Structural analysis of high-dimensional basins of attraction

Stefano Martiniani,[1, *] K. Julian Schrenk,[1] Jacob D. Stevenson,[2, 1] David J. Wales,[1] and Daan Frenkel[1]

[1]*Department of Chemistry, University of Cambridge, Lensfield Road, Cambridge, CB2 1EW, UK*
[2]*Microsoft Research Ltd, 21 Station Road, Cambridge, CB1 2FB, UK*

We propose an efficient Monte Carlo method for the computation of the volumes of high-dimensional bodies with arbitrary shape. We start with a region of known volume within the interior of the manifold and then use the multi-state Bennett acceptance-ratio method to compute the dimensionless free-energy difference between a series of equilibrium simulations performed within this object. The method produces results that are in excellent agreement with thermodynamic integration, as well as a direct estimate of the associated statistical uncertainties. The histogram method also allows us to directly obtain an estimate of the interior radial probability density profile, thus yielding useful insight into the structural properties of such a high dimensional body. We illustrate the method by analysing the effect of structural disorder on the basins of attraction of mechanically stable packings of soft repulsive spheres.

INTRODUCTION

In science we often face, and occasionally confront, the following question: "Can we estimate the *a priori* probability of observing a system in a very unlikely state?" An example is: "How likely is a given disordered sphere packing?", not to mention questions such as "How likely is life, or the existence of a universe like ours?" within the context of dynamical systems and of the multiverse. In a number of cases, where the states correspond to extrema in a high dimensional function, this question can be narrowed down to: "How large is the 'basin of attraction' of a given state?". In such cases, estimating the probability of observing a particular state is equivalent to computing the volume of the (high-dimensional) basin of attraction of this state. That simplifies the problem, but not by much [1, 2]: analytical approaches are typically limited to highly symmetric (often convex) volumes, whilst 'brute force' numerical techniques can deal with more complex shapes, but only in low-dimensional cases. Computing the volume of an arbitrary, high-dimensional body is extremely challenging. For instance, it can be proved that the exact computation of the volume of a convex polytope is a NP-hard problem [3–5] and, of course, the problem does not get any easier in the non-convex case.

Yet, the importance of such computations is apparent: the volume of the basin of attraction for the extrema of a generic energy landscape, be that of biological molecules [6], an artificial neural network [7–9], a dynamical system [10, 11], or even of a "string theory landscape" (where the minima corresponds to different de Sitter vacua [12, 13]), is essential for understanding the systems' behavior.

In high dimensions, simple quadrature and brute-force sampling fail [14] and other methods are needed. In statistical mechanics, the problem is equivalent to the calculation of the partition function (or, equivalently, the free energy) of a system, and several techniques have been developed to tackle this problem (see e.g [15]). The earliest class of techniques to compute partition functions is based on thermodynamic integration (TI) [15–17], which is based on the idea that a transformation of the Hamiltonian of the system can transform an unknown partition function into one that is known analytically. More recent techniques include histogram-based methods (Wang-Landau [18], parametric and non parametric weighted histogram analysis method (WHAM) [19]) or Nested Sampling [20, 21]. In essence, all these techniques reduce the computation of the partition function to the numerical evaluation of a one-dimensional integral.

Among the above methods Nested Sampling and Wang Landau are Monte Carlo algorithms in their own right, that produce the (binned) density of states as a by-product. On the other hand, TI can be identified as a particular Umbrella Sampling scheme [15], that outputs multiple sets of equilibrium states that can be analysed, either by numerical quadrature (e.g. see the Einstein crystal method [22]), or by WHAM and multi-state Bennet acceptance ratio method (MBAR). All the above methods can be used to compute high-dimensional volumes. However, the choice of the MBAR method [23] is an optimal one. Not only is MBAR non-parametric (no binning is required) and has the lowest known variance reweighting estimator for free energy calculations, but it also eliminates the need for explicit numerical integration of the density of states, thus reducing to a minimum the number of systematic biases.

One reason why brute force methods are not suited to estimate the volumes of high-dimensional bodies, is that for such bodies the volume of the largest inscribed hypersphere, quickly becomes negligible to the volume of the smallest circumscribed hypersphere – and most of the volume of the circumscribed hypersphere is empty. Hence, using a Monte Carlo 'rejection method' to compute the volume of the non-convex body as the fraction of volume contained in a hypersphere [24, 25], does not yield accurate results: the largest contribution should come from points that are barely sampled, if at all.

In this Letter we show that MBAR can be used, not only to arrive at an accurate estimate of a high-dimensional, non-convex volume, but that it also can be

used to probe the spatial distribution of this volume.

COMPUTING HIGH-DIMENSIONAL VOLUMES

Our aim is then to measure the volume of a $n-$ dimensional connected compact manifold $\Omega \subseteq \mathbb{R}^n$ with boundaries. We require this body to be "well guaranteed", i.e. it has both an inscribed and a circumscribed hypersphere [2]. To explore different parts of the non-convex volume, we use a spherically symmetric bias that either favors the sampling of points towards the center, or towards the periphery. We start by performing a series of $K + 1$ random walks under different applied bias potentials, similarly to the Einstein-crystal method [22]. We refer to each of the walkers as a "replica" R_i. Unlike TI, where biasing is always 'attractive' (i.e. it favors larger confinement), in MBAR we are free to choose both attractive and repulsive bias potentials (see SM for details of our implementation). Additionally MBAR uses the full posterior distribution (hence all moments) rather than just the average log-likelihood computed over the posterior, as for TI. The present method directly yields an estimate for the statistical uncertainty in the results that depends on the full distributions and is sensitive to their degree of overlap, thus making the method more robust to under-sampling. In contrast, TI would require an expensive resampling numerical procedure to achieve the same objective.

The Markov Chain Monte Carlo (MCMC) random walk of replica $i \in [0, K]$ will generate samples with unnormalised probability density $q_i(\mathbf{x})$, which for a standard Metropolis Monte Carlo walk is

$$q_i(\mathbf{x}) \equiv e^{-\beta_i U_i(\mathbf{x})} \tag{1}$$

with biasing potential $U_i(\mathbf{x})$ and inverse temperature β_i; from now on we assume $\beta_i = 1$ for all walkers R_i, without loss of generality. The normalised probability density is then

$$p_i(\mathbf{x}) = Z_i^{-1} q_i(\mathbf{x}) \tag{2}$$

with normalisation constant

$$Z_i = \int_{\mathbb{R}^n} q_i(\mathbf{x}) \, \mathrm{d}\mathbf{x}. \tag{3}$$

We require that the bias potential $U_i(\mathbf{x})$ can be factorised as

$$U_i(\mathbf{x}) = \chi_\Omega(\mathbf{x}) u_i(\mathbf{x}) \tag{4}$$

where u_i is the reduced potential function and $\chi_\Omega(\mathbf{x})$ is the "oracle" [2], such that for all choices of $u_i(\mathbf{x})$,

$$U_i(\mathbf{x}) = \begin{cases} u_i(\mathbf{x}) & \text{if } \mathbf{x} \in \Omega \\ \infty & \text{if } \mathbf{x} \notin \Omega \end{cases} \tag{5}$$

We thus have that the normalisation constant in Eq. (3) becomes an integral over the manifold Ω

$$Z_i = \int_{\mathbb{R}^n} e^{-U_i(\mathbf{x})} \, \mathrm{d}\mathbf{x} = \int_\Omega e^{-u_i(\mathbf{x})} \, \mathrm{d}\mathbf{x}. \tag{6}$$

If replica R_M is chosen to have bias $u_M = 0$, by definition Eq. (6) becomes the volume V_Ω. Hence if we can compute the partition function for the reduced potential function $u_M = 0$, we can compute the volume V_Ω.

The MBAR method [23] is a binless and statistically optimal estimator to compute the difference in dimensionless free energy for multiple sets of equilibrium states (trajectories) $\{\mathbf{x}\}_i$ obtained using different biasing potentials $u_i(\mathbf{x})$. The difference in dimensionless free energy is defined as

$$\Delta \hat{f}_{ij} \equiv \hat{f}_j - \hat{f}_i = -\ln \left(\frac{Z_j}{Z_i} \right) \tag{7}$$

which can be computed by solving a set of self-consistent equations as described in Ref. [23]. Note that only the differences of the dimensionless free energies are meaningful as the absolute values $\hat{f}_i$ are determined up to an additive constant and that the "hat" indicates MBAR estimates for the dimensionless free energies, to be distinguished from the exact (reference) values.

Let us define the volume $V_\omega = \pi^{n/2} r_\omega^n / \Gamma(n/2 + 1)$ of a n-ball $\omega \subseteq \Omega$ with radius r_ω centred on $\mathbf{x}_0$ and absolute dimensionless free energy $f_\omega = -\ln V_\omega$. For instance, when the volume of a basin of attraction in a potential energy landscape is to be measured, $\mathbf{x}_0$ is chosen to be the minimum energy configuration and $\omega \subseteq \Omega$ the largest n-ball centred at $\mathbf{x}_0$ that fits in Ω. We also define $\{\mathbf{x}\}_i$ to be the set of states sampled with biasing potential u_i and $\{\mathbf{x}\}_\omega = \cup_{i=0}^K \{\mathbf{x} : |\mathbf{x} - \mathbf{x}_0| \le r_\omega\}_i$ to be the set of states re-sampled within ω with reduced potential

$$u_\omega(\mathbf{x}) = \begin{cases} 0 & \text{if } |\mathbf{x} - \mathbf{x}_0| \le r_\omega \\ \infty & \text{if } |\mathbf{x} - \mathbf{x}_0| > r_\omega \end{cases} \tag{8}$$

In other words we augment the set of states with the additional reduced potential u_ω. Note that MBAR can compute free energy differences and uncertainties between sets of states not sampled (*viz.* with a different reduced potential function) without any additional iterative solution of the self-consistent estimating equations, see Ref. [23] for details.

Computing the free energy difference between the sets of equilibrium states $\{\mathbf{x}\}_\omega$ and $\{\mathbf{x}\}_M$, chosen to have reduced potentials $u_M = 0$ and u_ω, we find that the absolute free energy for the unbiased set of states $\{\mathbf{x}\}_M$ is

$$f_M = f_\omega + (\hat{f}_M - \hat{f}_\omega) \tag{9}$$

where the free energy difference $\hat{f}_M - \hat{f}_\omega$ is obtained by MBAR with associated uncertainty $\delta \Delta \hat{f}_{M\omega}$. The volume

of the manifold is then just $V_\Omega = \exp(-f_M)$ with uncertainty $\delta V_\Omega = V_\Omega \delta \Delta \hat{f}_{M\omega}$. Note that the set of biasing potentials u_i must be chosen so that there is sufficient overlap between each neighbouring pair of $p_i(\mathbf{x})$. For instance for the harmonic bias $u_i = k_i|\mathbf{x} - \mathbf{x}_0|^2/2$ we must choose a set of coupling constants k_i so that all neighbouring replicas have a sufficient probability density overlap.

Under an appropriate choice of biasing potential the present method may yield information such as the radial posterior probability density function, as an easy to compute by-product, details are discussed in the SM.

BASINS OF ATTRACTION IN HIGH DIMENSIONS

We define a basin of attraction as the set of all points that lead to a particular minimum energy configuration by a path of steepest descent on a potential energy surface (PES). Exploring a basin of attraction is computationally expensive because each call to the oracle function $\chi_\Omega(\mathbf{x})$ requires a full energy minimisation and equilibrating a MCMC on a high dimensional support is difficult [26–29]. For this reason little is known about the geometry of these bodies [27, 29–31].

Ashwin et al. [25], defined the basin of attraction as the collection of initial zero-density configurations that evolve to a given jammed packing of soft repulsive disks via a compressive quench. On the basis of 'brute-force' calculations on low-dimensional systems, Ashwin et al. suggested that basins of attraction tend to be "branched and threadlike" away from a spherical core region. However, the approach of ref. [25] breaks down for higher dimensional systems for which most of the volume of the basin is concentrated at distances from the 'minimum' where the overwhelming majority of points do not belong to the basin. The method that we present here allows us to explore precisely those very rarified regions where most of the 'mass' of a basin is concentrated.

In general the representation of *all* high dimensional *convex* bodies should have a hyperbolic form such as the one proposed in the illustration by Ashwin et al. due to the exponential decay in volume of parallel hypersections (slices) away from the median (or equator) [32]. This holds true even for the simplest convex bodies, such as the hypercube, and the underlying geometry need not be "complicated", as one would guess at first from the two-dimensional representation. For the simplest cases of the unit d-sphere and the unit d-cube it can be shown that most of the volume is contained within $\mathcal{O}(1/d)$ of the boundary and that at the same time the volume is contained in a slab $\mathcal{O}(1/\sqrt{d})$ and $\mathcal{O}(1)$ from the equator, irrespective of the choice of north pole, respectively [1, 33]. Hence, there is virtually no interior volume. Such phenomena of concentration of measure are ubiquitous
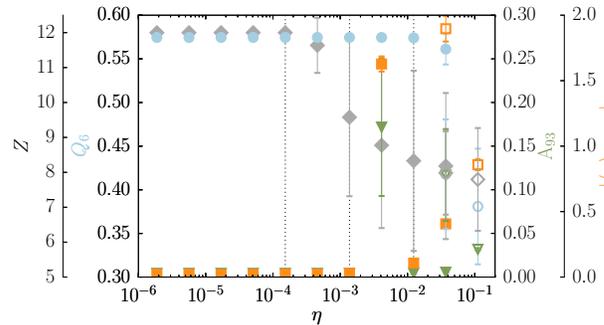
FIG. 1: Structural disorder as a function of polydispersity η is quantified by the average coordination number Z (grey diamonds) and the Q_6 bond orientational order parameter (blue circles); error bars correspond to one standard deviation of the distribution of values per particle. Basin shape is characterized by the asphericity factor A_d (green triangles) and the mean distance of the centre of mass from the minmum (orange squares); error bars correspond to the standard error. Filled and empty markers correspond to packings obtained starting from an fcc and a disordered arragement respectively. Dotted lines show the η after which, in order, Z, A_d and Q_6 change from the fcc value.

in high dimensional geometry and are closely related to the law of large numbers [33].

As we will show, the results presented by Ashwin et al. are, within the resolution available to their method, qualitatively consistent with those for a simple (unit) hypercube.

Effect of structural disorder on the basins of attraction of jammed sphere packings

We characterise the basins of attraction for a number of 32 hard-core plus soft-shell three-dimensional sphere packings, analogous to the ones described in Ref. [29]. The soft shell interactions are short ranged and purely repulsive, the full functional form of the potential and further technical details are reported in the SM. We systematically introduce structural disorder by preparing packings with (geometrically) increasing particle size polydispersity η, i.e. the (positive) radii are sampled from a normal distribution $\mathcal{N}(1, \eta)$. For each η we prepare ~ 10 packings at a soft packing fraction $\phi = 0.74148$ with a soft to hard-sphere radius ratio of $r_{SS}/r_{HS} = 1.12$. The particles are placed initially in a fcc arrangement $\mathbf{x}_{fcc}$ and then relaxed via an energy minimisation to a mechanically stable state $\mathbf{x}_0$. Thus, for the lowest polydispersities the packings remain in a perfect fcc structure and with increasing η they progressively move away into a disordered glassy state. For the largest polydispersity,
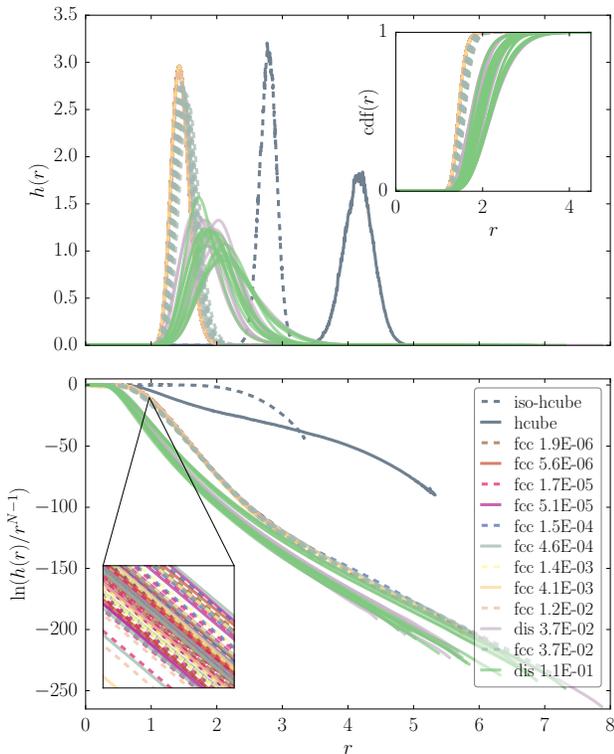
FIG. 2: Top plot shows the measured basin radial probability density function $h(r)$ (DOS) for packings at different polydispersities. The solid and dashed blue curves correspond to the DOS of a 93D hypercube, measured from the centre of mass ('iso-cube') and from a point in one of the corners. The top inset shows the cumulative distribution function for $h(r)$. The bottom panel shows the logarithm of the ratio of the DOS of the basin and of a 93D hyperball. The bottom inset shows the set of barely distinguishable overlapping curves measured for low polydispersities. Top and bottom plots share the x-axis.

for which hard-core overlaps do not allow an initial fcc arrangement, we sample a series of completely random initial states followed by an energy minimisation. Note that even for $\eta \approx 0$, due to the high packing fraction, starting from a completely random set of coordinates, an energy minimisation does not lead to the fcc crystal but rather to the closest glassy state (inherent structure). We are interested in the effect of structural disorder on the shape of the basin of attraction for the soft sphere packings.

We determine the amount of structural disorder in the packing by computing the Q_6 bond orientational order parameter [34] and the average number of contacts per particle Z, shown in Fig. 1. As the polydispersity of the system is increased, the coordination number Z decays monotonically from the close-packed value of 12 to a

value $Z_{\mathrm{fcc}} > Z > Z_{\mathrm{iso}}$, where $Z_{\mathrm{iso}} = 6$ is the average contact number at iso-staticity for a three-dimensional packing of frictionless spheres [35]. The Q_6 order parameter, computed using a solid-angle based nearest-neighbor definition [36], decays from its fcc value well after the contact number has dropped below the close-packed value of 12.

We start characterising the shape of the high dimensional basins of attraction associated with these packings by performing an unconstrained random walk within the basin and performing principal component analysis (PCA) on the trajectory thus obtained [14]. PCA yields a set of eigenvectors that span the d-dimensional configurational space with associated eigenvalues $\lambda_1, \ldots, \lambda_d$. If the basin posses d-dimensional spherical symmetry then all the eigenvalues are expected to be equal. A measure of the shape of a random walk is then the asphericity factor [37]

$$A_d = \frac{\sum_{i>j}(\lambda_i - \lambda_j)^2}{(d-1)\left(\sum_{i=1}^{d}\lambda_i\right)^2}, \qquad (10)$$

that has a value of 0 for a spherically symmetric random walk and of 1 for a walk that extends only in one dimension. Furthermore, we compute the distance of the centre of mass (CoM) position from the minimum energy configuration for the random walk, $|\langle \mathbf{x} \rangle - \mathbf{x}_0|$. This quantity reveals whether the basin is isotropic around the minimum or not. Both quantities, averaged over all packings, are plotted as a function of polydispersity in Fig. 1 along with the structural order parameters. Interestingly, we observe that for low η the basins are, on average, spherically symmetric and isotropic around the minimum. With the onset of structural disorder we observe a marginal increase in asphericity and in the CoM distance from the minimum. In order to observe a significant change however, we need to go to the fully disordered packings at higher polydispersity. With increasing polydispersity, we observe significant changes in the structural order parameters and in the asphericity factor A_d and CoM distance from the minimum.

The implementation details of the MBAR method that we have used are discussed in the SM. Using this method to compute the volume of the basins of attraction, we find excellent agreement with thermodynamic integration, see Fig. S2. As a natural by-product of the computation we are able to compute the radial probability density function (DOS), shown in Fig. 2 together with the logarithm of the ratio between the measured DOS, and that of a d-hypersphere. The log-ratio curves clearly show that all basins have a well-defined hyperspherical core region, where the curves are flat around 0, followed by a series of exponential decays at larger distances from the minimum. For $\eta < 10^{-4}$ the curves are mostly indistinguishable from one another with most of the probability mass concentrated between $1 < r < 3$, as it can be seen from the inset showing the corresponding cumulative distribu-

tion function (CDF). For higher polydispersity, the DOS curves have ever longer tails, as it is also shown by the systematic shift in the CDF.

Importantly, the curves show that a 'rejection' method to measure the basin volume will fail. In this method, the volume of the basin is determined by integrating the fraction of points on a hyper-shell with radius r that fall inside the basin. That fraction is the function shown in the bottom panel of Fig. 2. The most important contribution to the integral would come from the range of r values where $h(r)$ (top panel of Fig. 2) has a significant value. As can be seen from the figure, for disordered systems this happens for values of r where the fraction of hyper-sphere points within the basin is extremely small, in the example shown $\mathcal{O}(10^{-30})$. Hence, the dominant part of the integral would come from parts that are never sampled.

To interpret our results for the DOS curves, it is useful to compare with the corresponding result for a unit hypercube (see Fig. 2). In one instance we do so by placing the 'origin' of the hypercube at its CoM, and in another by placing the origin on one of the 2^d corners of the hypercube, to generate a DOS of a system with a very anisometric density distribution. Not surprisingly, moving the origin of the system from the center to the corner of a hypercube has a dramatic effect on the shape of the DOS, which is now much more similar to the curves for large η, with similar characteristic changes of slope observed for the basins. Again, this agrees with the observation that the CoM distance increases with increasing structural disorder. The effect of the basin asphericity, as measured by the asphericity factor A_d is difficult to infer from the DOS alone.

We thus observe that the structural isotropy and high degree of rotational symmetry in the crystal, as indicated by the Q_6 parameter, is reflected in the isotropy and spherical symmetry of the basin around the minimum, even for relatively large polydispersities when the average contact number has already dropped considerably from the close-packed value. Similarly, the structural disorder at larger η is reflected in the anisotropy and asphericity of the basin. Hence, changes in the basin structure, as indicated by the asphericity factor, the CoM and the density profile, occur before any observable changes occur in Q_6 and after the average contact number ($Z \lesssim 9$) has fallen well below the close-packed value of 12.

* sm958@cam.ac.uk

[1] K. Ball, Flavors of geometry **31**, 1 (1997).
[2] M. Simonovits, Mathematical programming **97**, 337 (2003).
[3] M. E. Dyer and A. M. Frieze, SIAM Journal on Computing **17**, 967 (1988).
[4] L. G. Khachiyan, Uspekhi Mat. Nauk **44**, 199 (1989).
[5] L. G. Khachiyan, Izvestia Akad. Nauk SSSR, Engineering Cybernetics **3**, 216 (1988).
[6] M. A. Miller and D. J. Wales, Journal of Chemical Physics **111**, 6610 (1999).
[7] L. Sagun, V. U. Guney, and Y. LeCun, arXiv:1412.6615 (2014).
[8] A. Ballard, J. D. Stevenson, R. Das, and D. J. Wales, The Journal of Chemical Physics , accepted (2016).
[9] A. Ballard, J. D. Stevenson, and D. J. Wales, in submission (2016).
[10] D. A. Wiley, S. H. Strogatz, and M. Girvan, Chaos: An Interdisciplinary Journal of Nonlinear Science **16**, 015103 (2006).
[11] P. J. Menck, J. Heitzig, N. Marwan, and J. Kurths, Nature Physics **9**, 89 (2013).
[12] J. Frazer and A. R. Liddle, Journal of Cosmology and Astroparticle Physics **2011**, 026 (2011).
[13] B. Greene, D. Kagan, A. Masoumi, D. Mehta, E. J. Weinberg, and X. Xiao, Physical Review D **88**, 026005 (2013).
[14] C. M. Bishop, *Pattern recognition and machine learning* (Springer, New York, 2009).
[15] D. Frenkel and B. Smit, *Understanding molecular simulation* (Academic Press, San Diego, 2002).
[16] J. G. Kirkwood, Journal of Chemical Physics **3**, 300 (1935).
[17] A. Gelman and X.-L. Meng, Statistical Science **13**, 163 (1998).
[18] F. Wang and D. P. Landau, Physical Review Letters **86**, 2050 (2001).
[19] M. Habeck, in *International Conference on Artificial Intelligence and Statistics* (2012) pp. 486–494.
[20] J. Skilling, AIP Conf. Proc. Bayesian inference and maximum entropy methods in science and engineering **735**, 395 (2004).
[21] S. Martiniani, J. D. Stevenson, D. J. Wales, and D. Frenkel, Physical Review X **4**, 031034 (2014).
[22] D. Frenkel and A. J. C. Ladd, Journal of Chemical Physics **81**, 3188 (1984).
[23] M. R. Shirts and J. D. Chodera, Journal of Chemical Physics **129**, 124105 (2008).
[24] S. Liu, J. Zhang, and B. Zhu, in *Computing and Combinatorics*, Lecture Notes in Computer Science, Vol. 4598, edited by G. Lin (2007) pp. 198–209.
[25] S. S. Ashwin, J. Blawzdziewicz, C. S. O'Hern, and M. D. Shattuck, Physical Review E **85**, 061307 (2012).
[26] N. Xu, D. Frenkel, and A. J. Liu, Physical Review Letters **106**, 245502 (2011).
[27] D. Asenjo, J. D. Stevenson, D. J. Wales, and D. Frenkel, Journal of Physical Chemistry B **117**, 12717 (2013).
[28] D. Asenjo, F. Paillusson, and D. Frenkel, Physical Review Letters **112**, 098002 (2014).
[29] S. Martiniani, K. J. Schrenk, J. D. Stevenson, D. J. Wales, and D. Frenkel, Phys. Rev. E **93**, 012906 (2016).
[30] N. Xu, V. Vitelli, A. J. Liu, and S. R. Nagel, EPL **90**,

56001 (2010).

[31] K. Wang, C. Song, P. Wang, and H. A. Makse, Physical Review E **86**, 011305 (2012).

[32] V. Milman, in *European Congress of Mathematics* (Birkhuser, Basel, 1998) pp. 73–91.

[33] V. Guruswami and R. Kannan, "Computer science theory for the information age," (2012), `https://www.cs.cmu.edu/~venkatg/teaching/CStheory-infoage/`.

[34] P. J. Steinhardt, D. R. Nelson, and M. Ronchetti, Physical Review B **28**, 784 (1983).

[35] C. S. O'Hern, L. E. Silbert, A. J. Liu, and S. R. Nagel, Physical Review E **68**, 011306 (2003).

[36] J. van Meel, L. Filion, C. Valeriani, and D. Frenkel, The Journal of Chemical Physics **136**, 234107 (2012).

[37] J. Rudnick and G. Gaspari, Science **237**, 384 (1987).

Supplemental material: Structural analysis of high-dimensional basins of attraction

Stefano Martiniani,[1, *] K. Julian Schrenk,[1] Jacob D. Stevenson,[2, 1] David J. Wales,[1] and Daan Frenkel[1]

[1]*Department of Chemistry, University of Cambridge, Lensfield Road, Cambridge, CB2 1EW, UK*
[2]*Microsoft Research Ltd, 21 Station Road, Cambridge, CB1 2FB, UK*

I.   VOLUME COMPUTATION

We choose a set of harmonic bias potential functions

$$u_i = \frac{1}{2}k_i|\mathbf{x} - \mathbf{x}_0|^2 \tag{S1}$$

with $k_{i \in [1,M]} = \{k_1, \ldots, k_{M-1}, 0\}$ and perform 10^{10} Hamiltonian Parallel Tempering steps as described in Ref. [1] [2]. Note that each Monte Carlo step is followed by a full energy minimisation to test whether the walker has stepped outside the basin of attraction. We choose half of the k's to be positive and the other half negative to accelerate equilibration as well as to increase the DOS resolution near the boundary of the basin. The distributions obtained from the replicas with negative coupling constants contribute to the final MBAR volume estimation, unlike for TI. We stress that the choice of biasing potential is arbitrary. Near the origin we sample the set of configurations $\{\mathbf{x}\}_0$ directly from a hypersphere centred at $\mathbf{x}_0$ with radius sampled from a Gaussian distribution with standard deviation $\sigma = \sqrt{\langle|\mathbf{x} - \mathbf{x}_0|^2\rangle_{k_1}}$, corresponding to a coupling constant $k_0 = 1/\sigma^2$. This choice of k_0 is such that there is sufficient overlap between the distributions of $\{\mathbf{x}\}_0$ and $\{\mathbf{x}\}_1$, as can be verified looking at the two leftmost curves in Fig. (S1) [3]. The corresponding bias potential function is

$$u_0 = (n-1)\log|\mathbf{x} - \mathbf{x}_0| + \frac{1}{2}k_0|\mathbf{x} - \mathbf{x}_0|^2, \tag{S2}$$

where the first term on the right-hand-side is the log-DOS for a n-ball, necessary to account for the greater entropy associated with the regions of space further away from the origin. For a system of N particles in d dimensions with fixed centre of mass we have $n = (N-1)d$ degrees of freedom. The overhead associated with this calculation is insignificant compared to the Hamiltonian Parallel Tempering since the samples thus drawn are completely uncorrelated.

We compute the reduced free energy differences between each of $1 + 31$ replicas with reduced potential functions given by Eqs. (S1)–(S2) using PyMBAR [4, 5]. As reference volume we choose ω to be the n-ball of radius r_ω centred on $\mathbf{x}_0$ with approximately $\mathcal{R} = 0.9$ of its volume contained within the basin Ω. We choose $\omega \not\subseteq \Omega$ to allow more samples with $|\mathbf{x} - \mathbf{x}_0| \leq r_\omega$ thus reducing the uncertainty in the MBAR estimate. For $\mathcal{R} \approx 1$ we can correct exactly for this by noting that

$$\mathcal{R} = \frac{1}{V_\omega}\int_\omega p_0(\mathbf{x})d\mathbf{x} \tag{S3}$$

and $\mathcal{R}$ can be computed directly by Monte Carlo. We thus rewrite Eq. (9) as

$$f_M = f_\omega - \log\mathcal{R} + (\hat{f}_M - \hat{f}_\omega). \tag{S4}$$

Note that the difference in reduced free energies computed using a reference sphere of radius $r_\omega/2$ or $2r_\omega$ is within the statistical uncertainty, hence the method is robust with respect to the choice of reference sphere. We also note that this method ought not be limited to the n-ball as the choice of reference volume, in fact any geometrical body $\omega \subseteq \Omega$ of known volume and surface (thus for which a similar expression to Eq. (S2) can be derived) is suitable, for instance a hypercube or a hyperellipsoid. If $\omega \not\subseteq \Omega$ then an accurate estimate of $\mathcal{R}$ must be available.

A.   Density of states

From the analysis of the posterior probability density functions, the present method may yield structural information, as an easy to compute by-product. Choosing a set of biasing potentials $u_i(r)$ that are a function of the distance

---

* sm958@cam.ac.uk
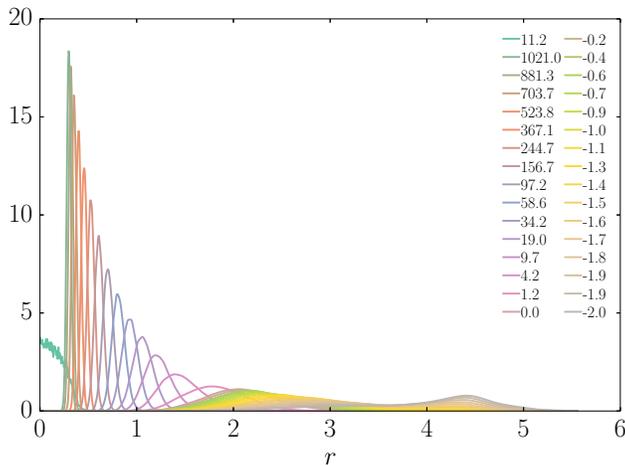
FIG. S1: Kernel density estimation of the distance sampled by a random walk within the basin, coupled to the minimum with decreasing coupling constant from left to right. The left-most curve was obtained by direct sampling as described in the main text. Replicas with negative coupling constants explore regions of the volume that would otherwise never be visited. This particular example is of a disordered packing with polidispersity $\eta = 0.037$.

from the origin $r = |\mathbf{x} - \mathbf{x}_0|$, we can compute the overall density of states (DOS) for the manifold as a function of r. From each of the $K + 1$ replicas' trajectories $\{\mathbf{x}\}_i$ we obtain a (binless) kernel density estimation (KDE) [6] of the probability density functions $h_i(r)$, see Fig. (S1) for an example, which must be unbiased and summed over all replicas to obtain the overall log-DOS function as

$$\log h(r) = \sum_{i=0}^{K} w_i(r) \left(\log h_i(r) + u_i(r) - \Delta \hat{f}_{0i} \right). \tag{S5}$$

where $w_i(r) = h_i(r) / \sum_{i=0}^{K} h_i(r)$ are normalised weights and $\Delta \hat{f}_{0i}$ are the free energy differences between replicas R_i and R_0.

B. Comparison to thermodynamic integration
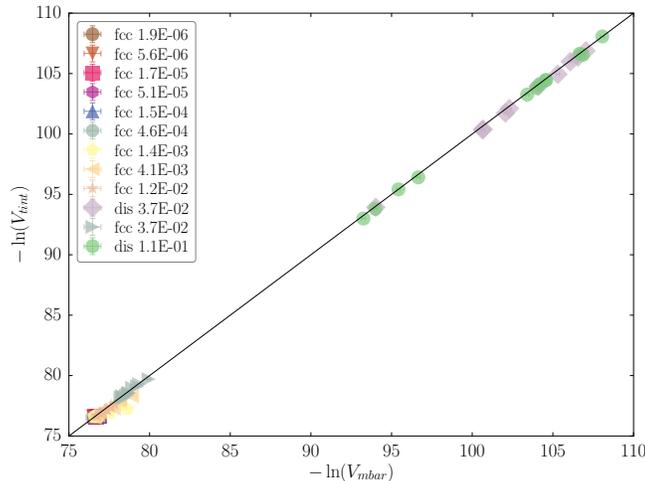
FIG. S2: Comparison of the volumes computed by thermodynamic integration, using only the replicas with positive coupling constant, and by MBAR following the protocol described in this work.
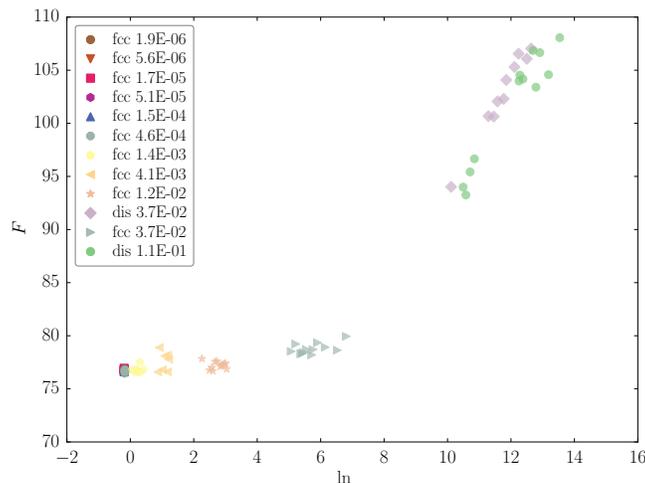
FIG. S3: Dimensionless free energy ($F \equiv -\ln V$) versus pressure of the mechanically stable states analysed in the main text. For details of the the strong correlation between pressure and volume refer to Martiniani et al.[1].

II. JAMMED PACKINGS OF POLYDISPERSE HS-WCA SPHERES

We draw $N = 32$ particle radii $\{r_{\mathrm{HS}}\}_N$ from a Gaussian distribution Normal$(1, \eta) > 0$, truncated at $r_{\mathrm{HS}} = 0$, set the box size to meet the target packing fraction of the hard sphere fluid ϕ_{HS} and then place the particles in a valid, either fcc or random, initial hard sphere configuration, as described in the main text.

Given these hard sphere configurations, we switch on a soft repulsive interaction to generate over-compressed jammed packings of the particles and relax the system to a mechanically stable state by energy minimization. The particles are inflated with a WCA-like potential [7] to reach the target soft packing fraction $\phi_{\mathrm{SS}} > \phi_{\mathrm{HS}}^{(\mathrm{RCP})} > \phi_{\mathrm{HS}}$. The hard spheres are inflated proportional to their radius, so that the soft sphere radius is

$$r_{\mathrm{SS}} = \left(\frac{\phi_{\mathrm{SS}}}{\phi_{\mathrm{HS}}} \right)^{1/d} r_{\mathrm{HS}}, \tag{S6}$$

where d is the dimensionality of the box, r_{SS} and r_{HS} the soft and hard sphere radii respectively. Clearly, this procedure does not change the polydispersity of the sample.

We define the WCA-like potential around a hard core as follows: consider two spherical particles with hard core distance r_{HS} and soft core contact distance $r_{\mathrm{SS}} = r_{\mathrm{HS}}(1 + \theta)$, with $\theta = (\phi_{\mathrm{SS}}/\phi_{\mathrm{HS}})^{1/d} - 1$. We can then write a horizontally shifted hard-sphere plus WCA (HS-WCA) potential as

$$v_{\mathrm{HS\text{-}WCA}}(r) = \begin{cases} \infty & r \leq r_{\mathrm{HS}}, \\ 4\epsilon \left[\left(\dfrac{\sigma(r_{\mathrm{HS}})}{r^2 - r_{\mathrm{HS}}^2} \right)^{12} \right. & \\ \left. - \left(\dfrac{\sigma(r_{\mathrm{HS}})}{r^2 - r_{\mathrm{HS}}^2} \right)^6 \right] + \epsilon & r_{\mathrm{HS}} < r < r_{\mathrm{SS}}, \\ 0 & r \geq r_{\mathrm{SS}} \end{cases} \tag{S7}$$

where $\sigma(r_{\mathrm{HS}}) = (2\theta + \theta^2)r_{\mathrm{HS}}^2/2^{1/6}$ guarantees that the potential goes to zero at r_{SS}. For computational convenience (avoidance of square-root evaluations), the potential in Eq. S7 differs from the WCA form in that the inter-particle distance in the denominator of the WCA potential has been replaced with a difference of squares.

Numerically evaluating this potential, we match the gradient and linearly continue the function $v_{\mathrm{HS\text{-}WCA}}(r)$ for $r \leq r_{\mathrm{HS}} + \varepsilon$, with $\varepsilon > 0$ an arbitrary small constant, such that minimisation is still meaningful if hard core overlaps do occur.

Energy minimisations are performed with the CG_DESCENT algorithm [8–10].

---

[1] Stefano Martiniani, K. Julian Schrenk, Jacob D. Stevenson, David J. Wales, and Daan Frenkel, "Turning intractable counting into sampling: Computing the configurational entropy of three-dimensional jammed packings," Phys. Rev. E **93**, 012906 (2016).

[2] Note that for convenience we used the same choice of positive k's as required for thermodynamic integration. For this particular method any choice of k's is appropriate, typically a geometric distribution, denser for small k and coarser near k_1 is also suitable.

[3] To do so we sample a direction from the surface of the unit sphere and the length of the displacement from a Normal$(0, \sigma)$.

[4] M. R. Shirts and J. D. Chodera, "Statistically optimal analysis of samples from multiple equilibrium states," Journal of Chemical Physics **129**, 124105 (2008).

[5] https://github.com/choderalab/pymbar.

[6] C. M. Bishop, *Pattern recognition and machine learning* (Springer, New York, 2009).

[7] John D Weeks, David Chandler, and Hans C Andersen, "Role of repulsive forces in determining the equilibrium structure of simple liquids," Journal of Chemical Physics **54**, 5237 (1971).

[8] William W Hager and Hongchao Zhang, "A new conjugate gradient method with guaranteed descent and an efficient line search," SIAM Journal on Optimization **16**, 170 (2005).

[9] William W Hager and Hongchao Zhang, "Algorithm 851: CG_DESCENT, a conjugate gradient method with guaranteed descent," ACM Transactions on Mathematical Software **32**, 113 (2006).

[10] https://github.com/smcantab/PyCG_DESCENT.