

Packing Pyramids: Is the Space Race Over?

By Barry A. Cipra

Among the many boneheaded assertions the great Greek philosopher Aristotle casually tossed into his voluminous writings, one stands out for its geometric errancy: Book 3, part 8, of *On the Heavens* opens with the observation that “there are only three plane figures which can fill a space, the triangle, the square, and the hexagon, and only two solids, the pyramid and the cube.”* Actually, four out of five isn’t bad. Plato’s student nailed the only three regular polygons capable of filling the plane, and he spotted the cube as a space-filling Platonic solid. But he was dead wrong about the pyramid—more properly called the tetrahedron.

As with much of classical proto-science, Aristotle’s observation went unquestioned for the better part of two thousand years. It was left to the 15th-century astrologer Regiomontanus to catch the error. Some subsequent scholar-apologists insisted that the philosopher never explicitly said he was talking about *regular* polyshapes, but that defense seems weak. (The early history is nicely summarized in a 1981 *Mathematics Magazine* article, “Which Tetrahedra Fill Space?,” by Marjorie Senechal of Smith College.)

In modern terms, the irrationality (as a multiple of π) of the solid angles at the vertices of the regular tetrahedron is a clear indication that copies of it cannot fit together to fill all of space. But this leaves open an enticing question: How much of space *can* they fill?

Geometers are still a long way from knowing—or even conjecturing—an exact answer. But recent research has led to dramatic progress: Several groups have found increasingly dense packings of Aristotle’s pyramid, leaping from a paltry proven 36.7% of space up to 85.6%.

Making Sphere Packing Look Easy

The juxtaposing of geometric objects has long brought intellectual joy and perplexity to mathematicians and mathematically minded crystallographers and materials scientists. The most famous example is the so-called Kepler conjecture—now better known as Hales’s theorem—that the closest packing of identical spheres is achieved by the “face-centered cubic” lattice packing, which fills $\pi/\sqrt{18} = 74.048\%$ of space. Posited by Johannes Kepler in 1611, this seemingly obvious assertion outwitted all who attempted to prove it, until Thomas Hales, now at the University of Pittsburgh, put the pieces in place for a theoretically grounded computational attack. Hales’s proof, first announced in 1998, was published in 2005, albeit with a caveat from the referees that they could not vouch for all the calculations. (Unsatisfied himself with the status of the proof, Hales launched a long-term project to develop computational tools that will provide a formal verification.)

When something as simple as a sphere proves so stubborn, it’s no wonder that polyhedra pose such difficulty: For a sphere packing, all you have to worry about is the locations of the centers of the spheres; with polyhedra, you have a continuum of orientations to consider as well. This makes it tricky even if you restrict yourself to “Bravais lattice” arrangements, in which all copies have the same orientation and are centered at the regularly repeating points of a parallelepiped tiling of space. For tetrahedra, this was settled only in the 1960s. In 1962, Helmut Groemer, now an emeritus professor at the University of Arizona, found such a packing with density $18/49$ (see Figure 1); in 1969, Groemer’s student Douglas Hoylman proved that his adviser’s packing was optimal among Bravais lattice arrangements.

At 36.7% of space, the Groemer-Hoylman arrangement isn’t much of a packing, but no one gave the matter much thought for the rest of the century. In 1999, Ulrich Betke of Universität Siegen and Martin Henk of Universität Magdeburg pretty well polished off the Bravais lattice part of the problem. They developed an efficient algorithm that finds the densest Bravais lattice packing for any given polyhedron, and they applied it to the five Platonic solids and the 13 Archimedean solids. (One of the latter, the truncated octahedron, joins the cube in tiling space. The rest are known not to do so, although, as with the tetrahedron, how much of space they’re capable of filling is unknown.)

The tetrahedron got its second wind in 2006, when John Conway and Salvatore Torquato of Princeton University re-examined its packing properties. They observed that you can get 20 tetrahedra to fill 85.68% of an icosahedron.[†] When they put this together with Betke and Henk’s optimal Bravais lattice packing of icosahedra, which fills 83.63% of space (the exact value involves the golden ratio and the root of a bizarre cubic polynomial), the density of tetrahedra jumped to 71.66%.

Conway and Torquato also found that because of the looseness of the 20 tetrahedra inside each icosahedron—the tetrahedron’s solid angle measures $\arccos(23/27) = .5513$ steradians, so there’s technically room for 22 of them to touch at a point—along with details of the way icosahedra abut one another in the Betke-Henk lattice packing, they could enlarge the tetrahedra slightly after jostling the contact points. One systematic jostling procedure got them to 71.7455%. This left them thinking that whatever the actual densest packing of tetrahedra turned out to be, it would be less than Kepler’s 74.08% for spheres. That would be of interest in part because in 1972, Stan Ulam had conjectured that the sphere had the poorest packing density of any convex 3-D object. Conway and Torquato thought that the honor might belong to the tetrahedron instead.

They were wrong.




Figure 1. *Home of the Bravais.* If identical, non-overlapping tetrahedra are centered at points of a lattice and all assigned the same orientation, they can fill at best 36.7% of space. Reprinted with permission from Figure Packing, Tiling, and Covering with Tetrahedra, by John Conway and Salvatore Torquato, Proceedings of the National Academy of Sciences, Vol. 103, No. 28, July 11, 2006, doi:10.1073/pnas.0601389103; copyright 2006 National Academy of Sciences, USA.

*Translation of J.L. Stocks.

[†]The exact fraction is $\tau^2/8$, τ being the golden ratio, $(1 + \sqrt{5})/2$. The height of a tetrahedron is $\sqrt{2/3} = .8165$ times the length of a side, whereas the perpendicular distance from the center of an icosahedron to a face is only $\tau^2/(2\sqrt{3}) = .7558$ times the length of a side; 20 tetrahedra can thus touch at the center of an icosahedron with their opposing faces on the equilateral faces of the icosahedron if their sides are scaled down to length $\tau^2/(\sqrt{8}) = .9256$.

Clusters, Disorder, and Dimers

The Conway–Torquato paper, which appeared in July 2006, was a surprise to Elizabeth Chen, a grad student at the University of Michigan, who had begun working on tetrahedra in 2005. Chen, too, had noticed the packing possibilities of 20 tetrahedra inside a repeating array of icosahedra. Her task now became to improve on what the Princeton pair had done.

She did, big time.

By the end of the year, Chen had shot past the sphere-packing threshold, with an arrangement of tetrahedra that filled 77.86% of space. Chen’s new packing ignores icosahedra; it starts instead with two sets of four tetrahedra, hinged to a pair of non-adjacent edges of a ninth tetrahedron. The dihedral angles of the tetrahedra, $\arccos(1/3) = 70.53^\circ$, permit some play in the arrangement. With the optimal values Chen found for these “swivel” parameters, a pair of nine-unit clusters can be nestled together and then tiled in a repeating pattern. Her results, which became her PhD dissertation, were published in 2008.

Chen’s breakthrough caught Torquato’s eye. Last year, he and Princeton grad student Yang Jiao made the cover of *Nature* with a paper that nudged the tetrahedron’s space-filling density to 78.2%. They developed a general optimization approach for polyhedral packing, which they call the “adaptive shrinking cell” method. Starting with a number of copies of whatever is to be packed inside a regularly repeating “fundamental cell,” ASC randomly jostles one of the copies, checks to make sure the jostled object doesn’t overlap any of its neighbors—a trivial computation for spheres but far and away the most demanding part of the process for polyhedra—and then checks to see if the jostling allows for the fundamental cell to shrink slightly. For the new tetrahedral record, they took eight copies of Chen’s nine-unit clusters, and let ASC take over.

Torquato and Jiao also ran the other Platonic solids through ASC, obtaining results in each case close to the density of Betke and Henk’s Bravais lattice results. This suggests that those may in fact be the densest possible packings. Indeed, Torquato and Jiao conclude, it seems likely that among the Platonic and Archimedean solids, the tetrahedron and truncated tetrahedron are the only ones for which the densest packing is not a Bravais lattice arrangement. Notably, these two (conjectured) exceptions are also the only ones that are not “centrally symmetric”—i.e., lacking an interior point that is the midpoint of every chord passing through it.

The pace has only picked up. Several groups have jumped on the packing bandwagon since the *Nature* paper appeared, and tetrahedral packings have vaulted from 78.2% to well beyond 80%. Running ASC on a “dilute” starting arrangement of 314 tetrahedra filling less than 1% of their fundamental cell, Torquato and Jiao made it to 82.26%. The ending arrangement was a disordered jumble of distorted, Chen-like clusters and individual tetrahedra within each fundamental cell, prompting the Princeton pair to speculate that the densest packing of tetrahedra might be truly disordered, approximable only for larger and larger numbers of tetrahedra in larger and larger fundamental cells.

Or maybe not.

The ink was barely dry on Torquato and Jiao’s disordered-packing paper (published in *Physical Review E*, October 5, 2009) when Yoav Kallus and Veit Elser of Cornell and Simon Gravel of Stanford posted a paper on arXiv presenting a denser packing for a mere four tetrahedra per fundamental cell. Their packing was essentially a repeating arrangement of two “dimers,” each dimer being a pair of tetrahedra glued face-to-face. Its simplicity allowed for an exact density value: $100/117 = 85.47\%$.

Torquato and Jiao responded in December with a four-unit arrangement of their own that fills $12,250/14,319 = 85.55\%$ of space. Chen, with co-authors Michael Engel and Sharon Glotzer of the University of Michigan, topped that in January, showing that a pair of tetrahedral dimers can be packed to fill $4000/4671 = 85.63\%$ of space (see Figure 2).

That may be the end of the race for denser packings. The Michigan trio ran Monte Carlo “compressions” with up to 16 tetrahedra per fundamental cell. When the number was a multiple of four, the numerical results always maxed out at the analytic result; otherwise, the densities were lower. But if pyramid packers have learned anything from Aristotle, it’s not to be too sure.

Barry A. Cipra is a mathematician and writer based in Northfield, Minnesota.




Figure 2. *Bravais New World.* If small clusters of tetrahedra are centered at points of a lattice and given the same orientation, the amount of space they can fill varies with the size of the cluster, but seems to be greatest when N is a multiple of 4. Shown here are the optimal cluster arrangements for $N = 1$ to 4, which allow for packings of density, 36.7%, 71.9%, 66.7%, and 85.6%, respectively. (The exact values are $18/49$, $(139+40\sqrt{10})/369$, $2/3$, and $4000/4671$.) Figure courtesy of Elizabeth Chen, Michael Engel, and Sharon Glotzer, University of Michigan, from *Dense Crystalline Dimer Packings of Regular Tetrahedra*, <http://arxiv.org/pdf/1001.0586v2>.