k, \lambda, 2^{-k})$ -patchable for some $\lambda = o(n^{0.3})$. Since $2^{-k} \ll \varepsilon$ and $\lambda \ll \mu$, the assumptions of Theorem 6.2 are satisfied with μ, λ, k and ε as above, and we obtain

$$|M - \mu| \leq 3\sqrt{\mu\lambda} = o(n^{0.4}),$$

with probability at least $1 - 3\varepsilon = 1 - \exp(-n^{0.02})$. ■

Thus it only remains to prove our lemma.

PROOF OF LEMMA 6.3. We are given a hypergraph H with $\rho(H) =: k$. That is, there is an $S \in \mathcal{S}_n^{(2)}$ such that $P := S - H$ satisfies $|P| = k$. We can assume without loss of generality that S is the disjoint union of P and H , since otherwise we may remove superfluous faces in H by replacing it with $H \cap S$. Our aim is show that there exists an H -patch with cost $O(k^{3/4}/n^{1/12}) = o(k)$ with high probability. That is, a P' such that $H \cup P'$ contains *some* sphere in $\mathcal{S}_n^{(2)}$ – not necessarily the same as S . Before detailing the formal proof, we present an overview. The proof proceeds in two steps:

1. We color the faces in H green, and the faces in P red. We use an asymmetric version of the Lipton–Tarjan planar separator theorem (Lemma 6.4 below) to remove a set Q of roughly s vertices from S (for some s with $\sqrt{s} \ll k \ll s \ll n$), such that the boundary of $S \setminus Q$ is a cycle C of length $O(\sqrt{s})$ (fig. 1(a)). Patch the hole in $S \setminus Q$ by arbitrarily triangulating C without interior vertices. Color the newly added triangles red. This gives us a sphere S' on $n - |Q|$ vertices, with at most $k + O(\sqrt{s}) \approx k$ of its faces being red and the rest green, as well as $|Q| \approx s$ free vertices (fig. 1(b)).
2. For any face xyz in S' , we can use one of the free vertices $v \in Q$ to replace it with a barycentric subdivision, i.e. with the three faces xyv , yzv and $z xv$. We do this for all red faces (fig. 1(b)), and for sufficiently many green faces to use up all vertices in Q (fig. 1(c)). We want to minimize the total cost of the added faces of such a matching between faces in S' and vertices in Q . We pick this matching greedily, and optimize the resulting cost over s .

STEP 1


We proceed with the formal proof. For the first step, we will use the following asymmetric version of the Lipton–Tarjan planar separator theorem.⁵

LEMMA 6.4 ([16, Lemma 3.8]): *Let $s > 0$ be a sufficiently large integer, and let G be a triangulation of \mathbb{S}^2 with $n > 2s$ vertices. Then G contains a cycle C of length at most $28\sqrt{s}$, such that $G \setminus C$ is disconnected into two sets of vertices, Q and Q' , with the smaller of them satisfying $s/14 \leq |Q| \leq s$.*


We apply this lemma with G being the 1-skeleton of S , and $s = s(k, n)$ some integer such that $\sqrt{s} \ll k \ll s \ll n$, whose precise value will be determined later. Let C and Q be the resulting cycle and smaller set of vertices respectively, and let $q := |Q|$, so that $s/14 \leq q \leq s$ and $|C| \leq 28\sqrt{s}$.

Let D' be the simplicial complex obtained from S by deleting the vertices in Q (and all simplices containing any of them), and note that D' is homeomorphic to a disc with boundary C .


⁵The condition that $n > 2s$ is added by us to the lemma; the proof in [16] implicitly assumes that s is not too large. In our application, $s \ll n$.




(a) A sphere with some faces marked red (if colour is shown), with a short planar separator C (dark green) with a set (Q) of 8 interior vertices.



(b) Replace the interior of C with a red triangulation without interior vertices, freeing up the vertices in Q . Match all red faces with free vertices.



(c) Match the remaining free vertices with green faces.



(d) Patched sphere: Consists mostly of the original green faces, and some additional gray faces.

Figure 1: ‘Patching’ a 2-sphere. The aim is to remove the red faces (and possibly some green faces) and add some new faces of low total cost, such that the resulting complex is a 2-sphere. We show that doing the matchings in (b) and (c) greedily gives sufficiently cheap faces in the barycentric subdivisions (gray).

Let D be an arbitrary triangulation of C without interior vertices. That is, D is a 2-manifold homeomorphic to a disc, with boundary C and whose 1-skeleton is an outerplanar graph. Easily, the number of triangular faces in D equals the number of edges in C minus 2. Color the faces of D red. Note that D and D' are both discs with the same boundary C , and so $S' := D \cup D'$ is a simplicial 2-sphere on $n - q$ vertices. Let R be the set of the red faces of S' , and note that $R \subseteq P \cup D^2$, where D^2 denotes the set of faces of D . Thus there are at most $|P| + |D^2| \leq k + 28\sqrt{s} = (1 + o(1))k$ red faces of S' . Let G be the set of green faces of S' , so that S' is the disjoint union of R and G .

STEP 2

For any face $\sigma = xyz$ in S' , we can replace σ with a barycentric subdivision using any vertex $v \in Q$. That is, replace σ with the faces xyv, yzv and zxv at cost

$$X_{\sigma v} := W_{xyv} + W_{yzv} + W_{zxv}.$$

Doing so for all red faces will result in a 2-sphere, possibly on fewer than n vertices because not all q vertices in Q have been used. We can additionally replace some green faces with their barycentric subdivisions until all vertices in Q have been used, finally obtaining a 2-sphere on n vertices.

For convenience, we construct an auxiliary graph: the complete bipartite graph K with left vertex set $R \cup G$ and right vertex set Q , with edge costs $X_{\sigma v}$. Since S' has $n - q$ vertices, the left set has $|R \cup G| = 2(n - q) - 4 =: \ell$ vertices, while the right has $|Q| = q$. By construction, $\ell \gg q$.

Let \mathcal{M} be the set of matchings in K containing every vertex in $Q \cup R$ (and hence such that most vertices in G are *not* matched). Using the convention of identifying an edge σv of K with the set of simplices $\{xyv, yzv, zxv\}$, any $M \in \mathcal{M}$ is an H -patch. The barycentric subdivisions described above correspond to picking a matching in \mathcal{M} .

The total cost of the matching M is $X_M := \sum_{\sigma v \in \mathcal{E}(M)} X_{\sigma v}$, and this is also the cost of the corresponding H -patch. We want to show that there exists an $M \in \mathcal{M}$ of low total cost, using a greedy exploration procedure.

However, this procedure would have been easier to analyse if the edge costs in K were independent. In our case, though, the edge costs $X_{\sigma v}$ and $X_{\sigma'v}$ are not independent if the triangular faces σ and σ' overlap in an edge: If $\sigma = xyz$ and $\sigma' = xyz'$, then both $X_{\sigma v}$ and $X_{\sigma'v}$ depend on W_{xyv} . Since this dependence is only between faces sharing an edge, consider the dual graph of the 1-skeleton of S' , where vertices corresponding to the triangular faces are connected if they share an edge. This graph is 3-regular, thus it is 4-colourable. Let R_1, R_2, R_3 and R_4 be the intersections of R with the four colour classes. Also let $G_1 \subseteq G$ be the largest intersection of a colour class with G , so that $|G_1| \geq |G|/4 = \Omega(n)$. At most $3|R| = o(n)$ faces in G_1 share an edge with a red face. Removing these faces leaves a set $G_0 \subseteq G_1$, with $|G_0| = \Omega(n)$. Crucially, the R_i 's and G_0 picked in this way are independent sets. Finally, let Q_1, Q_2, Q_3 and Q_4 be a partition of Q into four sets of similar sizes – each with at least $|Q|/5$ vertices, say. We will first greedily match all vertices in each R_i with vertices in the corresponding Q_i , and then

greedily match the vertices in Q that have not yet been matched, with vertices in G_0 . We will use the following Chernoff bound on the cost of these matchings.

CLAIM 2 (Chernoff bound): *Let U_1, U_2 and U_3 be i.i.d. random variables following a $U(0, 1)$ -distribution. Assume $\{Z_{ij}\}_{1 \leq i \leq k_0, 1 \leq j \leq n_0}$ are i.i.d. random variables such that $Z_{ij} \stackrel{d}{=} U_1 + U_2 + U_3$. Then for any $b \geq 20k_0$,*

$$\mathbb{P} \left(\sum_i^{k_0} \min_j Z_{ij} \geq bn_0^{-1/3} \right) \leq e^{-b/25}$$

PROOF. Observe that $\mathbb{P}(Z_{ij} < x + n_0^{-1/3} \mid Z_{ij} \geq x)$ is minimised when $x = 0$, at $1/(6n_0)$. It follows that $\mathbb{P}(\min_j Z_{ij} < x + n_0^{-1/3} \mid \min_j Z_{ij} \geq x) \geq 1 - \exp(-1/6) > 0.15$.

Consequently $\lceil n_0^{1/3} \min_j Z_{ij} \rceil$ is dominated by a $\text{Geo}(0.15)$ random variable, and these are independent for different i . Hence the probability that the sum is too large is at most the probability that a $\text{Bin}(b, 0.15)$ random variable takes a value less than $k_0 \leq 0.05b$, which by a standard Chernoff bound is less than $\exp(-b/25)$. ■

Any of the $k_0 = |R_i| \leq |R| < k + o(k)$ vertices in R_i has at least $n_0 = |Q_i| - k_0 = \Omega(s)$ vertices in Q_i to choose from. Apply Claim 2 with $b = 25tk$ for some $t \geq 1$. Then with probability at least $1 - e^{-tk}$, the greedy algorithm matches all vertices in R_i with vertices in Q_i with total cost at most $C_1 k/s^{1/3}$ for some constant $C_1 = C_1(t)$.

After the vertices in R have been matched with vertices in Q in this way, $k_0 = |Q| - |R| \leq s$ vertices in Q remain unmatched, and each of them have at least $n_0 = |G_0| - |Q| = \Omega(n)$ vertices in G_0 to choose between. Apply Claim 2, now with $b = 25ts$. Then with probability at least $1 - e^{-tk}$, the cost of this matching is at most $C_2 s/n^{1/3}$ for some constant $C_2 = C_2(t)$.

The total cost of the greedy matchings is then at most $4C_1 k/s^{1/3} + C_2 s/n^{1/3}$ with probability at least $1 - 5e^{-tk}$. Let $s = k^{3/4} n^{1/4}$, making both terms above $\Theta(k^{3/4}/n^{1/12})$ (this is close to optimal). The lemma follows by picking t sufficiently large compared to r . ■

REMARK 6.5: *In order to try to apply this proof method to spanning spheres in dimension $d \geq 3$, we would need something similar to the planar separator theorem (Lemma 6.4) in higher dimension. If we had such a theorem, then a weaker lower bound on $M(\mathcal{S}_n^{(d)})$ would suffice to prove that this random variable is sharply concentrated. Such a separator theorem might also help approach Gromov's aforementioned question [14] with a divide-and-conquer strategy.*

On the other hand, recall that the planar separator theorem was just used in order to gain access to a sufficient number (at least k) of free vertices. We could instead have considered the complete hypergraph $K_n^{(d+1)}$ on n vertices as a subgraph of the complete hypergraph $K_{n+k}^{(d+1)}$ on $n+k$ vertices, and used these additional k vertices to perform the matching of step 2. Instead of inequality (5), this would lead to an inequality relating $M(\mathcal{S}_{n+k}^{(d)})$ to the quantile functions of both $M(\mathcal{S}_n^{(d)})$ and $M(\mathcal{S}_{n+k}^{(d)})$. Although heuristically

$M(\mathcal{S}_{n+k}^{(d)})$ should only be slightly larger than $M(\mathcal{S}_n^{(d)})$, we neither know the size of their (typical) difference, nor even whether $\mathbb{E}M(\mathcal{S}_{n+k}^{(d)})$ is indeed larger than $\mathbb{E}M(\mathcal{S}_n^{(d)})$.

7 CONCENTRATION FOR $d \geq 3$

7.1 CONCENTRATION ASSUMING MODERATE GROWTH

We will prove the following theorem in the more general setting of an arbitrary family \mathcal{F} of subhypergraphs; we denote the median of $M(\mathcal{F})$ by μ .

THEOREM 7.1: *Let \mathcal{F} be a family of subhypergraphs of $K_n^{(d+1)}$ which for any $m \geq tn$ contains at most $K^m m^{\beta m}$ subgraphs with m hyperedges, for some $t, K > 0$ and $\beta < 1/2$, and which contains no subgraph with fewer than tn hyperedges.*

Then $M(\mathcal{F})$ is sharply concentrated: there exists $\delta > 0$ such that

$$\mathbb{P}(|M(\mathcal{F}) - \mu| > \mu^{1-\delta}) \leq \exp(-n^\delta)$$

for n sufficiently large.

Part (iii) of Theorem 1.2 is a special case of Theorem 7.1. To prove Theorem 7.1, we will need the following simple bound from [25].

LEMMA 7.2 (Special case with $q = 1$ of [25, Proposition 7.1]): *For any $t \geq 0$,*

$$\mathbb{P}(M(\mathcal{F}) > t\mu) \leq 2^{1-t}.$$

PROOF OF THEOREM 7.1. We will show this for the explicit value $\delta = (1/2 - \beta)/4 > 0$. Split \mathcal{F} into two families: \mathcal{F}' consisting of those sets $F \in \mathcal{F}$ of size at most some $m_1 \gg n$ (to be determined later), and $\mathcal{F}'' := \mathcal{F} - \mathcal{F}'$ consisting of the larger F 's. Note that $M(\mathcal{F}) = \min(M(\mathcal{F}'), M(\mathcal{F}''))$. We will show that $M(\mathcal{F}) \leq \mu^{1+4\delta} \ll M(\mathcal{F}'')$ w.h.p., from which it follows that $M(\mathcal{F}) = M(\mathcal{F}')$ w.h.p. We will then show that $M(\mathcal{F}')$ (and hence $M(\mathcal{F})$) is sharply concentrated.

Applying the first moment bound Proposition 5.4 to \mathcal{F} gives $\mu = \Omega((nt)^{1-\beta}) = \Omega(n^{1/2+4\delta})$. Similarly, applying Proposition 5.4 to \mathcal{F}'' gives $M(\mathcal{F}'') = \Omega(m_1^{1/2+4\delta})$.

The bound we will prove on the fluctuations of $M(\mathcal{F}')$ around its median will be roughly $\sqrt{m_1}$. We therefore set $m_1 := \mu^{2-4\delta}$. Observe that, since $0 < 4\delta < 1/2$, we have $(2 - 4\delta)(1/2 + 4\delta) > 1 + 4\delta$.

Using the first moment bound on $M(\mathcal{F}'')$ above, we obtain

$$M(\mathcal{F}'') = \Omega(\mu^{(2-4\delta)(1/2+4\delta)}) \gg \mu^{1+4\delta},$$

and similarly $m_1 = \Omega(n^{(1/2+4\delta)(2-4\delta)}) \gg n^{1+4\delta}$.

To upper bound $M(\mathcal{F})$, we use Lemma 7.2 with $t = \mu^{4\delta} \gg 1$, giving

$$\mathbb{P}(M(\mathcal{F}) \geq \mu^{1+4\delta}) \leq 2^{1-\mu^{4\delta}} = o(1).$$

We now focus on $M(\mathcal{F}')$, and show that it is sharply concentrated around its median. \mathcal{F}' is trivially (k, k, ε) -patchable for any k and ε : Simply put back the k elements that were removed, at cost at most k .⁶ We apply Theorem 6.2 for some $\varepsilon = \varepsilon(n)$ tending to zero not too fast ($\log(\varepsilon^{-1}) \leq \mu^{2\delta}$, say), and $\lambda = k := \sqrt{8 \log(\varepsilon^{-1}) m_1}$. We then have $\log(\varepsilon^{-1}) m_1 \leq \mu^{2-2\delta}$ and hence $\lambda \leq \mu^{1-\delta}$. Letting μ' be the median of $M(\mathcal{F}')$, Theorem 6.2 then gives us that

$$\mathbb{P}\left(|M(\mathcal{F}') - \mu'| \geq 3\mu^{1-\delta}\right) \leq 3\varepsilon.$$

Since $M(\mathcal{F}) = M(\mathcal{F}')$ w.h.p., this inequality also holds for $M(\mathcal{F})$ if we add an extra $o(1)$ term to the right-hand side. The right-hand side is strictly less than $1/2$, so we must furthermore have $|\mu' - \mu| \leq 3\mu^{1-\delta}$. By the triangle inequality, $|M(\mathcal{F}) - \mu| \leq 6\mu^{1-\delta}$ with high probability. ■

7.2 SMALL SUBFAMILIES AND LOCALLY CONSTRUCTIBLE SPHERES

In this section we prove Theorem 1.4, but before doing so let us discuss an important family of spheres to which it applies, namely the *locally constructible spheres*. (The reader may skip this discussion and proceed directly to the proof.)

Tutte's proof of eq. (1) involves a clever and intricate use of generating functions. It is however fairly easy to see that $B_{n,2n-4}^{(2)} \leq C^n$ for *some* constant C , and we include a sketch of that argument here.

PROOF SKETCH. Any 2-sphere can be constructed in a two-stage process. First, glue together $2n - 4$ triangular faces along edges in such a way that they form a disc without interior vertices. That is, the face-dual graph is a tree. This forms a triangulation of a $(2n - 2)$ -gon, and the number of such triangulations is given by the Catalan number $C_{2n-4} = o(4^n)$. Next, pair up the edges along the boundary with a planar perfect matching, and glue matched pairs of edges together. The number of planar matchings can also be counted by the Catalan numbers, so again grow exponentially with n . We can obtain any given isomorphism class of 2-spheres (in many different ways, as we have counted them), since we may base the first stage on any spanning tree of the dual graph. ■

In dimension $d = 3$, the argument above fails in the second stage. After constructing a 3-ball, we are no longer constrained to *planar* matchings when pairing up faces on its boundary.

The way the edges of the disc are glued together in dimension $d = 2$ can also be thought of as a sequence of 'local' glueings: First glue together two adjacent edges. This

⁶For this trivial patchability condition, Theorem 6.2 reduces to the Talagrand inequality, in particular the 'certifiability' corollary as found in [2]. It would have been slightly easier to apply this inequality directly instead of using Theorem 6.2, however we use the latter to avoid stating the inequality and definitions from [2].

shrinks the length of the boundary by two, and causes two previously non-neighbouring edges to become neighbours. Repeat this process until the boundary is eliminated.

Generalizing this notion of local glueings to higher dimensions, Durhuus & Jónsson [10] introduced the class of *locally constructible* manifolds (see Definition 7.3 below). They showed that there are only exponentially many locally constructible manifolds with m facets. As sketched above, all 2-spheres are locally constructible. Durhuus & Jónsson posed the question of whether this is also true for 3-spheres, but it was answered in the negative by Benedetti & Ziegler [7]. For more background on locally constructible spheres, as well as how they relate to shellable and collapsible simplicial complexes, we refer the reader to [7].

DEFINITION 7.3: *The locally constructible (LC) d -balls are those simplicial d -balls that can be constructed by a sequence of the following two moves, starting from a single $(d + 1)$ -simplex.*

1. *Glue a new $(d + 1)$ -simplex along a d -simplex to a d -simplex on the boundary.*
2. *Glue two d -simplices together, provided that they both lie on the boundary and already overlap in a $(d - 1)$ -simplex.*

A locally constructible sphere is the boundary of a locally constructible ball.

We now show that $S_{n,d}^*$ is a locally constructible sphere, and so Theorem 1.4 applies to the locally constructible spheres.

PROOF OF PROPOSITION 5.1. Let $\text{Cone} = \text{Cone}_z$ be the cone operator that takes a simplicial complex \mathcal{K} and adds to it a new vertex z and a simplex $K \cup \{z\}$ for every $K \in \mathcal{K}$.

Note that $B(C) = \text{Cone}^d(C) = \text{Cone}_{v_1}(\text{Cone}_{v_2}(\cdots \text{Cone}_{v_d}(C)\cdots))$. $\text{Cone}(C)$ is homeomorphic to the 2-ball (that is, the disc), and is clearly LC. Note that the cone over a k -ball is homeomorphic to the $(k + 1)$ -ball, and furthermore that the cone of an LC k -ball is also LC (by considering coning every step of the construction). Hence $\text{Cone}^d(C)$ is homeomorphic to the $(d + 1)$ -ball, its boundary is homeomorphic to the d -sphere, and both are LC.

If y, y' are the two neighbours of some vertex x in C , then $\{x, v_1, \dots, v_d\}$ lies in both ∂f_{xy} and $\partial f_{xy'}$, and hence they cancel out in $S_C = \partial B(C)$. On the other hand, any d -simplex containing both x and y (with $xy \in E(C)$) occurs only once, namely in ∂f_{xy} . Hence the facets are

$$\{f_{xy} - \{v_i\} : xy \in E(C), 1 \leq i \leq d\},$$

and there are $|E(C)| \cdot d = (n - d) \cdot d$ such facets. ■

Benedetti & Pavelka [6] extended Definition 7.3 to give the class of t -LC spheres, where the requirement to overlap in a $(d - 1)$ -simplex is relaxed to a $(d - t)$ -simplex. They showed that the larger class of 2-LC spheres is also exponentially bounded, and thus Theorem 1.4 also applies to this class. However, even for $d = 3$ (where these are also known as Mogami spheres), not all spheres are 2-LC [5].

PROOF OF THEOREM 1.4. The argument is similar to that of the proof of Theorem 7.1. First, split \mathcal{P} into two families: \mathcal{P}' , consisting of all spheres in \mathcal{P} with at most $m_1 := (d + \varepsilon)n$ facets, and $\mathcal{P}'' := \mathcal{P} - \mathcal{P}'$. For the upper bound, the sphere S_C in the proof of Proposition 5.2 is in \mathcal{P} by assumption. Furthermore, it has $d(n - d) < m_1$ facets, and hence it also lies in \mathcal{P}' . Thus $M(\mathcal{P}) \leq M(\mathcal{P}') \leq Cn^{1-\frac{1}{d}}$ with high probability, for some C .

Next, there are at most $2^{O(m)}n!$ spheres in \mathcal{P} with m facets, also by assumption. For m in the range $dn \leq m \leq m_0$, we have $n! \leq m^{m/d}$, while for $m \geq m_0$, we have $n! \leq m^{m/(d+\varepsilon)}$. Using Proposition 5.4 on \mathcal{P}' and \mathcal{P}'' with $\beta = 1/d$ and $\beta = 1/(d + \varepsilon)$ respectively therefore yields $M(\mathcal{P}') \geq c'n^{1-\frac{1}{d}}$ and $M(\mathcal{P}'') \geq c''n^{1-\frac{1}{d+\varepsilon}}$, with high probability, for some $c', c'' > 0$. Together with the upper bound on $M(\mathcal{P}')$ above, we deduce $M(\mathcal{P}'') \gg M(\mathcal{P}')$ w.h.p., and hence $M(\mathcal{P}) = M(\mathcal{P}')$ w.h.p.

The proof of concentration then proceeds just as in Theorem 7.1 by using the Talagrand inequality to show that $M(\mathcal{P}')$ is sharply concentrated. \blacksquare

The aforementioned question of Durhuus & Jónsson would have implied the following, if the answer had been positive, using the method of section 4 or 5.2.

CONJECTURE 7.4: *There is a constant $c > 0$ such that for all sufficiently large n ,*

$$\mu_n^{(3)} \geq cn^{2/3}.$$

Here $\mu_n^{(3)}$ denotes the median of $M(\mathcal{S}^{(3)})$ as in Theorem 1.3.

8 OUTLOOK

We conjecture that the upper bound on $\mu_n^{(d)}$ provided by Theorem 1.2 (i) is sharp:

CONJECTURE 8.1: *For every $d \geq 2$ there is a constant $c_d > 0$ such that*

$$\mu_n^{(d)} = c_d n^{1-1/d} + o(n^{1-1/d}).$$

We remark that this is open even for $d = 2$, in which case Theorem 1.1 bounds the possible values of c_2 within the interval $[\alpha, \frac{\varepsilon}{2}\alpha]$.

As mentioned in the introduction, we can consider a more general model where the cost of a triangulation in $\mathcal{S}_n^{(d)}$ influences its probability of being selected via a Boltzmann distribution. To make this precise, let $\lambda = \lambda_{d,n}$ denote the Lebesgue measure on the cube $Q := [0, 1]^{(d+1)}$, to be thought of as a random weighing of $K_n^{(d)}$, assigning independent $U(0, 1)$ -distributed costs to its hyperedges. For each triangulation $S \in \mathcal{S}_n^{(d)}$, and each weighing $w \in Q$, define the *energy* $H(S, w) := \sum_{e \in S} w(e)$, i.e. the total cost of S with respect to w . Finally, choose a random sphere $R_{\beta, n}^{(d)} \in \mathcal{S}_n^{(d)}$ with probability density

$$f_\beta(S) := \frac{e^{-\beta H(S, w)}}{Z_\beta},$$

where Z_β denotes the *partition function* $\sum_{S \in \mathcal{S}_n^{(d)}} \int_Q e^{-\beta H(S,w)} d\lambda(w)$, and $\beta \in [0, \infty)$ is the *inverse temperature* parameter. Thus the probability that $R_{\beta,n}^{(d)}$ coincides with a given $S \in \mathcal{S}_n^{(d)}$ is $\frac{1}{Z_\beta} \int_Q e^{-\beta H(S,w)} d\lambda(w)$.

All results and questions of this paper can be studied with $M(\mathcal{S}_n^{(d)})$ replaced by $R_{\beta,n}^{(d)}$. In particular, we can let $\mu_{n,\beta}^{(d)}$ be the median of $R_{\beta,n}^{(d)}$, and adapt Conjecture 8.1 as follows:

CONJECTURE 8.2: *For every $d \geq 2$ and $\beta \in [0, \infty)$, there are constants $e^{(d)}(\beta), c_{d,\beta} > 0$ such that*

$$\mu_{n,\beta}^{(d)} = c_{d,\beta} n^{e^{(d)}(\beta)} + o(n^{e^{(d)}(\beta)}).$$

Note that for $\beta = 0$ we recover the uniform n -vertex triangulation of the d -sphere. In particular, for $d = 2$, we have $\mu_{n,0}^{(2)} = \Theta(n)$. We believe that $e^{(2)}(\beta) = 1$ for every $\beta \in [0, \infty)$. On the other hand, our Theorem 1.1 says that $\mu_{n,\infty}^{(2)} = \Theta(n^{1/2})$ if we interpret $\beta = \infty$ appropriately. We know nothing about $e^{(d)}(\beta)$ in any dimension $d \geq 3$.

PROBLEM 8.3: *For $d = 1, 2, \dots$, determine $e^{(d)}(\beta)$. In particular, is $e^{(d)}(\beta)$ a constant function of $\beta \in [0, 1)$ for every d ?*

If Conjecture 8.2 is true, it would be interesting to study $c_{d,\beta}$ as a function of β , in particular, to decide its smoothness.

Recall that the uniform n -vertex triangulation $R_n^{(2)}$ of the 2-sphere converges, in the sense of Benjamini & Schramm [9], as proved by Angel & Schramm [3]. It would be very interesting to generalise this to other values of β (and dimensions), and to understand how the limit object varies with β :

CONJECTURE 8.4: *For every $d \geq 2$, and every $\beta \in [0, \infty]$, $R_{\beta,n}^{(d)}$ converges in the sense of Benjamini & Schramm as $n \rightarrow \infty$.*

There are various ways to formulate this conjecture precisely; in its simplest form, we ask about the convergence of the 1-skeleton of $R_{\beta,n}^{(d)}$, considered as a graph with the costs ignored.

There is another well-known convergence result about the uniform n -vertex triangulation of the 2-sphere: if we endow $R_n^{(2)}$ with the graph-metric, and rescale it appropriately (in such a way that its diameter stays bounded), then this metric space converges in distribution in the Gromov–Hausdorff sense as $n \rightarrow \infty$. The limit object is a random metric space, called the *Brownian map*, which is homeomorphic to \mathbb{S}^2 and yet has Hausdorff dimension 4 [26, 32, 4]. The Brownian map and similar models of random fractal geometry are being studied extensively due to connections with mathematical physics, see e.g. [15] and references therein. Our next question is how this generalises to our setup. Let $|R_{\beta,n}^{(d)}|$ denote the metric space obtained from $R_{\beta,n}^{(d)}$ by rescaling the metric by a factor of $1/\text{diam}(R_{\beta,n}^{(d)})$.

PROBLEM 8.5: Does $|R_{\beta,n}^{(d)}|$ converge in distribution in the Gromov–Hausdorff sense for every $d \geq 2$, and every $\beta \in [0, \infty]$? If yes, what is the Hausdorff dimension of the limit space as a function of β ? In particular, is this function monotone in β ?

All of the definitions and questions of this section can be reformulated, and are equally interesting, when the random costs are placed on the 1-cells rather than the d -cells of K_n^d , as in Theorem 4.1.

ACKNOWLEDGEMENT

We are grateful to Christoforos Panagiotis for a helpful discussion about Section 8.

REFERENCES

- [1] David J. Aldous. The $\zeta(2)$ limit in the random assignment problem. *Random Structures Algorithms*, 18(4):381–418, 2001.
- [2] Noga Alon and Joel H. Spencer. *The probabilistic method*. Wiley Series in Discrete Mathematics and Optimization. John Wiley & Sons, Inc., Hoboken, NJ, fourth edition, 2016.
- [3] Omer Angel and Oded Schramm. Uniform infinite planar triangulations. *Comm. Math. Phys.*, 241(2-3):191–213, 2003.
- [4] John C. Baez. The Brownian map. *Notices Amer. Math. Soc.*, 68(5):801–803, 2021.
- [5] Bruno Benedetti. Mogami manifolds, nuclei, and 3D simplicial gravity. *Nuclear Phys. B*, 919:541–559, 2017.
- [6] Bruno Benedetti and Marta Pavelka. 2-LC triangulated manifolds are exponentially many. *Ann. Inst. Henri Poincaré D*, 11(2):363–382, 2024.
- [7] Bruno Benedetti and Günter M. Ziegler. On locally constructible spheres and balls. *Acta Math.*, 206(2):205–243, 2011.
- [8] Itai Benjamini, Eyal Lubetzky, and Yuval Peled. Minimum weight disk triangulations and fillings. *Trans. Amer. Math. Soc.*, 374(5):3265–3287, 2021.
- [9] Itai Benjamini and Oded Schramm. Recurrence of distributional limits of finite planar graphs. *Electron. J. Probab.*, 6:no. 23, 13, 2001.
- [10] Bergfinnur Durhuus and Thórdur Jónsson. Remarks on the entropy of 3-manifolds. *Nuclear Phys. B*, 445(1):182–192, 1995.
- [11] A. M. Frieze. On the value of a random minimum spanning tree problem. *Discrete Appl. Math.*, 10(1):47–56, 1985.

- [12] Alan Frieze. On random symmetric travelling salesman problems. *Math. Oper. Res.*, 29(4):878–890, 2004.
- [13] Agelos Georgakopoulos, John Haslegrave, Richard Montgomery, and Bhargav Narayanan. Spanning surfaces in 3-graphs. *J. Eur. Math. Soc. (JEMS)*, 24(1):303–339, 2022.
- [14] Misha Gromov. Spaces and questions. *Geom. Funct. Anal.*, Special Volume (Part I):118–161, 2000.
- [15] Ewain Gwynne. Random surfaces and Liouville quantum gravity. *Notices Amer. Math. Soc.*, 67(4):484–491, 2020.
- [16] Sariel Har-Peled. A simple proof of the existence of a planar separator. Preprint, arXiv:1105.0103, 2011.
- [17] Freddie Illingworth, Richard Lang, Alp Müyesser, Olaf Parczyk, and Amedeo Sgueglia. Spanning spheres in Dirac hypergraphs. Preprint, arXiv:2407.06275, 2024.
- [18] Gil Kalai. Rigidity and the lower bound theorem. I. *Invent. Math.*, 88(1):125–151, 1987.
- [19] Gil Kalai. Many triangulated spheres. *Discrete Comput. Geom.*, 3(1):1–14, 1988.
- [20] Richard M. Karp. A patching algorithm for the nonsymmetric traveling-salesman problem. *SIAM J. Comput.*, 8(4):561–573, 1979.
- [21] Peter Keevash, Jason Long, Bhargav Narayanan, and Alex Scott. A universal exponent for homeomorphs. *Israel J. Math.*, 243(1):141–154, 2021.
- [22] Daniela Kühn and Deryk Osthus. Hamilton cycles in graphs and hypergraphs: an extremal perspective. In *Proceedings of the International Congress of Mathematicians—Seoul 2014. Vol. IV*, pages 381–406. Kyung Moon Sa, Seoul, 2014.
- [23] Andrey Kupavskii, Alexandr Polyanskii, István Tomon, and Dmitriy Zakharov. The extremal number of surfaces. *Int. Math. Res. Not. IMRN*, 2022(17):13246–13271, 2022.
- [24] Joel Larsson Danielsson. An entropic upper bound on the number of combinatorial 3-spheres with n vertices and m faces. Manuscript, in preparation.
- [25] Joel Larsson Danielsson. A concentration inequality for random combinatorial optimisation problems. Preprint, arXiv:2407.12672, 2024.
- [26] Jean-François Le Gall. Uniqueness and universality of the Brownian map. *Ann. Probab.*, 41(4):2880–2960, 2013.
- [27] Nathan Linial and Roy Meshulam. Homological connectivity of random 2-complexes. *Combinatorica*, 26(4):475–487, 2006.

- [28] Nathan Linial and Yuval Peled. On the phase transition in random simplicial complexes. *Ann. of Math. (2)*, 184(3):745–773, 2016.
- [29] Jason Long, Bhargav Narayanan, and Corrine Yap. Simplicial homeomorphisms and trace-bounded hypergraphs. *Discrete Anal.*, 2022(6):12pp, 2022.
- [30] Zur Luria and Ran J. Tessler. A sharp threshold for spanning 2-spheres in random 2-complexes. *Proceedings of the London Mathematical Society*, 119(3):733–780, 2019.
- [31] R. Meshulam and N. Wallach. Homological connectivity of random k -dimensional complexes. *Random Structures Algorithms*, 34(3):408–417, 2009.
- [32] Grégory Miermont. The Brownian map is the scaling limit of uniform random plane quadrangulations. *Acta Math.*, 210(2):319–401, 2013.
- [33] Vincent Rivasseau. Spheres are rare. *Europhys. Lett. EPL*, 102(6):61001, 2013.
- [34] W. T. Tutte. A census of planar triangulations. *Canadian J. Math.*, 14:21–38, 1962.
- [35] Johan Wästlund. Replica symmetry and combinatorial optimization. Preprint, arXiv:0908.1920, 2009.
- [36] Johan Wästlund. The mean field traveling salesman and related problems. *Acta Math.*, 204(1):91–150, 2010.
- [37] Johan Wästlund. Replica symmetry of the minimum matching. *Ann. of Math. (2)*, 175(3):1061–1091, 2012.