Geometrical frustration in two dimensions: Idealizations and realizations of a hard-disk fluid in negative curvature

Carl D. Modes and Randall D. Kamien

Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, Pennsylvania 19104-6396, USA

We examine a simple hard-disk fluid with no long-range interactions on the two-dimensional space of constant negative Gaussian curvature, the hyperbolic plane. This geometry provides a natural mechanism by which global crystalline order is frustrated, allowing us to construct a tractable, one-parameter model of disordered monodisperse hard disks. We extend free-area theory and the virial expansion to this regime, deriving the equation of state for the system, and compare its predictions with simulations near an isostatic packing in the curved space. Additionally, we investigate packing and dynamics on triply periodic, negatively curved surfaces with an eye toward real biological and polymeric systems.

PACS number(s): 64.10.+h, 02.40.Ky, 05.20.Jj, 61.43.−j

I. INTRODUCTION

The true nature of glasses and the glass transition has long remained inscrutable. Glasses belong to a wide class of disordered materials and systems which include granular materials like sand [1] or wheat grains [2], frustrated systems like spin glasses [3], and soft systems like foams [4] and emulsions [5]. Unfortunately, the analogies afforded by this classification have done little to lift the veil of mystery: from misunderstandings about apparent "flow" displayed by stained glass windows from the middle ages [6] to debate over the meaning of randomness [7], few aspects of the problem have escaped some measure of controversy. The prominent role of geometrical frustration in these systems, however, has long been agreed upon [8], and it is here that we turn our attention.

We consider an idealized colloidal glass, made of monodisperse spheres in three dimensions with no long-range interactions and a purely hard-core, short-range potential. The entropically favored, high-density fcc lattice (there is a virtual tie between fcc and hcp lattices and lattices with the stacking faults that relate them) for the system is well understood and highly stable, but nevertheless, the system can exhibit a glass transition to a dynamically arrested, amorphous state at a packing fraction significantly lower than that of the crystal. One way to understand this is to recognize that there is a competition at play between local and global properties of the spheres' arrangements. On the one hand, the best the colloids can hope to do, entropically speaking, is to arrange themselves into one of the close-packed, global crystal configurations. On the other hand, these crystal configurations are *not* the best local arrangements—a tetrahedral packing is locally preferable, but cannot be maintained beyond small distances without the introduction of defects.

In spite of this qualitative understanding of the importance of geometrical frustration, the full colloidal glass system remains difficult to tame, particularly near the transition [9]. In principle, one could imagine constructing a simpler system that shares enough of the qualities of the colloidal glass that a full analysis yields insights into the more complicated case. A time-honored tradition along these lines is the study of the same system in fewer dimensions [10]. In

this case, we would have monodisperse, hard disks with no long-range interaction. Notice, however, that we have lost geometrical frustration. The entropically-favored, global, hexagonal crystal is composed precisely of the locally optimal units, triangular configurations of disks. Indeed, two-dimensional, monodisperse disks do not exhibit a glass transition [9].

How, then, can we use a simplified system to say anything about the genuine article? We must find a way to modify the two-dimensional case so as to recover some measure of geometrical frustration. One method that has been well studied is the use of bi- or polydisperse disks [9], frequently with disk radii in a ratio of 1.4:1. The use of two or more disk sizes successfully reintroduces geometrical frustration and recovers the glass transition. Unfortunately, this method of turning on a glass transition in two dimensions is less than ideal. Phase transitions into locally monodisperse, crystalline regions are a constant worry along with the complications inherent to the explicit introduction of additional particle types. Even more troubling, however, is the nature of the parameter space for the system: polydispersity among the particles begets a *functional* parameter, the size-distribution function. Since there is not even a restriction on the continuity of this function, the parameter space is extraordinarily large, making a systematic exploration hopelessly untenable.

So, is there a way we can keep the two-dimensional system monodisperse and yet still manage to frustrate the local and global packings, all while suppressing the size of the relevant parameter space? Indeed, in the seminal work of Kléman and Sadoc [8] and Nelson [11], background curvature is introduced to the system, achieving the desired frustration. Even better, if we require that the system be isotropic and homogeneous, this background curvature must be uniform, leaving us with a simple, *one-dimensional* parameter space to explore. To picture how curvature may be used to frustrate the system, imagine that our two-dimensional space is a sphere and that the hard disks are embedded in it as spherical caps. Notice that there is no longer enough room around each spherical cap to fit six additional caps, as there would be in flat space. More precisely, the number of disks (or spheres or hyperspheres) that can simultaneously touch a central one, the kissing number n_{kiss}, varies continuously as a

function of the curvature K. In two dimensions, this relationship is given by [12]

$$n_{\text{kiss}} = \frac{2\pi}{\cos^{-1}\left(\frac{\cosh(2\sqrt{-K}r)}{1+\cosh(2\sqrt{-K}r)}\right)}, \qquad (1)$$

where r is the radius of the disks in question. In order for local triangular packings of disks to be commensurate with the best global packing, n_{kiss} must be an integer. This condition is satisfied only for flat space ($K=0 \rightarrow n_{\text{kiss}}=6$) and for isolated nonvanishing curvatures.

Historically, curved backgrounds in this context have been used to study crystalline states on curved surfaces [11,12], and the ordered defect patterns that emerge as a result, or to imagine the ideal, unfrustrated, ordered configurations in three dimensions with the hope of understanding how "decurving" might lead to the physical, disordered states we observe in glassy systems in the manner of the cholesteric blue phase [8,13]. In contrast to these approaches, we choose to study the disordered fluid regime of hard-core systems in uniform curvature. Positive curvature may, at first thought, seem more suitable due to its intuitive appeal and the fact that a sphere may be embedded in three-dimensional Euclidean space ($\mathbb{R}^3$) without distortion. However, since the surface area of the sphere is finite and smaller particles on the sphere's surface probe ever lower curvatures (see the Appendix), there is no way to extract a sensible thermodynamic limit at a given curvature.

Instead, we turn to the space of constant negative curvature, the hyperbolic plane $\mathbb{H}^2$, where the thermodynamic limit is not impossible, but remains subtle nonetheless. While the sphere suffers from not having enough room for a thermodynamic limit, on $\mathbb{H}^2$ there may actually be *too much* room. To see why this might be so, consider a circle of radius r on $\mathbb{H}^2$. Its area is $A(r) = -2\pi K^{-1}[\cosh(\sqrt{-K}r)-1] \approx \pi r^2 - \frac{\pi}{12}Kr^4$ for small r, consistent with the known $K=0$ result. Meanwhile, the circle's circumference is $C(r) = 2\pi \sinh(\sqrt{-K}r)/\sqrt{-K}$, which grows as fast as $A(r)$. Thus, the area within an arbitrary distance of the boundary of a circle, or any other fixed shape, is always a finite fraction of the entire region's area, no matter how large the region becomes. Since the traditional thermodynamic limit depends on the effect of the boundary vanishing at infinite size, we must tread carefully. Because our system has interactions that are purely local, the ideal gas provides a useful test case for managing the issue.

The textbook approach to the ideal gas law begins with particle-in-a-box eigenenergies and takes the large-area limit to calculate the partition function for a single particle, $Z_1 = V/(\lambda_T)^d$ where $\lambda_T = \sqrt{2\pi\hbar^2/(mk_BT)}$ is the thermal wavelength. We might expect that the partition function, even on a curved manifold, would be proportional to the proper hypervolume of the box. Indeed, since any partition sum is dominated by the infinite number of high-energy, short-wavelength terms, one would think that these modes are insensitive to the curvature as soon as their de Broglie wavelength is shorter than the radius of curvature. For concreteness, let us consider the partition function for an ideal gas on a hemisphere where the wave functions are forced to vanish

on the equator; the energies are the eigenvalues of the Laplacian on the hemisphere with Dirichlet boundary conditions. On the whole sphere, the spherical harmonics $Y_{\ell m}(\Omega)$ are the complete set of eigenfunctions with eigenvalue $\ell(\ell+1)$. The eigenfunctions on the hemisphere will be a subset of the spherical harmonics, with the boundary condition enforcing the constraint that $\ell+|m|$ be odd. Therefore, the partition function will be

$$Z_1 = \sum_{\ell=1}^{\infty} \ell \exp\left\{-\frac{\hbar^2\ell(\ell+1)}{2mk_BTR^2}\right\} = \sum_{\ell=1}^{\infty} \ell \exp\left\{-\frac{\lambda_T^2\ell(\ell+1)}{4\pi R^2}\right\}, \qquad (2)$$

where R is the radius of the sphere. We can re-express this partition function as an expansion in λ_T/R by employing the Euler-Maclaurin formula [14]

$$\sum_{n=i}^{j} f(n) \sim \int_{i}^{j} f(x)dx + \frac{f(i)+f(j)}{2} \\ + \sum_{k=1}^{\infty} \frac{b_{2k}}{(2k)!}[f^{(2k-1)}(j) - f^{(2k-1)}(i)], \qquad (3)$$

where b_{2k} are the even Bernoulli numbers. Substituting $f(n)=n\exp[\lambda_T^2 n(n+1)/4\pi R^2]$, $i=0$, and $j=\infty$ gives the desired expansion

$$Z_1 = \frac{2\pi R^2}{\lambda_T^2} - \frac{\pi R}{2\lambda_T} + \frac{1}{6} + O(\lambda_T/R). \qquad (4)$$

Notice that we get terms proportional to both the area and the perimeter, but also a constant piece that one may speculate is related to the topology of the manifold. Similarly, we could consider the partition function with Neumann boundary conditions along the equator—i.e., vanishing derivative. In this case the constraint is that $\ell+|m|$ be even and the appropriate sum to consider is

$$Z_1 = \sum_{l=1}^{\infty} (\ell+1)\exp\left\{-\frac{\hbar^2\ell(\ell+1)}{2mTR^2}\right\} \\ = \frac{2\pi R^2}{\lambda_T^2} + \frac{\pi R}{2\lambda_T} + \frac{1}{6} + O(\lambda_T/R). \qquad (5)$$

Again we see an area and perimeter term, the latter with the opposite sign, and the same constant piece.

To probe more exotic surfaces, one could attempt this kind of calculation on a case-by-case basis, but fortunately, the spectrum of the Laplacian on a general Riemann surface is well characterized. Originally, Weyl studied the asymptotic behavior of the large eigenvalues [15]. Kac, in asking, "Can one hear the shape of a drum?" [16], found our constant $\frac{1}{6}$ and, finally, McKean and Singer [17] developed the "Weyl expansion" for the partition function of a single particle on a two-dimensional domain D:

$$Z_1(T) = \frac{1}{\lambda_T^2} \int_D dA \pm \frac{1}{4\lambda_T} \int_{\partial D} d\ell + \frac{1}{12\pi} \int_D K \, dA$$

$$+ \frac{1}{12\pi} \int_{\partial D} \kappa_g d\ell + O(\lambda_T), \qquad (6)$$

where the sign of the first correction depends on Neumann (+) or Dirichlet (−) boundary conditions and κ_g is the geodesic curvature of the boundary considered relative to its outward-directed normal. Our calculations are, unsurprisingly, in complete agreement with their theorem, which, interestingly, works in any dimension. Returning to the question at hand—namely, the hyperbolic plane—we can see that the circumference directly enters the partition function at uncomfortably low order. What is to be done? Periodic boundary conditions provide an out—by eliminating the boundaries and considering sufficiently large periodic boxes we may recover some semblance of a thermodynamic limit. We hasten to point out that even this procedure is not free of its own subtleties—due to the inherent length scale provided by the curvature (Appendix), a simple rescaling of the periodic box is not allowed. In particular, by symmetry, the edges of the periodic box must have vanishing curvature. Thus the Gauss-Bonnet theorem, which relates the background curvature K, the geodesic curvature κ_g, and the jump angles at any vertices, α_i,

$$\int\int_T K \, dA = 2\pi - \sum_i \alpha_i - \int_{\partial T} \kappa_g ds, \qquad (7)$$

imposes a relation between the area contained within the box and the polygon's angles. At the same time, to preserve homogeneity after the periodic boundaries have been applied, any identified vertices must end up with 2π of angular space around them, thus fixing the box's area. The simplest topological considerations force the number of sides to be a multiple of 4 [18], and we are left with the familiar square cell in flat space, along with a specific-size octagon, dodecagon, and so forth on $\mathbb{H}^2$. Luckily, the area of these p-gons increases with p [18], and we are reassured that sufficiently large periodic boxes are achievable.

Recently, we took the first several steps in a comprehensive treatment of just such a system of hard disks on $\mathbb{H}^2$ [19]. Our examination included a study of ordered disk packings leading to the application of free-area theory, a calculation of the first several virial coefficients of the system, and a molecular dynamics realization of the fluid in the smallest possible periodic region. In this paper, we extend and detail these results by considering a constant-curvature recasting of the free-area theory and calculating the virial coefficients to higher precision while applying Padé approximants to get a qualitative look at the behavior of the divergence in ϕ as a function of the curvature. We expand our simulation to a larger polygon with more sides, allowing for better statistics through more simulated disks, and use it to probe the zero-curvature limit from below. In addition, we study entropically governed packings on surfaces with average negative curvature in $\mathbb{R}^3$, such as the triply periodic minimal surfaces found in diblock copolymers, lipids, or even the mitochon-

dria of *Chaos carolinensis* [20], among other biological examples.

This paper is organized as follows. In Sec. II, we review the theory of regular tesselations in two dimensions of arbitrary curvature and extend this theory to packings of hard disks with crystalline order. Section III develops a free-area theory from these ordered packings on curved space. In Sec. IV, we calculate the first five terms in the virial expansion as a function of curvature and comment on the analogy between high-dimensional flat space and the highly curved plane. Section V presents a molecular dynamics model of a hard-disk fluid on $\mathbb{H}^2$ and compares the simulation results with the predictions of free-area theory and the virial expansion developed in earlier sections. In Sec. VI, we consider a hard-disk system constrained to reside on a negatively curved triply periodic surface such as those found in the cell or among diblock copolymers. We summarize our results in Sec. VII.

II. TILINGS AND DISK PACKINGS

We begin our study of the hyperbolic plane, and curved space in general, with a discussion of the classical problem of tiling a plane with regular patterns. Regarding the hyperbolic plane in particular, we outline the features of our chosen model, the Poincaré disk model, in the Appendix.

The problem of understanding the possible coverings of a flat plane with uniform tiles is deceptive in its simplicity, and the different shapes the tiles may take or patterns they may be adorned with have long been pondered. In the last 100 years or so it has become known that there are precisely 17 ways of covering a flat plane with these tiles, corresponding to the 17 so-called wallpaper symmetry groups [21]. The Alhambra, in Granada, Spain, for example, is famous for its artistic representations of all 17 possible tilings [22].

A. Tesselating two-dimensional spaces

If we restrict ourselves to covering the plane with only *unadorned* polygons, as one would consider in a Voronoi tesselation, however, we reduce the number of possible tilings considerably. Furthermore, since crystalline order is a necessary starting point, we also restrict our tiles to regular polygons. We then begin with a regular lattice of Voronoi cells and consider a tesselation of F p-gons, with q meeting at each vertex. We allow for the possibility that F may be either finite, corresponding to a compact object, or infinite, corresponding to the Euclidean or hyperbolic plane. Each p-gon has p vertices, but each vertex is shared by q equivalent p-gons. Thus the total number of vertices is then pF/q. Likewise, each edge is shared by two p-gons and so the total number of edges is $pF/2$. The total number of faces is F. The Euler character χ of the surface constrains these three numbers and it follows that

$$\chi = V - E + F = pF\left[\frac{1}{q} + \frac{1}{p} - \frac{1}{2}\right]. \qquad (8)$$

For the flat periodic plane (the torus) $\chi = 0$, so $p^{-1} + q^{-1} = 1/2$ and we find the three tilings represented by the Schläfli symbol $\{p,q\} = \{6,3\}$, $\{3,6\}$, and $\{4,4\}$, hexagons, triangles,
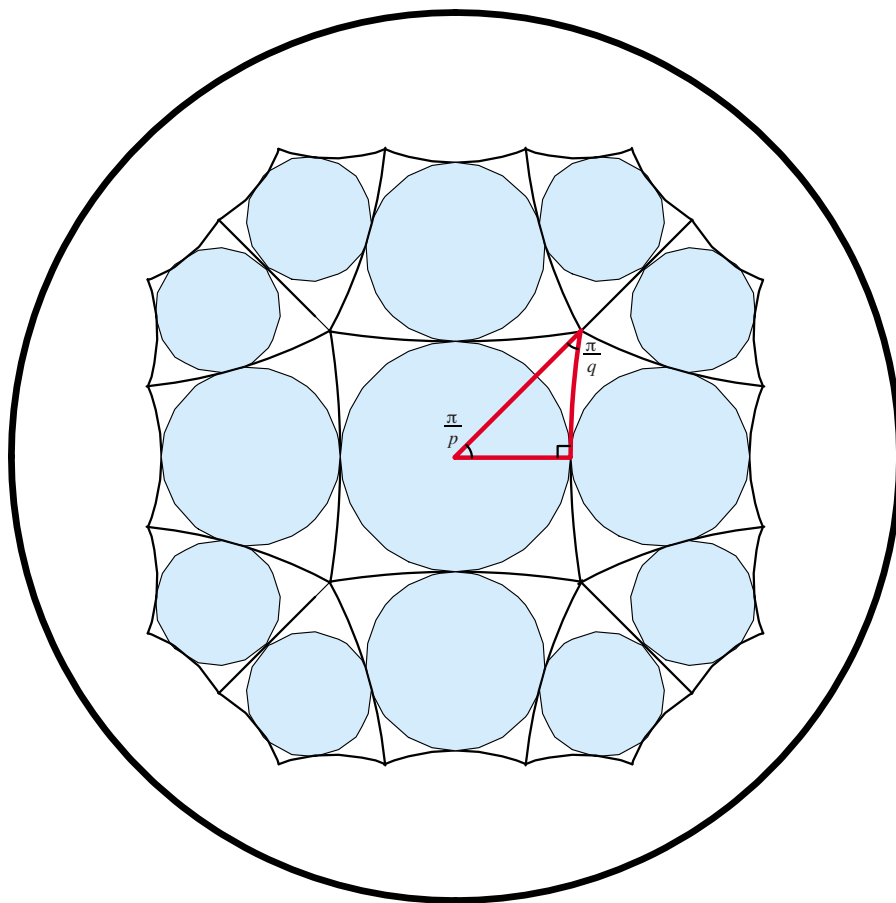
FIG. 1. (Color online) A {4,5} tiling of a segment of the hyperbolic plane, represented in the Poincaré disk model, with inscribed disks. Despite the apparent differences, the disks are of uniform size and the tiles are all equivalent regular polygons. The red triangle displayed in the central cell is used to calculate ϕ.

and squares, respectively. The spherical topology is the only one with positive χ (=2), admitting the five Platonic solids $\{p,q\}=\{3,3\}$, $\{3,4\}$, $\{3,5\}$, $\{4,3\}$, and $\{5,3\}$: the tetrahedron, octahedron, icosahedron, cube, and dodecahedron, respectively. Finally, we turn to the pertinent geometries with negative Gaussian curvature, for which $\chi<0$ and $p^{-1}+q^{-1}<\frac{1}{2}$. In this case there are an infinite number of integral pairs $\{p,q\}$ and an infinite number of corresponding regular tesselations—for example, $\{4,5\}$ (Fig. 1). As in the tilings of the Euclidean plane, the $\{p,q\}$ tesselation is dual to the $\{q,p\}$ tessellation—the vertices of one become the face centers of the other. However, this does not mean that their packing properties are the same, just as the packing of disks on the centers of hexagons is much more efficient than packing them on the centers of triangles.

B. Building ordered disk packings

We use these tesselations to produce close-packed, crystalline-ordered arrangements of disks simply by inscribing circles into their constituent regular polygons. Note that this procedure results in the familiar hexagonal crystal when applied to the $\{6,3\}$ tesselation of flat space. In general, those packings produced from tesselations with $q=3$ will result in the most efficient, densest crystals, as $q=3$ corresponds to local triangular close packing around each vertex. The rigidity or stability of a packing is another feature of interest that may be inferred from the parameters of the packing. In particular, a configuration is considered to be *isostatic* [23] if

there are exactly enough contacts to balance the degrees of freedom in the system; that is, the system is at the threshold of stability. In an N-particle packing in d dimensions there are dN degrees of freedom and $\frac{1}{2}\bar{z}N$ contacts, where $\bar{z}$ is the average number of contacts per particle, so the condition for isostaticity is $\bar{z}=2d$. Therefore, for an arrangement to be stable we must have $p \geq 4$, with isostaticity at $p=4$. Note that isostaticity is a topological issue and is independent of the curvature of the surface. We note that because we are forced to use periodic surfaces to achieve a thermodynamic limit, the usual isostatic counting should be revisited in a global context and is the subject of continuing work.

A central property of these crystalline arrangements is their area fraction, as it indicates which packings are favored entropically near a given curvature. We find the area fraction of the $\{p,q\}$ packing through application of the Gauss-Bonnet theorem and hyperbolic trigonometry. When there are F identical polygons, the area of each can be found from the Euler character χ:

$$A_{p\text{-gon}} = \frac{\int dA}{F} = \frac{-\int K\, dA}{KF} = \frac{-2\pi\chi}{KF}. \qquad (9)$$

The radius of the in circle can be determined via the dual law of hyperbolic cosines [24], $\cosh(\sqrt{-K}r) = \cos(\pi/q)/\sin(\pi/p)$ (Fig. 1). The area fraction for the $\{p,q\}$-packing is thus

$$\phi_{\{p,q\}} = \frac{\frac{\cos(\pi/q)}{\sin(\pi/p)} - 1}{p\left(\frac{1}{2} - \frac{1}{q} - \frac{1}{p}\right)}. \tag{10}$$

We recover the appropriate packing fractions $\pi/3\sqrt{3}$, $\pi/4$, and $\pi/2\sqrt{3}$ for the Euclidean tilings $\{3,6\}$, $\{4,4\}$, and $\{6,3\}$, respectively. As $p \to \infty$ for close-packed $q=3$ tesselations, we find $\phi = 0.9549$, the known packing fraction of the best packing on the hyperbolic plane at any curvature [25]. The lowest packing fraction for an isostatic configuration ($p=4$) is $\phi_{\{4,\infty\}} = \sqrt{2} - 1 \approx 0.4142$, an area fraction far below that for the Euclidean square lattice, $\pi/4 \approx 0.785$. Also of note is the broadening of the range in ϕ that supports stable configurations, from ~ 0.13 in two-dimensional flat space to ~ 0.54 on $\mathbb{H}^2$.

III. FREE-AREA THEORY IN A UNIFORMLY CURVED BACKGROUND

Now that we have explicitly constructed close-packed disk configurations with crystalline order, we have a foothold on the analysis of a thermodynamic fluid of hard disks residing in a curved space. Kirkwood's free-volume theory [26] is highly successful at modeling the equation of state for traditional crystals, particularly near their close-packed, maximum density [27,28]. Somewhat surprisingly, free-volume theory remains successful at lower densities as well; as such, it provides a useful tool for probing a thermodynamic system, so long as the high-density, high-correlation behavior is known.

Before we apply free-volume theory to our curved system, recall that in flat space the shapes of the Voronoi cell and the free-volume cell are equivalent and, thus, only the scaling of the free-volume cell is important to the equation of state. This scaling is dependent only on dimension, and we find a simple expression for the equation of state [29]:

$$P_{FV} = -\frac{k_B T \phi^2}{v}\frac{d}{d\phi}\ln\{[(\phi_{max}/\phi)^{1/D} - 1]^D\}. \tag{11}$$

Unfortunately, an analogously simple treatment does not exist in the presence of a curved background. The curvature itself, inducing geodesics to converge (positive K) or diverge (negative K), breaks the scale invariance enjoyed by all shapes on $\mathbb{R}^n$, and hence destroys the similarity relation between the Voronoi cell and the free-volume cell.

A. Idealized circular free-area cell

As an alternative, we consider an idealized case in which the symmetry of the system itself ensures that similarity is restored. If, for example, we have a two-dimensional fluid composed of hard disks and we posit that the free-area cells are *also* circular, then we may proceed in direct analogy with Euclidean space. The free area is simply the area of a disk whose radius is the difference between that of the Voronoi cell and that of the particle, $A_{free} = -\frac{2\pi}{K}\{\cosh[\sqrt{-K}(r_{cell} - r_{disk})] - 1\}$. The angle addition formula for the hyperbolic trigonometric functions, along with the recognition that ϕ is the ratio of A_{disk} to A_{cell}, allows us to write
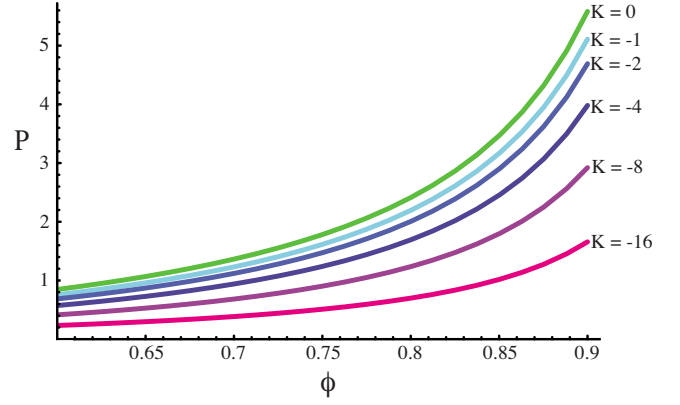

FIG. 2. (Color online) The equation of state for free-volume theory in a disk-shaped Voronoi cell. The top, green curve is for flat space; subsequent curves are at increasing negative curvature. Warmer colors correspond to more curvature.

$$
\begin{aligned}
A_{free} = -\frac{2\pi}{K}\Bigg(&\frac{1}{\phi}(\cosh[\sqrt{-K}r] + \phi - 1)\cosh[\sqrt{-K}r] \\
&- \frac{1}{\phi}\sqrt{(\cosh[\sqrt{-K}r] - 1)(\cosh[\sqrt{-K}r] + 2\phi - 1)} \\
&\times \sinh[\sqrt{-K}r] - 1\Bigg).
\end{aligned} \tag{12}
$$

Notice that we recover the expected Euclidean result of $\pi r^2(1/\phi^{1/2} - 1)^2$ as we take $K \to 0$. Furthermore, since, in general,

$$P_{FA} = -\frac{k_B T \phi^2}{a}\frac{d}{d\phi}\ln(A_{free}), \tag{13}$$

we find that the pressure diverges as a simple pole as $\phi \to 1$. Notice, however, that as we increase the background curvature the divergence at $\phi=1$ becomes sharper, apparent at ever higher area fractions (see Fig. 2), indicating a lower system pressure than is seen in a corresponding flat-space system. This suggests that both particle collisions and more complicated, many-particle arrangements become increasingly scarce as the curvature rises.

B. Free area for tesselation-derived packings

What if we were to consider a nonidealized system based on the ordered packings discussed in Sec. II, with polygonal Voronoi cells? In order to cope with the shape distortions caused by the background curvature, we divide the Voronoi cell into individual regions, as depicted in Fig. 3, and calculate the areas of the strip and corner regions separately, implying $A_{free} = A_{Voronoi} - pA_{strip} - pA_{corner}$. Without loss of generality we may set $K = -1$ since it is only the product of the disk size and $\sqrt{-K}$ which enter into all our expressions. We can access different curvatures by changing the size of the disks. Direct application of the Gauss-Bonnet theorem yields $A_{Voronoi} = p\pi - 2\pi - 2p\pi/q$. In order to find the other two areas, we first need s, δ, and α. By considering the triangle formed by the Voronoi cell's center and two adjacent vertices

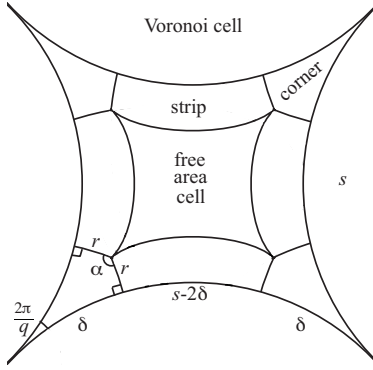FIG. 3. Partition of a $\{p,q\}$ Voronoi cell into its free-area cell, p corners, and p strips. Note that the boundaries of the free-area cell are not geodesics.

we may apply the dual hyperbolic law of cosines (Appendix A 2) to find

$$s = \cosh^{-1}\left[\frac{\cos\left(\frac{2\pi}{p}\right) + \cos^2\left(\frac{\pi}{q}\right)}{\sin^2\left(\frac{\pi}{q}\right)} \right]. \quad (14)$$

To proceed we construct two identical triangles out of the corner 4-gon by bisecting α. The dual hyperbolic law of cosines applied to the side of length r then gives

$$\alpha = \sin^{-1}\left[\frac{\cos\left(\frac{\pi}{q}\right)}{\cosh(r)} \right]. \quad (15)$$

Armed with an expression for α, applying the hyperbolic law of cosines again, this time to the side of length δ yields

$$\delta = \cosh^{-1}\left[\frac{\sqrt{\cosh^2(r) - \cos^2\left(\frac{\pi}{q}\right)}}{\cosh(r)\sin\left(\frac{\pi}{q}\right)} \right]. \quad (16)$$

Now A_{strip} must be the area within a given distance (namely, r) of a geodesic of length $s - 2\delta$. This area is $(s - 2\delta)\sinh(r)$ (see Appendix, Sec. 3). Given α, direct application of the Gauss-Bonnet theorem gives A_{corner}. Finally, we may eliminate r through ϕ:

$$r = \cosh^{-1}\left[\frac{p\phi}{2} - \frac{p\phi}{q} - \phi + 1 \right]. \quad (17)$$

Collecting terms and setting $\tau = 1 - \frac{\chi\phi}{F}$ and $\tau_{pq} = 1 - \frac{\chi\phi_{\{p,q\}}}{F} = \cos(\pi/q)/\sin(\pi/p)$, we have

$$A_{\text{free}} = 2p\,\sin^{-1}\left[\cos\left(\frac{\pi}{q}\right)\frac{\tau_{pq}}{\tau} \right] - 2\pi$$
$$- 2p\sqrt{\tau^2 - 1}\,\cosh^{-1}\left[\frac{\cos\left(\frac{\pi}{p}\right)}{\sin\left(\frac{\pi}{q}\right)} \right]$$
$$+ 2p\sqrt{\tau^2 - 1}\,\cosh^{-1}\left[\frac{\sqrt{1 - \frac{\tau_{pq}^2}{\tau^2}\sin^2\left(\frac{\pi}{p}\right)}}{\sin\left(\frac{\pi}{q}\right)} \right]. \quad (18)$$

As a check of the complex algebra, note that as $\phi \rightarrow \phi_{\{p,q\}}$, $\tau \rightarrow \tau_{\{p,q\}}$ and $A_{\text{free}} \rightarrow 0$. Likewise, as $\phi \rightarrow 0$, $\tau \rightarrow 1$ and
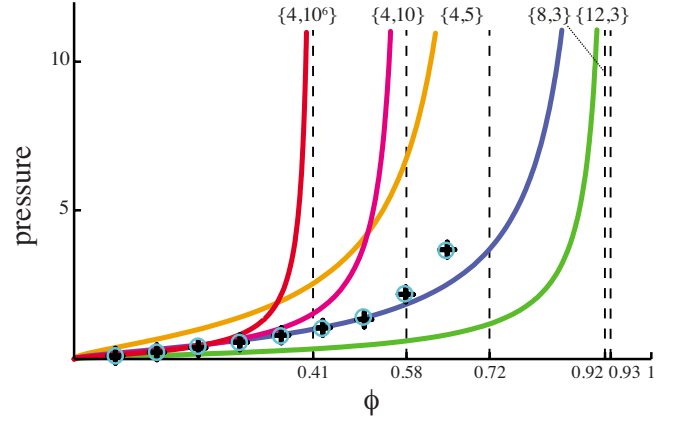

FIG. 4. (Color online) The equation of state for free-area theory for a range of $\{p,q\}$ tessellations. Note that the pressure curves diverge at their maximum packing fractions and that they differ even at low volume fractions ϕ. The vertical dashed lines indicate the maximum volume fraction for the tessellations listed on top of the figure. The (light blue) circles are the result of our molecular dynamics simulation for curvature just below the $\{4,5\}$ tessellation, and the (black) crosses are for curvature just above $\{4,5\}$. The number of particles ranges from 1 to 9.

$$A_{\text{free}} \sim 2p\left(\frac{\pi}{2} - \frac{\pi}{q}\right) - 2\pi = -2\pi p\chi/F = A_{\text{Voronoi}}. \quad (19)$$

Unfortunately, away from these limits, this expression is difficult to analyze, so we instead turn to Fig. 4, where we plot the free-area equations of state for various $\{p,q\}$ tessellations using Eq. (13)

It is important to note that by fixing the lattice of Voronoi cells we are forced to change the packing fraction by changing the disk size. Since this is equivalent to changing the curvature, our free-area equation of state samples different curvatures at different area fractions ϕ. If we wish to apply this free-area theory to a system of constant disk size to sample a constant background curvature, we must allow the Voronoi scaffolding to vary continuously. We will have to consider tessellations with nonintegral p and q; we can interpret these as the average number of edges per polygon and the average number of polygons meeting at a corner in a tessellation made of a variety of polygons. This is in the spirit of recent work on average polyhedra [30] which have a nonintegral number of sides. To maintain the curvature as we vary ϕ, we might fix p or q adjusting the other so as to keep r fixed in (17). We might also hold p/q fixed and set p to maintain r. The freedom afforded by choosing one's path through p-q parameter space is troubling—there seems to be no *a priori* way for one such path to be chosen over another, and we are forced to wonder how the behavior of the system is determined. Fortunately, upon considering specific cases, the problem disappears, as the fixed-radius equation of state displays insensitivity to this choice (Fig. 5), particularly near high density, indicating that perhaps there is a hidden, unaccounted-for symmetry among the isotropic Voronoi scaffoldings. We only find path dependence for very large values of q and p—there are no real solutions for p when q is large until we get close to the divergence in the pressure,
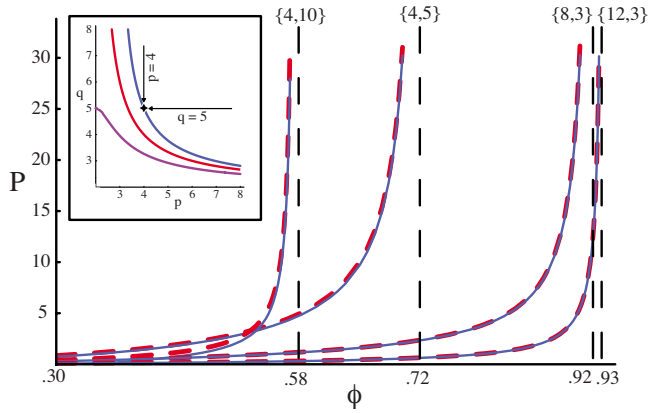
FIG. 5. (Color online) Fixed-radius equations of state with a sampling of $\{p,q\}$ tesselations as their divergent end points. The solid blue curves are for fixed q, and the dashed red curves are for fixed p. Note that in almost all cases these curves lie directly on top of one another. The inset displays the parameter space for the packing parameters $\{p,q\}$. The thick red line (center) is the allowed combinations of p and q in flat space and the purple line beneath it a representative example from positively curved space. The blue curve above gives the allowed combinations for $\mathbb{H}^2$ with $r=0.53$. The $\{4,5\}$ packing is identified and the paths through parameter space taken by holding either p or q constant are shown.

though there the divergence is insensitive to the details again.

IV. VIRIAL EXPANSION TO THE EQUATION OF STATE

Thus far, by considering the extension of free-area theory to uniformly curved spaces we have been working in the high-density regime, since free-area theory is the first-order term in an expansion about a close-packed state. On the other hand, we are also interested in the system's behavior when it is much more sparse. The expansion about low densities provided by the virial expansion to the equation of state [31] is suitable in this regime, although the connection between the two regimes is still not well understood in hard-core systems [32]. In contrast to free-area theory, the virial expansion's validity is dependent on the ability to take a proper thermodynamic limit. We therefore resort to the periodic spaces discussed earlier as our solution to this difficulty.

We note, however, that for hard disks, the virial expansion only includes local interactions among clusters of particles. Fortunately, as long as the local clusters under consideration do not wrap around the periodic space and self-interact, we may ignore the topology introduced by our usage of a periodic space and perform the cluster integrals on $\mathbb{H}^2$. In order to ensure that this is possible, we may continue to increase the area available to a cluster by constructing ever higher-genus manifolds [18]. Furthermore, for manifolds with constant $K<0$, $Z_1 = A[\lambda_T^{-2} + K/(12\pi)] > 0$ when $\lambda_T^2 < -12\pi/K$. In other words, there is an ample range of K in which to work where the curvature of the manifold does not probe the quantum regime of the gas. Hence, the pressure $P = N k_B T d \ln Z / dA$ is independent of the proportionality constant.

Just as in the last section, we take advantage of the lack of scale invariance on $\mathbb{H}^2$ to probe higher curvatures not by changing the ambient space, but simply by considering larger and larger disks. We consider disks of radius $r\sqrt{-K}$. Without loss of generality we again fix the Gaussian curvature to be $K = -1$. Dimensional analysis may always be employed to insert the appropriate factors of K. As in flat space, the second Virial coefficient $B_2(r)$ is given by half the excluded area of a single disk. To see that this is true, recall that, in general, $B_2(r)$ is given by

$$B_2(r) = -\frac{1}{2V}\int\int f_{12}\, d^2\mathbf{r}_1 d^2\mathbf{r}_2, \qquad (20)$$

where f_{12} is the Mayer function and the position integrals for both particle 1 and particle 2 are carried out over all space. Since $\mathbb{H}^2$ is isotropic and homogeneous, we may integrate out one particle position and pick up a factor of V. The remaining integration is over the separation between the disks. Meanwhile, the Mayer function for hard disks in the Poincaré disk model is

$$f_{ij} = \begin{cases} -1, & \text{if } \cosh^{-1}\left[1 + 2\dfrac{x^2}{1-x^2}\right] < 2r, \\ 0, & \text{otherwise,} \end{cases} \qquad (21)$$

as long as one of the disks is at the origin. The expression when both disks are away from the origin is more complicated, but follows from (A4). However, for the purposes of $B_2(r)$, we can always put one disk at the origin. Turning the step function into a condition on the bounds of the integral over x, inserting the metric from the Poincaré disk model (see the Appendix), and integrating out the angular part results in the following expression for $B_2(r)$:

$$B_2(r) = 4\pi \int_0^{\sqrt{[\cosh(2r)-1]/[\cosh(2r)+1]}} dx \frac{x}{(1-x^2)^2}. \qquad (22)$$

Notice that this is precisely the circularly-symmetric, hyperbolic, area integral needed to calculate the area of a disk of radius $2r$ on $\mathbb{H}^2$, and indeed, carrying out the integration leaves us with $B_2(r) = \pi[\cosh(2r) - 1]$, which is half the excluded area of a single disk, the classic result from the virial expansion. Further note that as $r \to 0$ and $B_2 \sim 2\pi r^2$, we recover the $d=2$ Euclidean result $B_2(r) = (4\pi)^{d/2}r^d/\Gamma(\frac{d}{2})d$.

We may carry out a similar procedure to calculate higher coefficients. The general relation for $B_3(r)$ is

$$B_3(r) = -\frac{1}{3V}\int\int\int f_{12}f_{23}f_{13}\, d^2\mathbf{r}_1 d^2\mathbf{r}_2 d^2\mathbf{r}_3, \qquad (23)$$

and we may still integrate out one of the disk positions. Unfortunately, regardless of which disk we set to the origin, we are now guaranteed to retain a Mayer function where neither of the disks whose interaction it represents have had their position integrated out. This means that we are forced to contend with an integral of the form
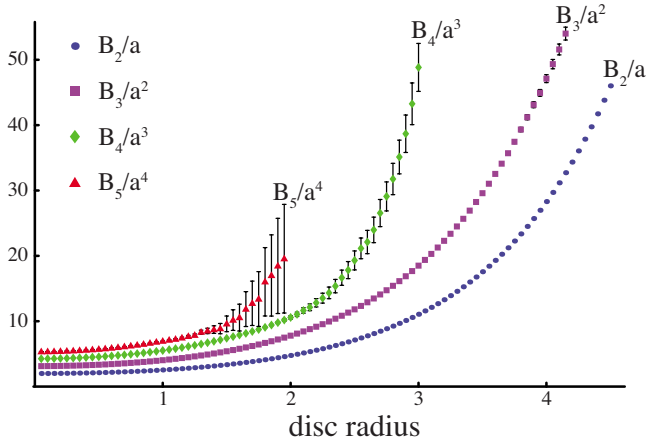
FIG. 6. (Color online) The second, third, fourth, and fifth virial coefficients in appropriate units of the area a as functions of disk size (for $K=-1$) r. The zero-size limit matches the flat-space coefficients. B_2 is exact, and error bars are indicated for all other coefficients. These errors are given by the estimated variance of the mean of the integrand and, for large η, scale like $\eta^{-1/2}$, where η is the number of Monte Carlo trials.

$$B_3 = -\frac{32\pi}{3}\int_0^b dx_1 \int_0^{2\pi} d\theta_{12} \int_0^b dx_2$$

$$\times \left[\frac{x_1}{(1-x_1^2)^2}\frac{x_2}{(1-x_2^2)^2}f_{12}(x_1,x_2,\theta_{12})\right], \qquad (24)$$

where b is the same upper limit as in Eq. (22) and the remaining Mayer function is written in terms of the coordinate distance of the two particles from the origin and the angle between them:

$$f_{12} = \begin{cases} -1, & \left[1+2\dfrac{x_1^2+x_2^2-2x_1x_2\cos\theta_{12}}{(1-x_i^2)(1-x_j^2)}\right] < \cosh(2r), \\ 0, & \text{otherwise.} \end{cases}$$

In order to evaluate this integral, and similar ones that appear in the higher virial coefficients, we turn to numerical Monte Carlo techniques. In addition, as is traditional in flat space [33], the higher coefficients B_n are reported in units of $a^{n-1} = [A(r)]^{n-1}$ to suppress size effects and generate the expansion for $P/(\rho k_B T)$ in the area fraction, $\phi = \rho a$:

$$P = \rho k_B T\left[1 + \sum_{n=2}^{\infty}\frac{B_n}{a^{n-1}}\phi^{n-1}\right]. \qquad (25)$$

We have checked the $r\to 0$ limit in all cases and find agreement with the known Euclidean results. We plot B_2/a, B_3/a^2, B_4/a^3, and B_5/a^4 in Fig. 6 as functions of r. To the accuracy of our calculation, over the range of curvatures considered, B_3, B_4, and B_5 all remain positive. There is evidence that, in higher-dimensional flat spaces, the virial coefficients alternate in sign [33,34] and, hence, the leading singularity that controls the radius of convergence of the expansion sits on the negative real axis [34]. The location of this singularity, and thus the range of applicability of the virial expansion itself, remains an open question in $D=2,3$. Though not con-

clusive, our results here suggest that the introduction of curvature does not *qualitatively* change the behavior or location of the leading singularity in the expansion.

Padé approximants and the radius of convergence

The quantitative behavior of the virial expansion's divergence as a function of the curvature, on the other hand, may be extracted from the coefficients that we have calculated, at least as a trend. To this end we have approximated the virial series by ratios of rational functions, the Padé approximants. We find the coefficients of two polynomials $a_m(\rho)$ and $b_n(\rho)$, of degree m and n, respectively so that

$$P = k_B T\left(\rho + \sum_{j=2}^{5} B_j\rho^j\right) = \frac{a_m(\rho)}{b_n(\rho)}. \qquad (26)$$

The left-hand side of (26) is fifth order, and so there are, in principle, six arbitrary coefficients. As a result, we can only study the Padé approximants with $n+m\leq 5$. We look for the divergences in these approximants to estimate where the virial expansion diverges—an estimate of the ultimate packing fraction if the virial expansion converges for all physical ρ. Our results are shown in Table I. Notice that, for the Padé approximants with only one pole (that is, of the form $[m,1]$), there is a clear trend toward divergence at lower ρ as the curvature increases. Furthermore, it is worth pointing out that among these approximants, the one that is the most sensitive, the $[3,1]$, also agrees best with the putative location of the divergence of the fluid branch in flat space, at around $\phi\approx 0.84$, as r goes to zero.

V. MOLECULAR DYNAMICS ON THE HYPERBOLIC PLANE

Equipped with these theoretical models for the system, particularly in the high- and low-density regimes, we turn to the full-blown hard-disk fluid. We choose to employ constant-energy molecular dynamics. There is some subtlety in handling the boundaries of the simulation region: recall that the area near the edge of a shape scales at the same rate of the area of the entire shape on $\mathbb{H}^2$. This means that the effect of hard boundaries cannot be minimized by considering larger regions. We thus choose to implement periodic boundary conditions, although there is subtlety here as well. Because of the scale dependence of the hyperbolic plane and the topological requirements of the periodic region, we are restricted to choosing for our simulation region the unit cell from tesselations of the form $\{4n,4n\}$ for n an integer greater than 1. This ensures both that the sides of the polygon may be identified in a systematic way to produce an orientable manifold and that, upon identification, there is not a conical singularity at the vertices of the polygon (so there are exactly 2π radians of space around each vertex).

We have chosen not to use the event-driven molecular dynamics traditional for pure hard-core systems, as projecting the next collision is extremely time intensive when the equations of motion are nontrivial as they are in a curved background. Instead, we evolve the system by time steps, checking for and resolving collisions as they occur. The col-

TABLE I. The leading divergences on the real axis in ϕ at various values of curvature and for different Padé approximants about the origin. Empty entries indicate that all poles for that approximant are located off the real axis. Though ϕ has physical meaning only in the range [0,1], a divergence at an unphysical value of ϕ may still be meaningful, if, for example, it encodes information about the radius of convergence of the virial expansion.

[m,n]			$\phi^*_{R_{[m,n]}}$		
	$r \to 0$	$r=0.5$	$r=1.0$	$r=1.5$	$r=2.0$
[0,4]	−3.444	4.335	—	—	—
[1,3]	2.720	−0.592	−8.695	0.797	0.650
[2,2]	—	—	—	0.798	0.695
[3,1]	0.798	0.797	0.794	0.770	0.676
[0,3]	−5.145	−45.09	0.634	0.409	0.258
[1,2]	—	—	0.888	0.821	0.829
[2,1]	0.735	0.733	0.732	0.730	0.737
[0,2]	0.737	—	—	—	—
[1,1]	0.639	0.637	0.630	0.620	0.611
[0,1]	0.500	0.470	0.393	0.298	0.210

lisions themselves require some care as well. The center-of-mass frame in a curved space is not necessarily inertial, even if the constituent particles are all following geodesics—imagine a sphere with a particle moving around the equator and another sitting motionless on the North pole—the center of mass follows the 45th parallel, which is not a geodesic. As a result, the textbook techniques and formulas for handling collisions do not apply. As an alternative, we decompose each particle's velocity into the component along the geodesic connecting the particle centers at contact and the component normal to this geodesic. Note that these directions will differ between the two colliding particles due to the complexity of parallel transport on $\mathbb{H}^2$. In order to resolve the collision, the components of velocity along the contact geodesic are swapped and the other components remain unchanged. Hence, collisions are guaranteed to conserve energy and momentum.

To calculate the equation of state from our simulation, we employ the collisional virial expression for the pressure [35]:

$$P = \frac{1}{A}\left(NkT + \frac{1}{2\tau}\sum_{\text{collisions}} d|\Delta\mathbf{p}| \right), \qquad (27)$$

where A is the area of the simulation region, τ is the time over which collisions are summed, d is the disk diameter, and $\Delta\mathbf{p}$ is the momentum transferred from one particle to another during the collision. Finally, we remind the reader that to fill out the equation of state with different values of ϕ it is necessary to change the number of particles simulated, N, as varying the disk size means changing the effective curvature.

A. Systems near the {4,5} tesselation

We probe curvatures $r=0.53$ and $r=0.531$, just *below* and just *above* that for the {4,5} tesselation, $r_{\{4,5\}}=0.5306$. At these curvatures no regular tesselation is allowed, and so noncrystalline arrangements give the best packings [36].

Thus we expect to probe disordered configurations even at high densities with this choice. At lower densities, however, the gas will still be sensitive to the {4,5} tesselation and should follow the free-area theory prediction for the {4,5}, failing only at higher area fractions. In particular, we expect that the system's pressure should be lower than that of the isostatic crystal for higher ϕ and diverge at larger ϕ. This is precisely the case in three dimensions where the liquid branch is lower in pressure than the free-volume crystalline pressure at low volume fractions. Indeed, these expectations are borne out by the data (Fig. 4).

Due to the fact that we have used a low number of particles in this regime, probing the spatial distributions of the particles is impractical. However, further future work will allow for the investigation of the pair correlation function. In particular, the effects of curvature on the split second peak [37] present in the flat two-dimensional system are of potential interest.

B. Approaching flat space from below: Monodispersity and disorder

In addition to investigating the system at packing fractions close to an isostatic crystal, we might use a limiting procedure to take the background curvature to zero. By looking at the fluid in low curvatures, we expect to gain some insight into hard disks in flat space. In particular, looking at curvatures near, but not at, zero guarantees that we will avoid crystallization, as no crystal packings are available in this curvature range. It is therefore possible that we may generate disordered, monodisperse configurations that can be carried into flat space with low probability of introducing overlaps or defects. Unfortunately, due to the broken scale invariance on $\mathbb{H}^2$, a very large number of simulated disks is needed to probe this regime, making the simulation numerically untenable, even for the high-speed digital computers of our day.

VI. PACKING AND DYNAMICS ON NEGATIVELY CURVED SURFACES IN FLAT SPACE

Thus far, we have used uniform, negative, background curvature as a way of constructing an idealized model in two dimensions for a more complicated three-dimensional system. Owing to the fact that no surface with uniform, negative curvature can be isometrically embedded into $\mathbb{R}^3$; however, this model is unphysical. But are there examples of real systems that are qualitatively close, for which the predictions of our model may be brought to bear directly? Triply periodic minimal surfaces, with their numerous physically advantageous properties [38], their prevalence in smectic liquid crystals, diblock copolymers, and even biology, and their status as infinite surfaces with negative curvature everywhere, are ideal for analysis with insights from the abstract hyperbolic plane. Further, recent work on the creation of toroidal bubbles whose interface is packed with colloids [39] indicates that the experimental creation of an ideal two-dimensional packing environment with negative curvature is within reach.

In lieu of a direct molecular dynamics simulation, as done on $\mathbb{H}^2$, we instead take advantage of the numerical insight provided by examining allowed and disallowed configurations [40]. In fact, we may obtain the pressure simply by randomly sampling particle configurations at a given density and determining the fraction that violate the no-overlap condition of the hard-core potential. If the probability of not finding an allowed state at density $\phi + \delta\phi$, given an allowed state at ϕ, is $J(\phi)\delta\phi$, then the following relation holds [40]:

$$J(\phi) = \frac{V}{\phi k_B T}[P(\phi) - \phi T/v],\tag{28}$$

where v is the area per disk, V the system volume, and T the temperature.

Hence, this approach gives us the equation of state. We have traded the numerical difficulty of dynamics of particles constrained to a complicated surface with the algorithmic difficulty of finding a way to generate points at random on an arbitrary surface such that the distribution is flat with respect to the local area.

A. Randomly distributed points on an arbitrary surface

We have employed a novel geometric algorithm for selecting points from an arbitrary surface randomly with even distribution in the area (as long as it is either compact or periodic in each dimension). Knowledge of the metric is not required, only a set of parametric equations that define the surface in $\mathbb{R}^3$. The algorithm works as follows. First, one repeat unit of the surface in whichever directions are periodic, and the entire surface in those directions for which the surface is compact, are considered to be inside a box. At this stage it is necessary to determine the maximum number of times a straight line can intersect this region of the surface—say, I_{max}—this can typically be done analytically, though for a very complicated surface numerics may prove more efficient. The box is inscribed in a sphere, and a great disk is chosen at random in this sphere. Next, a point is chosen at

random on this great disk and the line normal to the great disk is drawn through it. Finally, we determine the intersections of this line with the surface and randomly select one from among them. To ensure that the resulting distribution of points is flat with respect to area, it is necessary to begin the entire process anew if the line fails to intersect the surface. Furthermore, if the number of intersections is less than I_{max}, then an appropriate number of placeholders must be included when randomly selecting an intersection point. If a placeholder is selected, then the entire procedure is repeated. If an actual intersection point is selected, then that point is the output of the pass through the algorithm.

To prove that this procedure yields a distribution flat with respect to the local area, consider the shadow cast by the surface when light from an arbitrary direction strikes it. Now consider the contribution to this shadow from a differential piece of the surface. The area of the part of the shadow cast by this small patch will be directly proportional to the likelihood of selecting a point from it when the great disk chosen at random is normal to the direction of light, $P(dA) \propto dA|\cos(\theta)|$, where θ is the angle between the patch's surface normal and the light rays. But the algorithm specifies that we constantly choose, at random, new directions for the illumination. This has the effect of averaging over all incoming directions for the light, θ and ϕ, and removes the dependence on the angle, leaving the likelihood of selecting a point from any given patch of surface directly proportional to only the area of the patch. To see this, imagine a particular great disk parametrized by x and y. The surface has a shadow on the disk of area A_{shadow}. Ignoring, for the moment, multiple leaves, we may parametrize the surface in terms of x and y so that the probability of selecting any particular area element dA given that we pick a point in the shadow is

$$P(\ dA|\text{the point is in the shadow})$$

$$= \frac{dA|\cos\theta|}{\displaystyle\int_{\text{shadow}} dx\ dy} = \frac{dA|\cos\theta|}{A_{\text{shadow}}}.\tag{29}$$

Note that A_{shadow} depends on the choice of great disk, so this conditional probability does not pick area elements with equal weight. However, the probability of picking a point in the shadow is just $A_{shadow}/A_{great\ disk}$. Putting this together we see that $P(dA) = dA|\cos\theta|/A_{great\ disk}$. Integrating over all directions yields a flat distribution. Finally, if there are multiple leaves of the surface behind one another with respect to a given direction of incident light, randomly selecting from among them and multiple placeholders, as described above, guarantees that no patch is penalized or rewarded for its position relative to distant pieces of the surface.

B. Plumber's nightmare

The simplest triply periodic minimal surface, and the one that we consider first, is the Schwarz P-surface, sometimes referred to as the "Plumber's nightmare" for the impression it gives of a never-ending maze of pipe work. This surface appears in the study of lipid bilayers [41] and liquid crystals [42]. In lieu of using the exact minimal surface, however, we
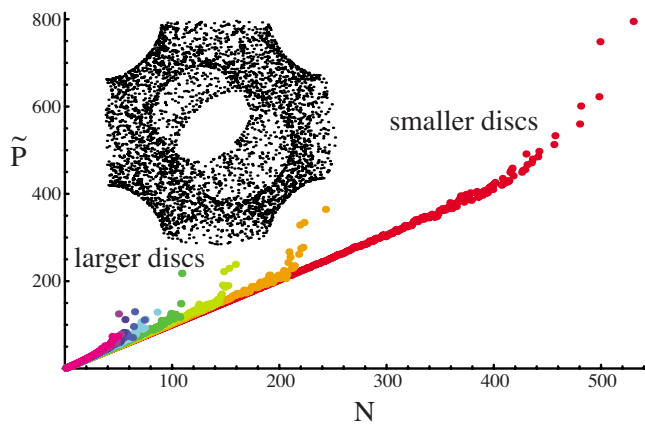
FIG. 7. (Color online) Dimensionless pressure as a function of *N* for disks of various sizes residing in the P-surface. The disk diameters range from 0.05 to 0.5 in units where the repeat length of the P-surface cell is 2π. Warmer colors correspond to smaller disks. A subset of the random points used is shown in the upper left corner.
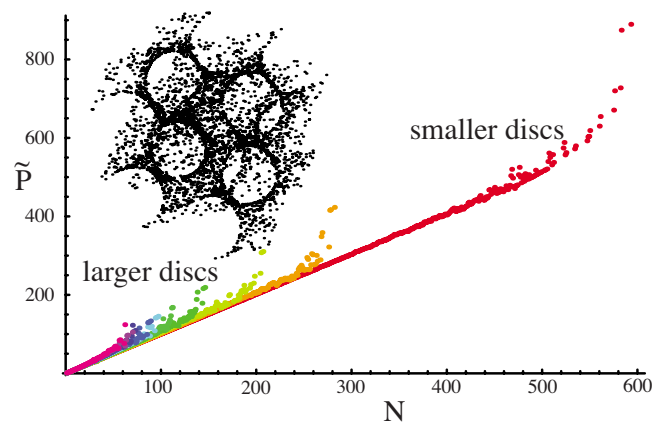

FIG. 8. (Color online) Dimensionless pressure as a function of *N* for disks of various sizes residing in the D-surface. The disk diameters range from 0.05 to 0.5 in units where the repeat length of the D-surface cell is 2π. Warmer colors correspond to smaller disks. A subset of the random points used is shown in the upper left corner.

use the trigonometric approximation, as its departure from minimality affects only the mean curvature, leaving the Gaussian curvature nearly unchanged. Our approximate Schwarz P-surface is defined by the following relation:

$$\cos(x) + \cos(y) + \cos(z) = 0. \tag{30}$$

Using this level-set approximation, we have used our algorithm to choose points as depicted in Fig. 7. We have checked our method by projecting these points along an axis and normalizing by the square root of the metric to see a uniform density. With this in hand we may now calculate the pressure using (28). We lay down *N* disks of a given size, one at a time, checking to see if they can be inserted without overlapping an existing disk. To accept or reject each disk, we check the distance between points along the surface and compare this to the disk diameter. The distances between disk centers along the surface are also estimated in a metric-free way by taking a small step along the line in $\mathbb{R}^3$ connecting the disk centers and finding the plane normal to this line at the location of the step. We then calculate the curve of intersection between the P-surface and the plane and find the point on this curve closest to the initial disk center. Tabulating this distance and then repeating the process with the location of the last step as the new initial point from which the next line in $\mathbb{R}^3$ is drawn completes the algorithm. Armed with an unambiguous state acceptance criterion, we have an avenue to calculate $J(\phi)$; however, our relevant control parameter is *N*. So we instead write

$$J(N) = \frac{V^2}{NT v}[P(N) - NT/V], \tag{31}$$

using $\phi = Nv/V$. Isolating a dimensionless pressure $\widetilde{P} = P(N)V/T$ allows us to wash out the hard-core system's insensitivity to the temperature and the overall scale invariance, leaving

$$\widetilde{P} = N\left(1 + \frac{v}{V}J(N)\right). \tag{32}$$

The dimensionless pressure calculated in this way is shown in Fig. 7 as a function of *N*. Notice that for all considered disk sizes, the pressure shows little deviation from an ideal gas state ($\widetilde{P}=N$) until near the divergence. This stands in excellent agreement with our predictions from free-area theory and meshes well with the observations made on hard-core systems in general by Kamien and Liu [40].

C. *Chaos Carolinensis* and the Schwarz D-surface

Another triply periodic minimal surface of interest is the Schwarz D-surface, or diamond surface. Lipid bilayer and liquid-crystalline systems also both exhibit this surface, though it appears in a more surprising context as well— biology. When the amoeba *Chaos Carolinensis* is deprived of food, its mitochondrial membrane transitions from an amorphous, unpatterned state to a D-surface, possibly to optimize its surface-to-volume ratio while cannibalizing its own membrane phospholipids for energy [20]. We again choose to use the trigonometric approximation; this time the defining relation is given by

$$\cos(x - y)\cos(z) + \sin(x + y)\sin(z) = 0. \tag{33}$$

As before, we generate points on this surface (see Fig. 8), and again we calculate the dimensionless pressure. For the full range of considered disk sizes there is little deviation from the ideal gas until near the divergence, similar to the P-surface, as shown in Fig. 8. Perhaps this is not surprising, as the fundamental cell of the P- and D-surfaces are made of the same local pieces patched together in different ways. Thus the purely local interactions are likely insensitive to this longer-range structure, leaving open the possibility that interactions on the scale of the fundamental cells might be sensitive to the changed topology.

VII. DISCUSSION

In summary, we have constructed a *one-parameter*, mono-disperse, two-dimensional system with hard-core repulsions by introducing negative background curvature to frustrate local close packing. This suppresses crystallization and brings the qualitative features of the system more in line with the frustrated three-dimensional flat case than with the unfrustrated two-dimensional system. We presented theoretical models of the system both at low and high density, with the Virial expansion and the enumeration of crystalline states along with a free-area theory based around these crystal end points. We simulated the system in a periodic piece of $\mathbb{H}^2$ and compared these results with our theory and with the well-known flat-space systems. In all cases, our results are qualitatively in line with the known behaviors in the flat systems, allowing us to conclude that this system will serve as a viable model for physical jamming problems that enjoys important advantages over the standard models in current use. In particular, our model has a vastly more manageable parameter space.

In addition to finding use as a model system for three-dimensional amorphous solids, our work may also be taken at face value as a direct investigation of two-dimensional systems in a curved background. We numerically analyzed such a system on a physically realizable curved surface with nonconstant negative curvature, laying the groundwork for future experimental work on triply periodic membranes and the like.

We emphasize, however, that these are only the first steps to a rich geometric framework that may serve as a promising alternative to poly- and bidisperse systems of hard or soft disks. Technological advances will allow a clean look at the system as the curvature is turned off and will make higher Virial coefficients attainable. More complicated particle shapes and interactions will allow a complete modeling of submicron scale packings constrained to curved membranes, as exist in lipids and diblock copolymers and abound in biology. In addition, it is worth considering $\mathbb{H}^2$ and the delicate balancing of factors that allow the thermodynamic limit to be trivially sensible in flat space and difficult to grasp anywhere else.

APPENDIX: CLASSICAL NON-EUCLIDEAN GEOMETRY

Here we review some properties of $\mathbb{H}^2$ and the well-known Poincaré model, and collect some useful classical identities and theorems on uniform curved geometry in general.

1. Poincaré disk model: Metric and distance

We employ the Poincaré disk model of $\mathbb{H}^2$, popularized by Escher [43], which conformally maps the hyperbolic plane with curvature K to the disk of radius $\sqrt{-K^{-1}}$ on $\mathbb{R}^2$. Using the coordinates (x,y), the metric is [18]

$$ds^2 = \frac{4(dx^2 + dy^2)}{[1 + K(x^2 + y^2)]^2}, \tag{A1}$$

where $x^2 + y^2 \leq -1/K$.

Geodesics in this model are arcs of circles which intersect the bounding circle normally; circles in the hyperbolic plane remain circles in this model, but their coordinate centers and radii vary: shifting outward from the apparent center and shrinking far from the origin, respectively. The radius of curvature sets a length scale and, consequently, any system on $\mathbb{H}^2$ is manifestly not scale invariant.

In addition, it is useful to understand that the Poincaré disk model arises via stereographic projection [44] in the same way as spherical mappings and, in particular, maps of the Earth. Define the following quadratic form: $Q(\mathbf{x}) = -x^2 - y^2 + z^2$. Then the collection of points with norm 1 with respect to this quadratic form define a two-sheeted hyperboloid. We carry out the stereographic projection by mapping points on the sheet at positive z onto the unit disk in the x-y plane via the collection of lines passing through the point $(0,0,-1)$. Given a point in the unit disk, (u_x, u_y) we now have the map to the hyperboloid:

$$(u_x, u_y) \rightarrow \left(\frac{2u_x}{1 - u^2}, \frac{2u_y}{1 - u^2}, \frac{1 + u^2}{1 - u^2} \right), \tag{A2}$$

where u is the length of the vector in the x-y plane. The distance function for the Poincaré disk model is then given by the symmetric bilinear form related to Q by polarization:

$$\cosh[d(\tilde{\mathbf{u}}, \tilde{\mathbf{v}})] = \langle \mathbf{u}, \mathbf{v} \rangle = \frac{1}{2}[Q(\mathbf{u} + \mathbf{v}) - Q(\mathbf{u}) - Q(\mathbf{v})], \tag{A3}$$

where we have denoted the two-dimensional vectors in the disk by $\tilde{\mathbf{u}}$ and $\tilde{\mathbf{v}}$ and their images on the hyperboloid in $\mathbb{R}^3$ by $\mathbf{u}$ and $\mathbf{v}$. Substituting the expressions for the images in terms of the model vectors and simplifying yields the distance function for the Poincaré disk model with $K = -1$:

$$d(\tilde{\mathbf{u}}, \tilde{\mathbf{v}}) = \cosh^{-1}\left[1 + 2\frac{|\tilde{\mathbf{u}} - \tilde{\mathbf{v}}|^2}{(1 - \tilde{u}^2)(1 - \tilde{v}^2)} \right]. \tag{A4}$$

Were we to write $\tilde{\mathbf{u}} = \tilde{\mathbf{v}} + d\tilde{\mathbf{u}}$, then we would find

$$ds^2 = [d(\tilde{\mathbf{u}}, \tilde{\mathbf{v}})]^2 = \left(\cosh^{-1}\left[1 + 2\frac{d\tilde{\mathbf{u}}^2}{(1 - \tilde{u}^2)(1 - \tilde{v}^2)} \right] \right)^2$$

$$\approx \frac{4d\tilde{\mathbf{u}}^2}{(1 - \tilde{u}^2)^2}, \tag{A5}$$

in agreement with (A1).

2. Hyperbolic law of cosines

Consider a triangle in the hyperbolic plane with side lengths a, b, and c and corresponding angles α, β, and γ.

Call the coordinate vertices A, B, and C. As before, we set $K = -1$ in order to simplify the notation. Now, because $\mathbb{H}^2$ is isotropic and homogeneous, we may assume that vertex A sits at the origin of the Poincaré disk model. The distance function (A4) then gives the following relations:

$$C = \sqrt{\frac{\cosh(b) - 1}{\cosh(b) + 1}} \quad \text{and} \quad B = \sqrt{\frac{\cosh(c) - 1}{\cosh(c) + 1}}. \quad \text{(A6)}$$

At the same time, applying the distance function in a similar way to the side of the triangle without a vertex on the origin gives

$$\cosh(a) = 1 + 2\frac{|\mathbf{B} - \mathbf{C}|^2}{(1 - B^2)(1 - C^2)}. \quad \text{(A7)}$$

But note that A, B, and C define a Euclidean, coordinate triangle as well. We apply the traditional law of cosines here:

$$|\mathbf{B} - \mathbf{C}|^2 = B^2 + C^2 - 2BC\cos(\alpha). \quad \text{(A8)}$$

Substituting (A6) and (A8) into (A7) eliminates B and C. Clearing denominators and repeated application of the hyperbolic trigonometric identities yields the hyperbolic law of cosines:

$$\cosh(a) = \cosh(b)\cosh(c) - \sinh(b)\sinh(c)\cos(\alpha). \quad \text{(A9)}$$

Alternatively, there is a dual form of this relation that one can obtain from the cyclic permutations of the hyperbolic law of cosines together with more trigonometric identities (hyperbolic or otherwise):

$$\cos(\alpha) = \cos(\beta)\cos(\gamma) + \sin(\beta)\sin(\gamma)\cosh(a). \quad \text{(A10)}$$

3. Area within a given distance of a geodesic

Consider a quadrangular region in the hyperbolic plane bordered on one side by a geodesic of length d and having geodesics of length r as the sides adjacent to the first. Finally, define the fourth side as the curve of constant geodesic curvature such that each point along it is a distance r from the length-d geodesic [45]. Recall the geodesic deviation equation

$$\frac{\partial^2 \xi}{\partial s^2} = -K\xi, \quad \text{(A11)}$$

where $\xi(s)$ is the separation between two geodesics as a function of the arclength parameter along each, s. Consider the two sides of length r in the context of this equation. The initial conditions are $\xi(0) = d$ and $\frac{\partial \xi}{\partial s}(0) = 0$ because the initial curve of length d is itself a geodesic and so first variations of its length vanish. We find $\xi(s) = d\cosh(\sqrt{-K}s)$. The area enclosed by our strip must then be given by

$$A = \int_0^r ds\ \xi(s) \quad \text{(A12)}$$

or

$$A = \frac{d\sinh(r\sqrt{-K})}{\sqrt{-K}}, \quad \text{(A13)}$$

as required.

---

[1] H. M. Jaeger, S. R. Nagel, and R. P. Behringer, Rev. Mod. Phys. **68**, 1259 (1996).

[2] I. Zuriguel, L. A. Pugnaloni, A. Garcimartin, and D. Maza, Phys. Rev. E **68**, 030301(R) (2003).

[3] K. Binder and A. P. Young, Rev. Mod. Phys. **58**, 801 (1986).

[4] K. Feitosa, O. L. Halt, R. D. Kamien, and D. J. Durian, Europhys. Lett. **76**, 683 (2006).

[5] H. Gang, A. H. Krall, H. Z. Cummins, and D. A. Weitz, Phys. Rev. E **59**, 715 (1999).

[6] E. D. Zanotto, Am. J. Phys. **66**, 392 (1998).

[7] S. Torquato, T. M. Truskett, and P. G. Debenedetti, Phys. Rev. Lett. **84**, 2064 (2000).

[8] M. Kléman and J. F. Sadoc, J. Phys. (France) Lett. **40**, L569 (1979); see also J. F. Sadoc and R. Mosseri, *Geometrical Frustration* (Cambridge University Press, Cambridge, England, 1999).

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

[10] E. A. Abbott, *Flatland: A Romance of Many Dimensions* (Seely and Co., London, 1884).

[11] D. R. Nelson, Phys. Rev. Lett. **50**, 982 (1983).

[12] M. Rubinstein and D. R. Nelson, Phys. Rev. B **28**, 6377 (1983).

[13] J. P. Sethna, Phys. Rev. Lett. **51**, 2198 (1983).

[14] G. A. Korn and T. M. Korn, *Mathematical Handbook for Scientists and Engineers* (Dover, Mineola, NY, 2000).

[15] H. Weyl, Rend. Circ. Mat. Palermo **39**, 1 (1915).

[16] M. Kac, Am. Math. Monthly **73**, 1 (1966).

[17] H. P. McKean, Jr. and I. M. Singer, J. Diff. Geom. **1**, 43 (1967).

[18] F. Sausset and G. Tarjus, J. Phys. A: Math. Theor. **40**, 12873 (2007).

[19] C. D. Modes and R. D. Kamien, Phys. Rev. Lett. **99**, 235701 (2007).

[20] Y. Deng and M. Mieczkowski, Protoplasma **203**, 16 (1998).

[21] E. S. Fedorov, Proc. St. Petersburg Mineral. Soc. **21**, 1 (1885).

[22] B. Grunbaum, Z. Grunbaum, and G. C. Sheppard, Comput. Math. Appl. **12**, 641 (1986); R. Perez-Gomez, *ibid.* **14**, 133 (1987).

[23] J. C. Maxwell, Philos. Mag. **27**, 250 (1864); J. C. Maxwell, Trans. - R. Soc. Edinbrgh **26**, 1 (1872).

[24] H. S. M. Coxeter, *Non-Euclidean Geometry*, 6th ed. (The Mathematical Association of America, Providence, RI, 1998).

[25] H. S. M. Coxeter, *Twelve Geometric Essays* (Southern Illinois University Press, Springfield, IL, 1968).

[26] J. G. Kirkwood, J. Chem. Phys. **18**, 380 (1950).

[27] Z. W. Salsburg and W. W. Wood, J. Chem. Phys. **37**, 798 (1962).

[28] A. Donev, S. Torquato, and F. H. Stillinger, Phys. Rev. E **71**, 011105 (2005).

[29] R. D. Kamien, in *Entropic Attraction and Ordering*, edited by G. Gompper and M. Schick (Wiley-VCH, Berlin, 2006).

[30] S. Hilgenfeldt, A. M. Kraynik, D. A. Reinelt, and J. M. Sullivan, Europhys. Lett. **67**, 484 (2004); M. E. Glicksman, Philos. Mag. **85**, 3 (2005).

[31] L. D. Landau and L. M. Lifshitz, *Course of Theoretical Physics*, 3rd ed. (Reed Educational and Professional Publishing, Oxford, 1980), Vol. 5.

[32] B. M. McCoy, e-print arXiv:cond-mat/0103556.

[33] M. Luban and A. Baram, J. Chem. Phys. **76**, 3233 (1982).

[34] N. Clisby and B. M. McCoy, Pramana J. Phys. **64**, 775 (2005); N. Clisby and B. M. McCoy, J. Stat. Phys. **114**, 1343 (2004).

[35] M. Robles, M. López de Haro, and Andrés Santos, J. Chem. Phys. **120**, 9113 (2004).

[36] L. Bowen and C. Radin, Discrete Comput. Geom. **29**, 23

[37] T. M. Truskett, S. Torquato, S. Sastry, P. G. Debenedetti, and F. H. Stillinger, Phys. Rev. E **58**, 3083 (1998).

[38] S. Torquato and A. Donev, Proc. R. Soc. London, Ser. A **460**, 1849 (2005).

[39] A. B. Subramaniam *et al.*, Nature (London) **438**, 930 (2005).

[40] R. D. Kamien and A. J. Liu, Phys. Rev. Lett. **99**, 155501 (2007).

[41] K. Larsson, J. Phys. Chem. **93**, 7304 (1989).

[42] B. A. DiDonna and R. D. Kamien, Phys. Rev. E **68**, 041703 (2003).

[43] M. C. Escher, *Circle Limit III & Circle Limit IV*.

[44] *B. Iversen, Hyperbolic Geometry, London Mathematical Society Student Texts* (Cambridge University Press, Cambridge, England, 1992), Vol. 25.

[45] T. H. Marshall and G. J. Martin, Proc. Edinburgh Math. Soc. **46**, 67 (2003).

(2002).