On Holyhedra*

J. Vinson

Department of Mathematics, Princeton University,
Princeton, NJ 08544, USA
jpvinson@math.princeton.edu

Abstract. We describe the construction of a *holyhedron*, which is a polyhedron so that the interior of every face is nonsimply connected. The holyhedron we construct has 78,585,627 faces and genus 60,380,421.

Introduction (by John Conway)

Some years ago I asked whether there could exist a polyhedron that “had a hole in every face,” and coined the name “holyhedron” for such things, should they exist.

I circulated the problem to a few friends on the Internet, and it rapidly acquired a certain notoriety. The precise definition of the problem (given below) was not always known or understood, and a number of people produced objects that would deserve the name “holyhedron” under weaker definitions. Others deliberately generalized or modified the problem with the aim of understanding where the difficulties lay, and so produced some more interesting examples. However, the problem as originally posed obstinately resisted solution until Jade Vinson produced the construction that he describes in the main part of this paper.

Superimpose two solid tetrahedra in such a way that a vertex of one pierces a face of the other. Then their union (Fig. 1) is a polyhedron for which one face is multiply connected; but this does not deserve to be called a “holyhedron” since the other faces are not.

Kepler’s “Stella Octangula” [1], [2] of Fig. 2 comes closer to being an example, if we regard it as having eight faces that are each composed of three triangles as in Fig. 3. These “faces” are certainly multiply connected (as closed sets), but since they are only connected by points, it seems more natural to disbar this example by regarding each of them as being three faces rather than one.

* The author was supported by an NDSEG fellowship.

The polyhedron of Fig. 4 is obtained by drilling square holes between each pair of opposite faces. We get an example having the same virtues and vices as the Stella Octangula by regarding it as having six faces of each of the two types shown in Fig. 5. Alternatively, regarding the second type as being four separate unpunctured faces, we could decide to puncture them in a similar way, so proceeding to the second stage of the Sierpinski sponge (Fig. 6). However no finite stage of this sponge is a true holyhedron.

Indeed, if we allow infinite polyhedra, there are more genuine examples. Figure 7 shows part of an infinite lattice of tetrahedra, with the property that each face of each tetrahedron is pierced by a vertex of another.

In the light of these examples, we give a precise statement of:

The Holyhedron Problem. Is there a polyhedron in Euclidean three-dimensional space that has only finitely many plane faces, each of which is a closed connected subset of the appropriate plane whose relative interior in that plane is multiply connected?

The rest of this Introduction discusses the necessity of most of the qualifications in this definition.

The Stella Octangula is an example in which the closed faces are multiply connected but their relative interiors are not. The lattice of interpenetrating tetrahedra is one with an infinite number of faces, and by rolling it up into a suitable three-dimensional flat torus, we can obtain an example in that space.

There are similar examples in three-dimensional spherical space. From the regular polytope 3,4,3, or “polyoctahedron,” we obtain a spherical tessellation that has 24 octahedral cells, and from its edges we can pick out the edges of 24 different cubes (they surround the 24 vertices). Each of these cubes contains two inscribed tetrahedra—choose one of these, and replace it by a smaller concentric tetrahedron on the sphere. In a similar way, each octahedral cell is inscribed in two different tetrahedra, again we replace one of these by a smaller one. Then there is an essentially unique way to make the choices so as to get a system of 48 tetrahedra on the 3-sphere that (if their sizes are suitably chosen) will interpenetrate in a manner analogous to Fig. 7. This yields a “spherical holyhedron” with 192 faces and genus 145. I have reduced these numbers to 72 and 55 by another construction.

These examples show that the Euclidean nature of the containing space was an important feature of the problem, and this is confirmed by the following proof that this method of taking unions of tetrahedra is doomed to fail in Euclidean 3-space.

Assume for contradiction the existence of a holyhedron in three-dimensional Euclidean space obtained by arranging a finite number of tetrahedra so that the face of each is pierced by the vertex of another and each vertex pierces at most one face. Taking a support plane to the holyhedron, we find a vertex which is not used to pierce anything. However, there were no vertices to spare: one of the faces is unpierced. Contradiction.

This negative result does not prohibit examples obtained by taking unions of other polyhedra, but shows that in them we must either somehow use one vertex to pierce a plurality of faces, or use polyhedra that have more vertices than faces. The first option is difficult: if a vertex goes all the way through another polyhedron and comes out the

Fig. 1. One face is connected but not simply connected.

Fig. 2. Eight “faces” which are not simply connected.

Fig. 3. The previous figure has eight “faces” like this one.

Fig. 4. A cube with square holes drilled in each side.

Fig. 5. The second face above is unacceptable.

other side, it introduces new faces which are simply connected (see Fig. 8). Vinson has chosen the second option, and used it to construct the holyhedron with 78,585,627 faces and genus 60,380,421 that is described in the rest of this paper. He and I both expect that there will be an example with less than 100 faces, and propose the problems of finding the minimal number of faces and the minimal genus to our readers.

1. Outline of Construction

We construct a holyhedron by taking the union of a large but finite number of convex polyhedra, arranged in such a way that every face of every polyhedron is pierced by the vertex of another. We use each vertex to pierce at most one face. Since a support plane to the holyhedron contains an unused vertex, our collection of convex polyhedra must contain more vertices than faces. An excess of vertices over faces is thus essential to construction.

This suggests a strategy: begin with a disjoint collection of polyhedra for which the difference $(\# \text{ vertices}) - (\# \text{ faces})$ is large. Introducing tetrahedra, we preserve the difference but are able to make the ratio $(\# \text{ unused vertices})/(\# \text{ unpierced faces})$ large. By choosing the initial collection properly, we can make this ratio as large as we wish.

Fig. 6. The Sierpinski sponge (second iteration).

Fig. 7. A lattice of tetrahedra: every face is pierced.

We then show how to trade four unused vertices in an inconvenient location for one in a more convenient location. After trading three times at the 4:1 ratio, we can reach any of the unpierced faces. If the ratio of ($\#$ unused vertices)/($\#$ unpierced faces) is large enough at the beginning of trading, then all of the unpierced faces can be pierced. The result is a holyhedron.

To create the initial excess of $(\# \text{ vertices}) - (\# \text{ faces})$, we construct a *honeycomb*. Asymptotically, as the number of cells increases, there are twice as many vertices as faces (see Fig. 9). This honeycomb is obtained by taking a regular tetrahedron and replacing one of the faces with a hexagonal lattice which has been bowed out to make it slightly convex. The other three sides of the honeycomb appear triangular, but one “side” of each “triangle” is actually a union of several line segments.

Although unnecessary for the construction, we note that for convex polyhedra the asymptotic ratio of two vertices for every face is optimal. Since each vertex is on at least three edges and each edge contains exactly two vertices, we have $E \geq \frac{3}{2}V$. Manipulating Euler’s formula for convex polyhedra, $V - E + F = 2$, we obtain the inequality $V \leq 2F - 4$. As a corollary, our method cannot hope to construct a holyhedron in which each face has two holes.

We now take two honeycombs and place them so that they oppose each other. In between the two honeycombs, we construct a forest of tetrahedra which, except for edge

Fig. 8. The piercing vertex goes all the way through its target, introducing three new unpierced faces.

Fig. 9. A honeycomb-shaped polyhedron.

effects, pierces each face of the two honeycombs. Except for edge effects, each face of each tetrahedron in the forest is pierced by a vertex from either a neighboring tetrahedron or from one of the two honeycombs.

Let M be the number of faces on one of the honeycombs. Then the number of unpierced faces is about $C_1\sqrt{M}$ —these are due to edge effects. Since introducing tetrahedra preserves the excess of unused vertices over unpierced faces, there are about $2M$ unused vertices. Half of these unused vertices will be used in our construction; we prove that the other half do not get in the way and forget about them.

We construct the forest of tetrahedra in such a way that there are corridors through the forest: empty cylinders of nonzero thickness independent of M which pass through the forest without intersecting any tetrahedra. There are about $C_2\sqrt{M}$ parallel corridors. It is possible for most of the M unused vertices to lean into a nearby corridor. Four nearby unused vertices leaning into the same corridor pierce the sides of a knife-shaped tetrahedron (a *knife*) which is very long, has very small height, and is much skinnier than it is tall. Three of the vertices of this knife are of no use. The fourth vertex which we call the *tip* is far away from the honeycombs and forest of tetrahedra and is far away roughly in the direction of the corridors.

It is possible for many knives to exit the forest through the same corridor, and so we have about $M/4$ tips far away. These tips are still slightly maneuverable, and it is possible to group them in fours and trade four tips for knives which return toward the honeycombs and forest of tetrahedra. By trading once more at this 4:1 ratio, we obtain about $M/64$ tips, each of which can reach any of the unpierced faces of the honeycombs or the forest of tetrahedra. Since there are at most $C_1\sqrt{M}$ unpierced faces, we can construct a polyhedron when M is sufficiently large.

2. Construction of Honeycombs

Let $\vec{u} = (0, -1)$ or $(0, -1, 0)$ depending on the context and $\vec{v} = (\sqrt{3}/2, 1/2)$ or $(\sqrt{3}/2, 1/2, 0)$. We begin with the vertices of a hexagonal lattice in the plane $z = 0$, and

then project these vertices to the parabola $z_0(x, y) = -\varepsilon(x^2 + y^2)$:

$$\{i\vec{u} + j\vec{v} + (0, 0, z_0(x, y)) \mid i + j \equiv 1 \text{ or } 2 \pmod{3}\}.$$

The vertices on the parabola corresponding to hexagonal cells turn out to be coplanar: these are the vertices of hexagonal faces. Take the region in R^3 below this multifaceted approximation of a parabola, and intersect it with the half-spaces

$$\begin{aligned} (x, y, z) \cdot \left(\frac{\sqrt{3}}{2}, \frac{1}{2}, -1 \right) &< \frac{3N+1}{2}, \\ (x, y, z) \cdot \left(-\frac{\sqrt{3}}{2}, \frac{1}{2}, -1 \right) &< \frac{3N+1}{2}, \\ (x, y, z) \cdot (0, -1, -1) &< \frac{3N+1}{2}. \end{aligned}$$

For ε sufficiently small (depending on N), the resulting polyhedron will resemble a tetrahedron, except that the top “face” will actually consist of $(3N+1)(3N+2)/2$ faces, of which $(3N-2)(3N-1)/2$ are hexagons. The bottom three faces appear triangular, but in fact the top “edge” of each is a union of $(3N+1)$ nearly colinear line segments.

We call the polyhedron just constructed the *bottom honeycomb* (see Fig. 10). We now construct the *top honeycomb* to oppose the bottom honeycomb. We shift the hexagonal lattice and project to the parabola $z_1(x, y) = 1 + \varepsilon(x^2 + y^2)$ instead of $z_0(x, y) = -\varepsilon(x^2 + y^2)$:

$$\{i\vec{u} + j\vec{v} + (0, 0, z_1(x, y)) \mid i + j \equiv 0 \text{ or } 2 \pmod{3}\}.$$

Fig. 10. Top view of the two honeycombs for $N = 2$ and $\varepsilon = 0$. The top honeycomb is drawn with dashed lines.

We intersect the region above the resulting faces with the half-spaces

$$\begin{aligned} (x, y, z - 1) \cdot \left(\frac{\sqrt{3}}{2}, \frac{1}{2}, 1 \right) &< \frac{3N}{2}, \\ (x, y, z - 1) \cdot \left(-\frac{\sqrt{3}}{2}, \frac{1}{2}, 1 \right) &< \frac{3N}{2}, \\ (x, y, z - 1) \cdot (0, -1, 1) &< \frac{3N}{2}. \end{aligned}$$

This polyhedron also resembles a tetrahedron, except that its bottom “face” really consists of $(3N)(3N + 1)/2$ faces, of which $(3N - 3)(3N - 2)/2$ are hexagons.

Figure 10 is a top view of the top and bottom honeycombs for $N = 2$ and $\varepsilon = 0$ (omitting the three “backside” faces of each). The bottom honeycomb is drawn with solid lines; the top honeycomb is drawn with dashed lines.

The parameters N and ε have yet to be specified.

3. Construction of a Forest

Since the faces of the honeycombs are simply connected, we must introduce polyhedra which pierce these faces. We pierce these faces with tetrahedra, which are driven so deeply into their target that their base triangle just barely remains above the surface. For the pairs $i + j \equiv 0 \pmod{3}$ which correspond to hexagonal faces on the bottom honeycomb, the tetrahedron which pierces that hexagonal face has coordinates

$$\begin{aligned} i\vec{u} + j\vec{v} + \left\{ (0, 0, \frac{-1}{2(1-\gamma)}), -\frac{2}{3} \left(1 + \frac{8\delta(1-\gamma)}{\gamma} \right) \vec{u} + \left(0, 0, \frac{4\delta}{\gamma} \right), \right. \\ \left. -\frac{2}{3} \left(1 + \frac{8\delta(1-\gamma)}{\gamma} \right) \vec{v} + \left(0, 0, \frac{4\delta}{\gamma} \right), \right. \\ \left. \frac{2}{3} \left(1 + \frac{8\delta(1-\gamma)}{\gamma} \right) (\vec{u} + \vec{v}) + \left(0, 0, \frac{4\delta}{\gamma} \right) \right\}. \end{aligned}$$

For those tetrahedra piercing hexagons on the top honeycomb (such hexagons correspond to pairs (i, j) with $i + j \equiv 1 \pmod{3}$), the coordinates of the vertices are

$$\begin{aligned} i\vec{u} + j\vec{v} + \left\{ \left(0, 0, 1 + \frac{1}{2(1-\gamma)} \right), \frac{2}{3} \left(1 + \frac{8\delta(1-\gamma)}{\gamma} \right) \vec{u} + \left(0, 0, 1 - \frac{4\delta}{\gamma} \right), \right. \\ \left. \frac{2}{3} \left(1 + \frac{8\delta(1-\gamma)}{\gamma} \right) \vec{v} + \left(0, 0, 1 - \frac{4\delta}{\gamma} \right), \right. \\ \left. -\frac{2}{3} \left(1 + \frac{8\delta(1-\gamma)}{\gamma} \right) (\vec{u} + \vec{v}) + \left(0, 0, 1 - \frac{4\delta}{\gamma} \right) \right\}. \end{aligned}$$

We call these *bottom red tetrahedra* and *top red tetrahedra*, respectively. See Figs. 11 and 12. The parameters δ and γ have yet to be specified. The reader may verify

that, for fixed N and $0 < \gamma < 1$, the red tetrahedra are disjoint and will pierce the appropriate hexagonal faces if ε and δ are sufficiently small. If necessary, we truncate the red tetrahedra so that they do not poke all the way through the honeycomb. Ignoring edge effects for now, all $2M$ of the honeycomb faces have been pierced.

Each time we pierce a hexagonal face with a red tetrahedron, we introduce four new faces which require piercing: three for the sides of the red tetrahedron and one for its base. To pierce these additional faces, we introduce *green* tetrahedra, the bases of which are pierced by vertices of the honeycomb corresponding to a pair (i, j) . The coordinates of the vertices of the bottom green tetrahedra, with $i + j \equiv 1 \pmod{3}$, are

$$i\vec{u} + j\vec{v} + \left\{ (0, 0, 1 + 3\delta), -\frac{2}{3}(1 + 3\delta)\vec{u} + \left(0, 0, z_0(x, y) + \frac{7\varepsilon}{18}\right), \right. \\ \left. -\frac{2}{3}(1 + 3\delta)\vec{v} + \left(0, 0, z_0(x, y) + \frac{7\varepsilon}{18}\right), \right. \\ \left. \frac{2}{3}(1 + 3\delta)(\vec{u} + \vec{v}) + \left(0, 0, z_0(x, y) + \frac{7\varepsilon}{18}\right) \right\}.$$

The coordinates for the top green tetrahedra, with $i + j \equiv 0 \pmod{3}$, are

$$i\vec{u} + j\vec{v} + \left\{ (0, 0, -3\delta), \frac{2}{3}(1 + 3\delta)\vec{u} + \left(0, 0, z_1(x, y) - \frac{7\varepsilon}{18}\right), \right. \\ \left. \frac{2}{3}(1 + 3\delta)\vec{v} + \left(0, 0, z_1(x, y) - \frac{7\varepsilon}{18}\right), \right. \\ \left. -\frac{2}{3}(1 + 3\delta)(\vec{u} + \vec{v}) + \left(0, 0, z_1(x, y) - \frac{7\varepsilon}{18}\right) \right\}.$$

Except for edge effects, all $8M$ of the red faces have been pierced. The side faces of the bottom red tetrahedra are pierced by the base vertices of bottom green tetrahedra. The base face of a bottom red tetrahedron is pierced by the tip vertex of a top green tetrahedron. The faces of the top red tetrahedra are pierced similarly by green tetrahedra. Since the bases of the green tetrahedra are pierced by vertices of the honeycomb, only $6M$ green faces remain to be pierced.

On each honeycomb, the vertices used to pierce the bases of green tetrahedra constitute only one of two cosets of vertices. Based on each vertex in the remaining coset, we add *blue* tetrahedra. The coordinates for the bottom blue tetrahedra, which correspond to $i + j \equiv 2 \pmod{3}$, are

$$i\vec{u} + j\vec{v} + \left\{ \text{tip vertex}, -\frac{2}{3}(1 - \delta)\vec{u} + \left(0, 0, z_0(x, y) + \frac{7\varepsilon}{18}\right), \right. \\ \left. -\frac{2}{3}(1 - \delta)\vec{v} + \left(0, 0, z_0(x, y) + \frac{7\varepsilon}{18}\right), \right. \\ \left. \frac{2}{3}(1 - \delta)(\vec{u} + \vec{v}) + \left(0, 0, z_0(x, y) + \frac{7\varepsilon}{18}\right) \right\}.$$

The coordinates for the top blue tetrahedra, also corresponding to $i + j \equiv 2 \pmod{3}$, are

$$i\vec{u} + j\vec{v} + \left\{ \begin{array}{l} \text{tip vertex, } \frac{2}{3}(1 - \delta)\vec{u} + \left(0, 0, z_1(x, y) - \frac{7\varepsilon}{18}\right), \\ \frac{2}{3}(1 - \delta)\vec{v} + \left(0, 0, z_1(x, y) - \frac{7\varepsilon}{18}\right), \\ -\frac{2}{3}(1 - \delta)(\vec{u} + \vec{v}) + \left(0, 0, z_1(x, y) - \frac{7\varepsilon}{18}\right) \end{array} \right\}.$$

We did not specify the coordinates of the tip vertices of the blue tetrahedra because there is a lot of flexibility in where we position the tip. We first say the properties we would like the blue tetrahedra to have, and then say what the admissible locations for the tip are. Since the desired properties and sufficient conditions are so similar for the top and bottom tetrahedra, we only list them for the bottom tetrahedra:

- (1) We would like the base faces of the bottom blue tetrahedra to be pierced by the honeycomb vertices.
- (2) We would like the base vertices of the bottom blue tetrahedra to pierce the sides of the bottom green tetrahedra.
- (3) We would like the side faces of the bottom blue tetrahedra to be pierced by the base vertices of the bottom red tetrahedra.

The first of these properties will hold if the tip vertex of the bottom blue tetrahedron is above the plane determined by its base. A sufficient condition for property (2) to hold when ε is sufficiently small is

$$\begin{aligned} (x, y, z) - i\vec{u} - j\vec{v} \cdot (-\vec{u} + (0, 0, -\tfrac{1}{3})) &< \tfrac{2}{3} - \delta, \\ (x, y, z) - i\vec{u} - j\vec{v} \cdot (-\vec{v} + (0, 0, -\tfrac{1}{3})) &< \tfrac{2}{3} - \delta, \\ (x, y, z) - i\vec{u} - j\vec{v} \cdot (\vec{u} + \vec{v} + (0, 0, -\tfrac{1}{3})) &< \tfrac{2}{3} - \delta. \end{aligned}$$

If we replaced the first inequality with equality, we would have the equation for the plane of the side face of the neighboring bottom green tetrahedron. Two of the base vertices of our bottom blue tetrahedron are on the “less than” side of this plane, the other base vertex is just barely on the “greater than” side. In order for this base vertex to pierce the side face of the bottom green tetrahedron, the tip vertex of the bottom blue tetrahedron must be on the “less than” side of the plane.

In our construction the first two properties of the blue tetrahedra are almost automatic, while the third requires care. A sufficient condition for the third property (that the red corners pierce the blue sides) to hold when ε is sufficiently small is

$$\begin{aligned} ((x, y, z) - i\vec{u} - j\vec{v}) \cdot \left(\vec{u} + \left(0, 0, \frac{16 - 17\gamma}{12}\right) \right) &> \frac{1 - \delta}{3}, \\ ((x, y, z) - i\vec{u} - j\vec{v}) \cdot \left(\vec{v} + \left(0, 0, \frac{16 - 17\gamma}{12}\right) \right) &> \frac{1 - \delta}{3}, \\ ((x, y, z) - i\vec{u} - j\vec{v}) \cdot \left(-\vec{u} - \vec{v} + \left(0, 0, \frac{16 - 17\gamma}{12}\right) \right) &> \frac{1 - \delta}{3}. \end{aligned}$$

The red corner with which we intend to pierce the blue face lies on one side or the other of the plane determined by the blue face: which side it is depends on whether the inequalities above are satisfied. Equivalently, the tip vertex of the bottom blue tetrahedron must lie in the intersection of three half-spaces.

Figure 11 is a top view of a portion of the forest of tetrahedra. The hexagonal lattice of the bottom honeycomb is shown in black. The solid colored triangles are the bases of the bottom tetrahedra. The dashed colored triangles are the bases of the top tetrahedra. The hexagonal lattice of the top honeycomb has been omitted for clarity.

A greatly simplified top view is shown in Fig. 12. Here we show only the bottom honeycomb and the bottom colored tetrahedra, with a particular choice of location for the tips of the bottom blue tetrahedra. The cross sections of these tetrahedra are shown at the two heights $z = 0$ and $z = 4\delta/\gamma$. The bottom level is drawn with solid lines; the top level is drawn with dashed lines. Observe that the red tetrahedron points down, so its cross section increases with height; the other tetrahedra point up. The slope of the sides of the red tetrahedra is almost twice that of the green tetrahedra; the slope of the blue tetrahedra is almost twice that of the red tetrahedra.

By examining Fig. 12 with a magnifying glass, one can just barely verify the following observations: the solid blue triangle pierces the solid green one, the solid green pierces the solid red, but the solid red does not pierce the solid blue. At the higher level (drawn with dashed lines), the dashed blue does not pierce the dashed green, the dashed green does not pierce the dashed red, but the dashed red does pierce the dashed blue. Although we have not indicated so in Fig. 12, the base triangle of the blue tetrahedron is slightly above (by an amount proportional to ε , which we choose to be much less than δ and γ) the base of the green tetrahedron. Thus the base vertices of blue tetrahedra pierce the sides of green tetrahedra; the base vertices of green tetrahedra pierce the sides of red tetrahedra; and the base vertices of red tetrahedra pierce the sides of blue tetrahedra.

Because the construction of the forest of tetrahedra is complicated, we refer the reader to the following table of what pierces what:

Type of face	Type of vertex which pierces it
Hexagon from honeycomb	Tip of red tetrahedron
Base of bottom red tetrahedron	Tip of top green tetrahedron
Side of bottom red tetrahedron	Base vertex of bottom green tetrahedron
Base of top red tetrahedron	Tip of bottom green tetrahedron
Side of top red tetrahedron	Base vertex of top green tetrahedron
Base of green tetrahedron	One coset of honeycomb vertices
Side of green tetrahedron	Base vertex of blue tetrahedron
Base of blue tetrahedron	Other coset of honeycomb vertices
Side of blue tetrahedron	Base vertex of red tetrahedron

Note that except for edge effects, all faces have been pierced. However, we now have $2M$ vertices which have not been used to pierce anything: the tips of the blue tetrahedra. We also have flexibility in where we place these tips. We only need half of these tips to

Fig. 11. Top view of a portion of the forest of tetrahedra, for $\delta = 0$. For clarity, only the bottom honeycomb and the bases of the colored tetrahedra are shown. The bases of the bottom colored tetrahedra are drawn with solid lines; the bases of the top colored tetrahedra are drawn with dashed lines. Refer to Figure 12 for the colors of the bottom tetrahedra in relation to the bottom honeycomb. Above bottom red tetrahedra are top green tetrahedra; above blue are blue; above green are red.

Fig. 12. Two cross sections of the bottom colored tetrahedra. The bottom level is drawn with solid lines; the top level with dashed lines. For this figure we used the parameters $\varepsilon = 0$, $\delta = 0.02$, and $\gamma = 0.3$. The central tetrahedron is red, the upper left tetrahedron is green, and the upper tetrahedron is blue.

continue: those of the bottom blue tetrahedra. To get the top blue tetrahedra out of the way, we chose their tips to make them as small as possible:

$$i\vec{u} + j\vec{v} + \left(0, 0, 1 - \frac{1 - \delta}{4 - 5\gamma}\right).$$

The parameters δ , γ , ε , and N have yet to be specified.

4. Existence of Corridors

There are corridors through the forest in the x -direction (there are also corridors in the $(-1/2, \sqrt{3}/2)$ and $(-1/2, -\sqrt{3}/2)$ directions, but these are not used in the construction). These corridors are apparent in Fig. 13, which is a projection of the forest of tetrahedra onto the yz plane. Also indicated are the allowed positions for the tips of the bottom blue tetrahedra. The figure shows that bottom blue tetrahedra on both sides of a corridor can lean into the corridor, and that there is a subcorridor whose entire cross section is available to the tips of bottom blue tetrahedra on both sides.

The only things which might block these corridors are the bottom blue tetrahedra; we must be careful when positioning these tips. In the next section we show how to position the tips inside the corridors to pierce the sides of knives. These knives will exit the forest through the corridors.

Fig. 13. A projection of the forest of tetrahedra onto the yz plane, with $\gamma = 0.025$ and $\delta = 0.0001$. The triangles with solid lines are the projections of the colored tetrahedra. The green tetrahedra between the two honeycombs resemble interlocking teeth; the red tetrahedra pierce into the honeycombs and point away from the intermediate region. The rectangular shaded areas lie above the bases of the bottom blue tetrahedra; the shaded regions above these bases outlined with dashed lines are the allowed locations for their tips. There is a corridor in the x -direction, indicated by the black spot, which is unobstructed and which is accessible to the tips of bottom blue tetrahedra from both the left and the right.

Fig. 14. A knife.

5. Knives Exiting Through Corridors

A *knife* is a tetrahedron which is very long, has very small height, and is much skinnier than it is tall (Fig. 14).

The following lemma describes sufficient conditions for constructing a knife in a cylinder and piercing it by the maneuverable tips of tetrahedra leaning into the cylinder. The reader may find it helpful to look at Fig. 15 while reading the statement of the lemma.

Lemma 5.1 (Knives Based in Cylinders). *Suppose a cylinder of finite length is divided into left and right ends by a plane perpendicular to the axis of the cylinder. Suppose that there are four triangles T_i in space (so adding a fourth vertex p_i to each would result in four tetrahedra). Suppose that there is flexibility in where these fourth vertices are placed: for each $i = 1, \dots, 4$ there is an entire cross section of the left end of the cylinder of allowable positions for p_i . Suppose that no matter how we choose each p_i from the i th cross section, the tetrahedra $T_i p_i$ are disjoint.*

Suppose also that the triangles T_i satisfy the slab separation conditions. By a slab we mean the region between two parallel supporting planes of the cylinder. Suppose that two of the triangles are on opposite sides of a slab, and that there is slab that all triangles are on the same side of.

Then for any specified point in the right end of the cylinder, we can construct a knife whose tip is at the specified point and whose remaining vertices are in the left end of the cylinder. We can find p_i so that each tetrahedron $T_i p_i$ pierces a distinct face of this knife.

Fig. 15. With respect to the cylinder, the four triangular bases satisfy the slab separation conditions required by Lemma 5.1.

Fig. 16. The knives in lower subcorridors are supported by bottom blue tetrahedra closer to the edge of the forest. For clarity, this diagram is not to scale: the corridor is not nearly so thick. The diagram shows only one blue tetrahedron supporting each knife when there are actually four.

Proof. See Fig. 15. □

It is possible for many knives to exit the forest through the same corridor; this is limited only by the number of bottom blue tetrahedra available to support the knives. When placing multiple knives in the same corridor, one must take care that knives do not intersect each other or the tetrahedra supporting other knives. One way to accomplish this is by dividing the corridor vertically into subcorridors, with one knife in each subcorridor. The knives in lower subcorridors are supported by tetrahedra closer to the edge of the forest. See Fig. 16.

By allowing several knives to use the same corridor, it is possible for $\lfloor (3N - 2)/2 \rfloor (\lfloor (3N - 2)/2 \rfloor + 1)/2$ knives to exit the forest. See Fig. 17.

Fig. 17. The bottom blue tetrahedra lean into nearby corridors (shown with dotted lines). We use each group of four (indicated by parallelograms) to pierce the four sides of a knife.

6. Knives Returning to the Forest

Now suppose that the tips of four knives can all be made to meet at a point p , and can in fact be maneuvered in a neighborhood of p . In the next lemma we describe conditions under which the tips of these four knives can be “traded” for a new knife based near the point p , and how maneuverable this new knife is.

Lemma 6.1. *Suppose that there are four convex cones, disjoint except that they share a vertex at $p_1 = p_2 = p_3 = p_4 = p$. Draw a small sphere around p which contains no part of the bases of any of the four cones. Then each cone can be associated with a closed convex region on the sphere.*

Let the unit vector $\vec{n}$ denote a direction. Suppose that there is a great circle through $\vec{n}$ so that all four regions are in the same hemisphere, and suppose that there is a great circle through $\vec{n}$ with two of the regions in opposite hemispheres.

Now suppose that each vertex p_i is allowed to move in a neighborhood of p . Then, for any $\lambda > 0$, there is a knife with its tip at $p + \lambda\vec{n}$ and with its base triangle in an arbitrarily small neighborhood of p . The tips p_i of the four cones can be maneuvered so that the four cones are disjoint and so that the four tips pierce the four sides of the new knife.

Proof. We construct a cylinder to which the previous lemma applies. We choose its axis to be the line determined by p and $p + \lambda\vec{n}$, and require that the cylinder be long enough to contain a neighborhood of both these points.

We can now pick $\tilde{p}_i = p + \lambda_i\vec{n}$ (all in the allowed neighborhood of p and all on the line segment from p to $p + \lambda\vec{n}$) so that the four cones are disjoint (Adrian Banner helped me find a proof that this is possible, which I omit). We then pick neighborhoods of the $\tilde{p}_i$ so that no matter how the p_i are chosen in these neighborhoods, the cones are disjoint.

The hemispheres in the assumptions of this lemma determine two planes, containing the axis of the cylinder, which respectively separate and unify the bases of the cones. The planes can only be thickened slightly without destroying the separating or unifying properties. The radius of the cylinder is also limited by requiring that the allowed neighborhoods around the $\tilde{p}_i$ include a full cross section of the cylinder.

We have shown that if the cylinder is sufficiently narrow, the hypotheses of Lemma 5.1 are satisfied. The picture-proof of Lemma 5.1 now applies. $\square$

We are mostly concerned with the case of trading four knife tips for a new knife, with the four knives coming from roughly the same direction. Consider those directions for the new knife to point which are in the same hemisphere as the source of the original knives. We may project this hemisphere from the center of the sphere onto the plane. We ask what points on the plane correspond to directions in which the new knife can point (see Fig. 18). Those points q are allowable which are outside of the convex hull of the regions and where there is a line through q which separates two of the regions. In this example, the set of excluded directions is very small, and so the tips of the new knives are very maneuverable.

Fig. 18. The region where the returning knife is allowed to point is lightly shaded. The regions corresponding to the four cones supporting it are the heavily shaded triangles.

7. Dealing with Edge Effects

In our construction of the honeycombs and the forest of tetrahedra, we claimed that all of the faces were pierced “except for edge effects.” Now that we have knives leaving the forest through corridors and are able to trade their tips for new knives, we are in a position to pierce the faces which, because of edge effects, have not been pierced yet.

First, we must know exactly which faces are yet unpierced. Please refer to Figs. 10 and 11. We claim that the following is the count of faces which need to be pierced:

$3(3N + 1)$	bottom red side faces,
$3(3N)$	bottom red base faces,
$3(3N)$	bottom green side faces,
$3(6N - 3)$	top blue side faces,
$3(3N - 1)$	top honeycomb faces (truncated hexagons),
6	backside faces of honeycombs.

All of these unpierced faces are visible from infinity. For this reason, they can be pierced by trading exactly twice for more maneuverable knives.

We now estimate how many of these maneuverable knives are available. Each such knife is the result of trading twice at the 4:1 ratio beginning with knives exiting the forest. We do not lose much by assuming that, for each maneuverable knife, the 16 original knives traded for it exited the forest through the same corridor.

Each successive corridor supports one knife more than its neighboring corridor. The biggest 16 corridors each support at least $\lfloor (3N - 2)/2 \rfloor - 15$ knives, which support $\lfloor \lfloor (3N - 32)/2 \rfloor / 16 \rfloor$ maneuverable knives. The next 16 biggest corridors each support $\lfloor \lfloor (3N - 32)/2 \rfloor / 16 \rfloor - 1$ maneuverable knives, and so on. The total number of

maneuverable knives is at least

$$8 \left\lfloor \frac{3N - 32}{32} \right\rfloor \left\lfloor \frac{3N}{32} \right\rfloor.$$

For N large enough, the number of knives available for piercing faces will exceed the number of faces which are unpierced because of edge effects. Using the estimates of this paper, it suffices to choose N so that

$$8 \left\lfloor \frac{3N - 32}{32} \right\rfloor \left\lfloor \frac{3N}{32} \right\rfloor \geq 54N - 3. \quad (1)$$

8. Doing the Steps in the Right Order

The construction of the honeycombs and the forest of tetrahedra depends on several parameters which we have introduced, and we must be careful about the order in which we specify these parameters. The parameters we introduced are: N , which is proportional to the sidelength of the honeycombs; ε , which describes the curvature of the honeycombs; δ , the amount by which the colored tetrahedra have been enlarged; and γ so that the red tetrahedra lean $2(1 - \gamma)$ times more steeply than the green tetrahedra.

We first choose $N = 800$, which is sufficiently large because of (1). As long as ε , δ , and γ are positive and sufficiently small, the honeycombs and the forest of tetrahedra will have the properties we claimed as far as what pierces what. For now we pick $\varepsilon = \delta = \gamma = 0$, so that the honeycombs and the forest of tetrahedra are degenerate.

Now observe that there are corridors through the forest. Pick δ and γ positive but sufficiently small that the corridors are not completely blocked. At this point the honeycombs and forest of tetrahedra are still degenerate because $\varepsilon = 0$.

We now draw line segments where we want the knives to go. Since each unpierced face (or each area which will become an unpierced face once $\varepsilon > 0$) is visible from infinity, have a ray go from just underneath the surface of the unpierced face to infinity without otherwise intersecting the honeycomb or the forest. Since we can perturb the angle slightly, we can assume that this ray is not in the positive or negative x -direction (this is the direction of the corridors through which knives will exit the forest). Now decide which rays correspond to which corridors (see Fig. 17), with the restriction that if k bottom blue tetrahedra are adjacent to a corridor, then at most $\lfloor k/64 \rfloor$ rays may be associated to that corridor.

For each corridor do the following: Order the rays that correspond to it. Count the rays and divide the column vertically into that many subcorridors. Associate 64 bottom blue tetrahedra to each subcorridor, with the bottom blue tetrahedra closest to the edge of the forest corresponding to the lowest subcorridors (as suggested by Fig. 16). For each subcorridor, we intend to obtain the corresponding ray by beginning with 16 line segments (each supported by four bottom blue tetrahedra) exiting the forest through the subcorridor and trading twice at the 4:1 ratio. We accomplish this by beginning with the ray and working backward:

For each ray, pick a point p on it far enough out so that its projection to the yz plane is outside the yz projection of the honeycomb and forest. Now pick four line segments from

p which do not intersect the honeycombs or forest and whose endpoints q_1, q_2, q_3, q_4 are four points in the appropriate subcorridor with distinct z -coordinates. Choose the q_i so that given pq_i , the direction of the ray is allowed in the sense of Lemma 6.1. Assume the q_i have distinct z -coordinates. Divide the subcorridor further vertically, with one q_i in each subsubcorridor, with 16 bottom blue tetrahedra for each subsubcorridor. As usual, bottom blue tetrahedra closest to the edge of the forest are assigned to the lowest subsubcorridors.

We now wish to pick for each q_i four points $r_{i,j}$ in the forest and in the subsubcorridor so that, given the directions $q_i r_{i,j}$, the direction $q_i p$ is allowed in the sense of Lemma 6.1. We associate to each line segment $q_i r_{i,j}$ a group of four bottom blue tetrahedra as in Fig. 17. We choose this correspondence so that the lower (i.e., smaller z -coordinate of $r_{i,j}$) line segments correspond to groups of tetrahedra closer to the edge of the forest. We also choose the $r_{i,j}$ so that the four bottom blue tetrahedra and a sufficiently skinny cylinder around the line segment $q_i r_{i,j}$ satisfy the hypotheses of Lemma 5.1 and that the x -coordinates of the $r_{i,j}$ are sufficiently large that they do not interfere with the bottom blue tetrahedra corresponding the next higher subsubcorridor.

It may be the case that some of the line segments we introduced intersect in an unintended way. If so, then perturb the line segments slightly so that they do not intersect. We now build cylinders around the $q_i r_{i,j}$ which are sufficiently small that none of these cylinders intersect except in a small neighborhood of the q_i and so that the hypotheses of Lemma 5.1 are still satisfied. Now construct knives in these cylinders and have bottom blue tetrahedra pierce their faces, but leave the tip free to move in a small neighborhood of q_i . Now fatten up cylinders around the $q_i p$ in a similar fashion, introducing no unintentional intersections between cylinders and so that the conditions of Lemma 6.1 are satisfied. Build knives, allowing the tips to move in a neighborhood of p and so that at point p the conditions of Lemma 6.1 are satisfied. Now build a knife from p which pierces the intended face.

Finally, choose ε sufficiently small that none of the piercings we constructed are destroyed and no unintended intersections are introduced. We are done.

Note that there is a little bit of freedom in the ordering of these steps. For instance, we could have chosen ε sufficiently small before fattening up the line segments into cylinders.

9. New Questions

The holyhedron we constructed has 78,585,627 faces:

Description	Number of polyhedra when $N = 800$	Number of faces when $N = 800$
$\frac{(3N+1)(3N+2)}{2} + 3$ faces on bottom honeycomb	1	2,883,604
$\frac{3N(3N+1)}{2} + 3$ faces on top honeycomb	1	2,881,203
$\frac{(3N+1)(3N+2)}{2}$ bottom red tetrahedra	2,883,601	11,534,404

$\frac{3N(3N+1)}{2}$ bottom green tetrahedra	2,881,200	11,524,800
$\frac{(3N-1)3N}{2}$ bottom blue tetrahedra	2,878,800	11,515,200
$\frac{(3N-3)(3N-2)}{2}$ top red tetrahedra	2,874,003	11,496,012
$\frac{(3N-2)(3N-1)}{2}$ top green tetrahedra	2,876,401	11,505,604
$\frac{(3N-1)3N}{2}$ top blue tetrahedra	2,878,800	11,515,200
$128 \left\lfloor \frac{3N-32}{32} \right\rfloor \left\lfloor \frac{3N}{32} \right\rfloor$ knives exiting forest	710,400	2,841,600
$32 \left\lfloor \frac{3N-32}{32} \right\rfloor \left\lfloor \frac{3N}{32} \right\rfloor$ first level knives	177,600	710,400
$8 \left\lfloor \frac{3N-32}{32} \right\rfloor \left\lfloor \frac{3N}{32} \right\rfloor$ second level knives	44,400	177,600
Grand total	18,205,207	78,585,627

We know that there are holyhedra with fewer faces: there are several places in the construction where we could improve the estimates at the expense of working harder and making the explanation less clear. For example, some of the unpierced faces (due to edge effects) could have been pierced after two trades instead of three.

A particularly clever improvement, suggested by Doug McIlroy, allows trading at the ratio 3:1 instead of 4:1. Instead of piercing a face with the tip of a knife, let a *horn* grow out of the face to be pierced. Like the knife, the horn is very long, has very small height, and width much less than its height. Unlike the knife, it has three faces that need to be pierced instead of four.

This begs the question: What is the minimum number of faces that a holyhedron can have? We would not be surprised if there is a holyhedron with less than 100 faces, but it almost certainly would not come from our construction. A related question is: What is the minimum genus of a holyhedron? The one we constructed has genus 60,380,421. (The genus is one minus the number of convex polyhedra plus the number of piercings; and the number of piercings is exactly the number of faces.)

We would have liked to build a three-dimensional model of a holyhedron, say out of posterboard. However, this is infeasible because of the number of faces our construction has. Even if we could build such a model, it would fail to be convincing [3]. There are many places in our construction where a slight perturbation would destroy the property of being a holyhedron. It would be interesting if someone were able to discover and build a holyhedron which not only has a small number of faces but for which a three-dimensional model is convincing.

Acknowledgment

I would like to thank Doug McIlroy for reading a draft of the paper suggesting ways to clarify the presentation, and John Conway for writing the Introduction.

References

1. Coxeter, H. S. M. *Regular Polytopes*, pp. 48–51. Dover, New York, 1973.
2. Kepler, J. *Harmonice Mundi*. Opera Omnia: Vol. 5. Frankfurt, 1864.
3. Lakatos, I. *Proofs and Refutations*. Cambridge University Press, Cambridge, 1976.

Received June 25, 1999, and in revised form February 21, 2000. Online publication May 8, 2000.