

A reprint from

American Scientist

the magazine of Sigma Xi, The Scientific Research Society

This reprint is provided for personal and noncommercial use. For any other use, please send a request Brian Hayes by electronic mail to bhayes@amsci.org.

The Science of Sticky Spheres

Brian Hayes

TAKE A DOZEN MARBLES, all the same size, and squeeze them into a compact, three-dimensional cluster. Now count the number of points where the marbles touch one another. What is the maximum number of contact points you can possibly achieve with 12 marbles? What geometric arrangement yields this greatest contact number? Is the optimal cluster unique, or are there multiple solutions that all give the same maximum?

When I first heard these questions asked, they did not seem overly challenging. For a cluster consisting of two, three, four or five equal-size spheres, I was pretty sure I knew the answers. But I soon learned that the problem gets harder in a hurry as the size of the cluster increases. Over the past three years, the maximum contact number has been determined for clusters of up to 11 spheres. Finding those answers required a variety of mathematical tools drawn from graph theory and geometry, as well as extensive computations and, at a few crucial junctures, building ball-and-stick models with a set of toys called Geomags. For clusters of 12 or more spheres, the answers remain unknown.

To be stumped by such simple questions about small clumps of spheres is humbling—but perhaps not too surprising. Sphere-packing problems are notoriously tricky. Some of them have resisted analysis for centuries.

Kepler and Newton

In 1611 Johannes Kepler declared that the densest possible packing of identical spheres is the arrangement seen in a grocer's pyramid of oranges. Any sphere in the interior of Kepler's lat-


*On the strange
attraction
of spheres
that like
to stick together*

tice touches 12 other spheres, and the fraction of space filled by the spheres is $\pi/\sqrt{18}$, or about 0.74. Kepler apparently believed that the superiority of this packing was so obvious that no proof was needed; as it happens, no proof was forthcoming for nearly 400 years. In 1998 Thomas C. Hales of the University of Pittsburgh finally showed that no other packing that extends throughout three-dimensional space can have a higher density.

Kepler's conjecture (and Hales's proof of it) apply to an infinite lattice


of spheres, but another centuries-old puzzle concerns finite clusters. The story begins with a dispute between Isaac Newton and his disciple David Gregory in the 1690s. According to one telling of the tale, Newton held that a central sphere could touch no more than 12 surrounding spheres of the same size, but Gregory thought there might be room for a 13th halo sphere. This problem of the "kissing number" was not resolved until 1953, when Kurt Schütte and B. L. van der Waerden proved that Newton was right—but just barely so. When a 13th satellite sphere is shoehorned into the assemblage, the diameter of the cluster increases by only about 5 percent.

Newton's kissing-number problem suggests a solar-system model of sphere packing, with a dozen planets all feeling the attraction of a central sun. The contact-counting problem has a more egalitarian character. There is no designated center of attraction; instead, all the spheres stick to one another, and the goal is to maximize



The kissing-number problem asks how many spheres can touch a central sphere of the same size. The question was debated by Isaac Newton and David Gregory in the 1690s but was answered only in the 1950s. At left 12 halo spheres make contact with the red central sphere, in a configuration derived from the hexagonal close-packed lattice of Johannes Kepler. At right a 13th halo sphere is added; as a result, all of the halo spheres lose touch with the central sphere, although the diameter of the cluster grows by only about 5 percent. A different sphere-packing problem of recent interest seeks to maximize the total number of contacts among all the spheres in a cluster. In the diagrams, any two spheres that touch have their centers connected by a rod.

Brian Hayes is senior writer for American Scientist. Additional material related to the Computing Science column appears at <http://bit-player.org>. Address: 11 Chandler St. #2, Somerville, MA 02144. E-mail: brian@bit-player.org



For clusters of up to five spheres, counting the maximum number of contacts is easy. A single sphere has no contacts, and two spheres can touch at only one point. Three spheres form an equilateral triangle with a total of three contacts; four spheres arranged in a tetrahedron have six contacts. In these four cases, each sphere touches every other sphere in the cluster, but no such configuration exists for five spheres. The best arrangement glues a fifth sphere to one face of a tetrahedron, creating a triangular dipyramid with nine contact points.

the overall number of contact points throughout the cluster.

Why is there so much interest in cramming spheres together? Kepler was trying to explain the symmetries of snowflakes, and much of the later work on sphere-packing has also been motivated by questions about the structure of solids and liquids. The recent focus on clusters with many sphere-to-sphere contacts arose from studies of colloids, powders and other physical systems in which particles are held together by extremely short-range forces.

The Well-Connected Cluster

In building clusters and counting contacts, it's convenient to work with *unit spheres*, which have a diameter of 1 and thus a radius of $1/2$. When two unit spheres are touching, the center-to-center distance is 1. The diagrams accompanying this column show a unit-length rod connecting the centers of spheres when they are in contact. Some physical models of clusters (including Geomags) keep this skeleton of connecting rods and omit the spheres altogether.

In any cluster of n spheres, let C_n denote the total number of contact points; then $\max(C_n)$ is the highest value of C_n found among all n -sphere clusters.

For the smallest values of n , finding $\max(C_n)$ is easy. The case of $n = 1$ is trivial: A single, isolated sphere doesn't touch anything, and so $\max(C_1) = 0$. Two spheres can meet at only one point, which means that $\max(C_2) = 1$. For three spheres the best solution puts the sphere centers at the vertices of an equilateral triangle; in this arrangement there are three contact points, and thus $\max(C_3) = 3$. A fourth sphere can be placed atop the triangle to create a regular tetrahedron with six pairwise contacts: $\max(C_4) = 6$.

Not only is it easy to construct these small clusters; it's also easy to prove that no other n -sphere configurations

could have a higher C_n . The reason is simply that in these clusters each sphere touches every other sphere, and so the number of contacts could not possibly be greater. In the terminology of graph theory, the cluster is a *clique*. The number of contacts in a cliquish cluster is $n(n-1)/2$. The sequence begins 0, 1, 3, 6, 10, 15, 21, ...

Going on to $n = 5$, cliquishness is left behind: In three-dimensional space there is no way to arrange five unit spheres so that they all touch one another. If such a five-sphere clique existed, it would have 10 contact points. The best that can actually be attained is $C_5 = 9$, which is the number of contacts formed when you attach a fifth sphere to any face of a tetrahedral cluster. The resulting structure is known as a triangular dipyramid.


At Sixes and Sevens

Up to this point, each value of n has had a unique cluster that maximizes C_n . Furthermore, in each case the best-connected cluster with $n+1$ spheres can be assembled incrementally by sticking a new sphere somewhere on the surface of the $\max(C_n)$ cluster. These properties come to an end at $n = 6$. With six spheres, two cluster


shapes both yield the same maximum contact number, $C_6 = 12$. (Note that a hypothetical six-sphere clique would have 15 contacts.) One of the $\max(C_6)$ clusters is built incrementally from the five-sphere triangular dipyramid. But the other $\max(C_6)$ cluster is a "new seed"—a structure that cannot be created simply by gluing a sphere to the surface of a smaller optimum cluster. The new seed is the octahedron (which might also be described as a square dipyramid).

Beyond $n = 6$, the problem of finding all the maximum-contact clusters becomes more daunting. For $n = 7$, the incremental approach of adding another sphere to the surface of an $n = 6$ cluster yields four solutions that have 15 contact points. Three of these $C_7 = 15$ clusters consist of four tetrahedra glued together face-to-face in various ways. The remaining product of incremental construction consists of an octahedron with a tetrahedron erected on one face. (One of the seven-sphere solutions has both left-handed and right-handed forms, but the convention is to count these "chiral" pairs as variants of a single cluster, not as separate structures.)

Finding this particular set of structures is not especially difficult. If you



Two distinct geometries both maximize the number of contacts in clusters of six spheres. The cluster at left is constructed incrementally from the optimal $n = 5$ cluster by attaching a sixth sphere to one face of the triangular dipyramid. At right is a "new seed": the octahedron. In both structures the six spheres have a total of 12 contact points.



The best-connected clusters of seven spheres have 15 contacts. Two of the structures above are assemblages of four tetrahedra glued together across various faces. The “tetratetrahedron” at left is highly symmetrical; the one in the middle is not only asymmetric but also chiral, meaning it has both left-handed and right-handed forms. The structure at right is an octahedron augmented with a tetrahedron erected on one face. All of these clusters can be built up incrementally by adding a single sphere to the surface of one of the two $n = 6$ optimal clusters.


spend some time playing with Geomags or some other three-dimensional modeling device, you are likely to stumble upon them. But having identified these four clusters with $C_7 = 15$, how do you know there aren’t more? And how do

you prove that no seven-sphere cluster has 16 or more contacts?



As it turns out, 15 is indeed the maximum contact number for seven spheres, but there *is* another $C_7 = 15$ cluster. It is a new seed, called a pen-

tagonal dipyramid. With its fivefold symmetry, it has no structural motifs in common with any of the smaller clusters. The novelty of this object again raises the question: How can we ever be sure there aren’t still more arrangements waiting to be discovered?

A successful program for answering such questions was initiated about five years ago by Natalie Arkus, who was then a graduate student at Harvard University. (She is now at Rockefeller University.) In a series of papers written with her Harvard colleagues Michael P. Brenner and Vinodhan N. Manoharan, she enumerated all the $\max(C_n)$ configurations for $n = 7$ through $n = 10$. The results were later extended to $n = 11$ by Robert S. Hoy, Jared Harwayne-Gidansky and Corey S. O’Hern of Yale University. (Hoy is now on the faculty of the University of South Florida.) All of the results I describe here come from the work of these two groups.



Two more seven-sphere clusters look superficially alike but actually have very different structures. At left is yet another way of assembling four tetrahedra by gluing faces together, yielding an almost-pentagonal form. It is *almost* pentagonal because the two spheres at the front are not touching. At right, a slight adjustment creates a pentagonal dipyramid, an object with true fivefold symmetry, but now the spheres above and below the five-member ring no longer touch.



The Geomags construction set offers another way to explore and visualize the structure of sticky-sphere clusters. The set consists of small steel balls linked by plastic-coated magnetic rods, which are all the same length. The steel balls represent the centers of spheres, which are touching if and only if a rod fits between the corresponding steel balls. In the two photographs above the yellow rods represent contacts between spheres. The red rods are not part of the skeleton of bonded spheres; they have been inserted to show where pairs of spheres that look like they might be in contact are actually too far apart for the rod to connect them.

Sticky Spheres


One way to solve sphere-packing problems is to view the spheres as particles subject to a physical force. Then, through mathematical analysis or computer simulation, you can try to find the geometric arrangement that minimizes the potential energy of the system. The force is usually defined as a smooth function of distance. When two particles are far apart, the force between them is negligible; at closer range the force becomes strongly attractive; at even smaller distances a “hard-core repulsion” prevents the spheres from overlapping. Under a force law of this kind, the particles settle into equilibrium at some small but nonzero separation.

The contact-counting problem can be translated into the language of forces and energy, but the physics of the system is rather peculiar. To begin with, the force law is not a smooth function of distance. Instead of hills and valleys representing gradual changes in energy, there is a sheer cliff, where the energy jumps abruptly. Imagine two spheres drifting through space. As long as they do not touch, there is no force acting between them—neither attraction nor repulsion. If the spheres happen to come in contact, however, they stick together; suddenly, the force becomes attractive. Yet any attempt to push them still closer is met by infinite resistance.

In this world of sticky spheres, the forces at work are not merely short-range but zero-range. (Martin Gardner once suggested a model for such systems: ping-pong balls coated with rubber cement.) Two spheres lower their total energy when they touch, and energy has to be supplied to pull them apart; but once they are separated, they have no further influence on one another. Minimizing the potential energy of the whole system means maximizing the number of contacts.

The discontinuous nature of the force law affects the choice of mathematical tools for solving the sticky-sphere problem. With a smooth force law, sphere-packing problems can be solved by optimization methods. An algorithm repeatedly attempts to reduce the total energy by making small adjustments to the particles' positions, continuing until no further progress is made. This scheme won't work with sticky spheres because there are no smooth gradients to guide the particles toward lower-energy configurations. For this reason, the sticky-spheres problem has seemed harder than most other sphere-packing tasks.

On the other hand, the discrete, all-or-nothing character of the sticky-spheres potential also brings an important advantage. Because each pair of spheres is either touching or not, the number of essentially different configurations is finite. In principle, you can examine all these possibilities and simply choose the one with the most contacts and hence the lowest potential energy. It was this insight—the idea that the problem can be solved by exhaustive enumeration—that led to the recent results of the Harvard and Yale groups.



Graph theory offers a convenient way of enumerating all sticky-sphere clusters with a given number of contacts. A mathematical graph is a set of vertices (usually drawn as dots) and a set of edges (drawn as lines between vertices). Two vertices are connected by an edge if and only if the corresponding spheres are in contact. The octahedral cluster at left is represented by the graph in the center, which ignores geometry but preserves the pattern of connectivity. The same information is recorded in the adjacency matrix at right. If i and j are any spheres in the cluster, the entry in row i and column j of the matrix is 1 if i touches j and otherwise is 0. Only the upper triangle (blue) needs to be specified; the lower triangle is its mirror image. All matrices with m 1s in the upper triangle are candidates for clusters with m contacts.


From Geometry to Graph Theory

All the essential facts about sphere-to-sphere contacts in a cluster can be captured in a graph—a collection of vertices and edges. Each sphere is represented by a vertex, and two vertices are connected by an edge if and only if the corresponding spheres are in contact. The same information can be encoded even more abstractly in an adjacency matrix. A cluster of n spheres becomes an n -by- n matrix of 0s and 1s. The matrix element at the intersection of row i and column j is 1 if sphere i touches sphere j and otherwise is 0.


Given a cluster of spheres, it's easy to construct the corresponding graph or adjacency matrix. But is it possible to go the other way—to start with an adjacency matrix and recover the full geometry of the cluster? In other words, with nothing more to go on than a table indicating which spheres are in contact,

can one determine the coordinates of all the spheres in three-dimensional space? The answer is: Not always. Consider the all-0s matrix, which reveals nothing about the locations of the spheres except that they're not touching. And the all-1s matrix describes a cluster that simply cannot exist when n is greater than 4. But for an important class of clusters the adjacency matrix does supply enough information to allow a full reconstruction. That class includes the clusters that maximize contact number. These facts suggest a direct problem-solving strategy: Generate all the candidate matrices and check to see which ones produce geometrically feasible clusters.

How many adjacency matrices need to be tested? Because contact is a symmetric relation—if i touches j , then j touches i —all the information in the matrix is confined to the upper triangle, which has $n(n-1)/2$ elements.



Geometric constraints determine which adjacency matrices describe clusters that can be realized in three-dimensional space. If two spheres are in contact (left), a third sphere can touch them both only if its center lies somewhere on a circle defined by the intersection of two "neighbor" spheres. With three spheres all in mutual contact (right), the available positions for adding a fourth sphere are reduced to two points, at the intersections of three circles.



Among 52 nine-sphere clusters that maximize the number of contacts, one configuration has a remarkable property: flexibility. Two sections of the structure can twist in opposite directions without breaking any contacts between spheres. The diagrams above show three phases in this oscillatory motion, with the neutral position in the middle. The source of the flexibility is a pair of square faces that share an edge.

Each element has two possible values, and so the total number of adjacency matrices is $2^{n(n-1)/2}$.

Sifting through $2^{n(n-1)/2}$ matrices would be a formidable task; the value of this expression is already beyond two million at $n=7$ and exceeds 10^{16} at $n=11$. But not all of those matrices are distinct; many of them represent mere relabelings of the same graph. When redundancies of this kind are eliminated, the number of distinct 7×7 matrices is reduced from 2,097,152 to just 1,044.

Arkus and her group took advantage of a further dramatic winnowing. For reasons that will be explained below, they examined only matrices that meet two criteria: Every column and row has at least three 1s, and the total number of 1s in the upper triangle is $3n-6$. For $n=7$, imposing these constraints reduces the number of can-

didate adjacency matrices from 1,044 to just 29. Among those 29 matrices, only five give rise to genuine three-dimensional sphere packings.

And Back to Geometry

How were these geometric structures determined? The key idea is to transform the adjacency matrix A into a distance matrix D . Whereas each element A_{ij} of the adjacency matrix is a binary value, answering the yes-or-no question “Do spheres i and j touch?,” the element D_{ij} is a real number giving the Euclidean distance between i and j .

As it happens, we already know some of those distances. Every 1 in the adjacency matrix designates a pair of unit spheres whose center-to-center distance is exactly 1; thus $A_{ij}=1$ implies $D_{ij}=1$. We even know something about the rest of the distances: A cluster is feasible only if every element of the distance matrix satisfies the constraint $D_{ij} \geq 1$. Any distance smaller than 1 would mean that two spheres were occupying the same volume.

To fully pin down the geometry of a cluster, we need to determine the x , y and z coordinates of all n spheres. A rule of elementary algebra suggests we would need $3n$ equations to determine these $3n$ unknowns, but in fact $3n-6$ equations are enough. The energy of the cluster depends only on the *relative* positions of the n spheres, not on the absolute position or orientation of the cluster as a whole. In effect, the locations of two spheres come “for free.” We can arbitrarily assume that one sphere is at the origin of the coordinate system and another is exactly one unit away along the positive x axis. In this way six coordinates become fixed. Then the $3n-6$ equations supplied by the 1s in the adjacency matrix are ex-

actly the number needed to locate the rest of the spheres.


Having just enough constraints to solve the system of equations is more than a convenient coincidence; it’s also a necessary condition for mechanical stability in a cluster. Specifically, having $3n-6$ contacts and at least three contacts per sphere gives a cluster a property called minimal rigidity. If any sphere had only one or two contacts, it could flap or wobble freely. Such a cluster cannot be a $\max(C_n)$ configuration because the unconstrained sphere can always pivot to make contact with at least one more sphere, thereby increasing C_n .

Each of the $3n-6$ equations has the form:

$$\sqrt{(x_i - x_j)^2 + (y_i - y_j)^2 + (z_i - z_j)^2} = 1,$$

defining the distance between the centers of spheres i and j . To recover the coordinates of all the spheres, this system of equations must be solved. To that end, Arkus first tried a technique called a Gröbner basis, which in recent years has emerged as a powerful tool of algebraic geometry. The method offers a systematic way to reduce the number of variables until a solution emerges. An implementation of the Gröbner-basis algorithm built into a computer-algebra system was able to solve the $n=7$ equations, but it became too slow for $n=8$.

Another approach relies on numerical methods that converge on a solution by successive approximation. The best-known example is Newton’s method of root-finding by refining an initial guess. Arkus found that the numerical techniques were successful and efficient, but she was concerned that they are not guaranteed to find *all* valid solutions.



A 10-sphere cluster with 25 contacts is the first to cross an important threshold. From $n=3$ through $n=9$, the maximum number of contacts in an n -sphere cluster is $3n-6$, which is the minimum required for mechanical stability. This 10-sphere cluster exceeds that requirement. It is built up from the flexible $n=9$ cluster by adding a sphere to one of the square faces (thereby eliminating the flexibility).

(Whenever the algorithm converges, the result is a correct solution, but failure to converge does not necessarily mean that no solution exists; it's also possible that the initial guess was in the wrong neighborhood.)

Setting aside the algebraic and numerical techniques, Arkus chose to rely on geometric reasoning both as a guide to assembling feasible clusters and as a means of excluding unphysical ones. A basic rule for unit spheres states that if i touches j , then j 's center must lie somewhere on a sphere of radius 1 centered on i —the “neighbor sphere.” If k touches both i and j , then k 's center must be somewhere on the circular intersection of two neighbor spheres. If l touches all three of i , j and k , the possible locations are confined to a set of two points. With a handful of rules of this general kind, it's always possible to solve for the unknown distances in a distance matrix—assuming that the adjacency matrix describes a feasible structure.

Other geometric rules can be applied to prove that certain classes of adjacency matrices cannot possibly yield a physical sphere packing. For example, if spheres i , j and k all touch one another, they must form an equilateral triangle. If the pattern of 1s in the adjacency matrix shows that more than two other spheres also touch i , j and k , then the cluster cannot exist in three-dimensional space. The unphysical matrices can be eliminated without even exploring the geometry of the clusters.

Arkus also made use of the Geomags construction set to check the feasibility of certain sphere arrangements. The Geomags set consists of polished steel balls and bar-magnet struts encased in colored plastic; all the struts are the same length, and so they can readily be assembled into a skeleton of unit-length bonds between sphere centers. Having a three-dimensional model you can hold in your hand is a great aid to geometric intuition.

Eight, Nine, Ten

The results of the survey of sticky-sphere clusters are summarized in the table on page 448. Arkus and her colleagues determined $\max(C_n)$ —and identified all clusters that exhibit these highest contact numbers—for all $n \leq 10$. Along the way, they discovered quite a few clusters with interesting quirks and personalities.

At $n = 8$ the maximum contact number is 18, and there are 13 distinct ways of achieving this bound. All but one of the clusters can be built up incrementally by attaching a new sphere to the surface of one of the $n = 7$ clusters.

Clusters of nine spheres have up to 21 contacts; there are 52 varieties, including four new seeds. In this crowd of sphere packings, one stands out from all the rest. It has a property not seen in any other $\max(C_n)$ cluster up to this point: flexibility. The structure can be twisted around one axis without breaking any bonds between spheres. This ability to wiggle may seem surprising, given that the adjacency matrices were designed with the explicit aim of ensuring “minimal rigidity.” But there's a reason this form of rigidity is called minimal. The requirement that every sphere make contact with at least three others implies that no individual sphere can move relative to the rest of the cluster without breaking at least one bond. But other modes of motion, in which larger groups of spheres flex or rotate, are not ruled out. In the flexible $n = 9$ cluster, two square faces joined by an edge can twist and deform slightly. The upper illustration on the opposite page shows three states of torsion, somewhat exaggerated for clarity. (The Web version of this column animates the motion.)



The $n = 10$ clusters cross another threshold: For the first time the number of contacts exceeds $3n - 6$ (which is 24 in this case). Some 259 clusters of 10 spheres have exactly 24 contacts, but another three clusters have 25 each. Again, it's mildly surprising that these objects find their way onto the list. The

search algorithm begins with a list of matrices that specify exactly $3n - 6$ adjacencies, so how can the search uncover a cluster with even more contacts? The explanation is that even if two spheres are not required to touch, they are not forbidden to do so. When you solve a system of equations that specifies 24 adjacencies, it can happen, as if by coincidence, that a 25th pair of spheres is also at a distance of exactly 1.

One of these happy accidents is shown in the lower illustration on the opposite page. The structure is derived from the flexible $n = 9$ cluster by adding an octahedral cap to one of the square faces. (The addition eliminates the flexibility.)

Eleven

To compile the catalog of 10-sphere clusters, Arkus and her colleagues had to examine more than 750,000 matrices for minimally rigid structures. The challenge of pushing the frontier out to $n = 11$ was taken up by Hoy, Harwayne-Gidansky and O'Hern at Yale. They relied on many of the same methods but adopted a different approach to streamlining the algorithms. For example, they took advantage of a curious fact proved by Therese Biedl, Erik Demaine and others: Any valid packing of spheres has a continuous, unbranched path that threads from one sphere to the next throughout the structure, like a long polymer chain. This fact implies that the rows and columns of the adjacency matrix can always be rearranged so that the superdiagonal (just above and to the right of the main diagonal) consists entirely of 1s. Confining attention only to these



The 11-sphere champion has 29 contacts, two more than the $3n - 6$ lower bound of 27. The structure is shown at left in a computer-generated space-filling diagram and at right in a photograph of a Geomags model. It can be built up from the flexible nine-sphere model by adding octahedral caps to both of the square faces.

matrices reduces the workload by a factor of 1,000 for $n = 11$.

The Yale group also devised simplified rules for excluding invalid packings. And they formulated more of the geometric rules in such a way that matrices could be tested without ever having to go through the time-consuming steps of computing sphere-to-sphere distances. For example, one such exclusion rule applies to clusters that have eight spheres arranged at the vertices of a cube. No ninth sphere can touch more than four of these corner spheres; to do so, the ninth sphere would have to lie somewhere inside the cube, but there's no room for it there. Violations of this rule can be detected merely by counting 1s in the adjacency matrix, a much quicker operation than calculating three-dimensional coordinates.

At $n = 11$ the Yale group identified 1,641 distinct clusters with $C_{11} \geq 27$. The vast majority of these structures have exactly 27 contacts (the $3n-6$ value), but there are 20 clusters with 28 contacts (equal to $3n-5$) and a single packing with 29 contacts ($3n-4$). This last object, shown on page 447, can be understood as a further elaboration of the "floppy" cluster described above. The flexible nine-sphere structure has two square faces exposed at the surface. One of those faces is capped to form an octahedron in the sole 10-sphere cluster with 25 contacts. Capping the other square

face to create a second octahedron leads to the unique 11-sphere cluster with 29 contacts.

Incidentally, there is an obvious way to add a 12th sphere to this cluster to produce a structure with 33 contacts, equal to $3n-3$. But whether or not 33 is the highest attainable C_{12} value remains a matter of conjecture, because no complete survey of 12-sphere clusters has been attempted.

Even for $n = 11$ it's possible to quibble over questions of completeness and certainty. Some of the Yale results rely on a numerical algorithm to solve the system of distance equations. As noted above, when this process fails to converge, it does not unequivocally prove that no solutions exists; even after many trials, there's always a possibility that one more run with a different initial guess might succeed. In practice, the chance that a valid sphere packing might have been missed is extremely slim; whether the question is worth worrying about is perhaps a matter of differing attitudes toward rigor in mathematics and the physical sciences.

Sticky Spheres in Action

Both the Harvard and the Yale groups were inspired to undertake this exercise by an interest in aggregations of material spheres rather than mathematical ones. Hoy, Harwayne-Gidansky and O'Hern discuss the mysteries of crystallization. Bulk materials tend to favor configurations that maximize density, such as the Kepler packing. But crystal growth must start from clusters of just a few atoms, where the configurations that minimize energy are not the same as those in the bulk. Some high-contact-number clusters exhibit motifs seen in the Kepler packing, but other clusters are incompatible with space-filling structures. For example, the $n = 7$ pentagonal dipyrmaid is not a pattern that can tile three-dimensional space.

Arkus is particularly interested in the self-assembly of nanostructures, both natural ones (such as viral capsids) and engineered materials. Guangan Meng of Harvard, working with Arkus, Brenner and Manoharan, has developed an experimental system that offers another way to explore small clusters of sticky spheres. The spheres are polystyrene beads one micrometer in diameter, suspended in water along with a large population of much smaller plastic nanoparticles. When two spheres come into contact, the nanoparticles are

excluded from the space between them; this phenomenon creates a short-range attractive force between the spheres. Hence the system is a good model of idealized sticky spheres.

In Meng's experiments the microspheres were spread over glass plates with thousands of cylindrical microwells, where they formed clusters with an average of about 10 spheres per well. The wells were scanned with a microscope to tabulate the relative abundance of various configurations. If potential energy were the sole criterion, then clusters with more contacts would be more common, but entropy also enters into this calculus: Structures that can be formed in many different ways are more probable. For the most part, results were broadly in accord with theoretical expectations. Entropy favored clusters with lower symmetry, and also enhanced the representation of nonrigid structures. But because extra contacts lower the potential energy, structures with more than $3n-6$ bonds were also overrepresented.

Apart from these physical applications of sticky spheres, the contact-counting model also evokes a celebrated open problem in pure mathematics. The question was raised by Paul Erdős in 1946: Given n points in d -dimensional space, how many pairs of points can be separated by the same distance? By scaling all distances appropriately, the repeated distance can always be set equal to 1, and so the problem is sometimes called the unit-distance problem. In three dimensions, the maximum-contact problem for unit spheres is equivalent to the Erdős unit-distance problem with the additional constraint that no distance is allowed to be less than 1. Thus the recent results on sticky spheres solve this restricted version of the problem for all $n \leq 11$.

What lies beyond $n = 11$? Arkus suggests that the main roadblock to enumerating maximum-contact clusters for higher n is not the geometric problem of solving for coordinates and distances but the combinatorial one of generating all appropriate adjacency matrices. Because so few of the matrices correspond to valid packings, the process becomes hideously wasteful. Arkus suggests a possible alternative approach, although it has not yet been successfully implemented. Through $n = 10$ she has shown that every cluster with $3n-6$ contacts can be converted into any other $3n-6$

n	$\frac{n(n-1)}{2}$	$3n-6$	C_n	multiplicity
1	0		0	1
2	1		1	1
3	3	3	3	1
4	6	6	6	1
5	10	9	9	1
6	15	12	12	2
7	21	15	15	5
8	28	18	18	13
9	36	21	21	52
10	45	24	24	259
			25	3
11	55	27	27	1,620
			28	20
			29	1

Results of searches by groups at Harvard and Yale summarize what is known about sticky-sphere clusters that maximize C_n , the number of contacts. Up to $n = 4$, $\max(C_n)$ is equal to $n(n-1)/2$, the number of contacts formed when every sphere touches every other sphere. Up to $n = 9$, C_n is equal to $3n-6$, the minimum for mechanical stability. The numbers in the multiplicity column give the number of clusters with each value of C_n .

cluster by some chain of simple transformations, in which a single bond is broken and another bond is formed. She conjectures that this property holds true for all n . If it does, the maximum-contact problem might be solved by generating any one structure with $3n-6$ contacts and then systematically traversing the tree of all single-bond-exchange transformations.

At some large enough value of n , the diversity of these curious geometric structures will necessarily begin to diminish, as all larger clusters come to look more and more like pieces of the Kepler packing. But we're not there yet, and there may still be oddities to discover.

Note: The Web edition of this column, at <http://www.americanscientist.org>, includes animated versions of many of the illustrations.

Bibliography

- Arkus, N., V. N. Manoharan and M. P. Brenner. 2009. Minimal energy clusters of hard spheres with short range attractions. *Physical Review Letters* 103:118303.
- Arkus, N., V. N. Manoharan and M. P. Brenner. 2011. Deriving finite sphere packings. *SIAM Journal on Discrete Mathematics* 25(4):1860–1901.
- Aste, T., and D. Weaire. 2008. *The Pursuit of Perfect Packing*, 2nd ed. New York: Taylor & Francis.
- Biedl, T. E., et al. 2001. Locked and unlocked polygonal chains in three dimensions. *Discrete and Computational Geometry* 26:269–281.
- Conway, J. H., and N. J. A. Sloane. 1999. *Sphere Packings, Lattices, and Groups*, 3rd ed. New York: Springer.
- Erdős, P. 1946. On sets of distances of n points. *American Mathematical Monthly* 53:248–250.
- Hales, T. C. 2005. A proof of the Kepler conjecture. *Annals of Mathematics* 162:1065–1185.
- Hoare, M. R., and J. McInnes. 1976. Statistical mechanics and morphology of very small atomic clusters. *Faraday Discussions of the Chemical Society* 61:12–24.
- Hoy, R. S., J. Harwayne-Gidansky and C. S. O'Hern. 2012. Structure of finite sphere packings via exact enumeration: Implications for colloidal crystal nucleation. *Physical Review E* 85:051403.
- Hoy, R. S., and C. S. O'Hern. 2010. Minimal energy packings and collapse of sticky tangent hard-sphere polymers. *Physical Review Letters* 105:068001.
- Kepler, J. 2010. *The Six-Cornered Snowflake: A New Year's Gift*. Philadelphia: Paul Dry Books.
- Meng, G., N. Arkus, M. P. Brenner and V. N. Manoharan. 2010. The free-energy landscape of clusters of attractive hard spheres. *Science* 327:560–563.
- Schütte, K., and B. L. van der Waerden. 1953. Das Problem der dreizehn Kugeln. *Mathematische Annalen* 125:325–334.
- Sloane, N. J. A., R. H. Hardin, T. D. S. Duff and J. H. Conway. 1995. Minimal-energy clusters of hard spheres. *Discrete and Computational Geometry* 14:237–259.