Spherical Codes for Gaussian Quantization

Jon Hamkins
Jet Propulsion Laboratory
May 28, 1997

Joint work with:
Ken Zeger, UCSD

Outline

1. Review of vector quantization (VQ) concepts
2. Vector quantization of a memoryless Gaussian source
   A. Properties of Gaussian vectors
   B. Shape-gain quantization using wrapped spherical codes
   C. Extensions

Main contribution:

Better low-complexity vector quantizers for memoryless Gaussian sources.

Qauntization

A communications system:

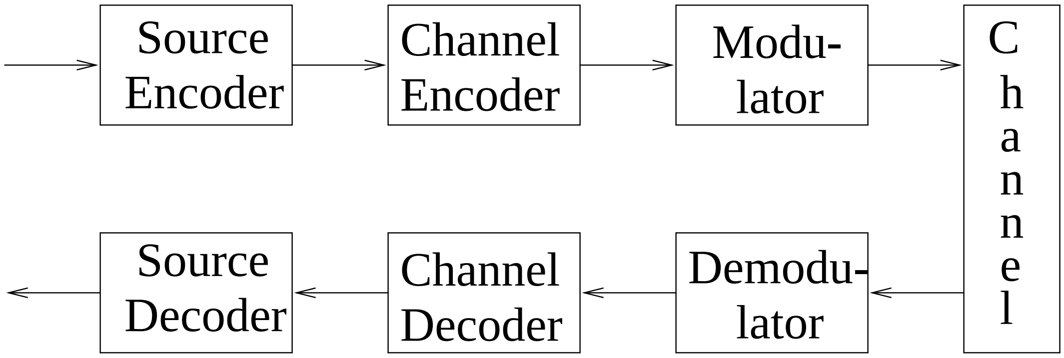

A simplified source coding model:

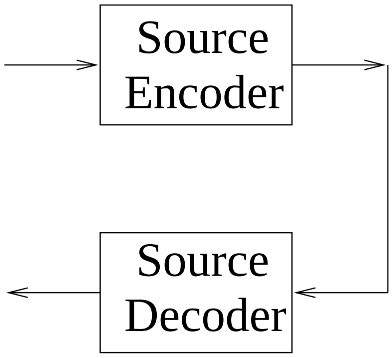

Ways to view a quantizer:

1. Analog to digital conversion

2. Signal/data compression

3. Function $Q : \mathbb{R} \to \mathbb{R}$

Quantizer Design

Try to minimize distortion:

$$D = E\|X - Q(X)\|^2$$

Example: uniform source

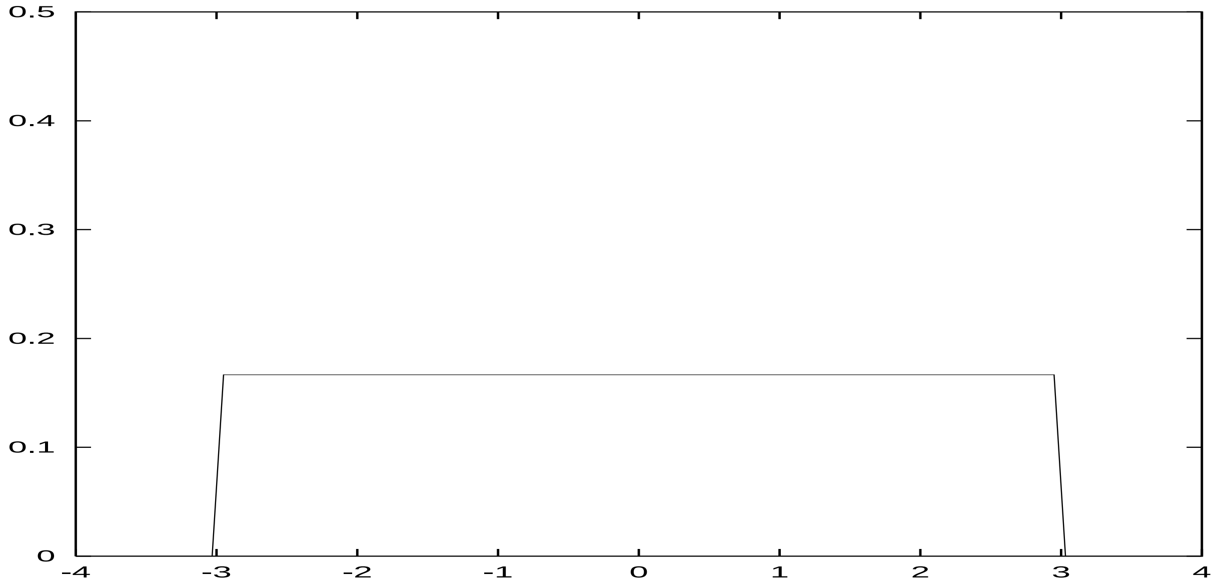

Lloyd-Max algorithm repeats the following:

1. Nearest neighbor (fix outputs, move cell boundaries)

2. Centroid condition (move outputs, fix cell boundaries)

This will converge to locally optimal quantizer.

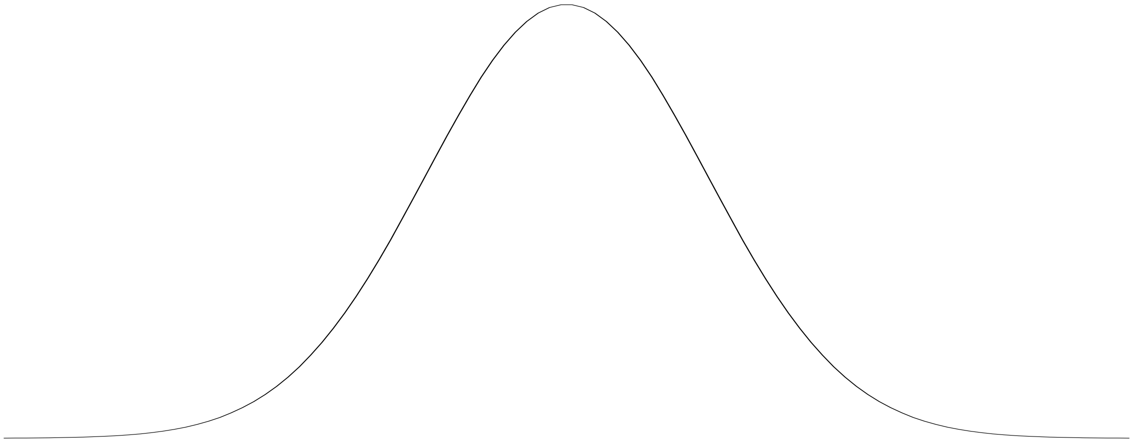

Vector Quantization

Group source samples into a vector, and quantize the whole vector.

Example: i.i.d. Uniform source, formed into 2D vectors

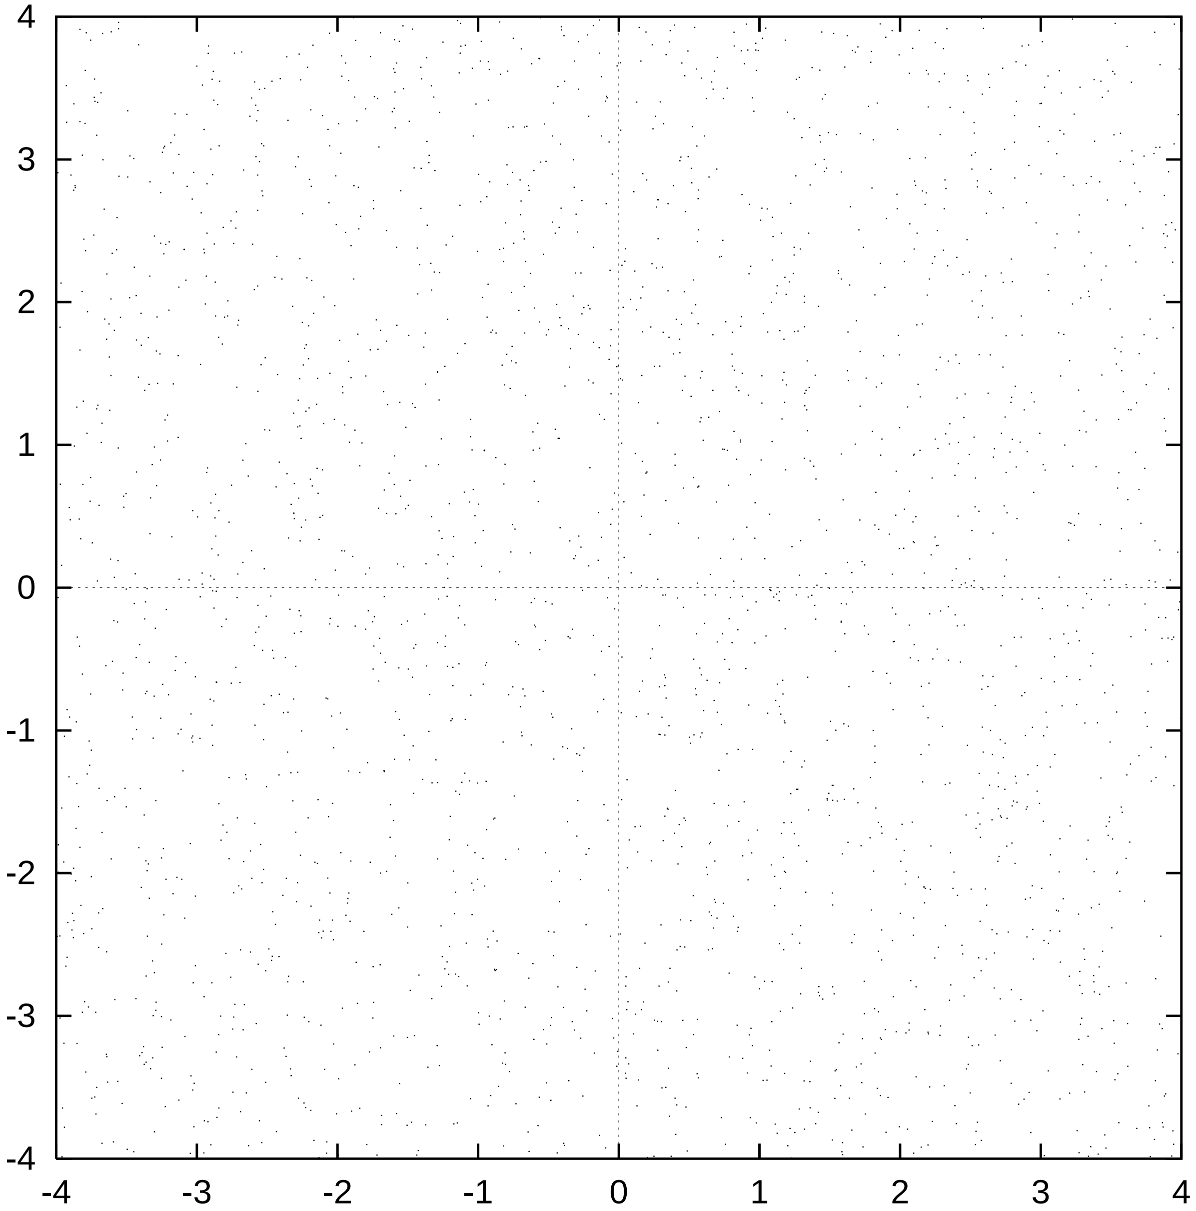

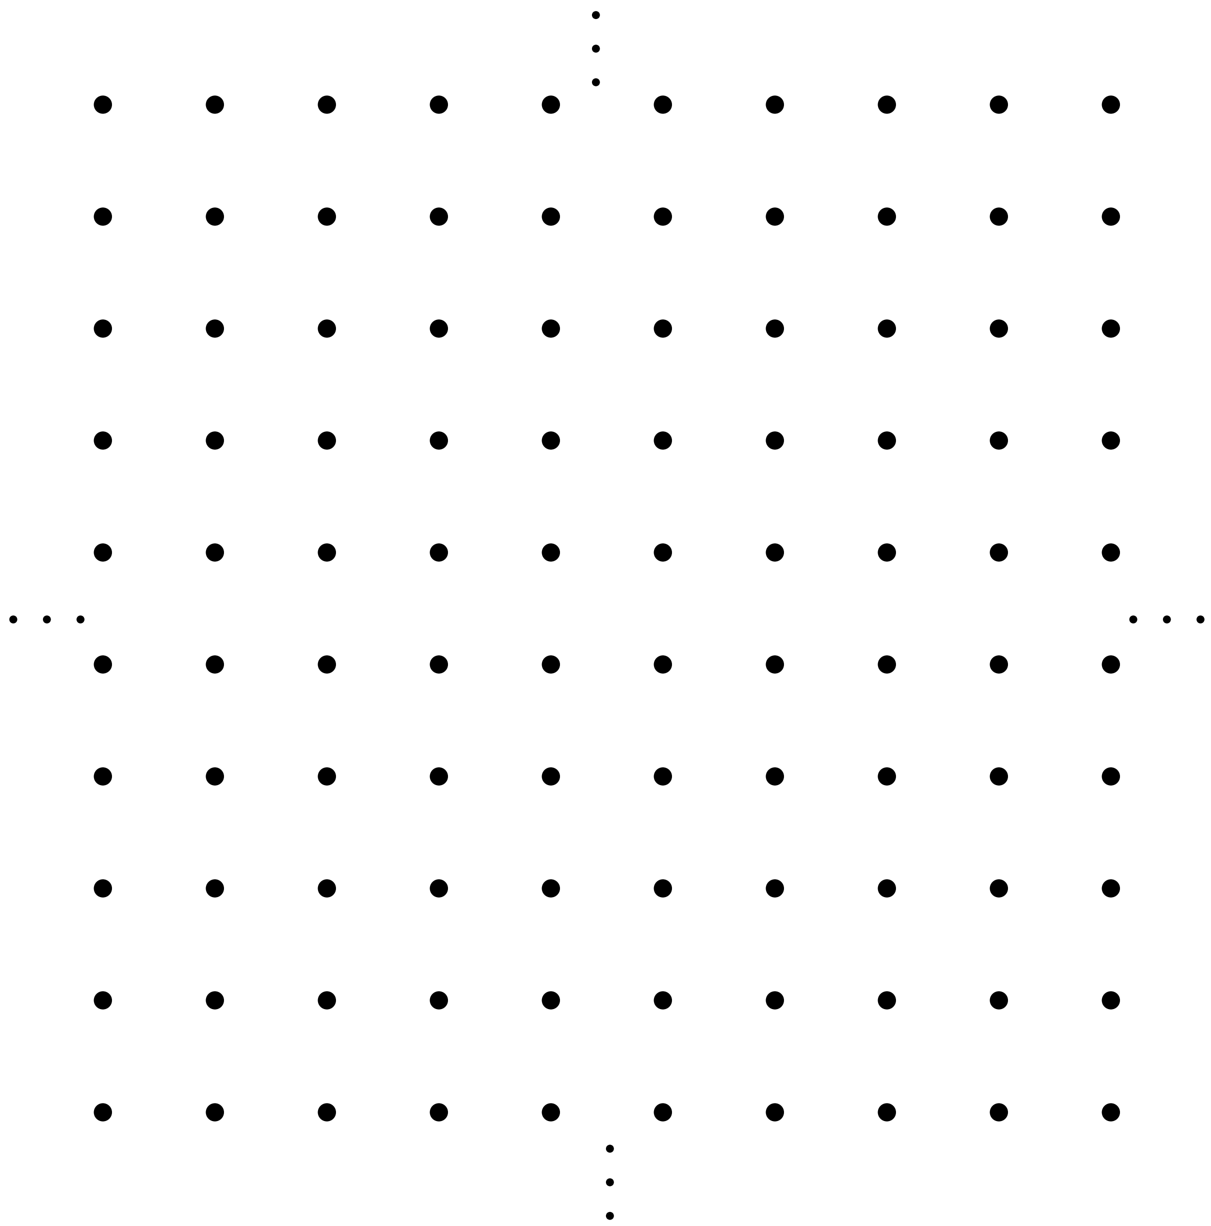

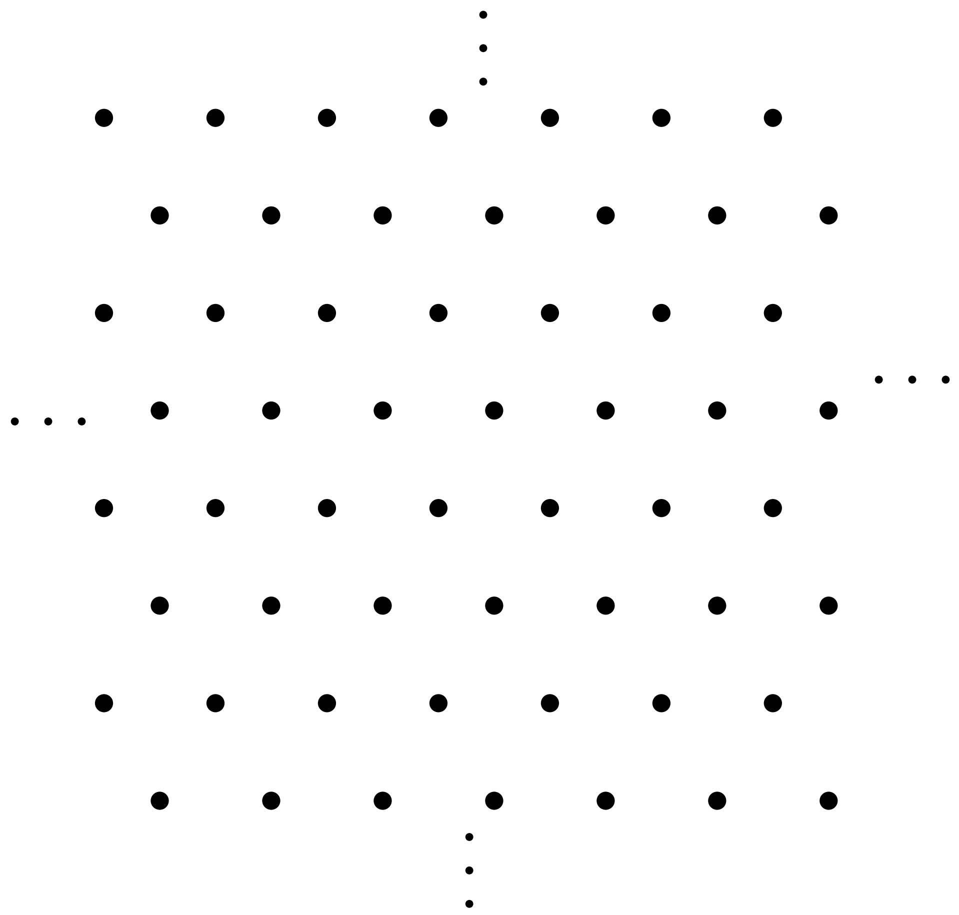

Quantization Coefficient

The quantization coefficient: average MSE per dimension for high rate quantization of a uniform source (scaled so as to be a dimensionless quantity).

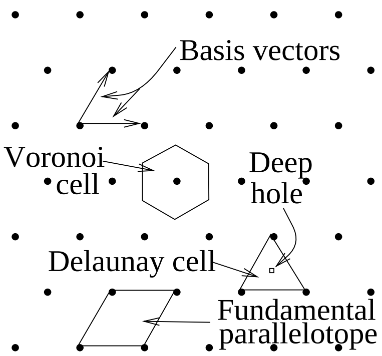

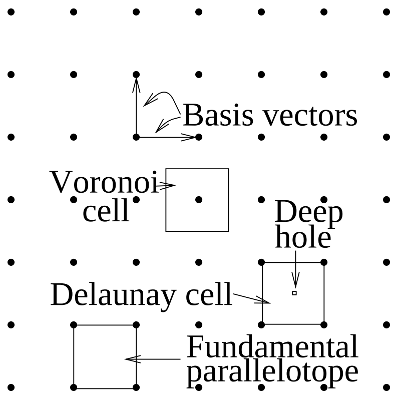

Lattice VQ's are good for uniform sources.

Dimension	Lattice	Quantization coefficient
k	$\mathbb{Z}^k$	$\frac{1}{12} \approx 0.0833$
2	A_2	$\frac{5}{36\sqrt{3}} \approx 0.0802$
3	A_3	0.0787
3	A_3^*	0.0785
24	Λ_{24}	0.0658
$\to \infty$		optimal: $\frac{1}{2\pi e} \approx 0.0586$

Note: maximizing packing density $\neq$ minimizing MSE

Dimension	Best known lattice quantizer	Densest known lattice
2	A_2^*	A_2
3	A_3^*	A_3
4	D_4^*	D_4
5	D_5^*	D_5
6	E_6^*	E_6
7	E_7^*	E_7
8	E_8^*	E_8
24	Λ_{24}^*	Λ_{24}

Implications of Shannon's Theorem

Good News:

Shannon's source coding theorem w.r.t. a fidelity criterion says that all k-dimensional quantizers $Q(\cdot)$ satisfy

$$\frac{1}{k}E\|X - Q(X)\|^2 \geq D(R)$$

and there is a sequence of quantizers such that

$$\frac{1}{k}E\|X - Q(X)\|^2 \to D(R)$$

as $k \to \infty$, where R is the rate of the quantizer ($\frac{1}{k}\log_2 M$), and where $D(R)$ is the distortion-rate function.

Bad News:

Computational complexity of encoding is

$$M = 2^{Rk},$$

i.e., exponential in both rate and dimension.

A Memoryless Gaussian Source

Let $X = (X_1, \ldots, X_k)$, where $X_i \sim N(0, \sigma^2)$.

In two dimensions ($k = 2$, $\sigma^2 = 1$):

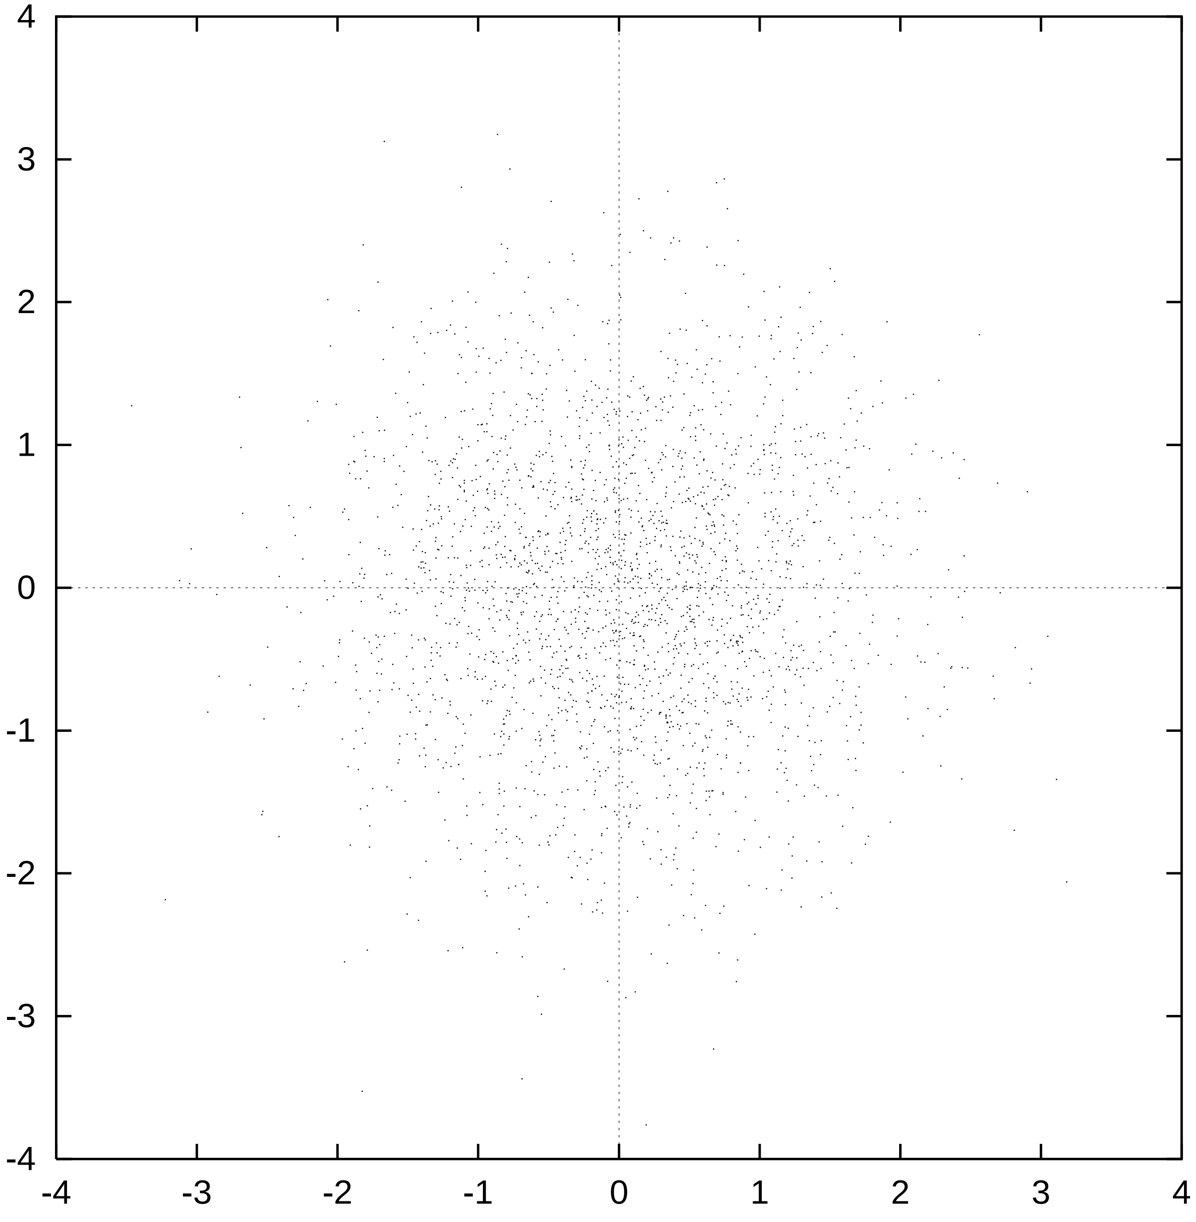

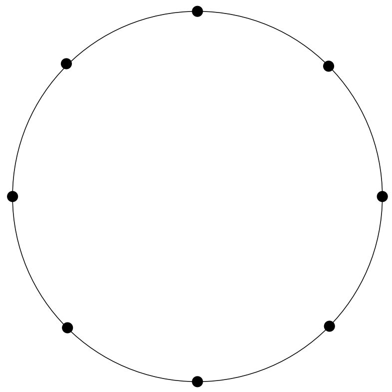

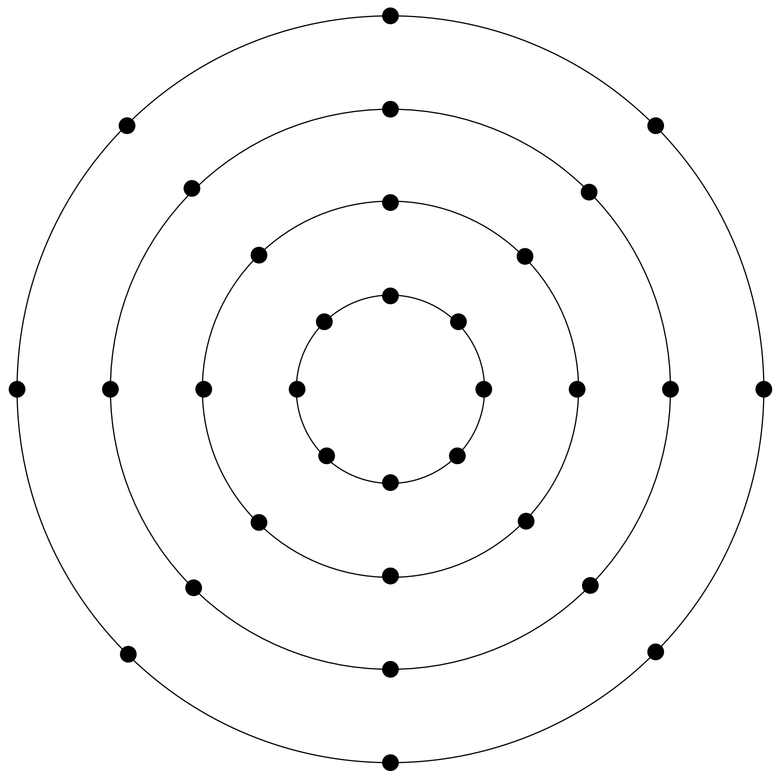

Properties of a Memoryless Gaussian Source

If $X = (x_1, \ldots, x_k)$, $x_i \sim N(0, \sigma^2)$, and if $Y = (y_1, \ldots, y_k)$, then

$$f_X(Y) = \prod_{i=1}^{k} \frac{\exp\left(\frac{-y_i^2}{2\sigma^2}\right)}{\sqrt{2\pi\sigma^2}} = \frac{\exp\left(\frac{-\|Y\|^2}{2\sigma^2}\right)}{(2\pi\sigma^2)^{k/2}}.$$

Lemma: The pdf of $g = \|X\|$ is

$$f_g(r) = \frac{2r^{k-1}\exp\left(\frac{-r^2}{2\sigma^2}\right)}{\Gamma(k/2)(2\sigma^2)^{k/2}}$$

Proof:

$$f_g(r) = \int_{\|Y\|=r} f_X(Y)\,dY = \frac{\exp\left(\frac{-r^2}{2\sigma^2}\right)}{(2\pi\sigma^2)^{k/2}} \cdot \underbrace{\int_{\|Y\|=r} dY}_{S_k r^{k-1}}. \blacksquare$$

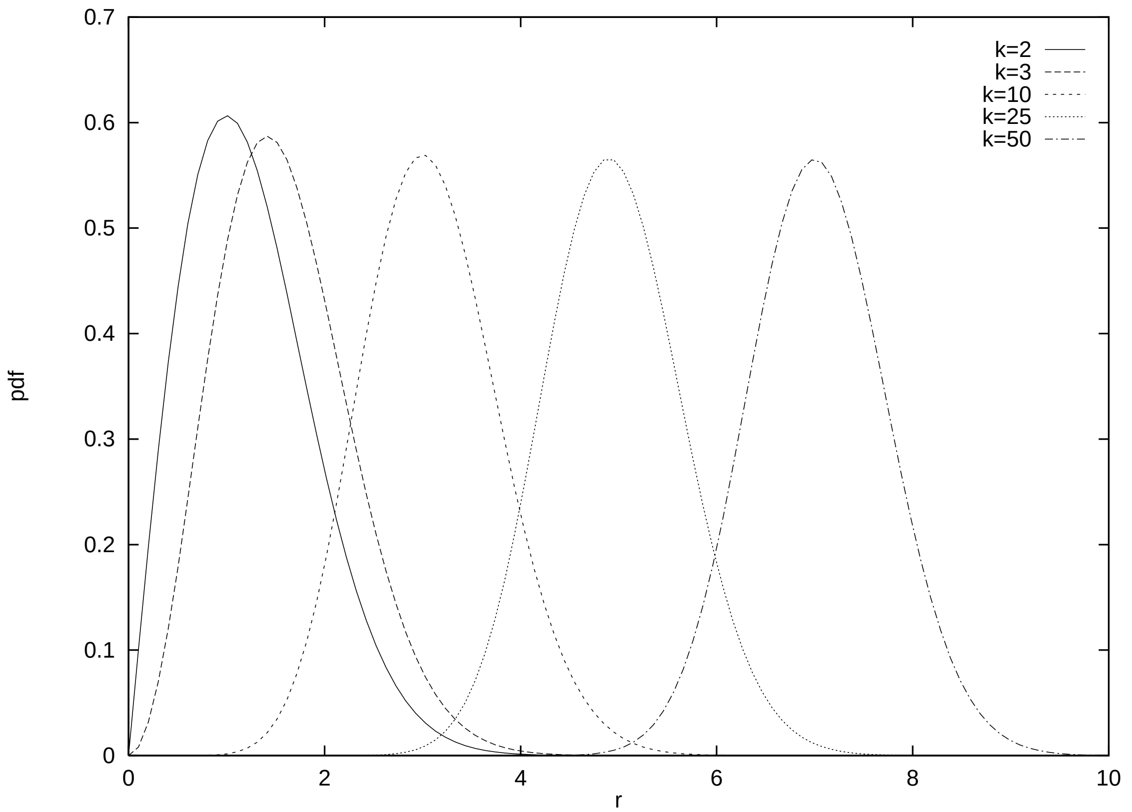

Variance of $\|X\|$

[Recall, $X = (X_1, \ldots, X_k)$, where $X_i \sim N(0, \sigma^2)$.]

$$\begin{aligned}
E\|X\| &= \frac{\sqrt{2\pi\sigma^2}}{\beta\left(\frac{k}{2}, \frac{1}{2}\right)} \approx \sigma\sqrt{k - (1/2)} \\
E\|X\|^2 &= k\sigma^2 \\
\text{var}\|X\| &= k\sigma^2 - \frac{2\pi\sigma^2}{\beta^2\left(\frac{k}{2}, \frac{1}{2}\right)}
\end{aligned}$$

Lemma: $\text{var}\|X\| < \sigma^2/2$ for all k.

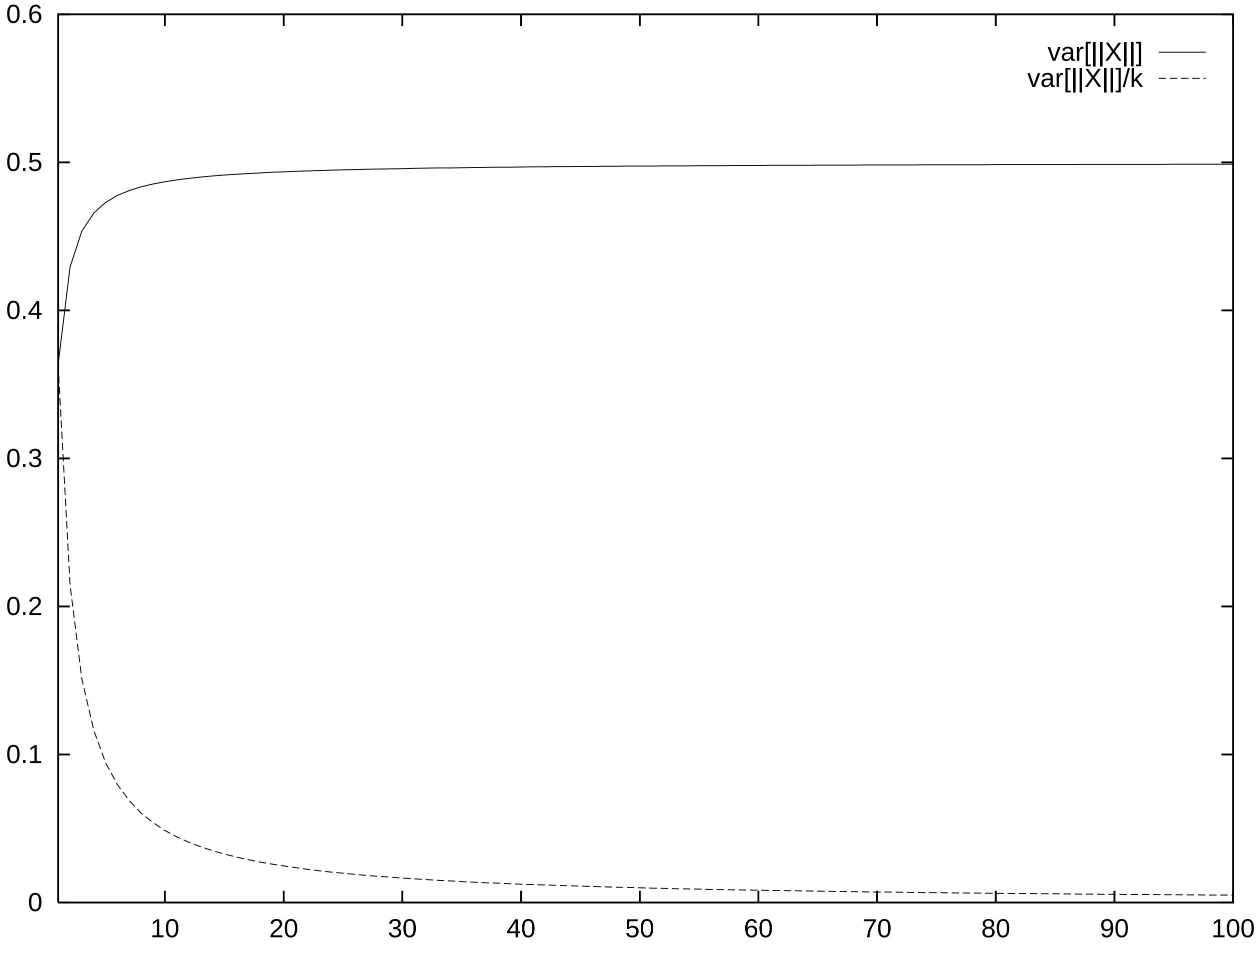

Motivation for Spherical Vector Quantization

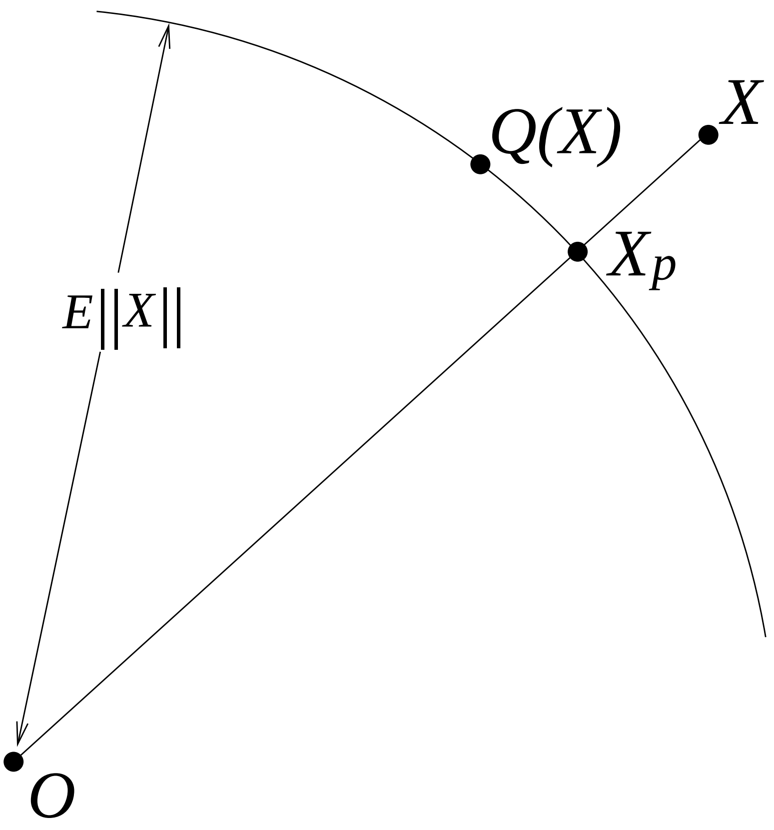

Sakrison showed (also stated as Eq. (6.6))

$$
\begin{aligned}
D &= \frac{1}{k}E\|X - Q(X)\|^2 \\
&= \frac{1}{k}E\|X_p - Q(X)\|^2 + \frac{1}{k}\underbrace{E\|X - X_p\|^2}_{\text{var}_{\|X\|}} \\
&= D_s + D_g \\
&\approx D_s, \text{ for large } k
\end{aligned}
$$

Spherical Vector Quantization

Previous work:

- Geometric development (Sakrison, '68)
- Polar Quantizers (Bucklew & Gallagher, '79, Wilson, '80)
- Spherical coordinates quant. (Swaszek & Thomas, '83)
- SVQ for speech coding (Adoul, Lamblin, & Leguyader, '84)
- Asymptotic Polar Quantizers (Swaszek & Wu, '86)
- Non-Gaussian distributions (Fischer, '86, '89)

Wrapped Spherical Vector Quantizer

Shape-gain approach: $g = \|X\|$, and $S = \frac{X}{g}$

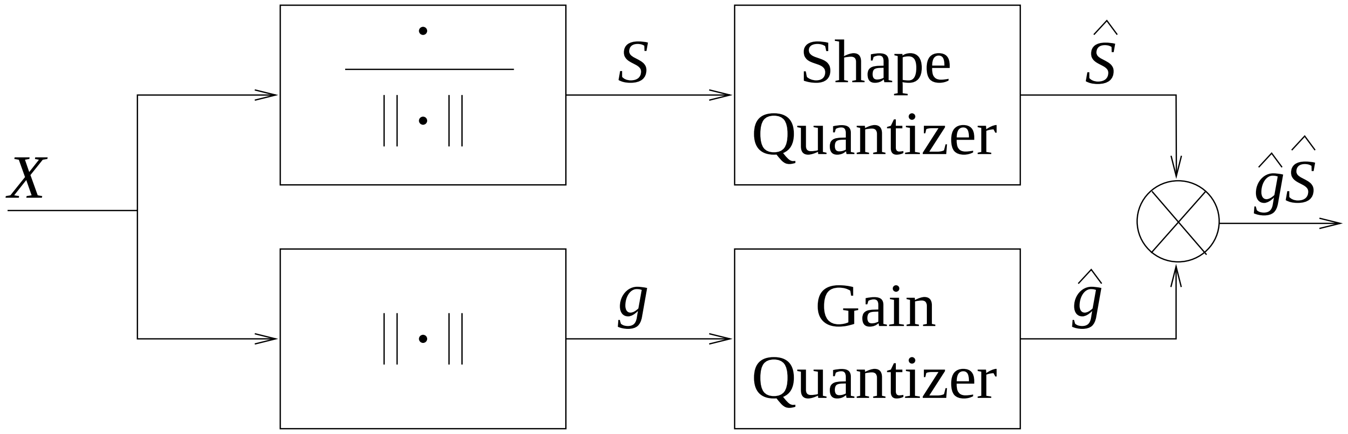

Gain codebook: scalar quantizer optimized by Lloyd-Max algorithm.

Shape codebook: wrapped spherical code, without buffer regions

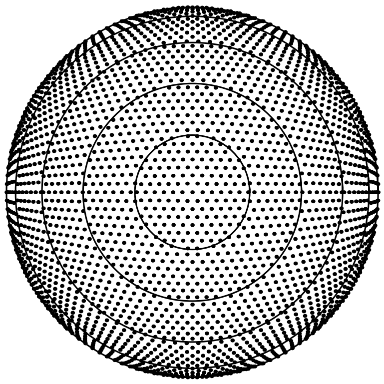

Definition of spherical code

Definition: A k-dimensional *spherical code* is a set of points which lie on the k-dimensional unit sphere.

Examples:

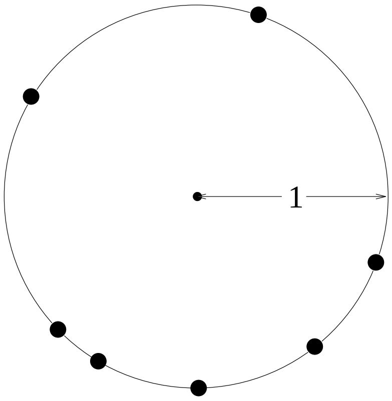

2 dimensions

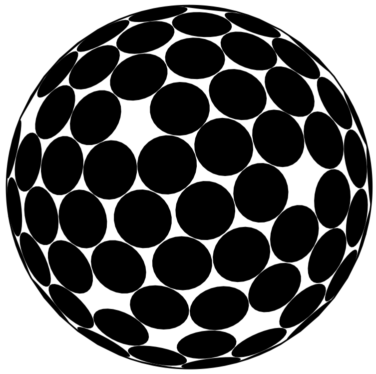

3 dimensions

In k dimensions:

$\mathcal{C} \subset \Omega_k \equiv \{(x_1, \ldots, x_k) : \sum x_i^2 = 1\}$.

Wrapped spherical code construction

Intuition: use Yaglom-type mapping, but with less distortion.

Yaglom's mapping: Project the best $(k-1)$-dimensional packing onto Ω_k.

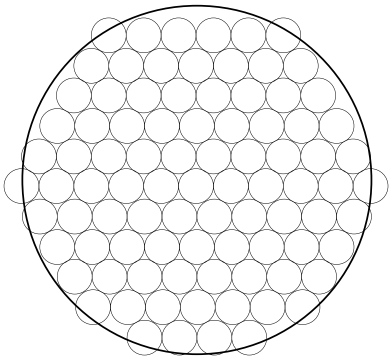

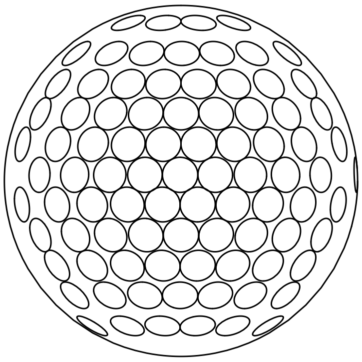

Partition Ω_k into annuli:

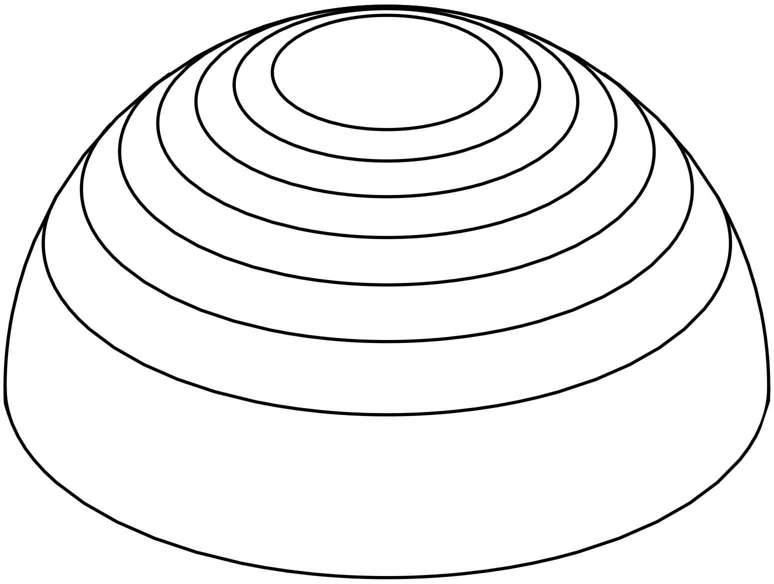

Wrapped spherical code construction

Within each annulus, introduce only small distortion:

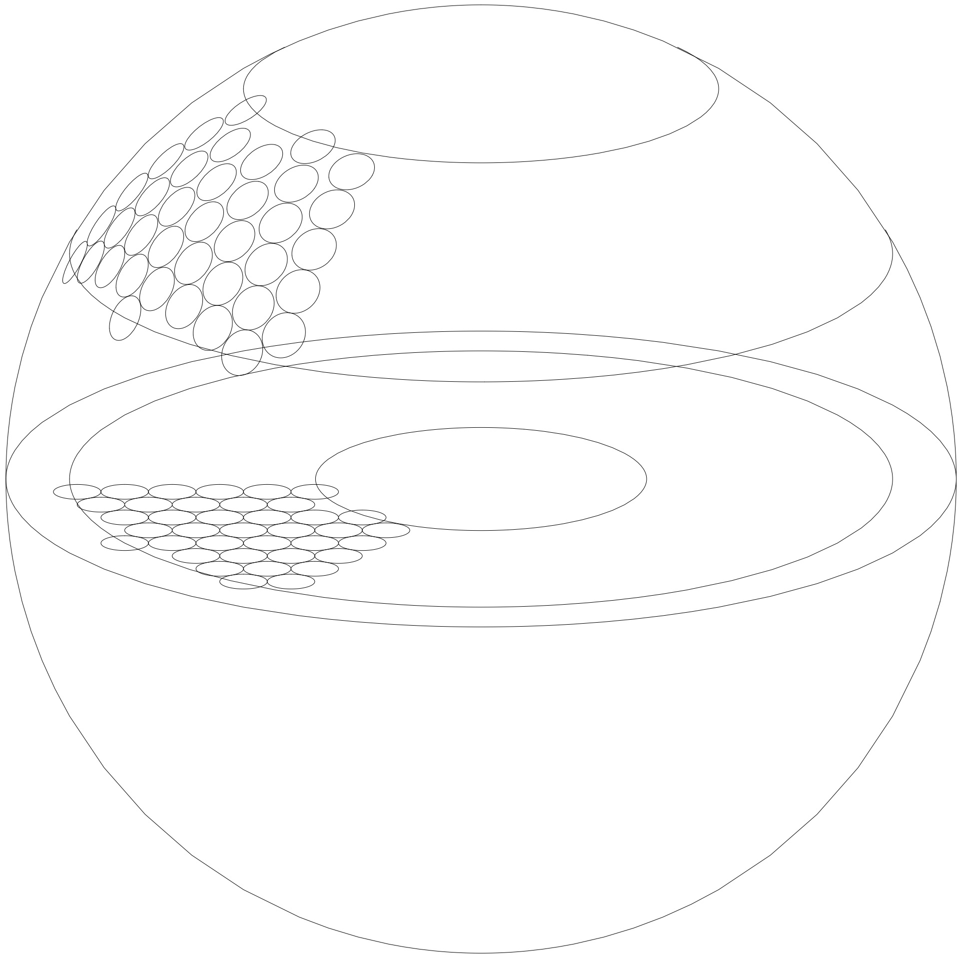

Theorem: The quantization coefficient of a wrapped spherical code is within $O(\sqrt{d})$ of the quantization coefficient of the underlying packing used to construct it.

Example of wrapped spherical code

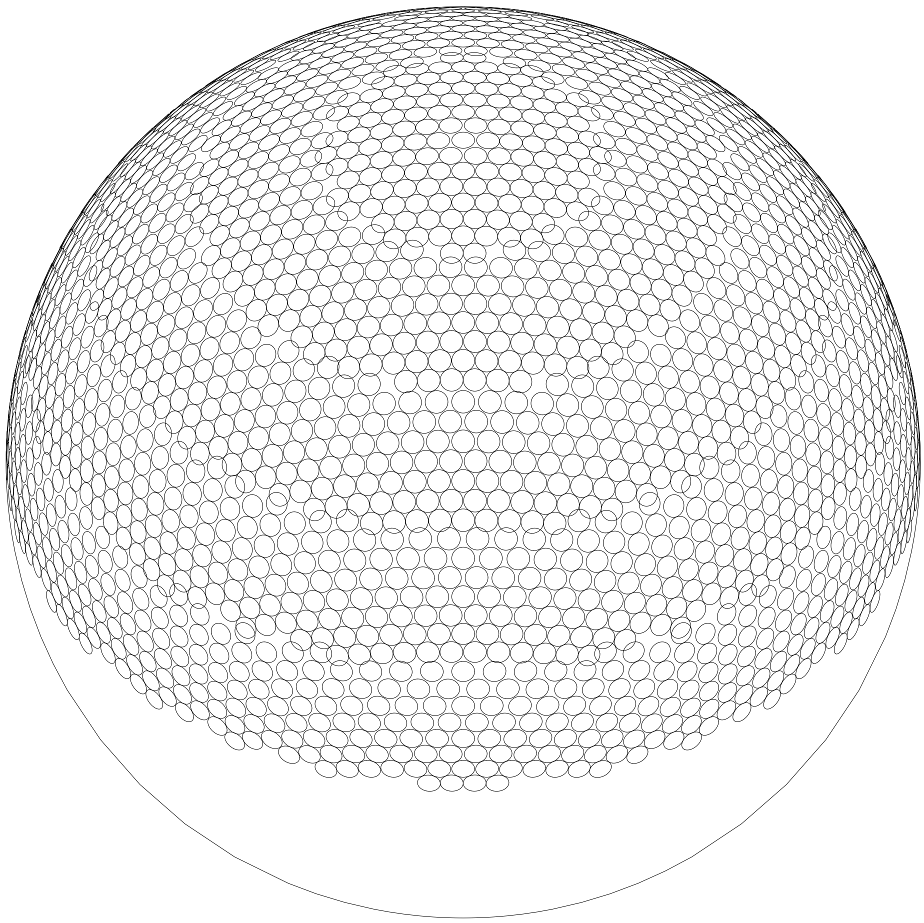

Shape-Gain Rate Allocation

$$
\begin{aligned}
R &= \frac{1}{k}\log_2\left[(\text{Gain CB size}) \times (\text{Shape CB size})\right] \\
&= \frac{1}{k}\log_2(\text{Gain CB size}) + \frac{1}{k}\log_2(\text{Shape CB size}) \\
&= R_g + R_s
\end{aligned}
$$

SQNR as a function of gain codebook size:

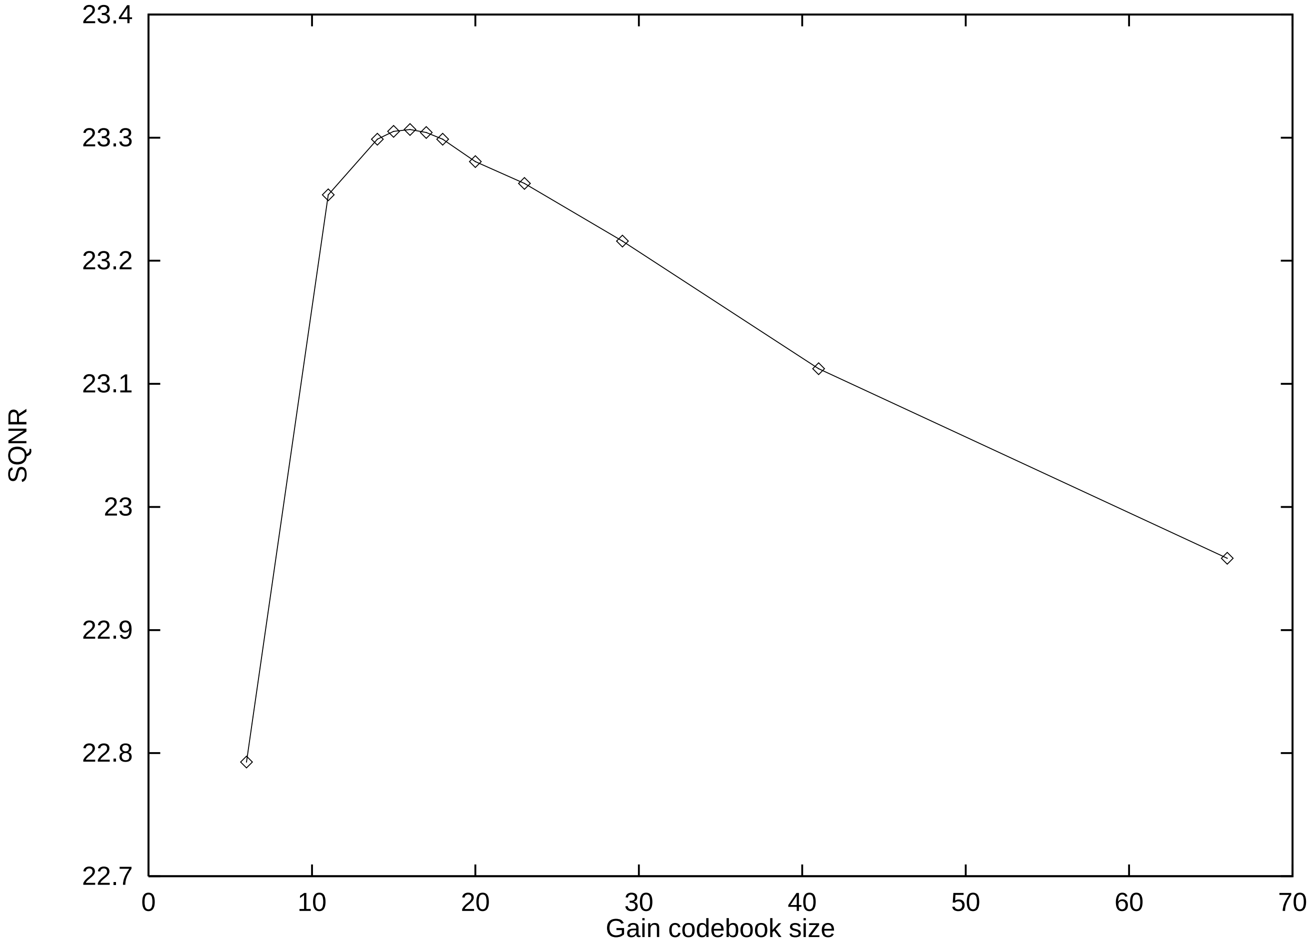

Optimize rate allocation by maximizing SQNR using numerical methods.

Index Assignment

Method to determine index, given $\widehat{g}\widehat{S}$:

1. Off-line, compute the number of codepoints in each annulus. Let P_a denote the number of codepoints of the ath annulus, and P be the total size of the shape codebook.

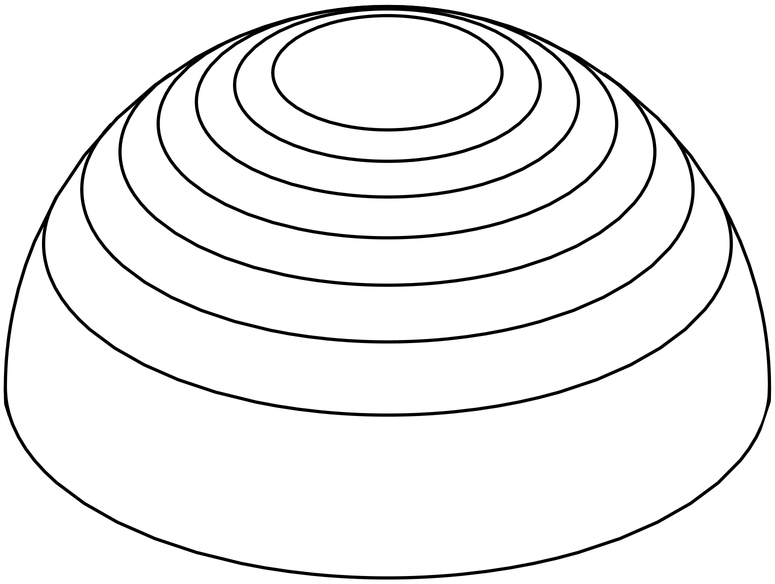

2. Map $\widehat{S}$ to lattice point $q \in \mathbb{R}^{k-1}$, using function from wrapped spherical code definition.

3. Identify index of q using lattice indexing algorithm. Suppose it has the lth lowest index among points in the annulus.

4. Let $\widehat{g}$ be the ith lowest gain output level.

5. The index is $iP + \left(\sum_{a=0}^{j-1} P_a\right) + l$.

Operation of Quantizer

Step 1.	Given k source samples, form vector $X \in \mathbb{R}^k$.	Encoder
Step 2.	Compute $g = \|X\|$ and $S = X/g$.	
Step 3.	Use gain codebook to quantize g as $\widehat{g}$.	
Step 4.	Map S to $f(S) \in \mathbb{R}^{k-1}$.	
Step 5.	Find nearest neighbor $\widehat{f}(S)$ to $f(S)$, using a nearest neighbor algorithm for Λ.	
Step 6.	Compute $f^{-1}(\widehat{f}(S))$ to identify quantized shape $\widehat{S}$.	
Step 7.	Compute the index of $\widehat{g}\widehat{S}$.	
Step 8.	Transmit index of $\widehat{g}\widehat{S}$ across (noiseless) channel.	Channel
Step 9.	Decode index to obtain $\widehat{g}\widehat{S}$	Decoder

Complexity: $7 + \frac{R+L+32}{k}$ arithmetic steps per sample, where

$$
\begin{aligned}
R &= \text{Rate} \\
k &= \text{Dimension} \\
L &= \text{Lattice encoding complexity}
\end{aligned}
$$

For Leech lattice, $L = 2955$, and the complexity is $127 + \frac{R}{25}$.

Performance Analysis

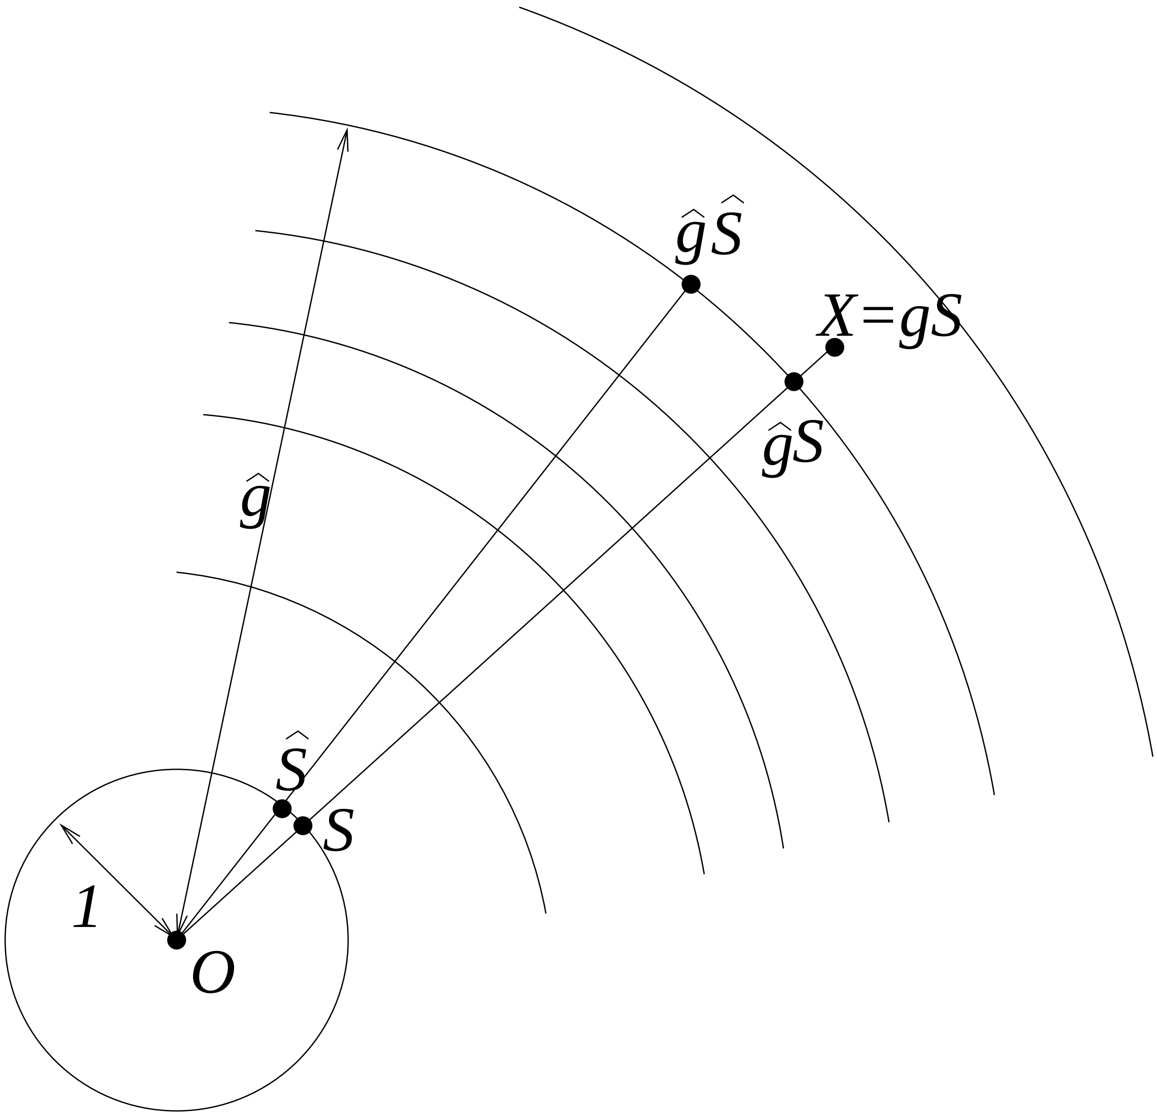

$$
\begin{aligned}
D &= \frac{1}{k} E\|X - \widehat{g}\widehat{S}\|^2 \\
&= D_g + D_s \\
&= \frac{1}{k} E(g - \widehat{g})^2 + \frac{1}{k} E\widehat{g}^2 E\|S - \widehat{S}\|^2
\end{aligned}
$$

Performance Analysis

In order to optimize rate allocation, $D = D_g + D_s$ must be estimated under differing shape and gain codebook sizes.

Evaluation of D_g

$$D_g = \frac{1}{k} E(g - \widehat{g})^2$$

Evaluate using $f_g(r)$ and table of $\widehat{g}$ outputs, which are known explicitly.

Compute the expectation numerically, if necessary.

Evaluation of D_S

$$D_s = \frac{1}{k} \underbrace{E\widehat{g}^2}_{1} \underbrace{E\|S - \widehat{S}\|^2}_{2}$$

1. $E\widehat{g}^2 = Eg^2 - E(g - \widehat{g})^2 \approx Eg^2$

2. $E\|S - \widehat{S}\|^2 \approx$ (MSE of underlying lattice, used as $(k-1)$-dimensional quantizer for Gaussian source.)

Simulations

Comparison of various quantization schemes for a memoryless Gaussian source. Values are listed as SQNR in decibels.

Method	Rate						
	1	2	3	4	5	6	7
D(R)	6.02	12.04	18.06	24.08	30.10	36.12	42.14
$W\Lambda_{24}$–SVQ	2.44	11.02	17.36	23.33	29.29	35.27	41.33
GLA (kR=8)		10.65		20.98			
Lloyd-Max Scalar	4.40	9.30	14.62	20.22	26.02	31.89	37.81
Uniform scalar	4.40	9.25	14.27	19.38	24.57	29.83	35.13
Entropy coded scalar	4.64	10.55	16.56	22.55	28.57	34.59	40.61
UPQ	4.40	9.63					
Fischer SVQ (estimated)	4.49	10.51	16.53	22.55	28.57	34.59	40.61
TCQ 256 state	5.56	11.04	16.64				
Wilson 128 state	5.47	10.87	16.78				
TB–SVQ (4 state)	5.14	11.11	16.77				
TB–SVQ (64 dim.,16 state)	5.49	11.28	17.05				

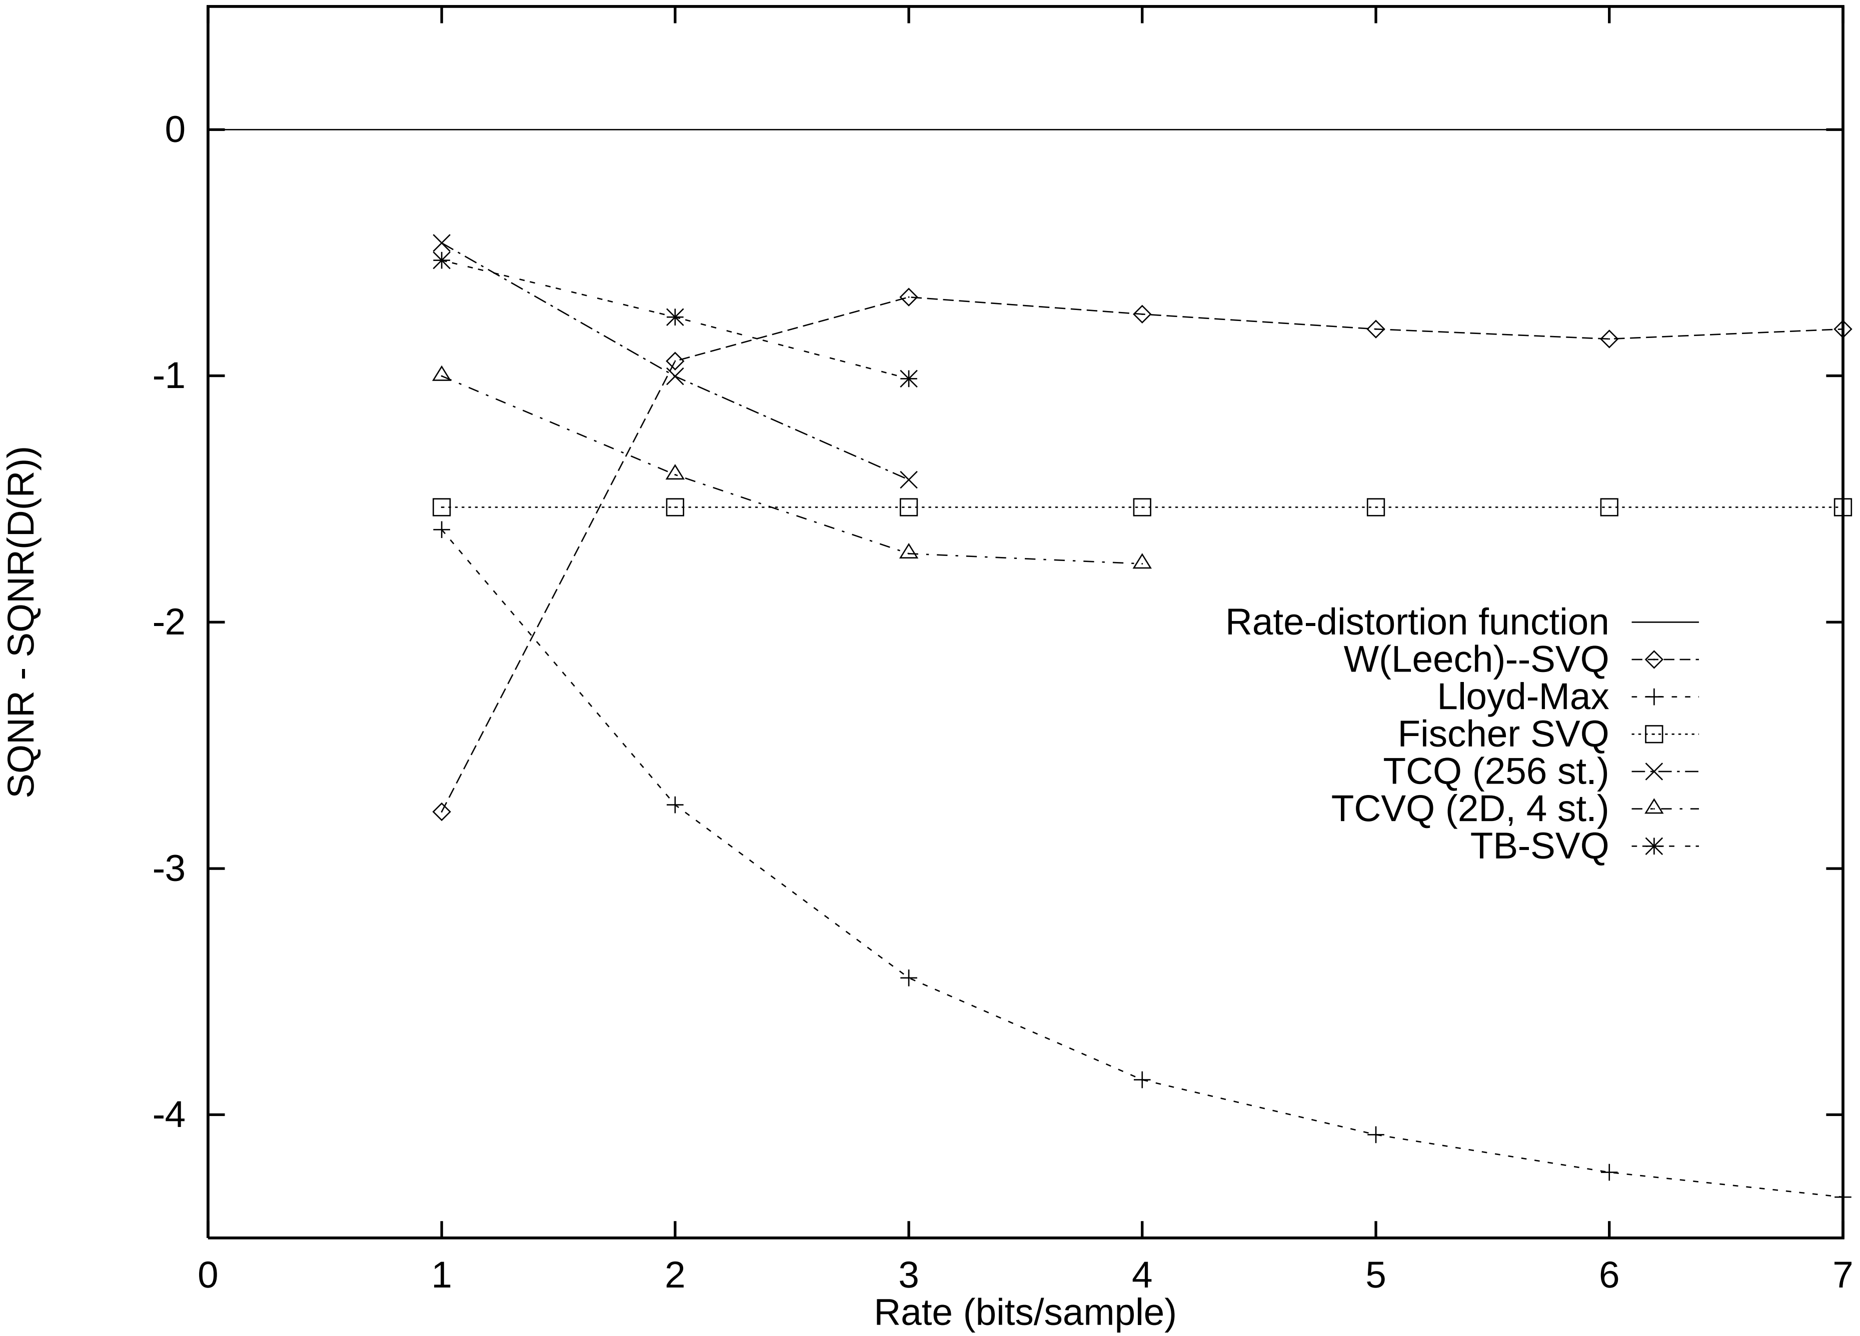

SQNR - SQNR(D(R))
0
-1
-2
-3
-4
0
1
2
3
4
5
6
7
Rate (bits/sample)
Rate-distortion function
W(Leech)--SVQ
Lloyd-Max
Fischer SVQ
TCQ (256 st.)
TCVQ (2D, 4 st.)
TB-SVQ

Improvements and Extensions

Vector quantization of the gain:

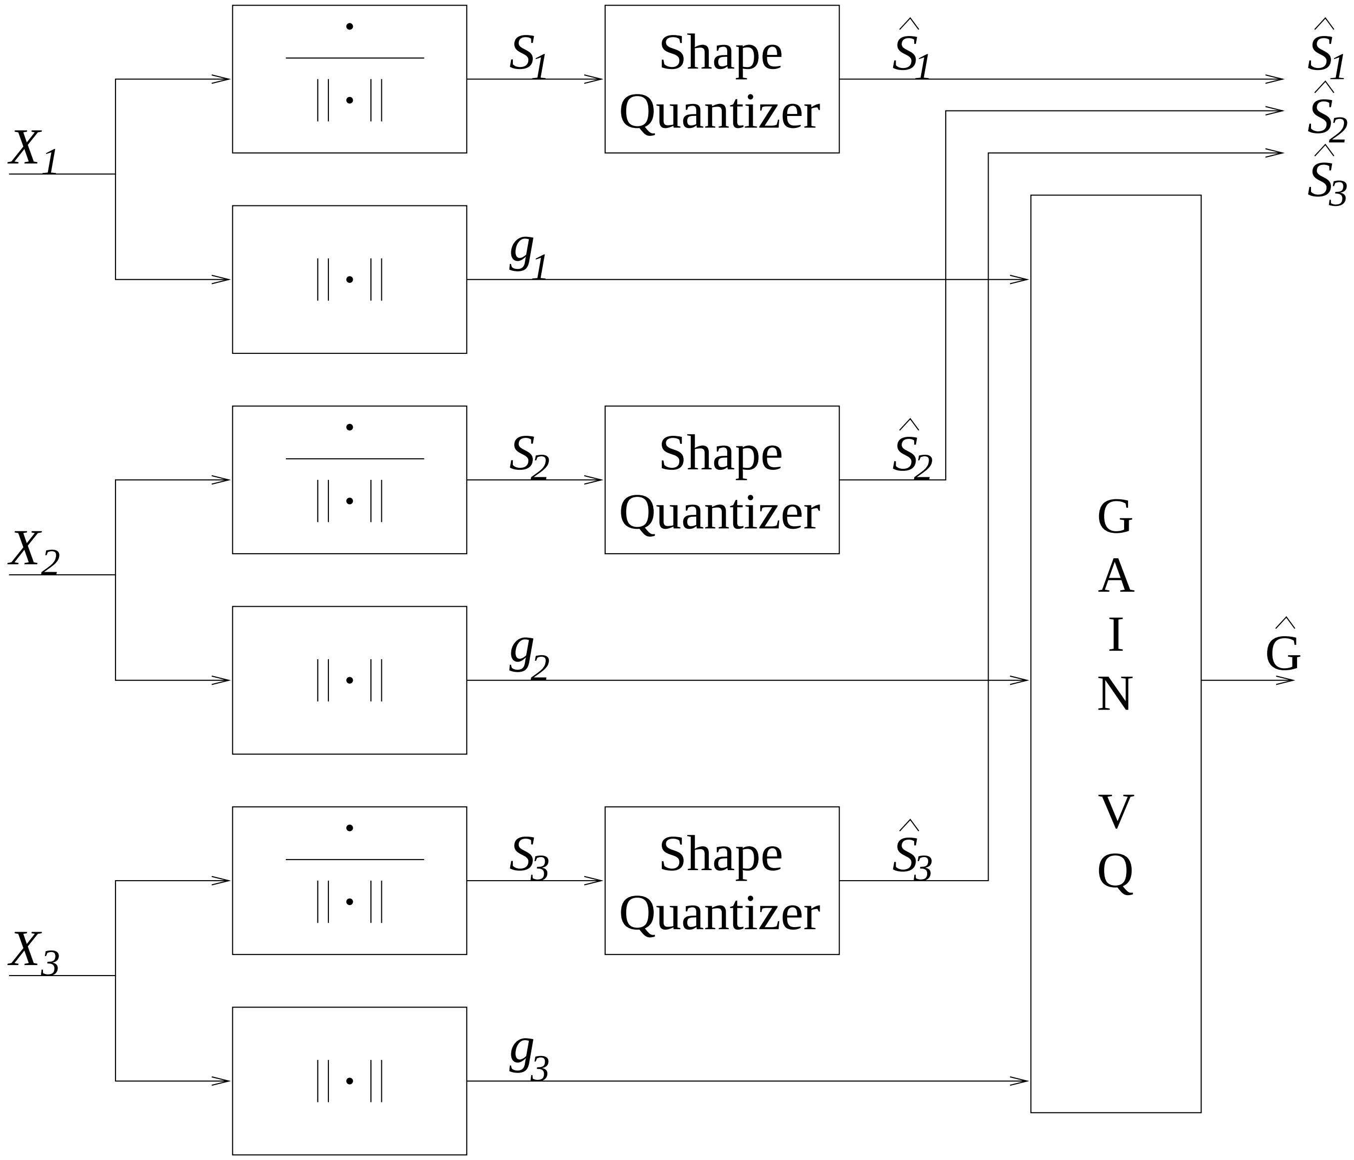

Improvements and Extensions

Shape classified quantization

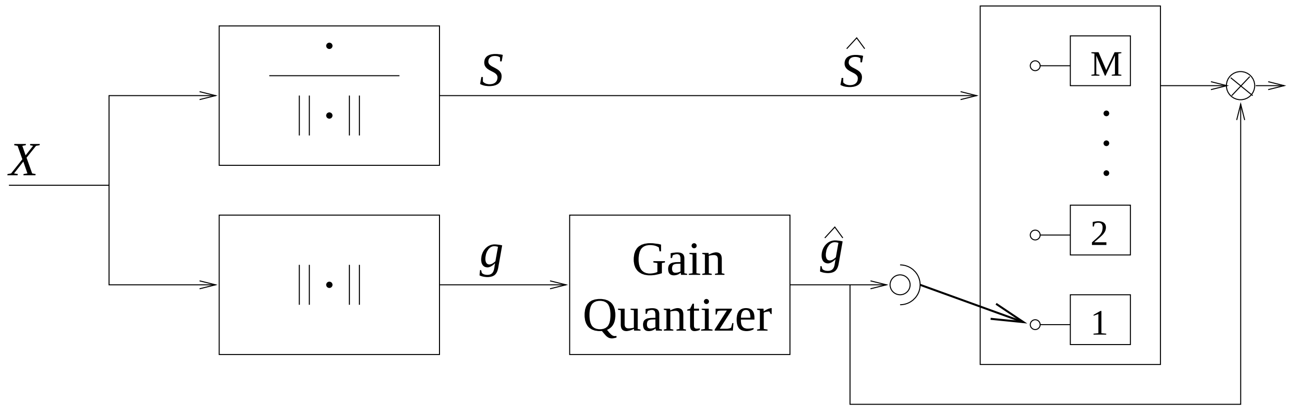

Universal Quantization

$$X \longrightarrow \boxed{H_m(\,)} \xrightarrow{\;Y\;} \boxed{Q} \xrightarrow{\;\widehat{Y}\;} \boxed{H_m^T(\,)} \xrightarrow{\;\widehat{X}\;}$$

Use intuition of central limit theorem.

$m \times m$ Hadamard matrix contains ± 1 and satisfies $H_m H_m^T = I$.

Example: $H_2 = \begin{pmatrix} 1 & -1 \\ -1 & -1 \end{pmatrix}$

Let $e = \widehat{Y} - Y$.

$$\begin{aligned} \widehat{X} &= H_m^T \widehat{Y} \\ &= H_m^T (Y + e) \\ &= H_m^T (H_m X + e) \\ &= H_m^T H_m X + H_m^T e \\ &= X + H_m^T e. \end{aligned}$$

The end-to-end distortion of this system is

$$\begin{aligned} E[\|\widehat{X} - X\|^2] &= E[(\widehat{X} - X)^T (\widehat{X} - X)] \\ &= E[(H_m^T e)^T (H_m^T e)] \\ &= E[e^T H_m H_m^T e] \\ &= E[e^T e] \\ &= E[(\widehat{Y} - Y)^T (\widehat{Y} - Y)] \\ &= E[\|\widehat{Y} - Y\|^2] \end{aligned}$$

Conclusion: we can transform distribution of X to a Gaussian distribution and use the usual wrapped SVQ without penalty.

Conclusions

- Performance and complexity of quantizers for memoriless Gaussian sources is equivalent to the performance and complexity of quantizers for the uniform source.

- Any future improvements to uniform source quantizers yield analogous improvements for Gaussian source quantizers.

Future work:

- Noisy channel vector quantization

- Trellis encoded spherical vector quantizers

Interesting Facts about Spheres

In 1885, L. Schläfli published an article containing a formula for the volume of the k-dimensional unit sphere Ω_k:

$$V_k = \frac{\pi^{k/2}}{\Gamma\left(\frac{k+2}{2}\right)} = \begin{cases} \frac{\pi^{k/2}}{(k/2)!} & \text{if } k \text{ even} \\ \frac{2^k \pi^{(k-1)/2}((k-1)/2)!}{k!} & \text{if } k \text{ is odd} \end{cases}$$

The surface area is given by

$$A_k = kV_k$$

Thus,

$$(V_2, V_3, V_4, \dots) = (\pi, \frac{4\pi}{3}, \frac{\pi^2}{2}, \dots) \approx (3.1, 4.2, 4.9, 5.3, \dots)$$

and

$$(A_2, A_3, A_4, \dots) = (2\pi, 4\pi, 2\pi^2, \dots) \approx (6.3, 12.6, 19.7, 26.3, \dots)$$

Amazing fact:

V_k and A_k each approach 0 as $k \to 0$.

V_k is maximum in dimension 5; A_k is maximum in dimension 7.

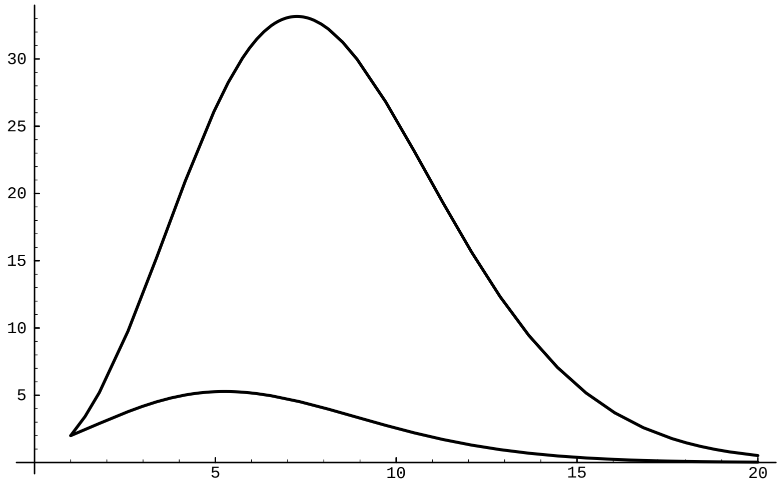

Interesting Facts about Spheres

A theorem of Archimedes* says that the surface area of a sphere is equal to the cylinder that contains it. In fact, a sphere 'slice' of thickness t has surface area $2\pi t$, regardless of where the slice is made.

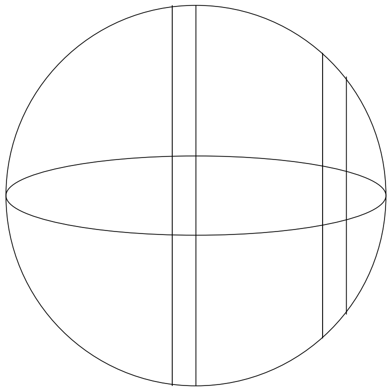

This implies the following:

Claim: Let P = (X, Y, Z) be a random vector uniformly distributed on the unit-radius sphere. Then the distribution of X is uniform on [-1,1].

* He was so proud of the theorem, he had a figure of it inscribed on his tombstone. The tombstone was lost but was rediscovered this century— they recognized the figure!

Interesting Facts about Spheres

The percentage of the surface area of Earth above 45° N is about 15% of the total surface area.

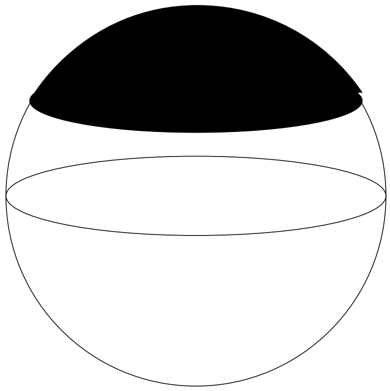

A similar 'polar cap' on a higher-dimensional sphere consists of points on the sphere whose last coordinate is at least $1/\sqrt{2}$ or higher.

The fraction of surface area such a polar cap occupies approaches 0 as the dimension tends to infinity. In dimension 25, it occupies less than two thousandths of a percent of the total surface area.

$$\text{fraction} = \frac{1}{S_k} \int_0^{\pi/4} S_{k-1} \sin^{k-2} \theta \, d\theta$$