The Poincaré Dodecahedral Space and the Mystery of the Missing Fluctuations

Jeffrey Weeks

Introduction

Because of the finite speed of light, we see the Moon as it was roughly a second ago, the Sun as it was eight minutes ago, other nearby stars as they were a few decades ago, the center of our Milky Way Galaxy as it was 30,000 years ago, nearby galaxies as they were millions of years ago, and distant galaxies as they were billions of years ago. If we look still deeper into space, we see all the way back to the final stages of the big bang itself, when the whole universe was filled with a blazing hot plasma similar to the outer layers of the modern Sun. In principle we see this plasma in all directions; it fills the entire background of the sky. So why don't we notice it when we look up at the night sky? The catch is that the light from it—originally visible or infrared—over the course of its 13.7 billion-year voyage from the plasma to us has gotten stretched out as part of the overall expansion of the universe. Specifically, the universe has expanded by a factor of about 1100 from then till now, so what was once a warm reddish glow with a wavelength around 10,000 angstroms is now a bath of microwaves with a wavelength of about a millimeter. So we cannot see the plasma with our eyes, but we can see it with a microwave antenna.

If our eyes were sensitive to microwaves as well as to visible light, a close-up view of the night sky might look something like Figure 1. In the foreground we see other galaxies as they were a few billion years ago. In the background we see the omnipresent plasma as it was 380,000 years after the big bang, a mere three one-thousandths of one percent of the universe's present age of 13.7 billion years. The plasma holds clues to the universe's birth, evolution, geometry, and topology. To harvest these clues, NASA launched the Wilkinson Microwave Anisotropy Probe (WMAP) on 30 June 2001. On 10 August 2001 WMAP reached its orbit about the so-called second Lagrange point, where the combined gravity of the Sun and the Earth are just right to keep the satellite orbiting the Sun in synchronization with the Earth, and WMAP began its four years of observations (Figure 2). The observed radiation from the plasma, known as the Cosmic Microwave Background (CMB) radiation, is extremely uniform across the sky. Nevertheless, it exhibits small temperature fluctuations on the order of 1 part in 10^5. These CMB temperature fluctuations result from fluctuations in the density (not temperature!) of the primordial plasma: photons arriving from denser regions do a little extra work against gravity and arrive slightly cooler, while photons arriving from less dense regions do a little less work against gravity and arrive slightly warmer. So in effect temperature fluctuations on the microwave sky reveal density fluctuations in the early universe.

Cosmologists' current standard model posits an essentially infinite Euclidean space created by inflation and containing density fluctuations on all scales. On small scales WMAP observed these fluctuations as predicted. However, on scales larger than about 60° degrees across the sky, the

Jeffrey Weeks is a MacArthur Fellow. He can be reached at `http://www.geometrygames.org/contact.html`.

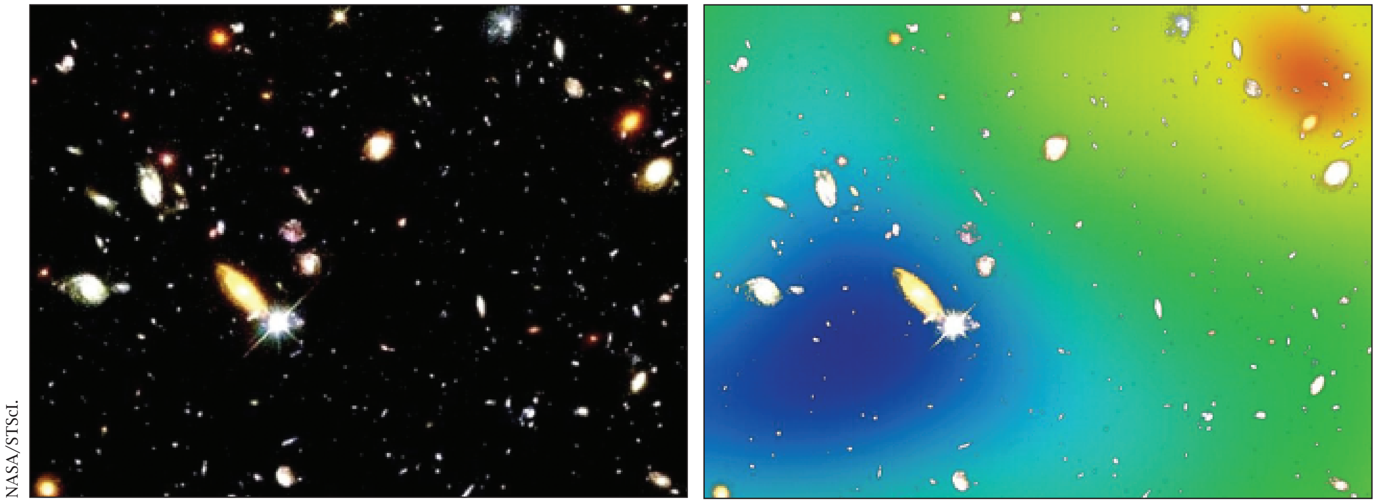

Figure 1. The primordial plasma provides a background for the whole sky. However, we see it not in the visible spectrum but in microwaves. The left panel shows a portion of the Hubble Deep Field image as we see it with our eyes. The right panel simulates what we might see if our eyes were sensitive to microwaves as well as to visible light.

fluctuations all but disappear,[1] leaving cosmologists with the Mystery of the Missing Fluctuations.

Finite Universe

One possible explanation for the missing large-scale fluctuations is that the universe is simply not big enough to support them. To take a simple 1-dimensional analogy, an infinite line supports waves of all sizes, while the circumference of a unit circle supports no wavelength longer than 2π. Similarly, if the real universe is a closed 3-manifold, it can support no waves longer than its own "circumference".

What 3-manifolds shall we consider? Observational evidence implies the observable universe is homogeneous and isotropic to a precision of one part in 10^4, so consider manifolds that locally look like the 3-sphere S^3, Euclidean space E^3, or hyperbolic space H^3. To construct a finite universe, take the quotient X/Γ of the simply connected space $X = S^3, E^3$, or H^3 under the action of a discrete fixed point free group Γ of isometries.

Observational data suggest the observable universe either is flat or has a small curvature that is more likely positive than negative. More precisely, on a scale where $\Omega < 1$ indicates a hyperbolic universe, $\Omega = 1$ indicates a flat universe, and $\Omega > 1$ indicates a spherical universe, analysis of the WMAP data yields $\Omega = 1.02 \pm 0.02$ at the 1σ level [2]. The parameter Ω measures the average mass-energy density of space, which by general relativity completely determines the spatial curvature, with low density yielding a hyperbolic universe and high density yielding a spherical universe. By definition Ω is the ratio of the actual density to the critical density that a flat universe would require.

The huge sphere on which we observe the primordial plasma is called the *horizon sphere*, and its radius in the modern universe is our *horizon radius*. If we live in a finite universe X/Γ, then our *injectivity radius* is the radius of the smallest Earth-centered sphere that "reaches all the way around the universe" and intersects itself. Equivalently, the injectivity radius is the radius of the largest Earth-centered sphere whose interior is embedded. Twice the injectivity radius is thus the minimal circumference of the universe, starting from Earth. If our horizon radius exceeds our injectivity radius, then we can trace two different lines of sight to the same distant region of space, meaning that in principle we see multiple images of the same astronomical sources (galaxies, quasars, plasma,…) in different parts of the sky, making detection of the universe's topology vastly easier.

A flat manifold E^3/Γ allows no a priori relationship between the horizon radius and the injectivity radius, because the latter is essentially arbitrary. Given any proposed Euclidean group Γ, we may easily stretch or shrink its translational components, via a similarity, to obtain any prescribed injectivity radius. We therefore have no reason to expect the injectivity radius to be comparable to our horizon radius. Thus successfully detecting a flat topology E^3/Γ would require a huge amount of luck. In spite of the long a priori odds against it, the possibility of a finite flat universe E^3/Γ continues to receive a fair amount of attention, because WMAP's $\Omega = 1.02 \pm 0.02$ result

[1] *Spergel et al. report the missing fluctuations in one of a series of papers released along with the first-year WMAP data [1]. Their Figure 16 compares predictions to observations, showing, in their words, "the lack of any correlated signal on angular scales greater than 60 degrees." More conservative observers point out that the signal is not totally missing, but merely very weak. Spergel et al. estimate the probability that such a weak signal could arise by chance to be either 0.0015 or 0.003, depending on which of their best-fit flat space models one compares to.*

Figure 2. A million miles from Earth, the Wilkinson Microwave Anisotropy Probe observes deep space from near the second Lagrange point, L2. In the (rotating) Sun-Earth coordinate system, the *Lagrange points* are the critical points of the effective gravitational potential. The second Lagrange point is a saddle point, unstable in the radial direction but stable in both the horizontal and vertical tangential directions. The satellite traces a gentle orbit about L2 in the vertical plane, with a slight nudge every few months to keep it from drifting towards or away from the Sun. The backward-facing solar panels support a large protective disk blocking microwave interference from the Sun, Earth, and Moon, giving the outward-facing microwave receivers an unobstructed view of deep space.

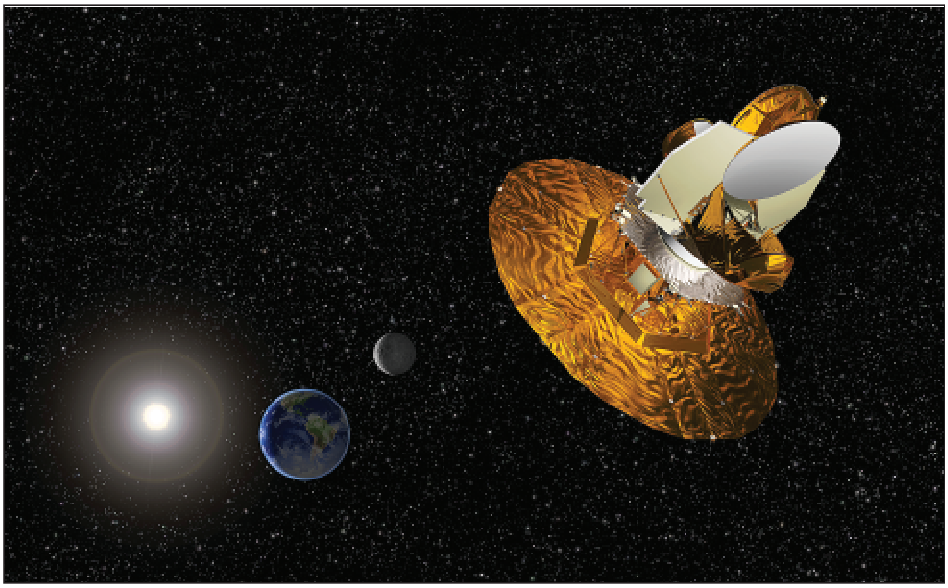

Courtesy of NASA/WMAP Science Team.

comfortably includes it and because the Mystery of the Missing Fluctuations still begs for an explanation.

Even though the 1σ estimate $\Omega = 1.02 \pm 0.02$ leaves plenty of room for a hyperbolic universe, a finite hyperbolic topology H^3/Γ would be difficult to detect. As the complexity of a hyperbolic group Γ increases, the volume of its quotient H^3/Γ increases as well, in contrast to spherical groups Γ whose quotient spaces S^3/Γ get smaller as the group gets larger. Moreover, even in the smallest hyperbolic manifolds, the injectivity radius R_{inj} is typically larger than the horizon radius R_{hor}. Astrophysicists know our horizon radius to be about 46 billion light-years.[2] Geometers, however, want to know our horizon radius not in light-years but in units of the *curvature radius* R_{curv}. In other words, geometers want to know the dimensionless ratio R_{hor}/R_{curv}, which tells the horizon radius in radians or "the usual hyperbolic units". Even though the horizon radius R_{hor} in light-years is known, the curvature radius R_{curv} depends strongly on Ω. If $\Omega = 0.98$, then the curvature radius R_{curv} is about 98 billion light-years, the geometer's horizon radius works out to $R_{hor}/R_{curv} = 0.47$, and randomly placed observers in the ten smallest known hyperbolic topologies H^3/Γ would have roughly a 50-50 chance of living at a point where their horizon radius exceeds the injectivity radius. However, if $\Omega = 0.99$, then the curvature radius R_{curv} increases to about 139 billion light-years, the dimensionless horizon radius drops to $R_{hor}/R_{curv} = 0.33$, and a randomly placed observer has only about a 10 percent chance of living at a point where the horizon radius exceeds the injectivity radius. As Ω approaches 1, the curvature radius goes to infinity, and the chances of detecting the nontrivial topology go to zero.

More promising from a purely topological point of view is the possibility of a spherical universe S^3/Γ. The curvature radius R_{curv}, which here is simply the radius in meters of the 3-sphere S^3 from which the universe S^3/Γ is constructed, comes out to 98 billion light-years when Ω takes the observed value of 1.02. Thus the astrophysicist's horizon radius $R_{hor} = 46$ billion light-years translates to the geometer's horizon radius $R_{hor}/R_{curv} = 0.47$ radians. In other words, at the nominal value of $\Omega = 1.02$, our horizon sphere's radius on the 3-sphere is 0.47 radians, meaning that we are seeing a modest yet nontrivial portion of the 3-sphere (Figure 3). Luckily, a horizon radius of 0.47 suffices to see the topology S^3/Γ for many of the simplest and most natural groups Γ, to be discussed below. Moreover, for more complicated groups Γ, the quotient S^3/Γ gets smaller, making the topology even easier to detect. This potential detectability, along with WMAP's observation of $\Omega \approx 1.02$, has fuelled considerable interest in the possibility of a spherical universe S^3/Γ.

Whether we consider hyperbolic, flat, or spherical manifolds, the question remains: which spaces best account for the Mystery of the Missing

[2] *Readers may wonder why the horizon has a 46 billion light-year radius when the universe is only 13.7 billion years old. The expanding universe provides the answer. The photons now reaching us from our horizon began their journey when the universe was 1100 times smaller than it is today. Thus the first light-year of space that a given photon traversed has since expanded to roughly 1100 light-years of space in the modern universe. In other words, the present day horizon radius is 46 billion light-years, but that same volume of space was much smaller, and more easily traversable, in the distant past.*

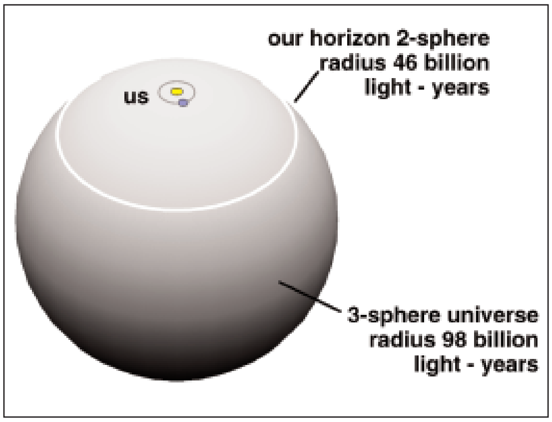

Figure courtesy of Bill Casselman.

Figure 3. At the nominal observed density $\Omega = 1.02$ the universe is positively curved with radius 98 billion light-years. Our horizon sphere has radius 46 billion light-years, which works out to 46/98 = 0.47 radians on the 3-sphere. Thus we see a modest but nontrivial portion of the 3-sphere.

Fluctuations? Surprisingly, not all small-volume universes suppress the large-scale fluctuations. As we will see below, some small-volume universes even elevate them!

Modes

A sustained musical tone may be expressed as the sum of a fundamental, a second harmonic, a third harmonic, and so on—in effect, its Fourier decomposition. This Fourier approach provides great insight into musical tones; for example, weak second and fourth harmonics characterize the sound of a clarinet. Similarly, any continuously defined field in the physical universe—for example, the density distribution of the primordial plasma—may be expressed as a sum of harmonics of 3-dimensional space. Technically these harmonics are the eigenmodes of the Laplace operator; intuitively they are the vibrational modes of the space, analogous to the vibrational modes of a 2-dimensional drumhead, so henceforth we will simply call them the *modes* of the space.

Just as the relative strengths of a clarinet's harmonics, its spectrum, characterize its sound, the relative strengths of the universe's modes characterize its physics. That is, just as the pressure and density fluctuations within the clarinet must conform to the clarinet's size and shape, the pressure and density fluctuations in the primordial plasma must conform to the size and shape of the universe. When we look out into space at our horizon, we see these density fluctuations. Of course we do not see the full 3-dimensional modes, but only their intersection with the 2-dimensional horizon sphere. Nevertheless, it is straightforward to calculate how a 3-manifold's modes restrict to 2-dimensional modes of the horizon sphere, ultimately allowing direct comparison to observations. Various sources of noise and other physical effects complicate the process but seem not to obscure the underlying topological and geometrical signatures. The strengths of the modes we observe on our 2-dimensional horizon sphere are called the *CMB power spectrum* (Figure 4). The CMB power spectrum's weak low-order terms conveniently quantify the Mystery of the Missing Large-Scale Fluctuations, just as the lack of low tones in a piccolo's spectrum reflect its small size.

The modes of a multiply connected space X/Γ, with $X = S^3$, E^3, or H^3 as before, lift in the obvious way to Γ-periodic modes of the simply connected space X. Conversely, each Γ-periodic mode of X

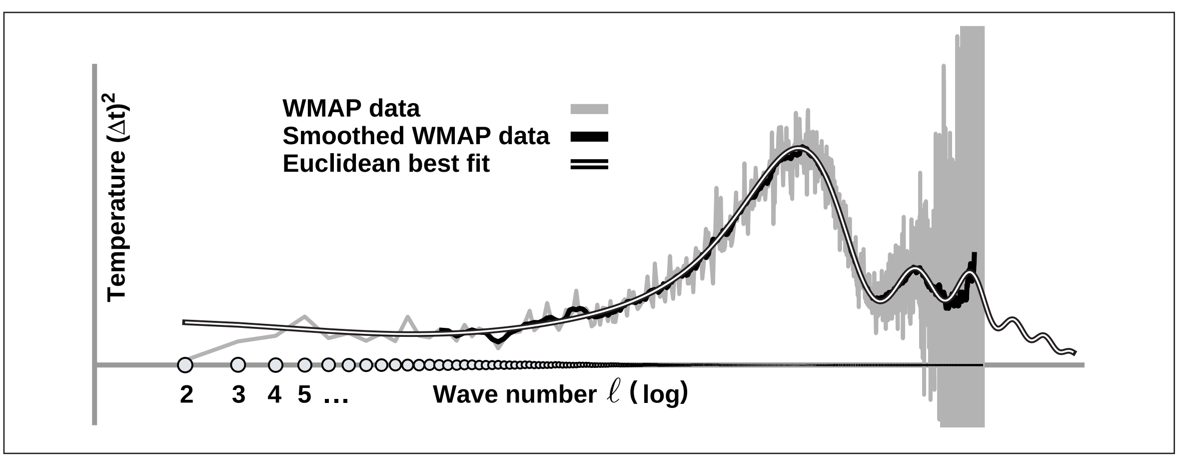

Figure courtesy of Bill Casselman.

Figure 4. Just as a musical tone splits into a sum of ordinary harmonics, the temperature fluctuations on our horizon sphere split into a sum of spherical harmonics. The resulting CMB power spectrum, shown here, tells much about the birth, evolution, geometry, and topology of the universe. The peaks in the spectrum fall more or less where expected, confirming physicists' theoretical understanding of the primordial plasma. The surprise lay in the weak lowest-order terms, which hint at a finite space.

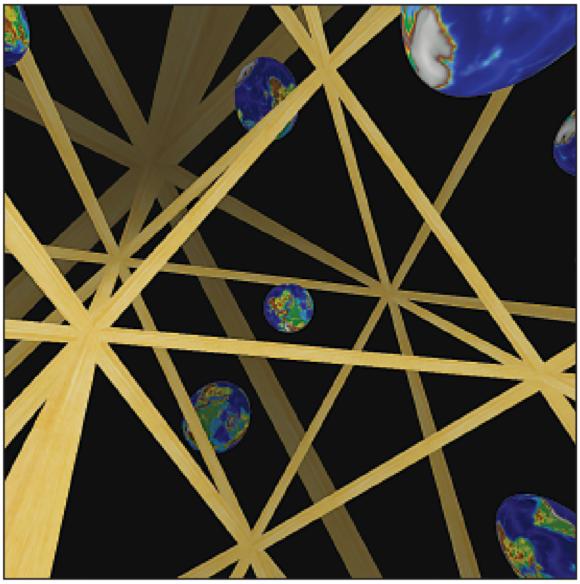

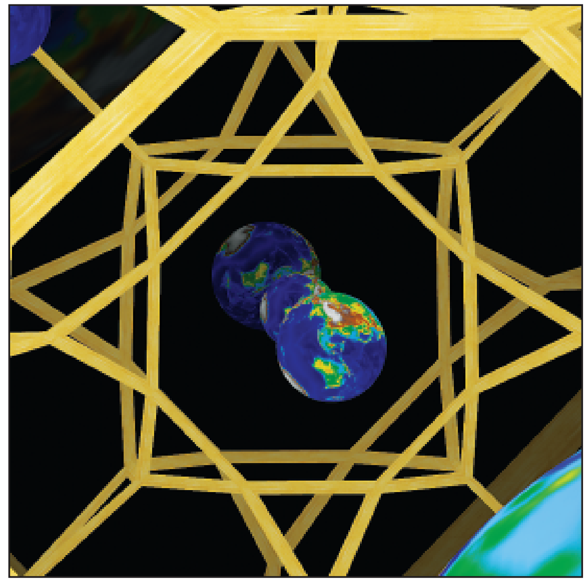

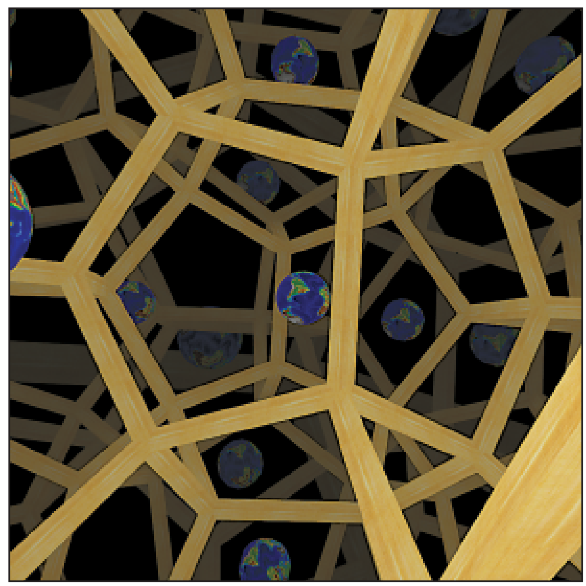

Figure 5. If the speed of light were infinite, inhabitants of the binary tetrahedral space S^3/T^* would see 24 images of every cosmological object (upper left); like atoms in a crystal the images repeat along a tiling of S^3 by 24 copies a fundamental octahedral cell. In the binary octahedral space S^3/O^* the images repeat along a tiling by 48 truncated cubes (upper right), and in the binary icosahedral space S^3/I^*, better known as the Poincaré dodecahedral space, the images repeat along a tiling by 120 octahedra (lower right). Because these still images provide only a weak understanding of the tiling, the reader is encouraged to fly around in them in real time using the free simulator available at `http://www.geometrygames.org/CurvedSpaces`**.**

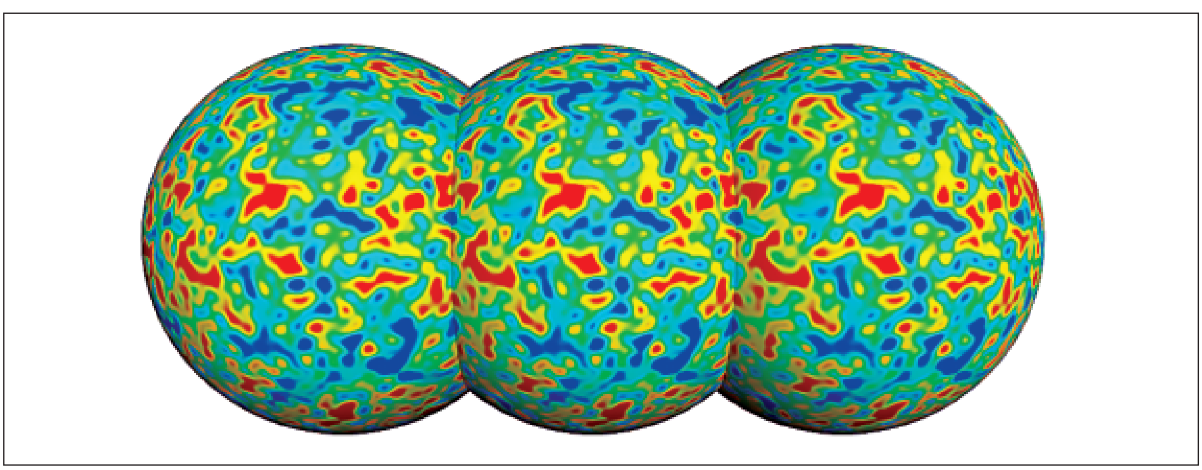

Figure 6. If our horizon radius exceeds our injectivity radius, the horizon sphere wraps all the way around the universe and intersects itself. Viewed in the universal cover, repeating images of the horizon sphere intersect. Observationally, we see the same circle of intersection on opposite sides of the sky.

projects down to a mode of the quotient X/Γ. Thus, even though we think of the modes conceptually as modes of X/Γ, in practice we invariably represent them as Γ-periodic modes of X.

In contrast to the irregular, hard-to-understand modes of hyperbolic space H^3, the modes of spherical space S^3 are regular and predictable. As a simple point of departure, consider the modes, or harmonics, of a circle S^1. Traditionally one writes them as $\{\{\cos\theta, \sin\theta\}, \{\cos 2\theta, \sin 2\theta\}, \dots\}$. Yet if we embed the circle as the set $x^2 + y^2 = 1$ in the xy-plane, we see that the transcendental functions $\cos\theta$ and $\sin\theta$ are completely equivalent to the linear functions x and y. Similarly, $\cos 2\theta$ and $\sin 2\theta$ become the quadratic polynomials $x^2 - y^2$ and $2xy$, and so on. In general, $\cos m\theta$ and $\sin m\theta$ become m^{th}-degree harmonic polynomials in x and y. By definition a polynomial $p(x_1, \dots, x_n)$ is *harmonic* if and only if it satisfies Laplace's equation $\nabla^2 p \equiv \frac{\partial^2 p}{\partial x_1^2} + \cdots + \frac{\partial^2 p}{\partial x_n^2} = 0$.

In perfect analogy to the modes of the circle S^1, the modes of the 2-sphere S^2 are precisely the homogeneous harmonic polynomials in x, y, and z, and the modes of the 3-sphere S^3 are precisely the homogeneous harmonic polynomials in x, y, z, and w. The only difference lies in the number of modes. On the circle the space of m^{th}-degree harmonic polynomials always has dimension 2, independent of m. On the 2-sphere the space of ℓ^{th}-degree harmonic polynomials has dimension $2\ell + 1$, and on the 3-sphere the space of k^{th}-degree harmonic polynomials has dimension $(k+1)^2$.

Cosmologists model the physics of a multiconnected spherical universe S^3/Γ using the Γ-periodic modes of S^3. For each degree k, the Γ-periodic modes form a subspace of the full $(k+1)^2$-dimensional mode space of S^3. Finding an orthonormal basis for that subspace requires, in principle, nothing more than a simple exercise in sophomore-level linear algebra. In practice the linear algebra works great for single-digit values of k but quickly bogs down as the size of the $(k+1)^2$-dimensional function space grows. Even simple cubic-time numerical matrix operations require $O(k^6)$ time on $(k+1)^2$-by-$(k+1)^2$ matrices, and efforts to work directly with the polynomials slow down even more dramatically. Adding insult to injury, accumulating round-off errors in floating point computations often render unusable the results of those computations that can be carried out within a reasonable time.

Thus for the past two years the main bottleneck for understanding and simulating the physics of multiconnected spaces X/Γ has been the efficient and accurate computation of the modes, the underlying local physics being already well understood. **News Flash:** Jesper Gundermann of the Danish Environmental Protection Agency has overcome the bottleneck and extended the CMB power spectra for various spherical spaces S^3/Γ from $\ell_{max} = 4$ out to $\ell_{max} = 15$. As this article goes to press, he is completing a rigorous statistical analysis, with results expected soon. His tentative results show an excellent fit.

Spherical Spaces

To recognize a spherical universe S^3/Γ by its imprint on the CMB power spectrum, we must first know which such spaces are possible. Fortunately 3-dimensional spherical spaces were classified by 1932 [3]. The possible groups Γ turn out to bear a close relationship to the symmetry groups of an ordinary 2-sphere! The easiest way to see the tight relationship between the symmetries of a 2-sphere and the symmetries of a 3-sphere is via the quaternions. Recall that the quaternions provide a noncommutative algebraic structure on $\mathbb{R}^4$ analogous to the commutative algebraic structure that the complex numbers provide on $\mathbb{R}^2$. Specifically the quaternions are spanned by $\{\mathbf{1}, \mathbf{i}, \mathbf{j}, \mathbf{k}\}$ subject to the rules $\mathbf{i}^2 = \mathbf{j}^2 = \mathbf{k}^2 = -\mathbf{1}$, $\mathbf{ij} = -\mathbf{ji} = \mathbf{k}$, $\mathbf{jk} = -\mathbf{kj} = \mathbf{i}$ and $\mathbf{ki} = -\mathbf{ik} = \mathbf{j}$. The length of a generic quaternion $a\mathbf{1} + b\mathbf{i} + c\mathbf{j} + d\mathbf{k}$ is just its obvious Euclidean length $\sqrt{a^2 + b^2 + c^2 + d^2}$. Just as the complex numbers restrict to a multiplicative group on the unit circle S^1, the quaternions restrict to a multiplicative group on the unit 3-sphere S^3.

Visualize the 3-sphere S^3 as the multiplicative group of unit length quaternions. A given quaternion $\mathbf{q} \in S^3$ can act on S^3 in two ways: by multiplication and by conjugation. When $\mathbf{q}$ acts by multiplication, taking each point $\mathbf{x} \in S^3$ to $\mathbf{qx}$, the result is a fixed-point free rotation of S^3. Any finite group Γ of such fixed-point free rotations defines a spherical space S^3/Γ.

When $\mathbf{q}$ acts by conjugation, taking each point $\mathbf{x} \in S^3$ to $\mathbf{qxq}^{-1}$, the result is a rotation with fixed points. Indeed, the point $\mathbf{1} \in S^3$ is fixed by all such rotations, because $\mathbf{q1q}^{-1} = \mathbf{1}$ for all $\mathbf{q}$, so in effect conjugation by $\mathbf{q}$ defines a rotation of the equatorial 2-sphere, which is the intersection of S^3 with the 3-dimensional subspace of purely imaginary quaternions $b\mathbf{i} + c\mathbf{j} + d\mathbf{k}$.

We now have a way to transfer symmetries from the 2-sphere to the 3-sphere. Start with a finite group of symmetries of S^2, for example, the tetrahedral group T consisting of the twelve orientation-preserving symmetries of a regular tetrahedron. Represent T as a set G of quaternions acting by conjugation. Now let the same set G act on S^3 by multiplication. Voilá! There is our group Γ of fixed-point free symmetries of the 3-sphere. The only catch is that each of the original symmetries of S^2 is realized by two different quaternions, $\mathbf{q}$ and $-\mathbf{q}$, so the group G has twice as many elements as the original group. In the present example, with the original group being the tetrahedral group T, the final group Γ is the *binary tetrahedral group* T^*, of order 24.

The finite symmetry groups of S^2 are well known:

- The *cyclic groups* Z_n of order n, generated by a rotation through an angle $2\pi/n$ about some axis.
- The *dihedral groups* D_m of order $2m$, generated by a rotation through an angle $2\pi/m$ about some axis as well as a half turn about some perpendicular axis.
- The *tetrahedral group* T of order 12 consisting of all orientation-preserving symmetries of a regular tetrahedron.
- The *octahedral group* O of order 24 consisting of all orientation-preserving symmetries of a regular octahedron.
- The *icosahedral group* I of order 60 consisting of all orientation-preserving symmetries of a regular icosahedron.

Transferring those groups from S^2 to S^3, as explained above, yields the single action symmetry groups of S^3:

- The cyclic groups Z_n of order n.
- The binary dihedral groups D_m^* of order $4m$, $m \geq 2$.
- The binary tetrahedral group T^* of order 24.
- The binary octahedral group O^* of order 48.
- The binary icosahedral group I^* of order 120.

The corresponding quotients S^3/Γ are the *single action spaces* (Figure 5). To fully understand these spaces, the reader is encouraged to fly around in them using the free simulator available at `www.geometrygames.org/CurvedSpaces`.

The full classification of finite fixed-point free symmetry groups of S^3 is more complicated, but only slightly. Geometrically, the mapping $\mathbf{x} \to \mathbf{qx}$ acts as a right-handed corkscrew motion. If we switch to right multiplication $\mathbf{x} \to \mathbf{xq}$, we get left-handed corkscrew motions instead of right-handed ones, and the quotient space S^3/Γ becomes the mirror image of what it had been. (To understand why changing $\mathbf{qx}$ to $\mathbf{xq}$ reverses the chirality of the corkscrew motion, multiply out the two products

$$(\cos\theta\mathbf{1} + \sin\theta\mathbf{i})(a\mathbf{1} + b\mathbf{i} + c\mathbf{j} + d\mathbf{k})$$

and

$$(a\mathbf{1} + b\mathbf{i} + c\mathbf{j} + d\mathbf{k})(\cos\theta\mathbf{1} + \sin\theta\mathbf{i})$$

and observe the results.) Two groups, Γ and Γ', may act on S^3 simultaneously, one by left multiplication and the other by right multiplication. The resulting quotient space is called a *double action space* (for details see [4]). So far cosmologists have largely neglected double action spaces and the related *linked action spaces*. One reason for the neglect is that they are more complicated and often require unrealistically large groups Γ. A more fundamental difficulty, however, is that a double action space is *globally inhomogeneous*, meaning its geometry and therefore the expected CMB power spectrum look different to observers sitting at different locations within the same space, while a single action space is *globally homogeneous*, meaning its geometry and the expected CMB fluctuations look the same to all observers. Obviously the single action spaces are far easier to simulate, because a single simulation suffices for all observers.

Encouraging News

Resolving the Mystery of the Missing Fluctuations turned out to be harder than expected. The principle that a small universe cannot support broad fluctuations was clear enough; the challenge lay in accounting for the particular CMB power spectrum that WMAP observed. In reality the $\ell = 2$ term, corresponding to a quadratic polynomial and generally called the *quadrupole*, was anomalously low (at the 1-in-100 level), while the $\ell = 3$ term, corresponding to a cubic polynomial and generally called the *octopole*, was somewhat low as well (at the 1-in-5 level). These 1-in-500 odds of weak low-ℓ fluctuations beg for an explanation, even though simple random chance cannot be excluded.

The reader may wonder what became of the lower-ℓ terms. The $\ell = 0$ constant term is simply the average CMB temperature. The $\ell = 1$ linear term, the *dipole*, is swamped by the far stronger dipole induced by the solar system's 300 km/sec motion relative to the CMB and is therefore unavailable.

Returning to the mysteriously low quadrupole and octopole, the simplest possible explanation, a flat 3-torus universe made by identifying opposite faces of a cubical block of space, failed to account for the observed power spectrum. A sufficiently small cubic 3-torus can of course suppress the quadrupole ($\ell = 2$) as strongly as desired, but not without suppressing other low-order modes ($\ell = 3, 4, 5$) along with it, contrary to observations. Making the 3-torus from a more general parallelepiped offers greater flexibility. The full 6-parameter space of parallelepipeds has yet to be fully explored, but initial investigations show that while noncubic rectangular 3-tori suppress the low-ℓ portion of the spectrum, they suppress the high-ℓ portion even more, leaving the low-ℓ portion *relatively* elevated [5], contrary to observation.

Among the spherical topologies, the lens spaces were considered first. A lens space $L(p, q)$ is the quotient S^3/Z_p of S^3 under the action of the cyclic group generated by

$$\begin{pmatrix} \cos\frac{2\pi}{p} & -\sin\frac{2\pi}{p} & 0 & 0 \\ \sin\frac{2\pi}{p} & \cos\frac{2\pi}{p} & 0 & 0 \\ 0 & 0 & \cos\frac{2\pi q}{p} & -\sin\frac{2\pi q}{p} \\ 0 & 0 & \sin\frac{2\pi q}{p} & \cos\frac{2\pi q}{p} \end{pmatrix}.$$

Because the group acts in only one direction, the quotient behaves roughly like a rectangular 3-torus that is narrow in one direction but wide in two others. Like a noncubic 3-torus, a lens space $L(p,q)$ suppresses the high-ℓ portion of the spectrum more heavily than the low-ℓ portion, in effect elevating the quadrupole and contradicting observations for all but the smallest choices of p.

The failures of both the rectangular 3-tori and the lens spaces to account for the low quadrupole combined to teach a useful lesson: a low quadrupole requires a *well-proportioned space*, with all dimensions of similar magnitude. Only in a well-proportioned space is the quadrupole suppressed more heavily than the rest of the power spectrum. This insight, combined with WMAP's hints of slight positive curvature, led researchers to consider the binary polyhedral spaces S^3/T^*, S^3/O^*, and S^3/I^*, all of which have both positive curvature and well-proportioned fundamental domains, namely a regular octahedron, a truncated cube, and a regular dodecahedron, respectively, as shown in Figure 5.

Of the binary polyhedral spaces, S^3/I^* seemed the most promising candidate, for the simple geometrical reason that its fundamental domain's inradius of $\pi/10 \approx 0.31$ fits easily within the horizon radius $R_{hor}/R_{curv} = 0.47$ corresponding to $\Omega = 1.02$. Topologists know S^3/I^* as the *Poincaré dodecahedral space.* Curiously, Poincaré himself never knew his namesake manifold could be constructed from a dodecahedron. Rather, he discovered the manifold in a purely topological context as the first example of a multiply connected homology sphere. A quarter century later Weber and Seifert glued opposite faces of a dodecahedron and proved that the resulting manifold was homeomorphic to Poincaré's homology sphere.

Lacking explicit formulas for the modes of the binary polyhedral spaces, the author and his colleagues computed the modes numerically. Unfortunately, in the case of the Poincaré dodecahedral space S^3/I^*, accumulating numerical errors limited the computation to the modes $k \le 24$ of the 3-dimensional space, in turn limiting the reliable portion of the predicted CMB power spectrum to $\ell = 2, 3, 4$. Nevertheless, the results were delightful: the predicted quadrupole ($\ell = 2$) and octopole ($\ell = 3$) matched observations [6]! (The $\ell = 4$ term was used to set the overall normalization.) Moreover, the best fit occurred in the range $1.01 < \Omega < 1.02$, comfortably within WMAP's observation of $\Omega = 1.02 \pm 0.02$.

Several factors made this result especially elegant. First and foremost was the near total lack of free parameters. Unlike the 3-torus, which can be made from an arbitrary parallelepiped (six degrees of freedom), the dodecahedral space can be made only from a perfectly regular dodecahedron (zero degrees of freedom). Second, the dodecahedral space is globally homogeneous. Unlike in a typical 3-manifold where the observer's position affects the expected CMB power spectrum (three degrees of freedom), the dodecahedral space looks the same to all observers (zero degrees of freedom). The only free parameter in our simple initial study was the density parameter Ω. Amazingly, with only one parameter to vary, the model correctly accounted for three independent observations: the quadrupole ($\ell = 2$), the octopole ($\ell = 3$), and the observed density itself. While far from a proof, such results were most encouraging.

Discouraging News

The dodecahedral model makes three testable predictions:

1. the weak large-scale CMB fluctuations,
2. matching circles in the sky (to be discussed momentarily), and
3. a slight curvature of space.

The WMAP satellite had already observed the weak large-scale CMB fluctuations, so there is no problem with the first prediction.

As for the third prediction, current measurements ($\Omega = 1.02 \pm 0.02$) fail to distinguish flat space ($\Omega = 1$) from the slight curvature that the dodecahedral model requires ($\Omega \approx 1.02$). Fortunately, upcoming data may suffice, either within a year or two by combining the WMAP results with other data sets to narrow the error bars on Ω, or by the end of the decade if we wait for more precise CMB measurements from the European Space Agency's Planck satellite.

This leaves the second prediction, currently the most controversial one. The basic insight is as follows. If the fundamental dodecahedron is smaller than our horizon sphere, then the horizon sphere will "wrap around the universe" and intersect itself. This is most conveniently visualized in the universal covering space (Figure 6) where repeating images of the horizon sphere intersect their neighbors. From our vantage point on Earth, at the center of our horizon sphere we can see the same circle of intersection sitting on opposite sides of the sky. If the observed CMB temperature fluctuations depended only on plasma density fluctuations, then the two images, one in front of us and one behind us, would display identical temperature patterns. Locating such pairs of matching circles would conclusively prove the universe is finite and reveal its topology.

When the news appeared last October that the dodecahedral model accounts for the Mystery of the Missing Fluctuations, another group of researchers was simultaneously searching for matching circles. Their massive computer search was expected to take months on a cluster of PCs, but by October their program, while not yet finished with the whole search, had finished checking for the diametrically opposite circle pairs predicted by the dodecahedral model and had found none [7].

The failure to find matching circles disappointed everyone, but might not have dealt a fatal blow to the idea of a finite universe. The observed CMB temperature fluctuations depend not only on plasma density fluctuations but also on other factors such as the Doppler effect of the plasma's motion and the gravitational influences the CMB photons experience over the course of their 13.7 billion-year journey from the plasma to us. For the most part the circle searchers carefully accounted for these various sources of contamination. However, they neglected residual contamination from foreground sources within our own Milky Way Galaxy. Total foreground contamination is intense, and the WMAP team devoted considerable effort to subtracting as much of it as possible before releasing their cleaned CMB sky maps. Nevertheless, the remaining contamination is easily visible to the naked eye along the galactic equator and is likely strong enough to disrupt circle matching along a wider swath. Future work should reveal the extent to which the residual foreground contamination may or may not obscure matching circles after resolving questions about how to model it correctly.

If the technical details get resolved and the matching circles really are not there, would we conclude that the universe is infinite? Not at all! First there is the possibility that the universe is finite but much larger than our horizon, in which case we could not detect its topology. More practically, there is the possibility that the universe is comparable to, or slightly larger than, the horizon. Such a universe would not generate detectable circles but might still account for the Mystery of the Missing Fluctuations. All participants in recent discussions—circle searchers as well as dodecahedral space modelers—agree this is the obvious Plan B. But how can one hope to detect topology lying beyond our horizon? It may not be as hard as it seems. Just as one may deduce the full length of a guitar string merely by observing vibrations on its middle 80 percent (with the endpoints hidden from view!), one may in principle deduce the topology of the universe by observing density fluctuations within a limited volume.

Conclusion

Where will the conflicting pieces of evidence lead? To a "small" dodecahedral universe lying wholly within our horizon? To a somewhat larger universe lying just beyond our view but within our experimental grasp? Or to a presumed but unconfirmed infinite flat space? Over the course of the decade improved measurements of curvature should provide the decisive clue. If Ω is sufficiently close to 1, then the dodecahedral model is dead. However, if Ω is found to be near 1.02—and bounded away from 1—then the dodecahedral model or some variation of it will almost surely prove correct, whether or not matching circles are found. For now, the Missing Fluctuations remain a Mystery.

Added in Proof

Studies of the first-year WMAP data find that certain features of the low-order CMB harmonics align with the ecliptic plane at roughly the 99.9% confidence level. Such alignments call into question the presumed cosmic origins of the low-order harmonics, suggesting instead either some hitherto unknown solar system contribution to the CMB, or perhaps some error in the collection and processing of the data. The second-year WMAP data, originally expected by February 2004 but delayed due to unexpected surprises in the results, may soon shed additional light on these anomalies. If the true cosmological low-order harmonics, after subtracting off any solar system effects and/or processing errors, turn out to be even weaker than previously believed, then the Mystery of the Missing Fluctuations will deepen. However, the case for the dodecahedral topology would vanish, forcing researchers to reconsider other topologies—and perhaps other explanations—in light of the revised data.

Acknowledgment

I thank the MacArthur Foundation for its support.

References

[1] D. Spergel et al., First year Wilkinson Microwave Anisotropy Probe (WMAP 1) observations: Determination of cosmological parameters, *Astrophys. J. Suppl. Ser.* **148** (2003), 175–194.

[2] C. Bennett et al., First year Wilkinson Microwave Anisotropy Probe (WMAP 1) observations: Preliminary maps and basic results, *Astrophys. J. Suppl. Ser.* **148** (2003), 1–27.

[3] W. Threlfall and H. Seifert, Topologische Untersuchung der Diskontinuitätsbereiche endlicher Bewegungsgruppen des dreidimensionalen sphärischen Raumes, *Math. Ann.* **104** (1930), 1–70; and **107** (1932), 543–586.

[4] E. Gausmann, R. Lehoucq, J.-P. Luminet, J.-P. Uzan, and J. Weeks, Topological lensing in elliptical spaces, *Classical Quantum Gravity* **18** (2001), 5155–5186.

[5] J. Weeks, J.-P. Luminet, A. Riazuelo, and R. Lehoucq, Well-proportioned universes suppress CMB quadrupole, www.arxiv.org/astro-ph/0312312, to appear in *Monthly Notices of the Royal Astronomical Society*.
[6] J.-P. Luminet, J. Weeks, A. Riazuelo, R. Lehoucq, and J.-P. Uzan, Dodecahedral space topology as an explanation for weak wide-angle temperature correlations in the cosmic microwave background, *Nature* **425** (2003), 593–595.
[7] N. Cornish, D. Spergel, G. Starkman, and E. Komatsu, Constraining the topology of the universe, www.arxiv.org/astro-ph/0310233, to appear in *Phys. Rev. Lett.*

About the Cover

Mapping the Universe with HEALPixels

In Jeff Weeks's article, the data that go into his analysis of the cosmic microwave background radiation have been collected by the WMAP satellite, which scans different directions in the sky and records temperatures. Data are sampled and amalgamated according to a scheme called **HEALPix** (**H**ierarchical **E**qual **A**rea iso-**L**atitude **Pix**elization), which has become the standard tool for such measurements—for the WMAP satellite as well as for its successor, Planck, scheduled to be launched in 2007. According to **HEALPix** the sphere is first divided into twelve large zones (outlined in gray on the cover), and each of these in turn is parceled into N^2 smaller "pixels" where $N = 2^n$. For WMAP $N = 512$, which means $3,145,728$ pixels altogether, and Planck will have many more, so there is a lot of calculation to be done in dealing with these data.

There is no canonical way to partition a sphere into uniform small regions, since there are only five regular solids. Any scheme used must therefore choose among various tradeoffs. The principal characteristic feature of **HEALPix** is that its pixels are all of equal area, more or less in the shape of rhombi, as the cover shows. Different formulas are used to generate the shapes in the polar and equatorial regions, but both depend on Archimedes' Theorem that cylindrical projection is area preserving.

Several considerations in addition to efficiency go into devising such a scheme: regions of equal area minimize the effect of noise, among other things. The distribution of regions along parallels of latitude allows using the fast Fourier transform in decomposing the data into spherical harmonics. **HEALPix** is probably close to being as fast as possible in facilitating spherical harmonic analysis of astrophysical data. Faster known mathematical techniques are not practical in this context, requiring as they do a partition of the sphere that does not deal well with noise.

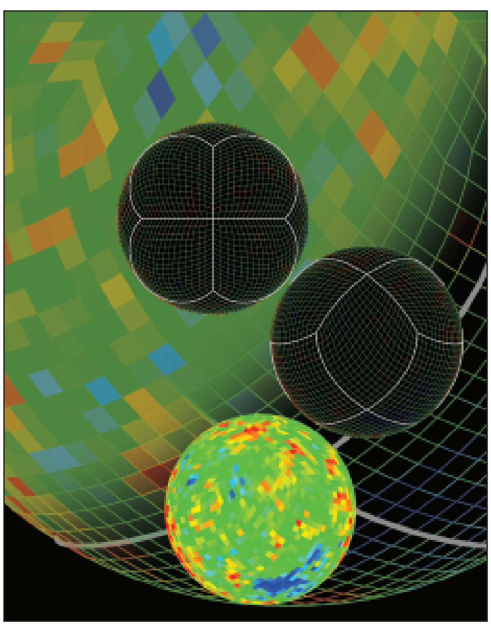

The **HEALPix** suite of programs, written in FORTRAN, is publicly available from http://www.eso.org/science/healpix/. This package includes a number of useful peripheral tools as well as core routines. The documentation explaining how to use it is good, although to a mathematician interested in understanding internals it will likely seem that the algorithms it incorporates are not well explained nor the programs themselves easy to follow.

The data for the cover are obtained by degradation from those in the file of internal linear combination data available at http://lambda.gsfc.nasa.gov/product/map/m_products.cfm.

I wish to thank William O'Mullane for making available to me his Java versions of portions of the **HEALPix** code, and Krzysztof Górski, the originator of **HEALPix**, for his spirited explanation of its role.

—Bill Casselman
Graphics Editor
(notices-covers@ams.org)