

Sampling-Efficient Mapping of Spherical Images

John Snyder
Microsoft Research

Don Mitchell
Microsoft Research

Abstract

Many functions can map images to the sphere for use as environment maps or spherical panoramas. We develop a new metric that asymptotically measures how well these maps use a given number of samples to provide the greatest worst-case frequency content of the image everywhere over the sphere. Since this metric assumes perfect reconstruction filtering even with highly anisotropic maps, we define another, conservative measure of sampling efficiency that penalizes anisotropy using the larger singular value of the mapping's Jacobian. With these metrics, we compare spherical maps used previously in computer graphics as well as other mappings from cartography, and propose several new, simple mapping functions (dual equidistant and polar-capped maps) that are significantly more efficient and exhibit less anisotropy. This is true with respect to either efficiency metric, which we show agree in the worst case for all but one of the spherical maps presented. Although we apply the metrics to spherical mappings, they are useful for analyzing texture maps onto any 3D surface.

Additional Keywords: cartography, environment map, Fourier transform, signal processing, singular value, surface parameterization, texture.

1. Introduction

Spherical images representing the radiance field at a point are useful to simulate reflections on shiny surfaces (environment maps) [2][7][9][20] and to produce arbitrary views from a point (spherical panoramas) [15][5]. To enhance realism, we expect more widespread use of such spherical images, including the use of multiple spherical and hemispherical images to approximate radiance at many points throughout an environment [11].

Rectangular arrays of samples (i.e., texture maps) are ubiquitous in graphics systems, providing advantages including locality of reference, simplicity of texel addressing, and ease of filtering for reconstruction. We omit schemes like spherical wavelets [17] which exploit local differences in frequency content, but are more complicated to implement in hardware. Still, many functions have been used to map samples from a 2D texture domain to the sphere, including the cube map [7][15][20], OpenGL map [8][13], polar coordinate (latitude/longitude) map [2], and dual stereographic map [9] (called "dual parabolic" in [9]). How should these mappings be compared? Is there anything better?

To answer these questions, we propose the following criteria:

1. *Sampling efficiency:* the mapping function should support the greatest worst-case frequency content in the spherical image using the fewest samples. We prefer a worst-case metric assuming blurry spots or directions in the mapping are undesirable, even if they allow greater fidelity elsewhere. Minimizing texture samples is important to conserve graphics memory and reduce memory bandwidth.
2. *Anisotropy:* isotropic filtering methods (e.g., trilinear interpolation within a MIPMAP) [21] penalize maps having significant anisotropy, since they cause excessive blur in the locally stretched

Map Name	Sampling Req.	Maximum Anisotropy	Map Comps.
OpenGL	∞	∞	1
Cube	24	1.73	6
Dual Stereographic	32	1	2
Lat/Long	19.7	∞	1
Dual Equidistant*	19.7	1.57	2
Low Distortion Equal Area*	19.7	3.45	1
Polar-Capped* (stretch invariant)	14.8	1.41	3
Polar-Capped* (conformal)	16.5	1	3
Polar-Capped* (hex. reparam.)	13.5	1.73	3
Optimal Isometric**	12.57	1	∞
Optimal**	10.9	1.73	∞

Figure 1: Summary of Spherical Map Properties. Single-starred are new maps described in this paper; double-starred are theoretical limits rather than actual maps. Sampling requirement is proportional to texture area required for a desired frequency content in the worst case (Section 2.2). OpenGL's infinite value means its number of samples grows faster than quadratically as frequency content increases. Smaller anisotropy values are better and represent degree of local conformality (Section 2.3).

direction. Even hardware having anisotropic filtering capability can not tolerate *unbounded* anisotropy, though anisotropy values up to 2.0 are handled with hardware implementations such as Nvidia's TNT and GeForce [10].

3. *Simplicity of projection function:* texture coordinate generation by the graphics system should be quickly computable.
4. *Ease of geodesic interpolation:* simple (e.g., linear) interpolation over the domain should closely approximate geodesic interpolation over the sphere.
5. *Ease of creation:* ideally, mappings should be suitable for dynamic creation using the rendering hardware. Note though that maps can be reparameterized using an additional rendering pass over a textured, tessellated sphere.

Also important is the use of map components, such as the six faces of the cube map or two hemispheres in the dual stereographic map. With enough components, or "pieces" in a piecewise map, sampling efficiency can be increased to an asymptotic limit. But more components create more complexity in evaluating the mapping function, reduced locality of reference, and greater difficulty in MIPMAP construction. Moreover, efficiency gains may be lost through area wasted in packing many small maps into the texture rectangle. For these reasons, only maps having few components are practical.

This paper formally defines sampling efficiency for texture mapping functions over regular 2D lattices by applying results from signal processing [14] and crystallography [3] to spatially-varying mappings. We define two metrics; one considering only frequency content and another, strictly larger, that takes into account anisotropy by considering the largest singular value of the Jacobian over the domain. Using these metrics, we analyze several types of spherical maps used previously in computer graphics as well as other mappings from cartography [16][19], and show that the two metrics agree in the worst case. We propose several new mappings, including the polar-capped map and dual equidistant map, having superior sampling efficiency and reduced anisotropy. We also provide theoretical limits on the efficiency attainable by any piecewise-differentiable map to the sphere. The results are

summarized in Figure 1. The new maps have simple-to-evaluate projection functions that can be implemented using existing hardware to reduce storage and bandwidth costs and avoid anisotropic blur of spherical images.

2. Sampling Efficiency and Other Local Mapping Properties

The *spherical mapping function*, $S(u, v)$, maps (u, v) points in the domain, D , to 3D unit-length vectors. Existing graphics systems use rectangular domains and sampling patterns (see Figure 1(a)). Each (u, v) sample from the tabulated spherical image represents the radiance associated with vector $S(u, v)$.

Analysis of sampling efficiency is first developed for regular lattices in the plane (Section 2.1). We then extend this analysis to spatially varying mappings by locally projecting the sampling grid onto the tangent plane at each point and applying the same planar analysis (Section 2.2). Section 2.3 presents mathematics for the analysis of local distortion (anisotropy).

2.1 Sampling in the Plane

Following [14][6], a regular sampling pattern in the plane can be represented as a *sampling matrix*, $V \equiv (v_1 \ v_2)$, where v_1 and v_2 are linearly independent vectors to two nearest-neighbor sample locations. For example, isotropic rectangular and hexagonal sampling have the following sampling matrices:

$$V_{rect} \equiv \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \delta, \quad V_{hex} \equiv \begin{bmatrix} \sqrt{3}/3 & \sqrt{3}/3 \\ 1 & -1 \end{bmatrix} \delta$$

where δ represents the sample spacing. The corresponding sample geometries are shown in Figure 2. Sample locations are derived using $t = Vn$ where $n \equiv (n_1, n_2)$ is a vector of integers (prime denotes transpose).

Denote the continuous signal to be sampled as $\tilde{f}(t)$. The corresponding sampled version is $f(n) \equiv \tilde{f}(t) = \tilde{f}(Vn)$. Taking the Fourier transform of \tilde{f} and then its inverse, we obtain

$$\begin{aligned} \tilde{F}(\Omega) &\equiv \int_{-\infty}^{+\infty} \tilde{f}(t) \exp(-i\Omega' t) dt \\ \tilde{f}(t) &= \frac{1}{4\pi^2} \int_{-\infty}^{+\infty} \tilde{F}(\Omega) \exp(i\Omega' t) d\Omega \end{aligned}$$

where $\Omega \equiv (\Omega_x \ \Omega_y)$ is in units of radians per unit length. Doing the same for the discrete signal, f , we get

$$\begin{aligned} F(\omega) &\equiv \sum_n f(n) \exp(-i\omega' n), \\ f(n) &= \frac{1}{4\pi^2} \int_{-\pi}^{+\pi} F(\omega) \exp(i\omega' n) d\omega \end{aligned} \quad (1)$$

where $\omega \equiv (\omega_x \ \omega_y)$ is in units of radians. So

$$f(n) = \tilde{f}(Vn) = \frac{1}{4\pi^2} \int_{-\infty}^{+\infty} \tilde{F}(\Omega) \exp(i\Omega' Vn) d\Omega.$$

Substituting $\omega = V'\Omega$ yields

$$f(n) = \frac{1}{4\pi^2} \int_{-\infty}^{+\infty} \frac{1}{|\det V|} \tilde{F}(V'^{-1}\omega) \exp(i\omega' n) d\omega.$$

The double integral in the $(\omega_x \ \omega_y)$ -plane can be broken into an infinite sum of integrals each covering a square area of $4\pi^2$:

$$f(n) = \frac{1}{4\pi^2} \sum_{k_x} \sum_{k_y} \int_{-\pi+2\pi k_x}^{\pi+2\pi k_x} \int_{-\pi+2\pi k_y}^{\pi+2\pi k_y} \frac{1}{|\det V|} \tilde{F}(V'^{-1}\omega) \exp(i\omega' n) d\omega$$

Replacing ω_x by $\omega_x - 2\pi k_x$ and ω_y by $\omega_y - 2\pi k_y$, simplifying, and comparing to (1) implies that




Figure 2: Example Plane Samplings

$$F(\omega) = \frac{1}{|\det V|} \sum_k \tilde{F}(V'^{-1}(\omega - 2\pi k)) \quad \text{or}$$

$$F(V'\Omega) = \frac{1}{|\det V|} \sum_k \tilde{F}(\Omega - Uk)$$

where $U \equiv 2\pi V'^{-1}$, and the integer vector $k \equiv (k_x, k_y)$. Thus the discrete signal's Fourier transform has a periodicity matrix related to the original sampling matrix V via its inverse transpose.

To eliminate aliasing, we need to bandlimit $\tilde{f}(t)$ such that its Fourier transform, $\tilde{F}(\Omega)$, is zero outside a finite region in frequency space. The region is chosen so as to have no overlap between neighboring periodically repeated *tiles* of $F(V'\Omega)$. For the rectangular and hexagonal sampling matrix examples, we have

$$U_{rect} \equiv \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix} \frac{\pi}{\delta}, \quad U_{hex} \equiv \begin{bmatrix} \sqrt{3} & \sqrt{3} \\ 1 & -1 \end{bmatrix} \frac{\pi}{\delta}$$

with frequency space tiles shown in Figure 3.

For such a bandlimited function, inside the periodic tile containing the origin, called the *baseband*, B , the continuous and discrete Fourier transformed functions are related via

$$F(V'\Omega) = \frac{1}{|\det V|} \tilde{F}(\Omega)$$

implying we can reconstruct such bandlimited functions exactly from the discrete samples. In fact, the reconstruction is given by

$$\tilde{f}(t) = \frac{|\det V|}{4\pi^2} \sum_n f(n) \int_B \exp(i\Omega'(t - Vn)) d\Omega.$$

Given a sampling matrix V , we find the radius of the largest inscribed circle in B ; that is, the maximum frequency content in all directions that a circularly bandlimited periodic signal sampled using V can support. This is calculated by computing $U \equiv 2\pi V'^{-1} = (u_1 \ u_2)$ and finding half the minimum distance of

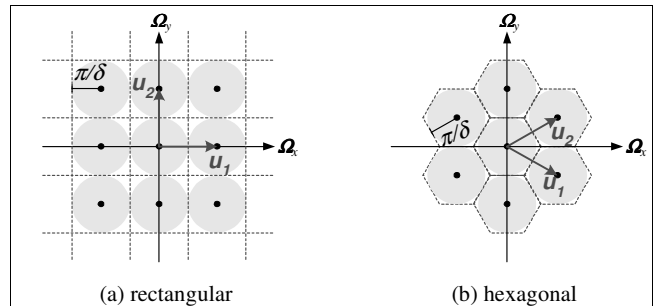


Figure 3: Fourier transform of discrete signals from example sampling matrices. Both rectangular and hexagonal sampling from Figure 2 provide for frequencies up to π/δ in all directions without aliasing—a frequency radius of π/δ is the largest for an inscribed circle in each periodic tile. But hexagonal sampling requires 13.4% fewer samples to accomplish this, since the circles are packed more tightly in each hexagonal tile, while the square tiles waste space on outside diagonal frequencies.

the origin to the vectors $n_1\mathbf{u}_1 \pm n_2\mathbf{u}_2$, for integers n_1 and n_2 not both zero. We thus define the *sampling spectral radius*, Δ , of \mathbf{V} :

$$\begin{aligned} \Delta &\equiv 1/2 \min_{(n_1, n_2) \neq (0,0)} (\|n_1\mathbf{u}_1 - n_2\mathbf{u}_2\|) \\ &= (\pi/|\det\mathbf{V}|) \min_{(n_1, n_2) \neq (0,0)} (\|n_1\mathbf{v}_1 - n_2\mathbf{v}_2\|) \end{aligned} \quad (2)$$

For the rectangular and hexagonal example samplings, $\Delta = \pi/\delta$. It can be shown [14] that among all sampling matrices having a given sampling spectral radius, the most efficient is always isotropic hexagonal sampling, in that it covers the greatest area with the fewest samples. Sampling density can be measured by $d \equiv 1/|\det\mathbf{V}|$, for which our example sampling matrices yield $d_{rect} = \delta^2$ and $d_{hex} = 2\sqrt{3}/3 \delta^2$.

2.2 Mapped Sampling

We are interested in frequency content on the sphere, not the plane. We therefore locally apply the previous section's analysis of regular sampling patterns in 2D by approximating the spherical samples in the neighborhood of $\mathbf{S}(u, v)$ by 2D points projected to the tangent plane there.

Let $\mathbf{S}_u(u, v) \equiv \partial\mathbf{S}/\partial u(u, v)$ and $\mathbf{S}_v(u, v) \equiv \partial\mathbf{S}/\partial v(u, v)$. Then an orthonormal basis for the tangent space is given by $(\mathbf{T}_1 \mathbf{T}_2)$

$$\mathbf{T}_1(u, v) \equiv \frac{\mathbf{S}_u}{\|\mathbf{S}_u\|}, \quad \mathbf{T}_2(u, v) \equiv \frac{(\mathbf{S}_u \square \mathbf{S}_v) \mathbf{S}_v - (\mathbf{S}_v \square \mathbf{S}_u) \mathbf{S}_u}{\|\mathbf{S}_u\| \tau}$$

where $\tau(u, v) \equiv \sqrt{(\mathbf{S}_u \square \mathbf{S}_u)(\mathbf{S}_v \square \mathbf{S}_v) - (\mathbf{S}_u \square \mathbf{S}_v)^2}$ represents the differential area at (u, v) . Letting $\mathbf{J}(u, v) \equiv (\mathbf{S}_u \mathbf{S}_v)$,

$$\mathbf{S}((u, v) + (\dot{u}, \dot{v})) - \mathbf{S}(u, v) \approx \mathbf{J}(\dot{u}, \dot{v})' \quad (3)$$

is an approximation good for small $\|(\dot{u}, \dot{v})\|$. Since the Jacobian \mathbf{J} maps perturbations in any domain direction to the tangent plane, it represents a "local" sampling matrix. To derive a 2D sampling matrix, we project it into the plane using the orthonormal basis $(\mathbf{T}_1 \mathbf{T}_2)$, a pure rotation which preserves the spectral properties. Defining \mathbf{K} as the resulting transformed Jacobian:

$$\mathbf{K}(u, v) \equiv (\mathbf{T}_1(u, v) \mathbf{T}_2(u, v))' \mathbf{J}(u, v) = \frac{1}{\|\mathbf{S}_u\|} \begin{bmatrix} \mathbf{S}_u \square \mathbf{S}_u & \mathbf{S}_u \square \mathbf{S}_v \\ 0 & \tau \end{bmatrix}$$

For rectangular sampling with spacing δ , the *local sampling matrix* mapped by \mathbf{S} is then given by $\mathbf{V} = \delta \mathbf{K}$.

We can now analyze the spectral radius determined by \mathbf{V} at any point (u, v) as if \mathbf{S} were everywhere equal to the local approximation in (3). By selecting the minimum sampling spectral radius for any point in \mathbf{D} , we determine the highest permissible frequency, Δ^* , in a circularly bandlimited function rectangularly sampled with spacing δ . Substituting the local sampling matrix into the definition of sampling spectral radius (2) and simplifying yields the *local sampling spectral radius* $\Delta(u, v)$:

$$\Delta(u, v) \equiv \frac{\pi}{\delta \Gamma(u, v)} \quad \text{where} \quad \Gamma(u, v) \equiv \frac{\tau(u, v)}{\min_{(n_1, n_2) \neq (0,0)} (\|n_1\mathbf{S}_u + n_2\mathbf{S}_v\|)}$$

where n_1 and n_2 are integers. We call Γ the *local sampling spectral stretch* of the mapping. The minimum local sampling spectral radius over the parameter domain, Δ^* , is then defined via

$$\Delta^* \equiv \min_{(u, v) \in \mathbf{D}} \Delta(u, v) = \frac{\pi}{\delta \Gamma^*} \quad \text{where} \quad \Gamma^* \equiv \max_{(u, v) \in \mathbf{D}} \Gamma(u, v)$$

As expected, for a given mapping function, we can make Δ^* arbitrarily large by reducing δ , essentially, by adding more samples.

Mapping efficiencies can be compared by fixing the sampling spectral radius and determining the number of samples required to generate that desired frequency content. The number of samples required to sample the domain \mathbf{D} is given by

$$N \equiv \frac{A}{\delta^2} = A / \left(\frac{\pi}{\Gamma^* \Delta^*} \right)^2 = (A \Gamma^{*2}) \left(\frac{\Delta^*}{\pi} \right)^2$$

where A is the area of \mathbf{D} . Defining the (*spectral*) *sampling efficiency* of a mapping, $\Psi^{-1} \equiv 1/(A \Gamma^{*2})$, then the number of samples required to achieve a minimum spectral radius Δ^* is $N = (\Delta^*/\pi)^2 / \Psi^{-1}$. Higher efficiency requires fewer samples for a given spectral radius everywhere in \mathbf{D} . We call the reciprocal of sampling efficiency the (*spectral*) *sampling requirement*, defined

$$\Psi \equiv A \Gamma^{*2}$$

2.3 Local Distortion Analysis

A mapping function locally transforms an infinitesimal circle into a general ellipse, with eccentricity and rotation that can be determined from the Jacobian of the mapping. The lengths of the major and minor axes of this ellipse, λ_1 and λ_2 , are given by the singular values of the Jacobian

$$\lambda_1^2(u, v) \equiv 1/2 \left((a+c) + \sqrt{(a-c)^2 + 4b^2} \right)$$

$$\lambda_2^2(u, v) \equiv 1/2 \left((a+c) - \sqrt{(a-c)^2 + 4b^2} \right)$$

where $a(u, v) \equiv \mathbf{S}_u \square \mathbf{S}_u$, $b(u, v) \equiv \mathbf{S}_u \square \mathbf{S}_v$, and $c(u, v) \equiv \mathbf{S}_v \square \mathbf{S}_v$. So λ_1 represents the maximum local stretch or length of the longest vector mapped from the set of unit tangent vectors in the domain, and λ_2 the maximum compression or length of the shortest mapped vector. The singular values are related to the differential area via $\tau = \lambda_1 \lambda_2$. We also define the *anisotropy* of the mapping, $\alpha \equiv \lambda_2/\lambda_1$, $\alpha \geq 1$, whose magnitude is useful for measuring the severity of artifacts from MIPMAP filtering.

Mapping functions can be categorized by their local properties as follows:

$$\begin{aligned} \text{isometric} &\Rightarrow \lambda_1(u, v) = \lambda_2(u, v) = \eta \\ \text{conformal} &\Rightarrow \lambda_1(u, v) = \lambda_2(u, v) \\ \text{equivalent or area-preserving} &\Rightarrow \lambda_1(u, v) \lambda_2(u, v) = \eta \\ \text{stretch-preserving} &\Rightarrow \lambda_1(u, v) = \eta \end{aligned}$$

where η is a constant and the properties are for all $(u, v) \in \mathbf{D}$.

The sampling spectral stretch of a mapping Γ is closely related to its largest singular value λ_j . $\Gamma \leq \lambda_1$ because

$$\begin{aligned} \Gamma &= \frac{\tau}{\min_{(n_1, n_2) \neq (0,0)} (\|n_1\mathbf{S}_u + n_2\mathbf{S}_v\|)} = \frac{\lambda_1 \lambda_2}{\min_{(n_1, n_2) \neq (0,0)} (\|n_1\mathbf{S}_u + n_2\mathbf{S}_v\|)} \\ &= \lambda_1 \frac{\min_{\|x, y\|=1} (\|x\mathbf{S}_u + y\mathbf{S}_v\|)}{\min_{(n_1, n_2) \neq (0,0)} (\|n_1\mathbf{S}_u + n_2\mathbf{S}_v\|)} \leq \lambda_1 \end{aligned}$$

and since $\|(n_1, n_2)\| \geq 1$ for integers not both 0. In the case of an orthogonal parameterization (in which $b \equiv \mathbf{S}_u \square \mathbf{S}_v = 0$), $\Gamma = \lambda_1$.

We also have $\Psi \leq \rho/k$ where ρ is the surface area of the output of the mapping function over \mathbf{D} and $k \equiv 2\sqrt{3}/3$. This is because isotropic hexagonal sampling is most efficient [14], for which $\mathbf{K}(u, v) = \mathbf{V}_{hex}$. For an isometric map, $\Psi = \rho$.

In practice, graphics hardware properly reconstructs maps having only limited anisotropy, something ignored by the Ψ metric. For example, the 2D shearing transformation $\mathbf{S}(u, v) = (u, v + tu)$ has $\Gamma = 1$ and thus constant sampling requirement for all t (i.e., no matter how much it shears), while $\lambda_1^2 = 1/2(2 + t^2 + \sqrt{t^4 + 4t^2})$ which increases as t increases. We therefore define another measure of sampling efficiency which models hardware without any anisotropic reconstruction filtering; i.e., only performing isotropic reconstruction filtering based on largest stretch. The *stretch sampling requirement*, \mathcal{I} , is defined

$$\mathcal{I} \equiv A \lambda_1^{*2} \quad \text{where} \quad \lambda_1^* \equiv \max_{(u, v) \in \mathbf{D}} \lambda_1(u, v)$$




Figure 4: Azimuthal Projections: Spherical circles having a radius of 3.5° and arranged at intervals of 12.5° in latitude and 22.5° in longitude are projected back into parameter space using the four azimuthal projections. The red circle represents the extent of a hemisphere. For the gnomonic map, which is unable to represent the entire hemisphere in a finite domain, a single face from the cube map is shown in red. Note that greater mapping stretch is indicated by smaller projected circles in the domain since we are projecting constant size spherical circles back into parameter space.

Then $\lambda_1 \geq \Gamma \Rightarrow \Upsilon \geq \Psi$. The smaller singular value is ignored assuming the source texture has been properly bandlimited to account for the worst case stretch (λ_1^*) by prefiltering more over the compressed direction corresponding to λ_2 . Optimizing with respect to this metric then reserves the hardware's limited anisotropic filtering capability to handle distortions arising from the perspective transformation and/or reflections rather than due to the parameterization.¹

3. Simple Spherical Maps

We analyze the sampling efficiency and anisotropy of simple projections from cartography [16], some of which have appeared in computer graphics. Where possible, we retain the terminology of cartography and relate it to that used in computer graphics.

3.1 Azimuthal Projections

Azimuthal projections transform the sphere into a tangent or intersecting plane such that parallels (lines of constant latitude on the earth) are projected to circles. Mathematically, these projections may be modeled as

$$S(u, v) = \left(\frac{u}{r} \sin \theta(r), \frac{v}{r} \sin \theta(r), \cos \theta(r) \right) \quad (4)$$

where $r = \sqrt{u^2 + v^2}$, $u, v \in [-r^*, r^*]$, $\theta(r) \in [0, \pi]$ arbitrarily reparameterizes the parallel spacing, and $\theta = 0, \pi$ represent the poles. Figure 5 summarizes the four most important azimuthal projections from cartography: gnomonic, stereographic, Lambert equal area, and equidistant, which are illustrated in Figure 4.

The *gnomonic* map projects the sphere onto a plane tangent to it, using a perspective transformation that looks directly at the point of tangency. It is the projection used in the “cube map” spherical image [7][15], for each of its six faces. The gnomonic map projects great circles on the sphere to straight lines in the map domain, an advantageous property for texture coordinate interpolation. Gnomonic maps can also be directly produced using the perspective projection of the rendering hardware. The *stereographic* map is a conformal map that also has the property that circles on the sphere project to circles in the domain. [9] proposes dual stereographic maps, one for each hemisphere, to parameterize environment maps. The *Lambert equal area* map is an area-

	Equidistant	Gnomonic	Stereographic	Lambert Equal Area
$\theta(r)$	$(\pi/2)r$	$\cos^{-1}(\sqrt{1/(r^2+1)})$	$\cos^{-1}(1-r^2/(1+r^2))$	$\cos^{-1}(1-r^2)$
properties	stretch-preserving	great circles project to lines	conformal, circles project to circles	area-preserving
r^* covers hemisphere	$[0, 1], [0, 2]$	$[0, \infty], -$	$[0, 1], [0, \infty]$	$[0, 1], [0, \sqrt{2}]$
$\lambda_1(\theta)$	$\pi/2$	$\cos \theta$	$1 + \cos \theta$	$2/\sqrt{1 + \cos \theta}$
$\lambda_2(\theta)$	$(\pi/2)\text{sinc} \theta$	$\cos^2 \theta$	$1 + \cos \theta$	$\sqrt{1 + \cos \theta}$
$\lambda_1^*(\theta)$	$\pi/2$	1	2	$2/\sqrt{1 + \cos \theta}$
$\Psi(\theta) = \Upsilon(\theta)$	$4\theta^2$	$4 \tan^2 \theta$	$16 \tan^2(\theta/2)$	$16 \tan^2(\theta/2)$
inverse map	$f = \frac{\pi}{2} \text{sinc}(\cos^{-1} z)$ $u = x/f$ $v = y/f$	$u = x/z$ $v = y/z$	$u = x/(1+z)$ $v = y/(1+z)$	$u = x/\sqrt{1+z}$ $v = y/\sqrt{1+z}$

Figure 5: Table of Mathematical Properties of Azimuthal Projections. We define $\text{sinc} \theta \equiv \sin \theta / \theta$. Note that the local properties do not vary as a function of the direction of the vector (u, v) , only as a function of its length r or equivalently, θ . The maximum larger singular value, λ_1^* , and the sampling requirement, Ψ , are taken over the portion of the sphere from $[0, \theta]$, thus allowing analysis of parts of the sphere from the pole to any parallel, such as the hemisphere ($\theta = \pi/2$). The inverse maps implement the texture coordinate generation required by graphics systems.

preserving map also called the “gazing ball” or OpenGL map [13]. Finally, the *equidistant* map is a stretch-preserving map that also preserves distances to the pole. Although it has not been used in computer graphics, its sampling efficiency exceeds that of the other azimuthal projections, as we will demonstrate shortly.

The local distortion properties of the four maps can be derived from (4) using the respective definition of $\theta(r)$ from the table. To derive the sampling efficiency of these maps, first note that the singular values of the projections are invariant over any circle in the domain centered at the origin, $D_r \equiv \{(u, v) \mid u^2 + v^2 = r^2\}$. It can also be shown that the metric tensor entries are given by

$$a = u^2 d + e, \quad c = v^2 d + e, \quad |b| = uvd$$

where $d \equiv (\lambda_1^2 - \lambda_2^2)/r^2$, $e \equiv \lambda_2^2$, so $d, e \geq 0$. Then

$$\Gamma^2 = \frac{\lambda_1^2 \lambda_2^2}{\min_{(n_1, n_2) \neq (0,0)} (n_1^2 a - 2n_1 n_2 |b| + n_2^2 c)}$$

Maximizing Γ^2 over a domain circle D_r requires minimizing the denominator above, since the numerator is invariant over D_r ; denote this maximum as Γ_r^2 . Examining the denominator, which we denote γ^2 ,

¹ Unfortunately, this model does not fully account for difficulties from anisotropy unless the hardware can isolate the parameterization distortion from the “real” (projection) distortion and compensate for prefiltering by a larger kernel in the λ_2 direction. This is why we also analyze each map's anisotropy as well as its sampling requirement.

$$\begin{aligned} \gamma^2 &\equiv \min_{(n_1, n_2) \neq (0,0)} (n_1^2 a - 2n_1 n_2 |b| + n_2^2 c) \\ &= \min_{(n_1, n_2) \neq (0,0)} ((n_1^2 + n_2^2) e + (n_1 u - n_2 v)^2 d) \end{aligned}$$

It can be seen that the minimum value of γ^2 over D_r , which we denote γ_r^2 , occurs at $(n_1, n_2) = (0, 1)$, $(u, v) = (r, 0)$, since $d, e \geq 0$ and the integer factors can not both be zero. Thus, $\gamma_r^2 = e = \lambda_2^2(r)$ so $\Gamma_r^2 = \lambda_1^2(r)$; in words, for any azimuthal projection, the maximum spectral stretch over a circle centered at the pole is the same as the larger singular value anywhere on the circle. So the spectral and stretch sampling requirements agree.

The sampling requirement tabulated in Figure 5 is then given by the domain area times the square of the maximum λ_1 ; i.e.,

$$\Psi(\theta) = \Gamma(\theta) = 4r^2(\theta) \lambda_1^{*2}(\theta) \quad \text{where} \quad \lambda_1^*(\theta) \equiv \max_{\phi \in [0, \theta]} \lambda_1(\phi)$$

While this domain is properly a disk in 2D, we take its area as $4r^2$ rather than πr^2 to calculate efficiency, assuming the disk samples must be embedded in a square to allow a practical indexing scheme. Because the equidistant map is stretch-invariant, its maximum spectral stretch does not depend on θ , and thus has optimal sampling efficiency among all azimuthal projections.

Figure 7 graphs sampling requirements of the various azimuthal maps for portions of the sphere up to a hemisphere. A single gnomonic map can't encompass the entire hemisphere; its sampling requirement is unbounded. The equidistant map has $\Psi(\pi/2) = 4(\pi/2)^2 \approx 9.87$ while both the stereographic and equal area maps have $\Psi(\pi/2) = 16 \tan^2(\pi/4) = 16$. To cover the sphere, it is more efficient to use two maps each covering a hemisphere (called a *dual projection*) rather than a single one covering the entire sphere. The dual equidistant map has sampling requirement $2\Psi(\pi/2) \approx 19.74$.

The sampling efficiency of a set of gnomonic maps, one for each face of a polyhedron circumscribing the unit sphere, is given by the surface area of that polyhedron. This is because the maximum spectral stretch of the gnomonic map is 1 and occurs at the center (pole) of the projection within each face; the domain area is the same as the area of the face. The following table lists the sampling requirements, Ψ , and maximum anisotropies, $\alpha^* \equiv \max_{(u,v) \in D} \alpha(u,v)$, of gnomonic map sets generated from the platonic solids to cover the sphere. It assumes triangular faces can be represented without wasting texture area:

Solid	Ψ	α^*	Components
tetrahedron	41.57	3.00	4
cube	24	1.73	6
octahedron	20.78	1.73	8
dodecahedron	16.65	1.26	12
icosahedron	15.16	1.26	20

Note that the sampling efficiency of dual equidistant maps is still

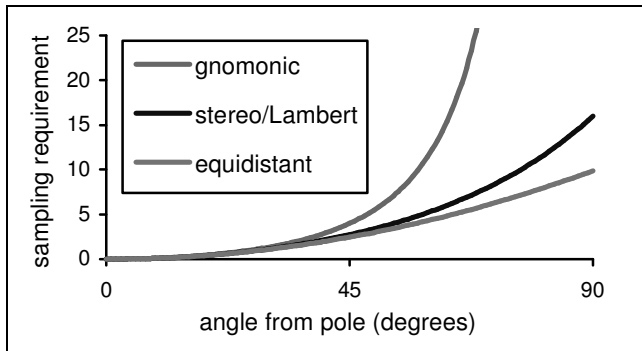


Figure 7: Sampling Requirements of Azimuthal Maps as a Function of Angular Coverage.

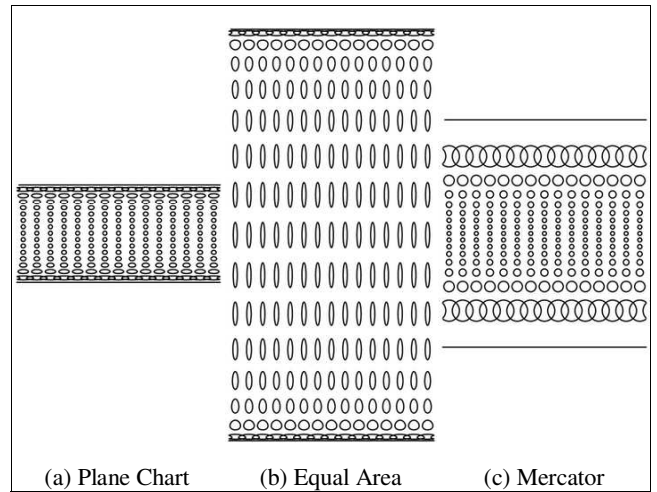


Figure 6: Cylindrical Projections: Spherical circles are distributed as defined in Figure 4, and projected into the parameter space of the three cylindrical mappings. u values are charted horizontally and v values vertically. Note that the Mercator projection has an unbounded domain in v , but only the spherical area very close to the poles is greatly stretched: the 3.5° spherical circles surrounding the poles map to the horizontal lines shown at the top and bottom.

better than the octahedral gnomonic map set. By using gnomonic maps derived from tessellations of more and more faces, we can approach but never attain the sampling requirement of an isometric map of 4π .

Considering the anisotropy of the azimuthal mappings, the stereographic map is best since it is conformal, but the equidistant map exhibits little anisotropy for angles less than about 45° , where it has $\alpha(\pi/4) = \sqrt{2}\pi/4 \approx 1.1$. This can be compared to its still reasonable anisotropy at the equator, $\alpha(\pi/2) = \pi/2 \approx 1.57$. Gnomonic maps have worst-case anisotropy at the greatest angular distance from the pole. For example, the cube map has $\alpha^* = \sqrt{3} \approx 1.73$ at the cube vertices.

3.2 Cylindrical Projections

Cylindrical projections transform the sphere to a tangent or intersecting cylinder such that parallels are projected to straight lines, with the model

$$S(u, v) = (\cos(2\pi u) \cos \theta(v), \sin(2\pi u) \cos \theta(v), \sin \theta(v))$$

where $u \in [-1/2, 1/2]$, $\theta(v) \in [-\pi/2, \pi/2]$, $\theta(v)$ arbitrarily reparameterizes the parallel spacing, and where $\theta = \pm\pi/2$ represent the poles, and $\theta = 0$ the equator. Three important cylindrical projections, the plane chart, equal area, and Mercator, are illustrated in Figure 6 with mathematical properties in Figure 8.

The *plane chart* is the standard latitude/longitude parameterization of the sphere. It is also stretch-preserving. As its name implies, the *equal area* cylindrical projection preserves area. Finally the *Mercator* projection is a conformal projection useful for navigation since it projects *loxodromes* on the sphere, or curves making a constant angle with the meridians, into straight lines. Since all cylindrical projections are orthogonal, their spectral stretch is identical to their larger singular value, so the two sampling metrics agree.

To analyze sampling efficiency, we consider the portion of the sphere from the equator to an angle $\theta \leq \pi/2$. The sampling requirement tabulated in Figure 8 and graphed in Figure 9 is

$$\Psi(\theta) = \Gamma(\theta) = v(\theta) \lambda_1^{*2}(\theta) \quad \text{where} \quad \lambda_1^*(\theta) = \max_{\phi \in [0, \theta]} \lambda_1(\phi)$$

	Plane Chart	Equal Area	Mercator
$\theta(v)$	$2\pi v$	$\sin^{-1} v$	$\sin^{-1}(\tanh(2\pi v))$
properties	stretch-preserving	area-preserving	conformal
v covers sphere	$[-1/4, 1/4]$	$[-1, 1]$	$[-\infty, \infty]$
$\lambda_1(\theta)$	2π	$\max(1/\cos\theta, 2\pi\cos\theta)$	$2\pi\cos\theta$
$\lambda_2(\theta)$	$2\pi\cos\theta$	$\min(1/\cos\theta, 2\pi\cos\theta)$	$2\pi\cos\theta$
$\lambda_1^*(\theta)$	2π	$\max(1/\cos\theta, 2\pi)$	2π
$\Psi(\theta) = \mathcal{I}(\theta)$	$2\pi\theta$	$\max(1/\cos^2\theta, 4\pi^2)\sin\theta$	$2\pi\tanh^{-1}(\sin\theta)$ $= \pi\ln((1+\sin\theta)/(1-\sin\theta))$
inverse map	$u = (\text{atan } 2(y,x))/(2\pi)$ $v = (\sin^{-1} z)/(2\pi)$	$u = (\text{atan } 2(y,x))/(2\pi)$ $v = z$	$u = (\text{atan } 2(y,x))/(2\pi)$ $v = \tanh^{-1} z/(2\pi)$ $= \ln((1+z)/(1-z))/(4\pi)$

Figure 8: Table of Mathematical Properties of Cylindrical Projections. The maximum singular value and sampling requirement are taken over the part of the sphere from $[0, \theta]$; i.e., from the equator to the θ parallel. Local properties of the cylindrical projections are invariant with respect to u , depending only on v or inverting, θ .

since the domain area is equal to $v(1/2 - -1/2) = v$. The plane chart is most sampling-efficient because it is stretch-invariant, just as the equidistant map is best among azimuthal projections. Covering the hemisphere using the plane chart has sampling requirement $\Psi(\pi/2) = 2\pi(\pi/2) \approx 9.87$, identical to that of the azimuthal equidistant map. The other two maps have unbounded sampling requirement for hemispherical coverage. The cylindrical equal area map has very poor sampling efficiency for any angular coverage, while the Mercator projection, at least for angles not more than about 45° , is only slightly worse than the plane chart, with the advantage of conformality.

At small angles from the equator, the plane chart has little anisotropy. But its anisotropy increases without bound near the pole making it a poor choice for MIPMAP texturing. The area-preserving map has poor anisotropy as well as sampling efficiency.

3.3 Low-Distortion Equal-Area Map

We also present an unusual sphere mapping developed for use in a stochastic ray tracer. This mapping was designed to project stratified and other specialized sampling patterns onto the sphere and to project spherical samples into 2D histogram bins. To solve those problems, the mapping had to be a bijection between the unit square and the unit sphere, area-preserving, and not severely anisotropic. Its projection, $(u, v) = S^{-1}(x, y, z)$, is defined by the composition of three area-preserving bijections. The first is a mapping from a hemisphere to a disk: $(u, v) = (x, y)/\sqrt{1+z}$. The second mapping is from a disk to a half disk: $(r', \theta') = (r, \theta/2)$.

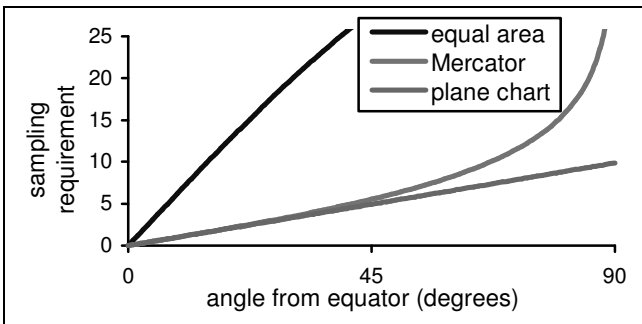


Figure 9: Sampling Requirements of Cylindrical Maps as a Function of Angular Coverage.

Thus the two halves of a sphere are converted into half disks, which are joined to form a single disk. The third mapping is Shirley's area-preserving bijection between the disk and the unit square [18]. This leaves us with an image of the sphere on the unit square, where the north and south pole are both mapped to the center of the square. If we roll the mapping halfway, the north pole will be on one edge of the square, the south pole on the opposite edge, and the interior of the square will be C_0 continuous. See Figure 11d for the mapping applied to the image of the earth.

Numerical analysis of this mapping shows the spectral sampling requirement to be $\Psi \approx 19.7$, about the same as the plane chart and dual equidistant mappings. But its stretch sampling requirement is much worse at $\mathcal{I} \approx 43.9$, the only map described where these metrics disagree. Its worst-case isotropy is about 0.29.

4. Polar-Capped Maps

Azimuthal and cylindrical projections are complementary, in that azimuthal projections tend to have better sampling efficiency and less anisotropy near the pole, while cylindrical projections are better near the equator. We can therefore improve both sampling efficiency and anisotropy with a three-component map set that uses an azimuthal projection for each of the two poles and a cylindrical projection near the equator, called a *polar-capped map*.

To create an optimal stretch-invariant polar-capped map, we examine the sampling requirement of a two-component map covering the hemisphere, containing a plane-chart projection near the equator and an azimuthal equidistant projection near the pole. The sampling requirement of such a map is given by

$$M_{\text{capped}}(\theta) \equiv M_{\text{equi}}(\theta) + M_{\text{plane}}(\pi/2 - \theta) = 4\theta^2 + 2\pi(\pi/2 - \theta)$$

where θ is the angle from the pole where the equidistant projection transitions to the plane chart. We seek θ minimizing the sampling requirement, which occurs at $\theta^* \equiv 45^\circ$ for which the sampling requirement is $3/4\pi^2 \approx 7.40$.

To cover the entire sphere, we can therefore continue the equatorial map to the southern hemisphere and add a third equidistant map covering the south pole to produce a polar-capped map with sampling requirement of roughly 14.80 and maximum (worst-case) anisotropy of $\alpha^* \approx 1.41$ (occurring at the 45° parallel of the plane chart). This represents an improvement of 54% over the dual stereographic map, 38% over the cube map (but with half as many map components and much less anisotropy), and 25% over the dual equidistant or plane chart maps (but with much less anisotropy). It is even a slight improvement (2.4%), over the unwieldy 20-map icosahedral gnomonic set! Moreover, no map can save more than 26% over its sampling requirement.

The domain structure for the stretch-invariant polar-capped map is




Figure 10: Stretch-Invariant Polar-Capped Map: The rectangle E represents the plane chart equatorial map, whose vertical resolution matches that of the two azimuthal equidistant polar caps, labeled N and S. Such a map with sample spacing δ produces a minimum sampling spectral radius of $\Delta = 2/\delta$, for a sampling requirement of $(3/2)\pi^2 \approx 14.8$.

shown in Figure 10. Polar-capped maps can also be defined that are conformal (to avoid anisotropy) and hexagonally reparameterized (to improve sampling efficiency slightly at the expense of greater anisotropy) [22].

5. Results

Figure 11 illustrates five texture maps of the earth with identical texture area: stretch-invariant polar-capped (a), gnomonic cube (b), dual stereographic (c), low distortion area preserving (d), and Lambert equal area (e). Note particularly the sizes of various features such as the polar ice caps. For example, it can be seen that the gnomonic cube (Figure 11a) shrinks Antarctica, and thus will sample that area less effectively than the polar-capped map (Figure 11b). The dual stereographic map (Figure 11c) reduces Antarctica still further.

Figure 12 compares the maps using a high-frequency test pattern on the sphere. This pattern is first sampled into texture maps of identical area for each of the 5 example maps. We then generate orthographic views of the sphere textured with each of these results, shown in the top row. We chose a view where the south pole has been rotated towards the viewer by 45° so that both polar and equatorial regions are visible. The row below zooms in on the south pole of the row above which represents a “bad spot”, or most undersampled region, for each of the maps.²

Our sampling efficiency metric is based on worst-case frequency preservation using the principle that all parts of the sphere must be sampled well to avoid visual artifacts. The bottom row of Figure 12 validates the mathematical ranking of the maps which decrease in sampling efficiency from left to right. Discriminating between the low distortion equal area (LDEA) map and the gnomonic cube is difficult. This is probably because we applied the spherical texture using bilinear filtering to simulate typical graphics hardware, which penalizes the more highly anisotropic LDEA map. The top row shows the polar-capped map’s sharpness uniformity over the entire sphere. The other maps have noticeable patches of blurriness near the south pole.

6. Conclusions and Future Work

Mapping functions used to represent spherical images in graphics systems are less than optimal in terms of sampling efficiency. This paper defines the notion of sampling efficiency and analyzes existing and new maps in terms of sampling efficiency and local anisotropy. We introduce pieces of projections used in cartography to form two- and three-component maps that have better sampling efficiency than the best map used previously, with little or no anisotropy. Implementation involves replacing the texture coordinate generation function with a simple alternative.

We believe the sampling efficiency metric can be extended to Monte Carlo integration for rendering. Stratified stochastic sampling seeks to scatter as few samples as possible to arrive at an integral estimate with greatest confidence, and often uses global mapping functions from simple domains [1][12]. While area-preserving mappings as in [1] assure that all samples contribute equally to the integral estimate, maps with much local stretch, a typical consequence of area-preserving mapping to surfaces with curvature, increase variance (i.e., reduce the effectiveness of stratification). How should these considerations be balanced? A surface integral of spectral stretch, rather than a simple maximum, may be the right comparison metric for mappings used for integrating rather than tabulating functions over surfaces.

7. References

- [1] Arvo, J., “Stratified Sampling of Spherical Triangles,” *Siggraph ’95*, August 1995, 437-438.
- [2] Blinn, J.F., and M.E. Newell, “Texture and Reflection in Computer Generated Images”, *Communications of the ACM*, 19(10):542-547, October 1976.
- [3] Brillouin, L., *Science and Information Theory*, Academic Press, 105-111, 1956.
- [4] Cabral, B., M. Olano, and P. Nemecek, “Reflection Space Image Based Rendering,” *Siggraph ’99*, August 1999, 165-170.
- [5] Chen, Shenchang, “QuickTime VR – An Image-Based Approach to Virtual Environment Navigation”, *Siggraph ’95*, August 1995, 29-38.
- [6] Dudgeon, Dan, and R. Mersereau, *Multidimensional Digital Signal Processing*, Prentice Hall, 1984.
- [7] Greene, N., “Environment Mapping and other Applications of World Projections,” *Computer Graphics and Applications*, 6(11):21-29, November 1986.
- [8] Haerberli, P., and M. Segal, “Texture Mapping as a Fundamental Drawing Primitive,” in *Fourth Eurographics Workshop on Rendering*, June 1993, 259-266.
- [9] Heidrich, Wolfgang, and H.P. Seidel, “Realistic, Hardware-Accelerated Shading and Lighting,” *Siggraph ’99*, August 1999, 171-178.
- [10] Hüttner, T., and W. Straßer, “Fast Footprint MIPmapping,” in *Proceedings 1999 Eurographics/Siggraph Workshop on Graphics Hardware*, Aug. 1999, 35-43.
- [11] Miller, G., M. Halstead, and M. Clifton, “On-the-fly Texture Computation for Real-Time Surface Shading,” *IEEE Computer Graphics and Applications*, March 1998, 44-58.
- [12] Mitchell, D., “Consequences of Stratified Sampling in Graphics,” *Siggraph ’96*, August 1996, 277-280.
- [13] *OpenGL Reference Manual*, Addison Wesley, 1992.
- [14] Peterson, D.P., and D. Middleton, “Sampling and Reconstruction of Wave-Number Limited Functions in N-Dimensional Euclidean Spaces,” *Information and Control*, 5(1962), 279-323.
- [15] Regan, M. and R. Pose, “Priority Rendering with a Virtual Reality Address Recalculation Pipeline”, *Siggraph ’94*, 155-162.
- [16] Robinson, Arthur, *Elements of Cartography*, John Wiley & Sons, New York, 1960.
- [17] Schröder, P., and W. Sweldens, “Spherical Wavelets: Efficiently Representing the Sphere,” *Siggraph ’95*, August 1995, 161-172.
- [18] Shirley, Peter and Kenneth Chiu, “A Low Distortion Map Between Disk and Square”, *Journal of Graphics Tools*, Vol 2, No 3, pp 45-52, 1997.
- [19] Snyder, John P., *Flattening the Earth: Two Thousand Years of Map Projections*, The University of Chicago Press, London, 1993.
- [20] Voorhies, Douglas, and Jim Foran, “Reflection Vector Shading Hardware,” *Siggraph ’94*, July 1994, 163-166.
- [21] Williams, Lance, “Pyramidal Parametrics,” *Siggraph ’83*, July 1983, 1-11.
- [22] Sampling-Efficient Mapping of Spherical Images, TR to be published, July 2000.


² The polar-capped map has no bad spot. The other maps have bad spots at both poles or at the south pole (Lambert).




(a) Stretch-invariant Polar Capped




(b) Gnomonic Cube (Cube Map)



(c) Dual Stereographic



(d) Low Distortion Equal Area



(e) Lambert Equal Area (OpenGL)

Figure 11: Texture Maps of the Earth

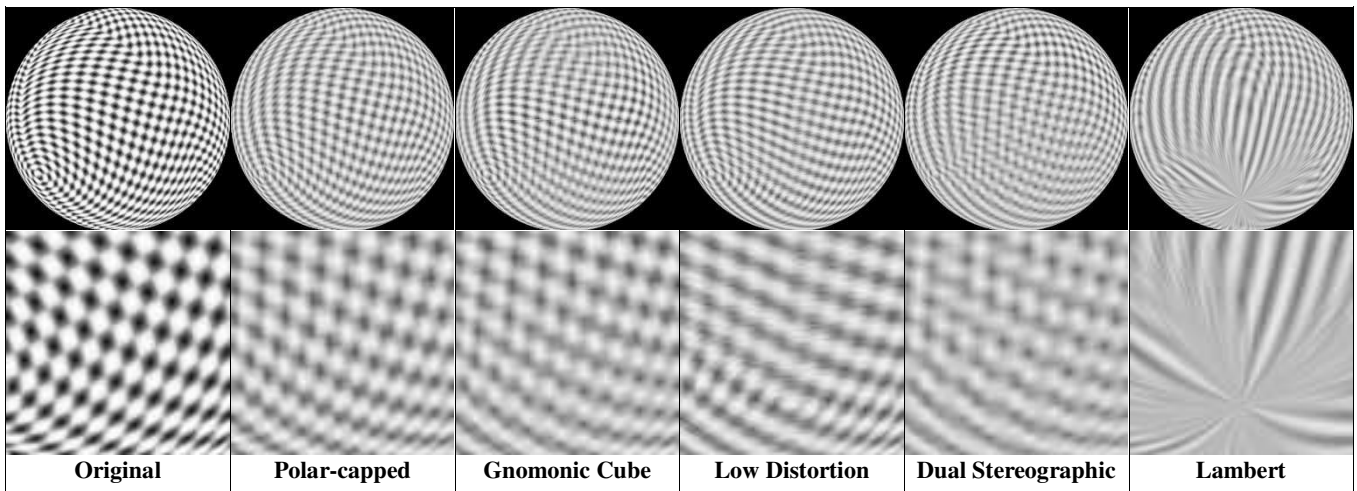


Figure 12: Results on Spherical Test Pattern