Quantization

Input-output characteristic of a scalar quantizer

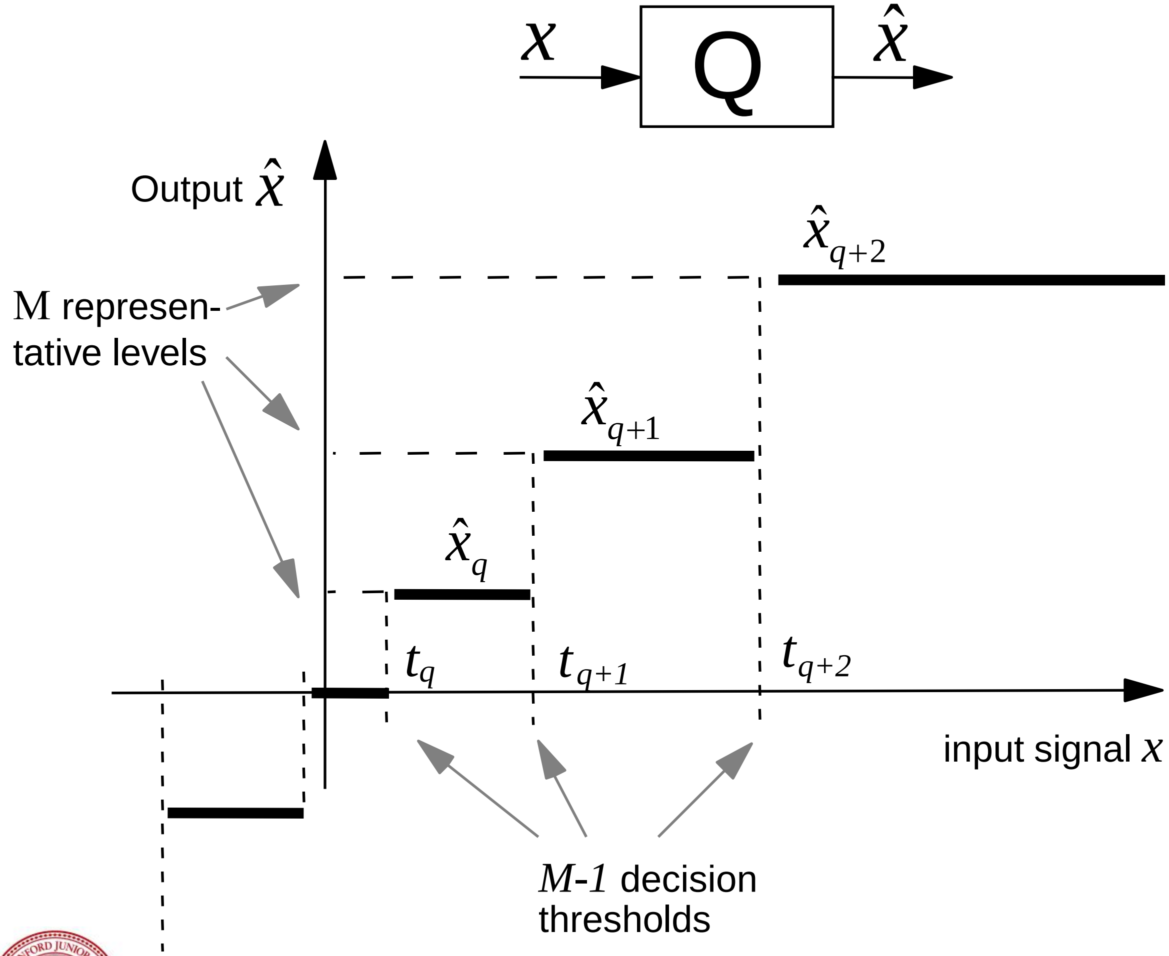

Sometimes, this convention is used:

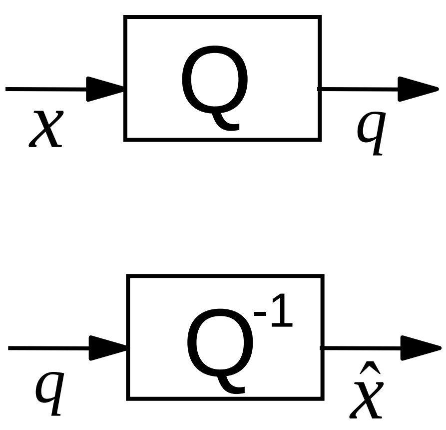

Example of a quantized waveform

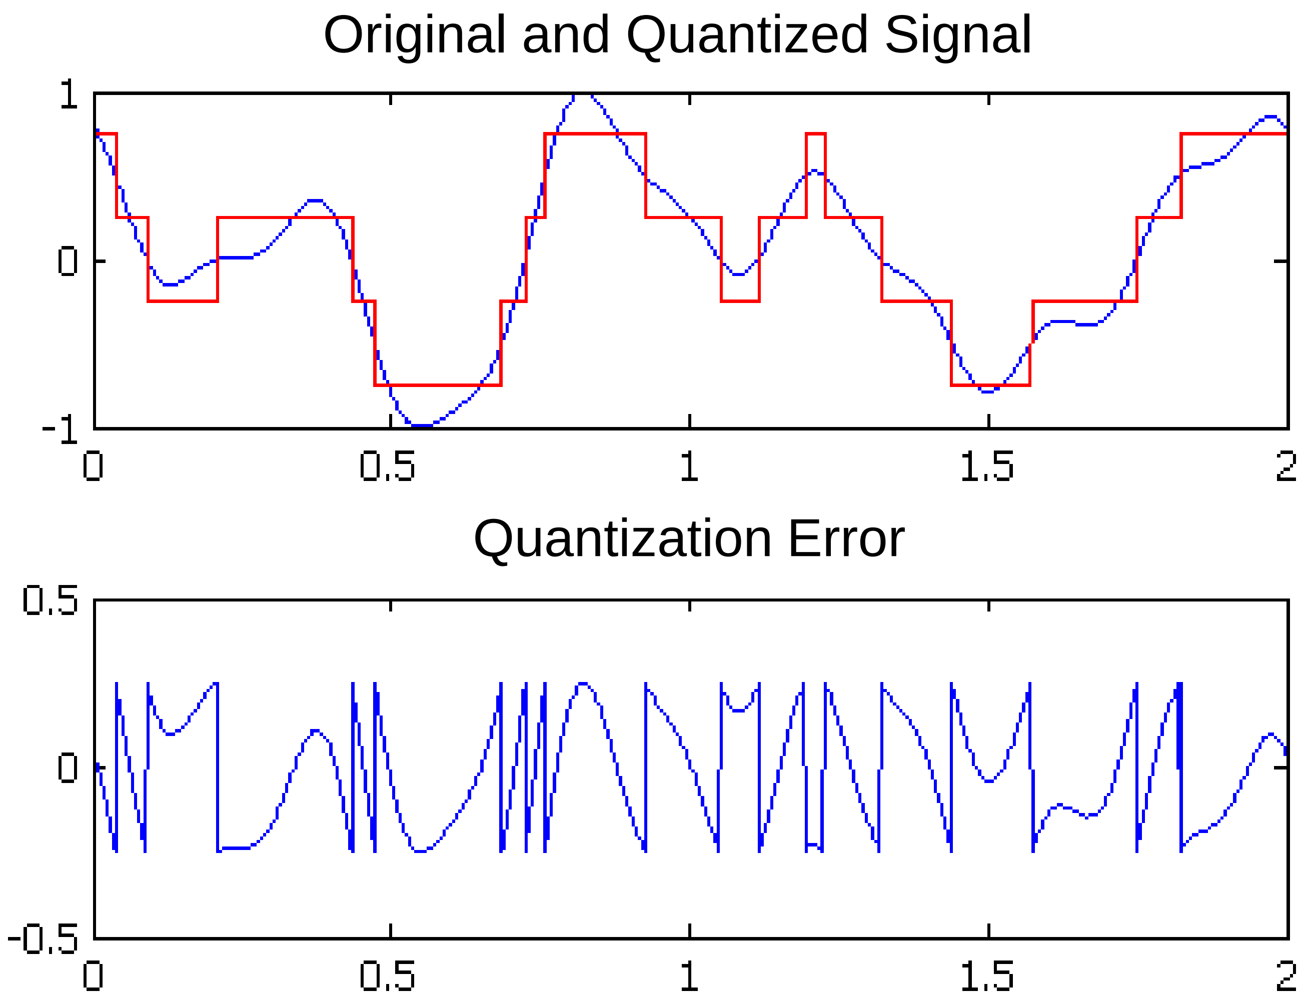

Lloyd-Max scalar quantizer

- Problem : For a signal x with given PDF $f_X(x)$ find a quantizer with M representative levels such that

$$d = MSE = E\left[\left(X - \hat{X}\right)^2\right] \rightarrow \min.$$

- Solution : Lloyd-Max quantizer
  [Lloyd,1957] [Max,1960]
  - M-1 decision thresholds exactly half-way between representative levels.
  - M representative levels in the centroid of the PDF between two successive decision thresholds.
  - Necessary (but not sufficient) conditions

$$t_q = \frac{1}{2}\left(\hat{x}_{q-1} + \hat{x}_q\right) \quad q = 1, 2, \ldots, M\text{ -1}$$

$$\hat{x}_q = \frac{\int_{t_q}^{t_{q+1}} x \cdot f_X(x)dx}{\int_{t_q}^{t_{q+1}} f_X(x)dx} \quad q = 0, 1, \ldots, M\text{ -1}$$

Iterative Lloyd-Max quantizer design

1. Guess initial set of representative levels $\hat{x}_q \quad q = 0, 1, 2, \ldots, M-1$
2. Calculate decision thresholds

$$t_q = \frac{1}{2}\left(\hat{x}_{q-1} + \hat{x}_q\right) \quad q = 1, 2, \ldots, M-1$$

3. Calculate new representative levels

$$\hat{x}_q = \frac{\int_{t_q}^{t_{q+1}} x \cdot f_X(x)\,dx}{\int_{t_q}^{t_{q+1}} f_X(x)\,dx} \quad q = 0, 1, \ldots, M-1$$

4. Repeat **2.** and **3.** until no further distortion reduction

Example of use of the Lloyd algorithm (I)

- X zero-mean, unit-variance Gaussian r.v.
- Design scalar quantizer with 4 quantization indices with minimum expected distortion D^*
- Optimum quantizer, obtained with the Lloyd algorithm
  - Decision thresholds -0.98, 0, 0.98
  - Representative levels –1.51, -0.45, 0.45, 1.51
  - D^*=0.12=9.30 dB

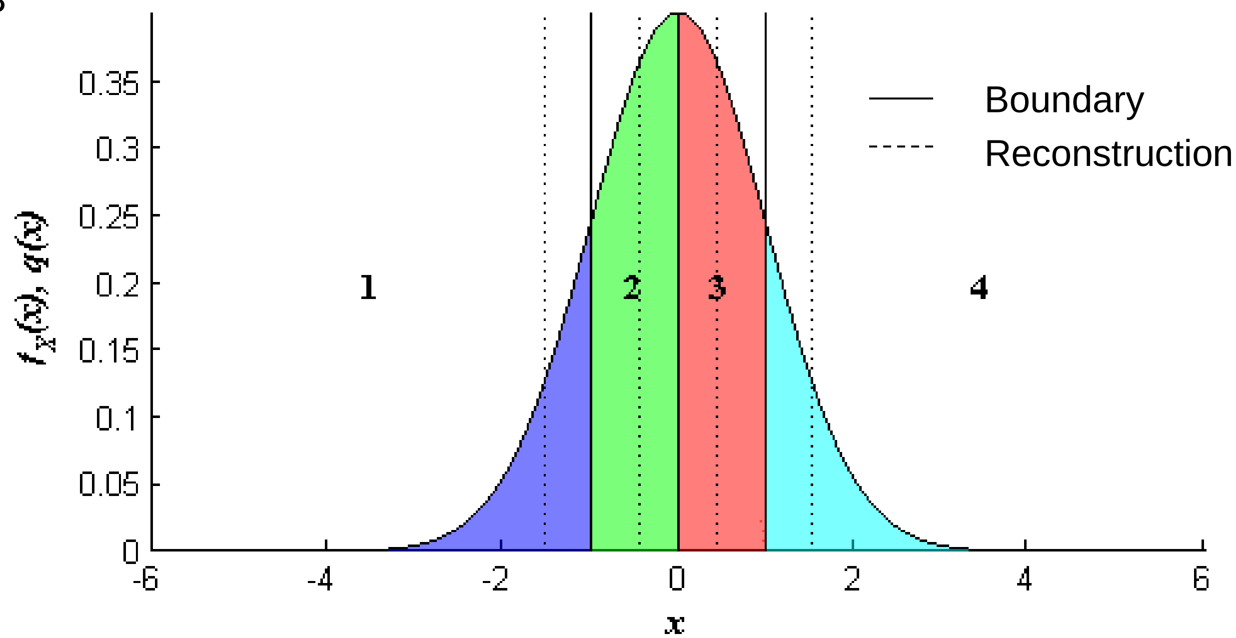

Example of use of the Lloyd algorithm (II)

- Convergence
  - Initial quantizer A:
    decision thresholds –3, 0 3

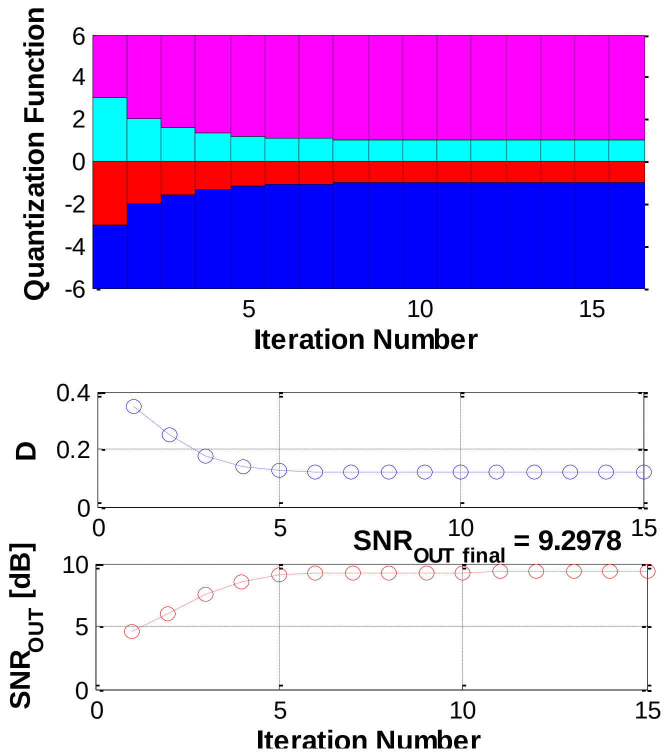

  - Initial quantizer B:
    decision thresholds –½, 0, ½

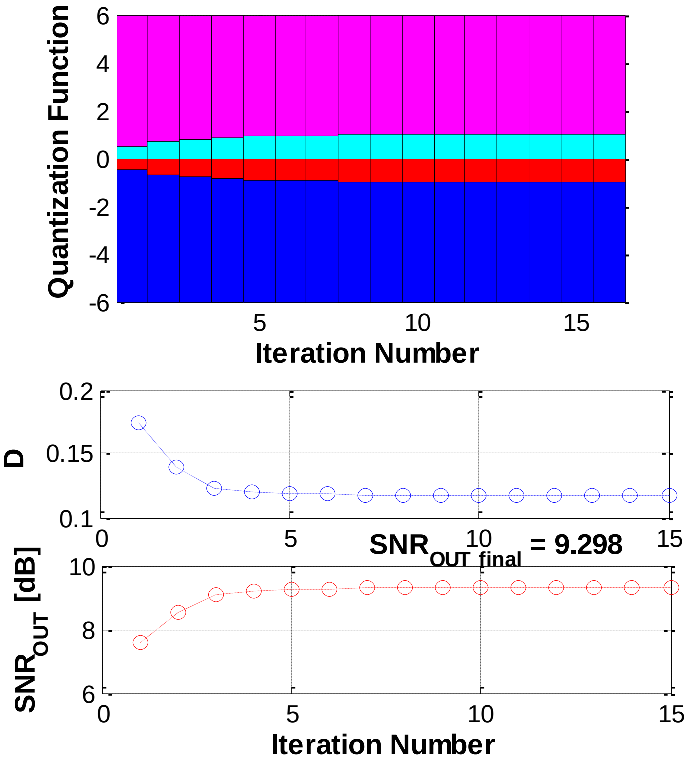

  - After 6 iterations, in both cases, $(D\text{-}D^*)/D^* < 1\%$

Example of use of the Lloyd algorithm (III)

- X zero-mean, unit-variance Laplacian r.v.
- Design scalar quantizer with 4 quantization indices with minimum expected distortion D^*
- Optimum quantizer, obtained with the Lloyd algorithm
  - Decision thresholds -1.13, 0, 1.13
  - Representative levels -1.83, -0.42, 0.42, 1.83
  - D^*=0.18=7.54 dB

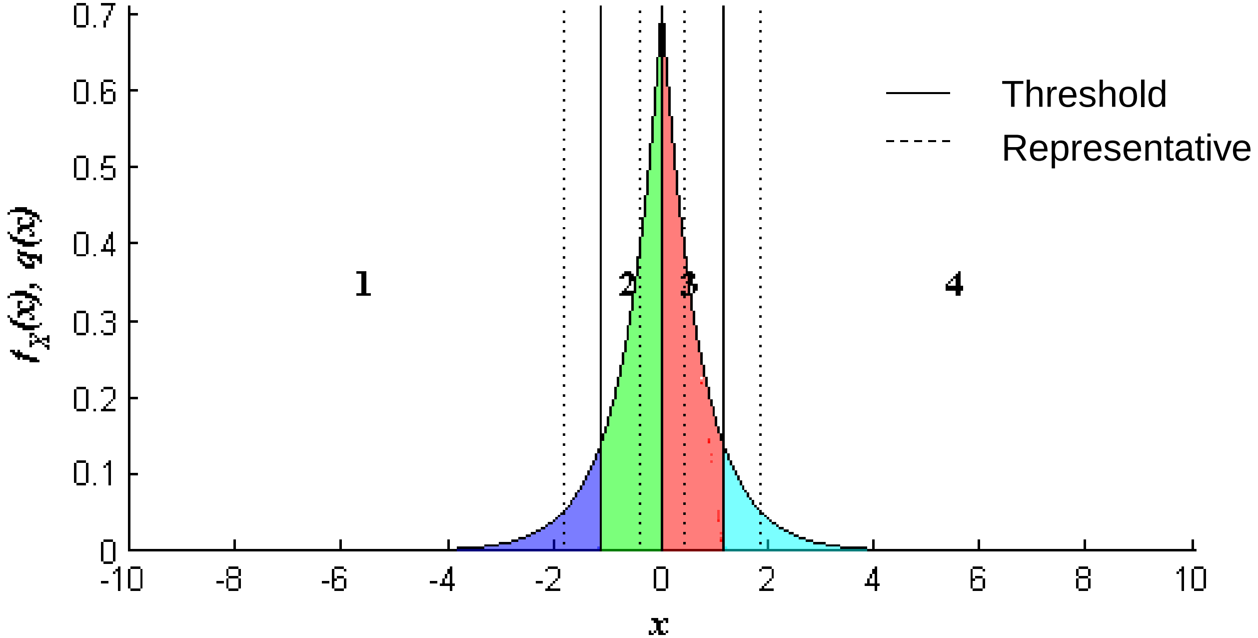

Example of use of the Lloyd algorithm (IV)

- Convergence
  - Initial quantizer A,
    decision thresholds –3, 0 3

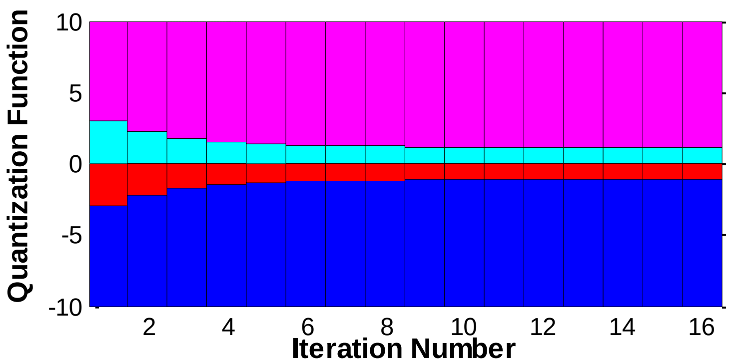

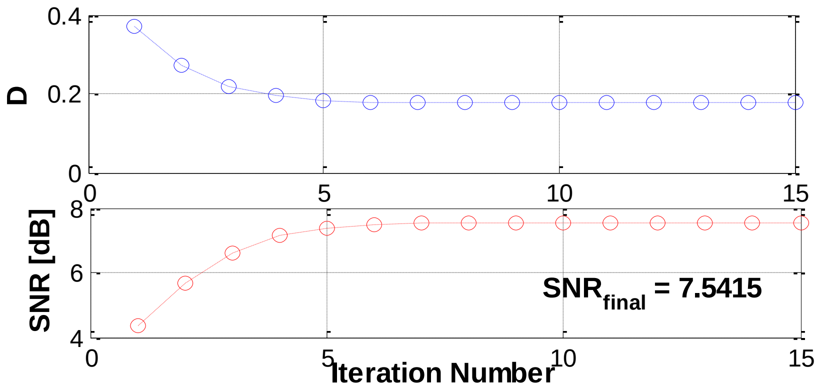

  - Initial quantizer B,
    decision thresholds –½, 0, ½

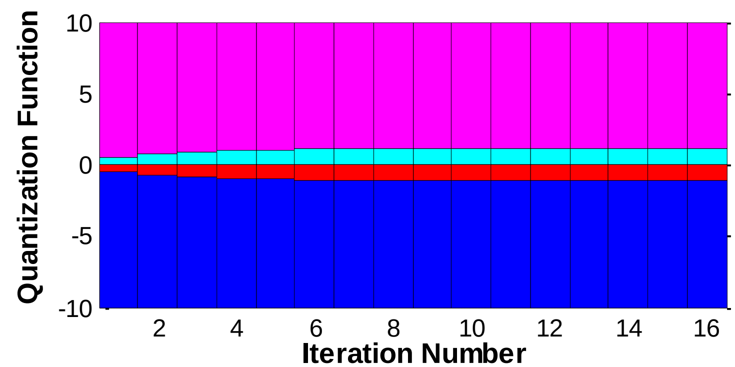

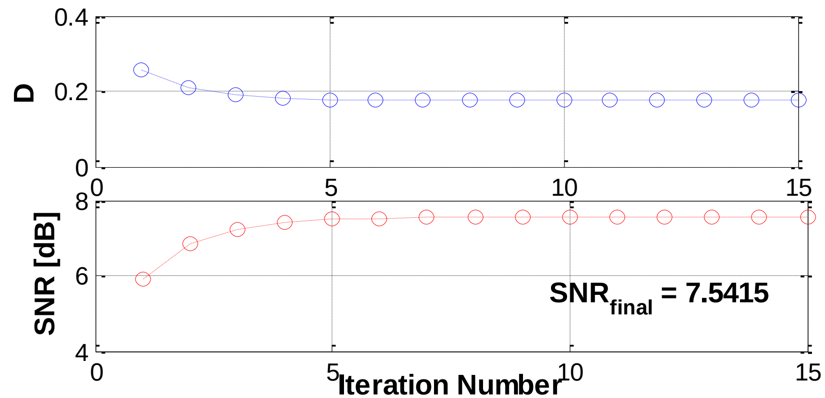

  - After 6 iterations, in both cases, $(D\text{-}D^*)/D^*<1\%$

Lloyd algorithm with training data

1. Guess initial set of representative levels $\hat{x}_q;\ q = 0,1,2,\ldots,M-1$
2. Assign each sample x_i in training set $\boldsymbol{T}$ to closest representative $\hat{x}_q$

$$B_q = \left\{ x \in \boldsymbol{T} : Q(x) = q \right\} \qquad q = 0,1,2,\ldots,M-1$$

3. Calculate new representative levels

$$\hat{x}_q = \frac{1}{\left\| B_q \right\|} \sum_{x \in B_q} x \qquad q = 0,1,\ldots,M-1$$

4. Repeat **2.** and **3.** until no further distortion reduction

Lloyd-Max quantizer properties

- Zero-mean quantization error

$$E\left[X - \hat{X} \right] = 0$$

- Quantization error and reconstruction decorrelated

$$E\left[\left(X - \hat{X} \right) \hat{X} \right] = 0$$

- Variance subtraction property

$$\sigma_{\hat{X}}^2 = \sigma_X^2 - E\left[\left(X - \hat{X} \right)^2 \right]$$

High rate approximation

- Approximate solution of the "Max quantization problem," assuming high rate and smooth PDF *[Panter, Dite, 1951]*

$$\Delta x(x) = const \frac{1}{\sqrt[3]{f_X(x)}}$$

Distance between two successive quantizer representative levels

Probability density function of x

- Approximation for the quantization error variance:

$$d = E\left[\left(X - \hat{X}\right)^2\right] \approx \frac{1}{12M^2}\left[\int_x \sqrt[3]{f_X(x)}dx\right]^3$$

Number of representative levels

High rate approximation (cont.)

- High-rate distortion-rate function for scalar Lloyd-Max quantizer

$$d(R) \cong \varepsilon^2 \sigma_X^2 \, 2^{-2R}$$

$$\text{with} \quad \varepsilon^2 \sigma_X^2 = \frac{1}{12} \left[\int_x \sqrt[3]{f_X(x)} dx \right]^3$$

- Some example values for ε^2

uniform	1
Laplacian	$\frac{9}{2} = 4.5$
Gaussian	$\frac{\sqrt{3}\pi}{2} \cong 2.721$

High rate approximation (cont.)

- Partial distortion theorem: each interval makes an (approximately) equal contribution to overall mean-squared error

$$\Pr\{t_i \le X < t_{i+1}\} E\left[\left(X - \hat{X}\right)^2 \middle| t_i \le X < t_{i+1}\right]$$
$$\cong \Pr\{t_j \le X < t_{j+1}\} E\left[\left(X - \hat{X}\right)^2 \middle| t_j \le X < t_{j+1}\right] \quad \text{for all } i, j$$

[Panter, Dite, 1951], [Fejes Toth, 1959], [Gersho, 1979]

Entropy-constrained scalar quantizer

- Lloyd-Max quantizer optimum for fixed-rate encoding, how can we do better for variable-length encoding of quantizer index?
- Problem : For a signal x with given pdf $f_X(x)$ find a quantizer with rate

$$R = H\left(\hat{X}\right) = -\sum_{q=0}^{M-1} p_q \log_2 p_q$$

such that

$$d = MSE = E\left[\left(X - \hat{X}\right)^2\right] \rightarrow \min.$$

- Solution: Lagrangian cost function

$$J = d + \lambda R = E\left[\left(X - \hat{X}\right)^2\right] + \lambda H\left(\hat{X}\right) \rightarrow \min.$$

Iterative entropy-constrained scalar quantizer design

1. Guess initial set of representative levels $\hat{x}_q;\ q = 0,1,2,\ldots,M\text{-}1$ and corresponding probabilities p_q
2. Calculate $M\text{-}1$ decision thresholds

$$t_q = \frac{\hat{x}_{q\text{-}1} + \hat{x}_q}{2} - \lambda \frac{\log_2 p_{q-1} - \log_2 p_q}{2\left(\hat{x}_{q\text{-}1} - \hat{x}_q\right)} \qquad q = 1,2,\ldots,M\text{-}1$$

3. Calculate M new representative levels and probabilities p_q

$$\hat{x}_q = \frac{\int_{t_q}^{t_{q+1}} x f_X(x)\,dx}{\int_{t_q}^{t_{q+1}} f_X(x)\,dx} \qquad q = 0,1,\ldots,M\text{-}1$$

4. Repeat **2.** and **3.** until no further reduction in Lagrangian cost

Lloyd algorithm for entropy-constrained quantizer design based on training set

1. Guess initial set of representative levels $\hat{x}_q;\ q = 0,1,2,\ldots,M-1$ and corresponding probabilities p_q
2. Assign each sample x_i in training set ***T*** to representative $\hat{x}_q$ minimizing Lagrangian cost $J_{x_i}(q) = \left(x_i - \hat{x}_q\right)^2 - \lambda \log_2 p_q$

$$B_q = \left\{x \in \boldsymbol{T} : Q_\lambda(x) = q\right\} \qquad q = 0,1,2,\ldots,M-1$$

3. Calculate new representative levels and probabilities p_q

$$\hat{x}_q = \frac{1}{\left\|B_q\right\|} \sum_{x \in B_q} x \quad q = 0,1,\ldots,M-1$$

4. Repeat **2.** and **3.** until no further reduction in overall Lagrangian cost

$$J = \sum_{x_i} J_{x_i} = \sum_{x_i} \left(x_i - Q(x_i)\right)^2 - \lambda \log_2 p_{q(x_i)}$$

Example of the EC Lloyd algorithm (I)

- X zero-mean, unit-variance Gaussian r.v.
- Design entropy-constrained scalar quantizer with rate $R \approx 2$ bits, and minimum distortion D^*
- Optimum quantizer, obtained with the entropy-constrained Lloyd algorithm
  - 11 intervals (in [-6,6]), almost uniform
  - D^*=0.09=10.53 dB, R=2.0035 bits (compare to fixed-length example)

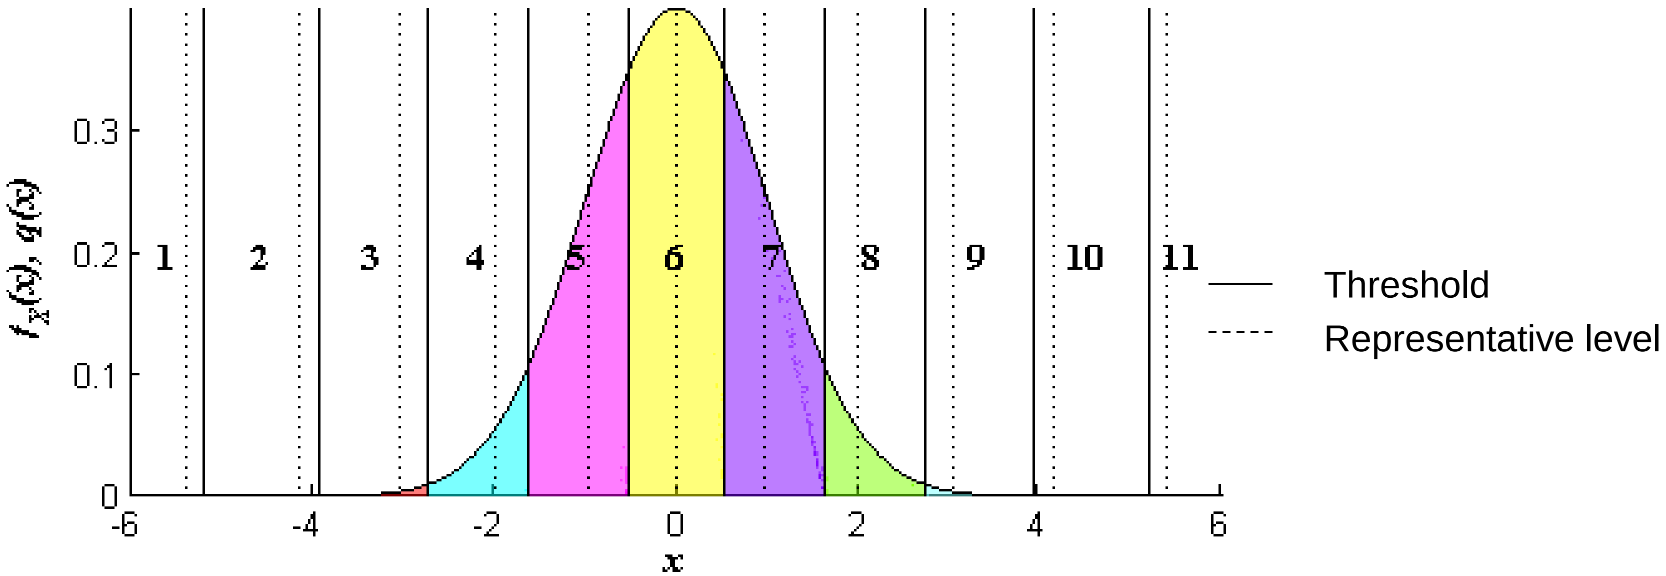

Example of the EC Lloyd algorithm (II)

- Same Lagrangian multiplier λ used in all experiments
  - Initial quantizer A, 15 intervals (>11) in [-6,6], with the same length
  - Initial quantizer B, only 4 intervals (<11) in [-6,6], with the same length

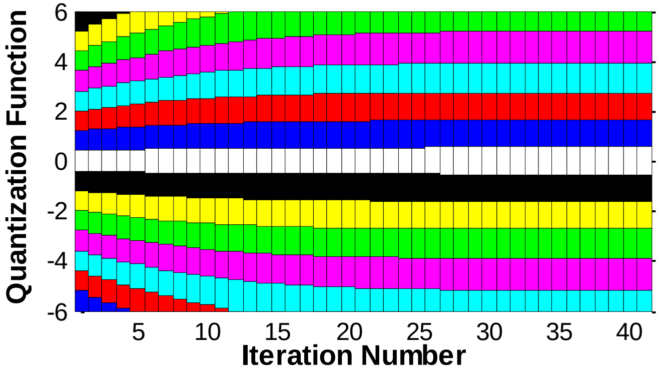

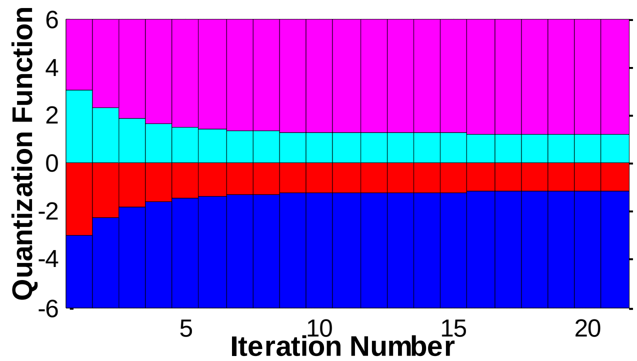

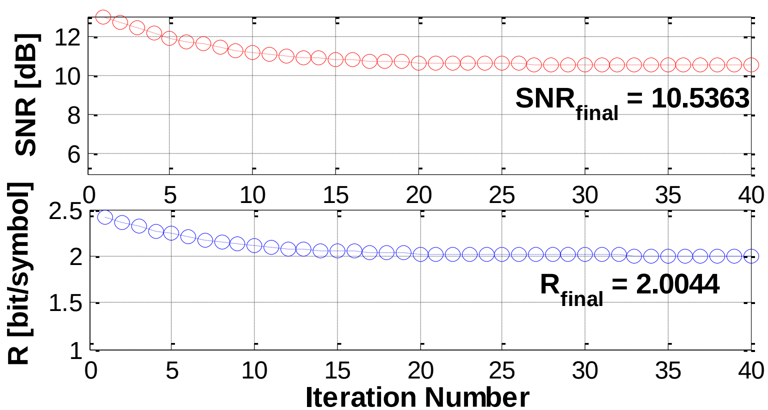

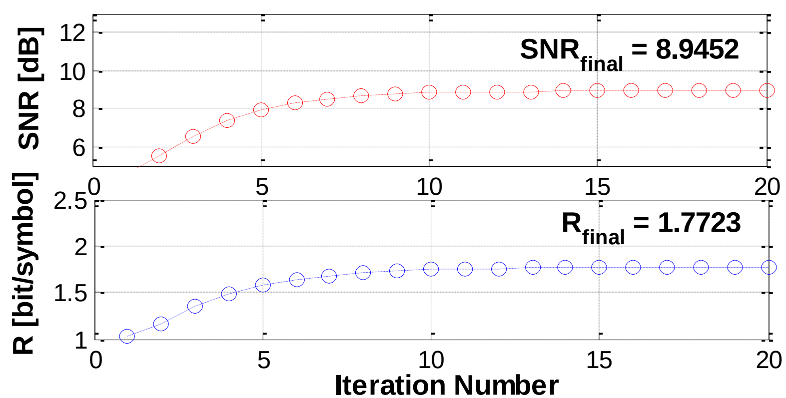

Example of the EC Lloyd algorithm (III)

- X zero-mean, unit-variance Laplacian r.v.
- Design entropy-constrained scalar quantizer with rate $R \approx 2$ bits and minimum distortion D^*
- Optimum quantizer, obtained with the entropy-constrained Lloyd algorithm
  - 21 intervals (in [-10,10]), almost uniform
  - D^*= 0.07=11.38 dB, R= 2.0023 bits (compare to fixed-length example)

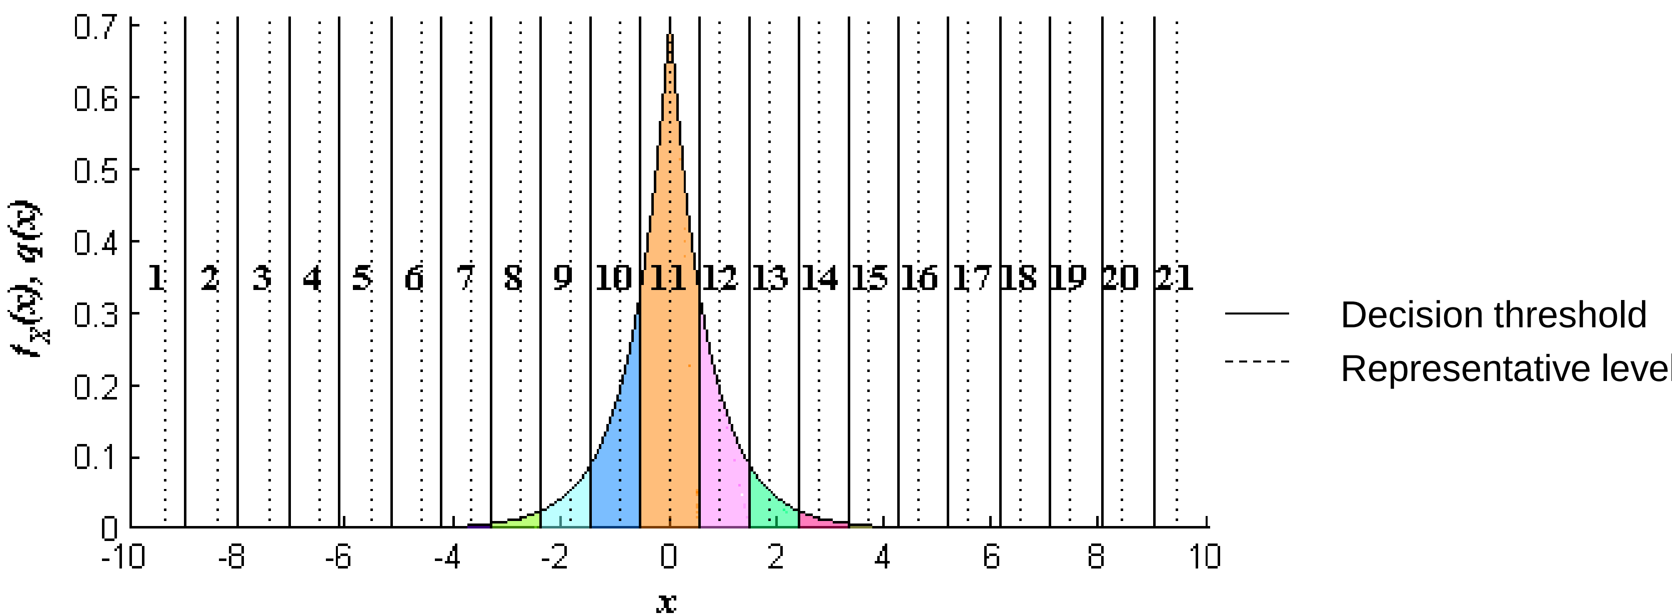

Example of the EC Lloyd algorithm (IV)

- Same Lagrangian multiplier λ used in all experiments
  - Initial quantizer A, 25 intervals (>21 & odd) in [-10,10], with the same length

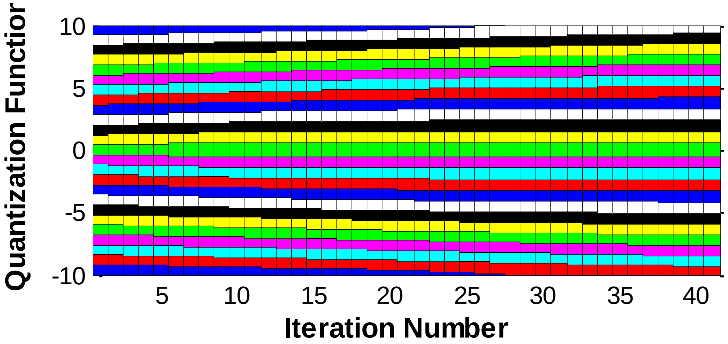

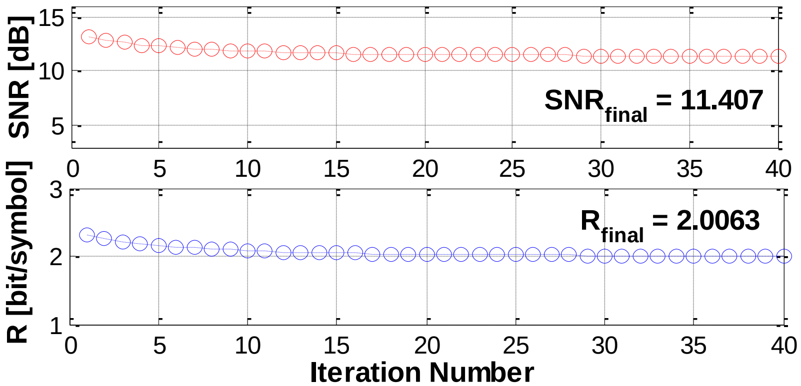

  - Initial quantizer B, only 4 intervals (<21) in [-10,10], with the same length

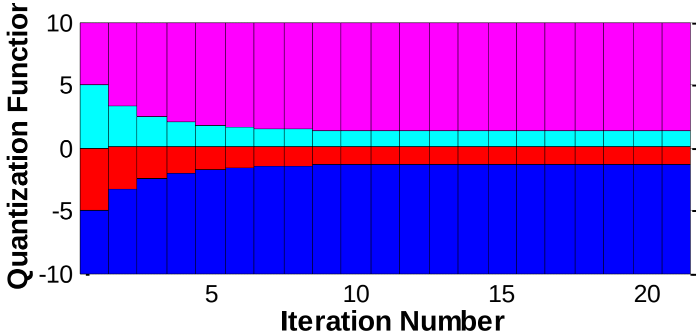

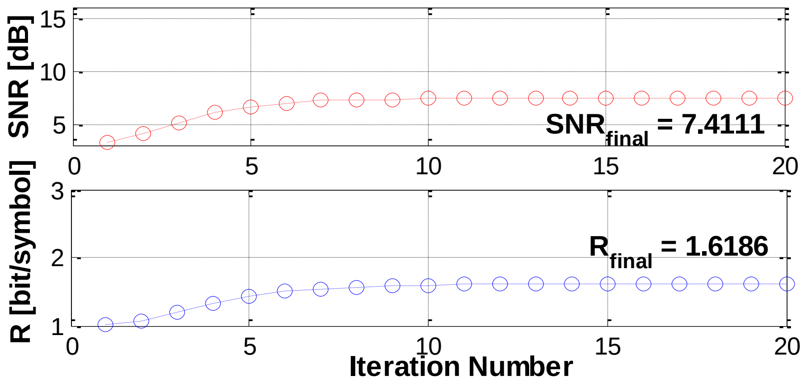

  - Convergence in cost faster than convergence of decision thresholds

High-rate results for EC scalar quantizers

- For MSE distortion and high rates, uniform quantizers (followed by entropy coding) are optimum *[Gish, Pierce, 1968]*
- Distortion and entropy for smooth PDF and fine quantizer interval Δ

$$d \cong \int_{-\Delta/2}^{\Delta/2} \varepsilon^2 d\varepsilon = \frac{\Delta^2}{12}$$

$$H\left(\hat{X}\right) \cong h\left(X\right) - \log_2 \Delta$$

- Distortion-rate function

$$d\left(R\right) \cong \frac{1}{12} 2^{2h(X)} 2^{-2R}$$

is factor $\frac{\pi e}{6}$ or 1.53 dB from Shannon Lower Bound

$$D\left(R\right) \geq \frac{1}{2\pi e} 2^{2h(X)} 2^{-2R}$$

Comparison of high-rate performance of scalar quantizers

- High-rate distortion-rate function

$$d(R) \cong \varepsilon^2 \sigma_X^2 \, 2^{-2R}$$

- Scaling factor ε^2

	Shannon LowBd	Lloyd-Max	Entropy-coded
Uniform	$\frac{6}{\pi e} \cong 0.703$	1	1
Laplacian	$\frac{e}{\pi} \cong 0.865$	$\frac{9}{2} = 4.5$	$\frac{e^2}{6} \cong 1.232$
Gaussian	1	$\frac{\sqrt{3}\pi}{2} \cong 2.721$	$\frac{\pi e}{6} \cong 1.423$

Deadzone uniform quantizer

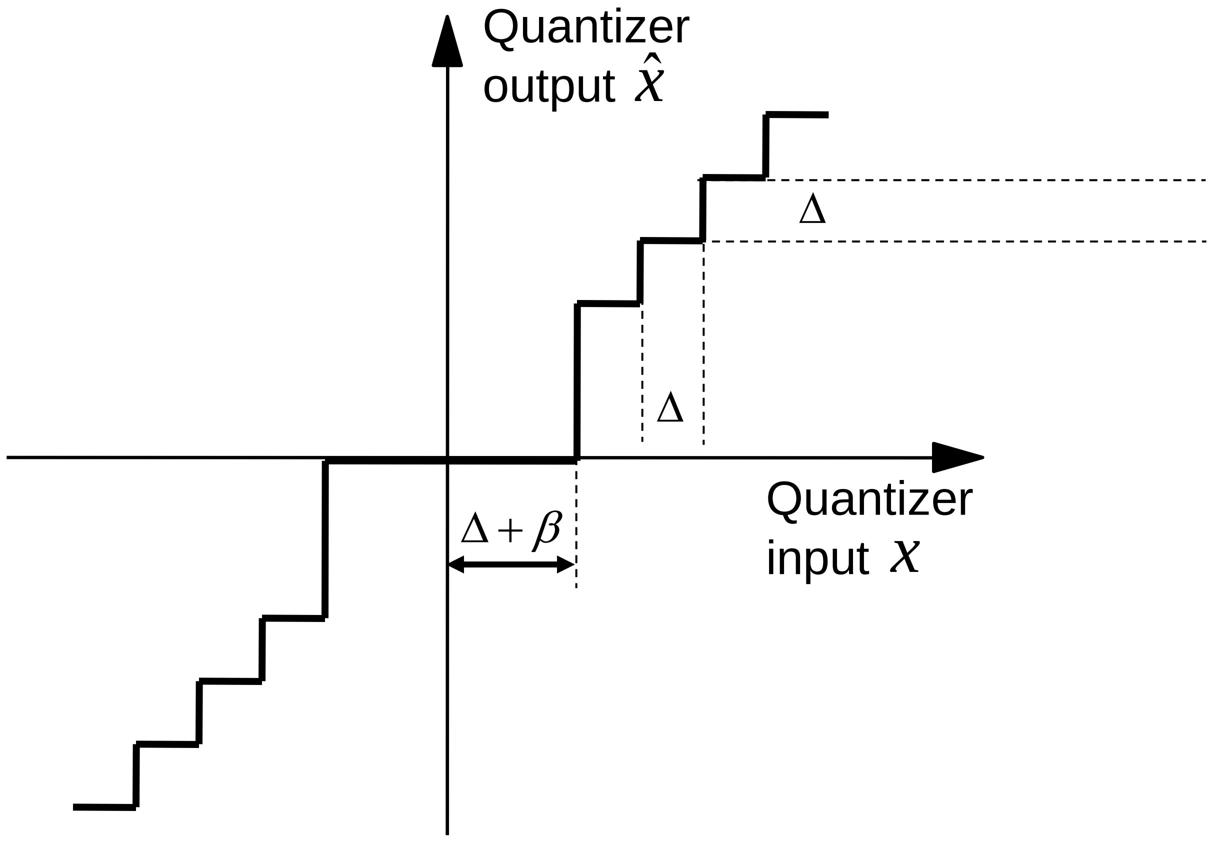

Embedded quantizers

- Motivation: “scalability” – decoding of compressed bitstreams at different rates (with different qualities)
- Nested quantization intervals

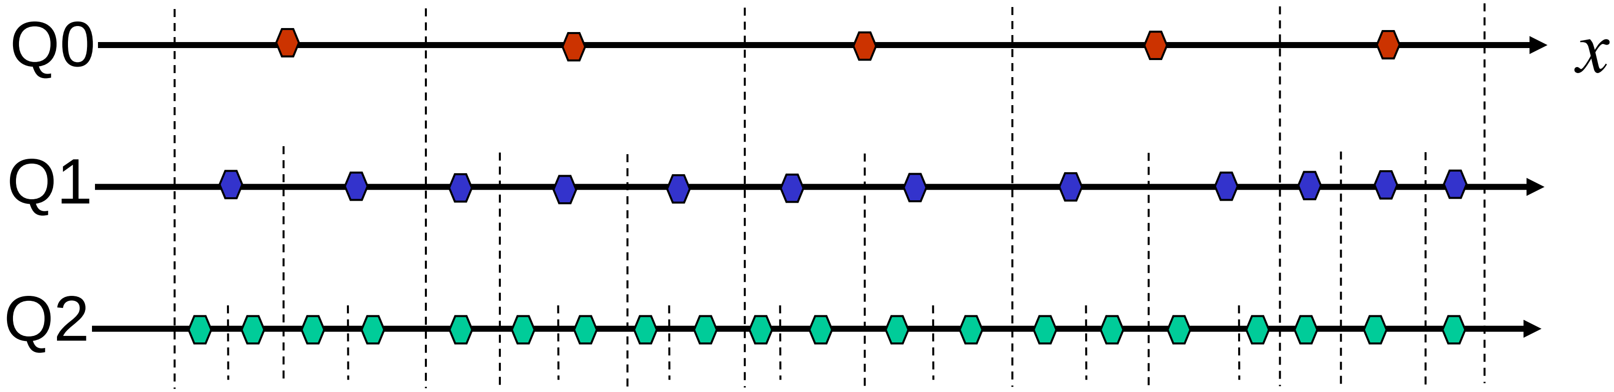

- In general, only one quantizer can be optimum (exception: uniform quantizers)

Example: Lloyd-Max quantizers for Gaussian PDF

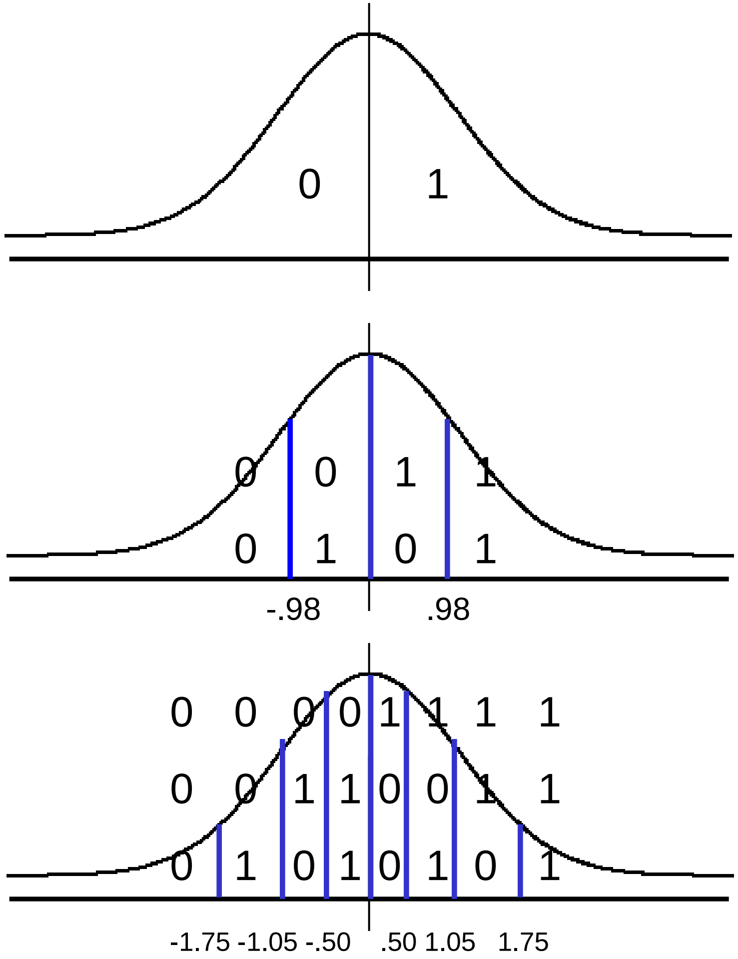

- 2-bit and 3-bit optimal quantizers not embeddable
- Performance loss for embedded quantizers

Information theoretic analysis

- "Successive refinement" – Embedded coding at multiple rates w/o loss relative to R-D function

$$E\left[d(X,\hat{X}_1)\right] \le D_1 \qquad I(X;\hat{X}_1) = R(D_1)$$

$$E\left[d(X,\hat{X}_2)\right] \le D_2 \qquad I(X;\hat{X}_2) = R(D_2)$$

- "Successive refinement" with distortions D_1 and $D_2 \le D_1$ can be achieved ***iff*** there exists a conditional distribution

$$f_{\hat{X}_1,\hat{X}_2|X}(\hat{x}_1,\hat{x}_2,x) = f_{\hat{X}_2|X}(\hat{x}_2,x)\, f_{\hat{X}_1|\hat{X}_2}(\hat{x}_1,\hat{x}_2)$$

- Markov chain condition

$$X \leftrightarrow \hat{X}_2 \leftrightarrow \hat{X}_1$$

[Equitz,Cover, 1991]

Embedded deadzone uniform quantizers

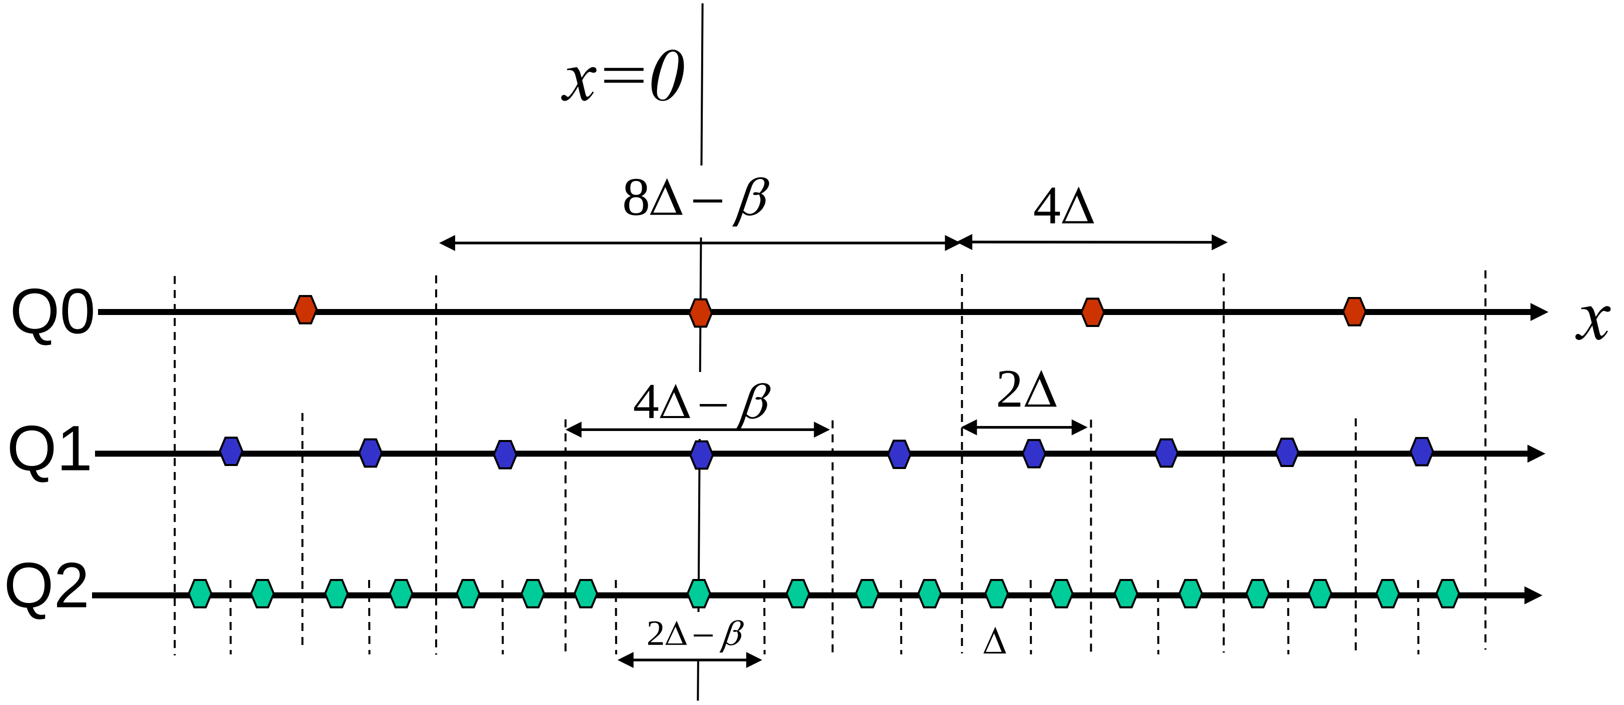

Supported in JPEG-2000 with general β for quantization of wavelet coefficients.

Vector quantization

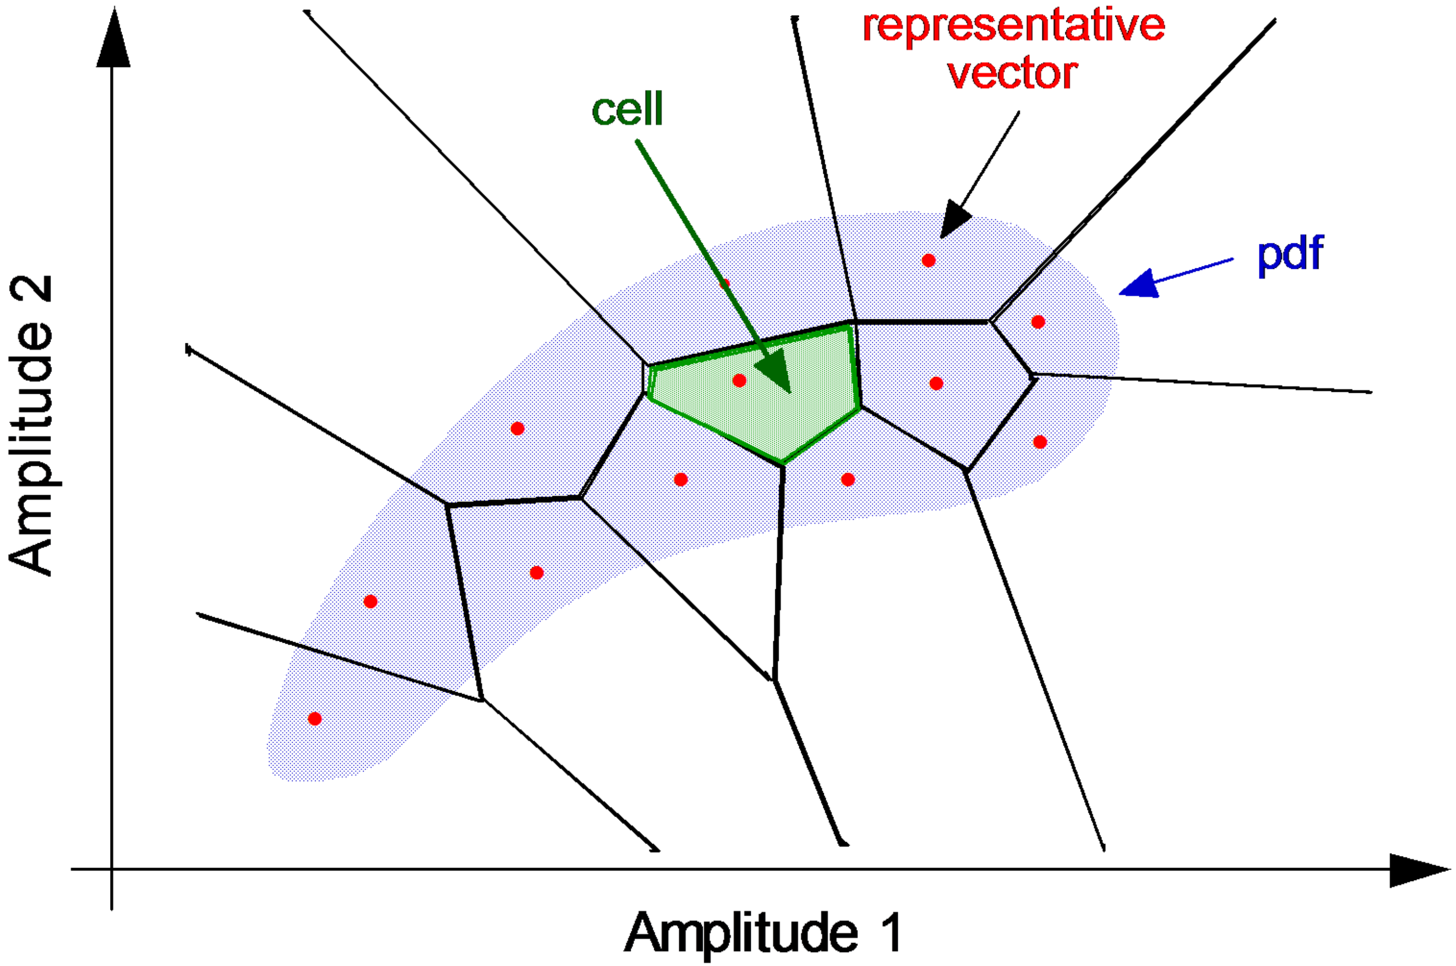

LBG algorithm

- Lloyd algorithm generalized for VQ *[Linde, Buzo, Gray, 1980]*

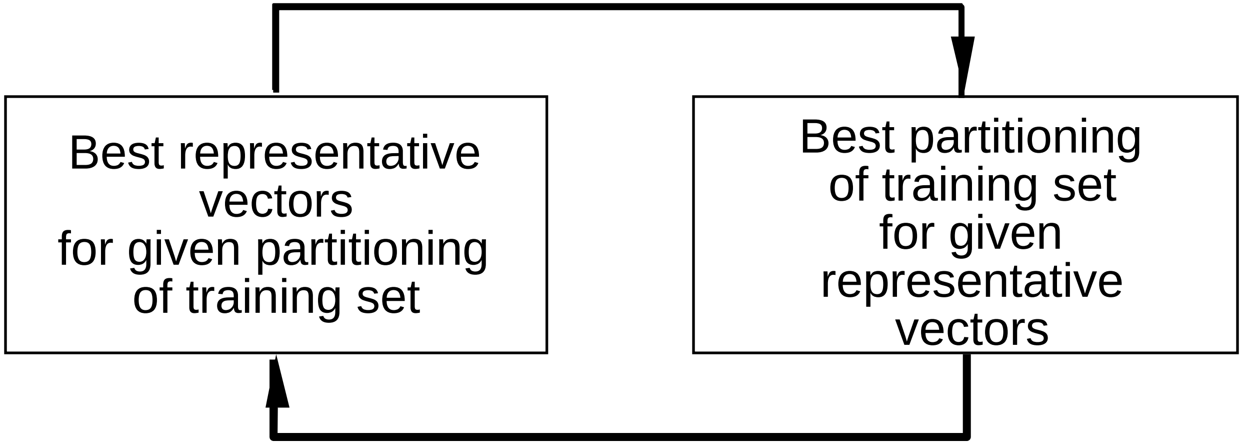

- Assumption: fixed code word length
- Code book unstructured: full search

Design of entropy-coded vector quantizers

- Extended LBG algorithm for entropy-coded VQ
  [Chou, Lookabaugh, Gray, 1989]
- Lagrangian cost function: solve unconstrained problem rather than constrained problem

$$J = d + \lambda R = E\left[\left\| X - \hat{X} \right\|^2\right] + \lambda H\left(\hat{X}\right) \rightarrow \min.$$

- Unstructured code book: full search for

$$J_{x_i}(q) = \left\| x_i - \hat{x}_q \right\|^2 - \lambda \log_2 p_q$$

The most general coder structure:
Any source coder can be interpreted as VQ with VLC!

Lattice vector quantization

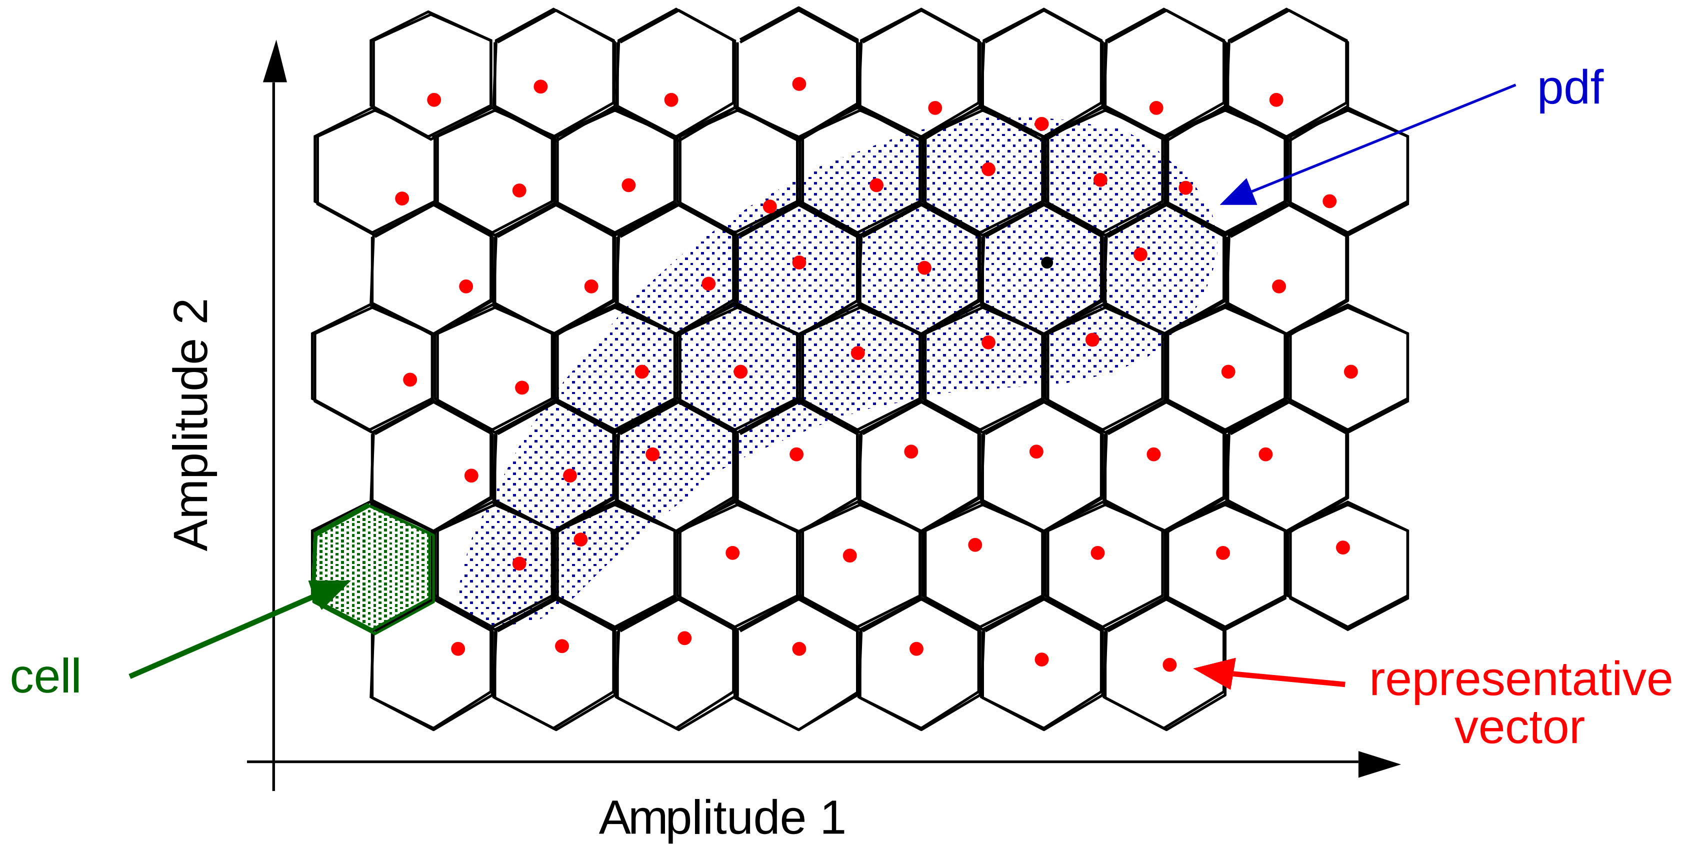

8D VQ of a Gauss-Markov source

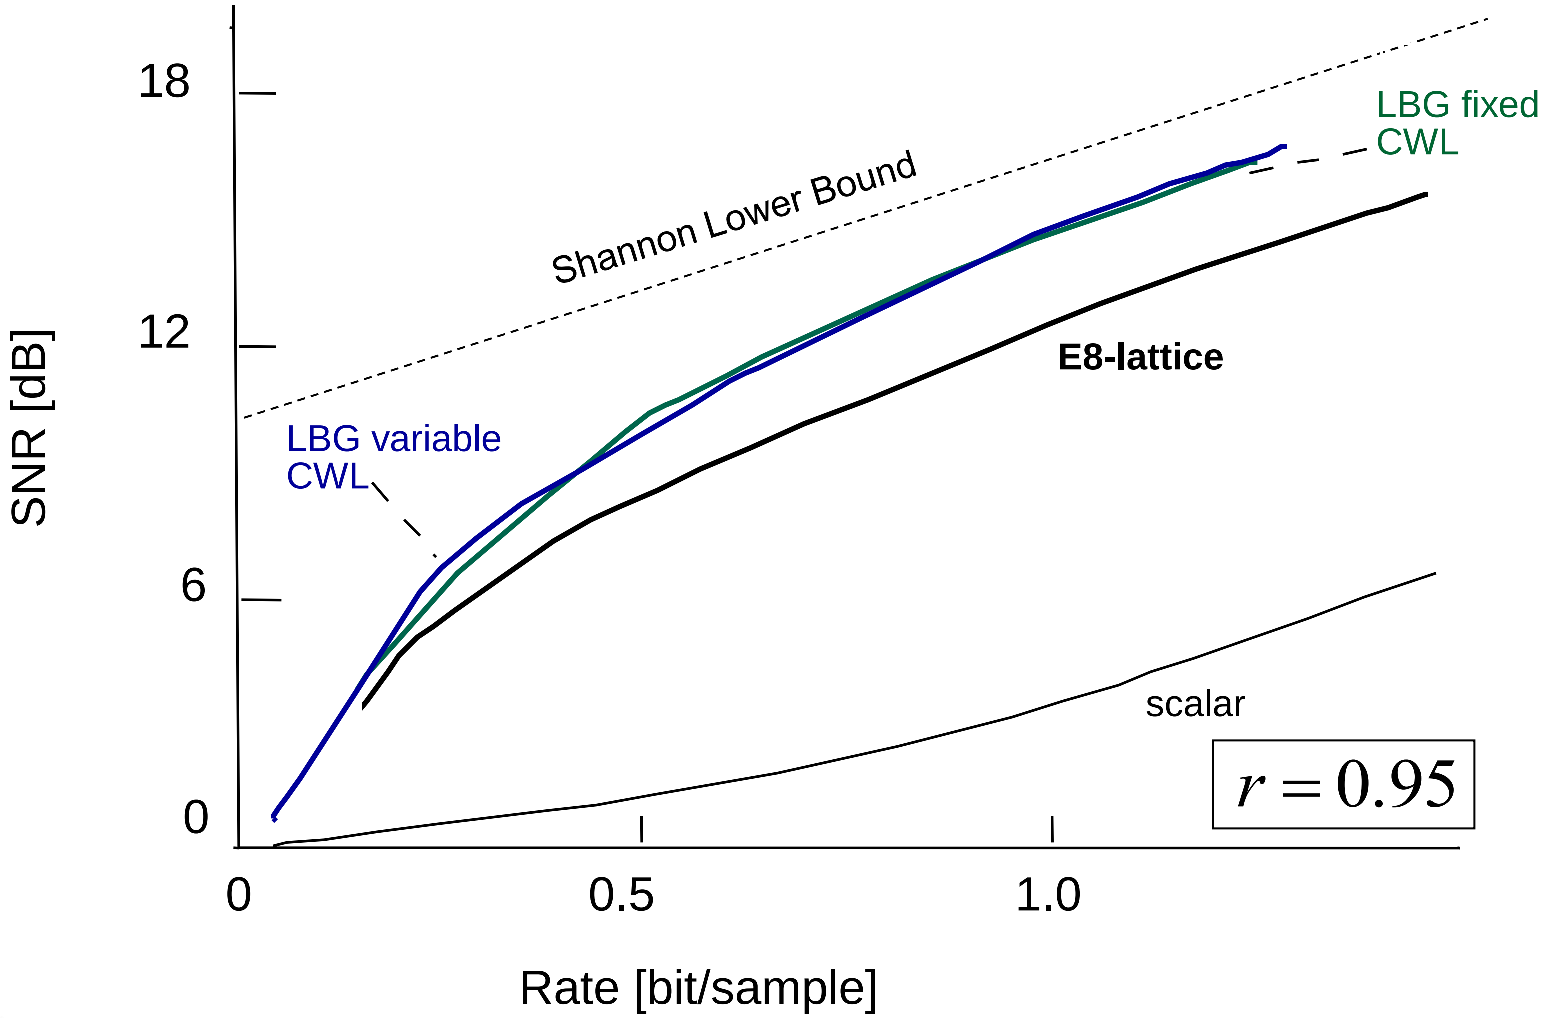

8D VQ of memoryless Laplacian source

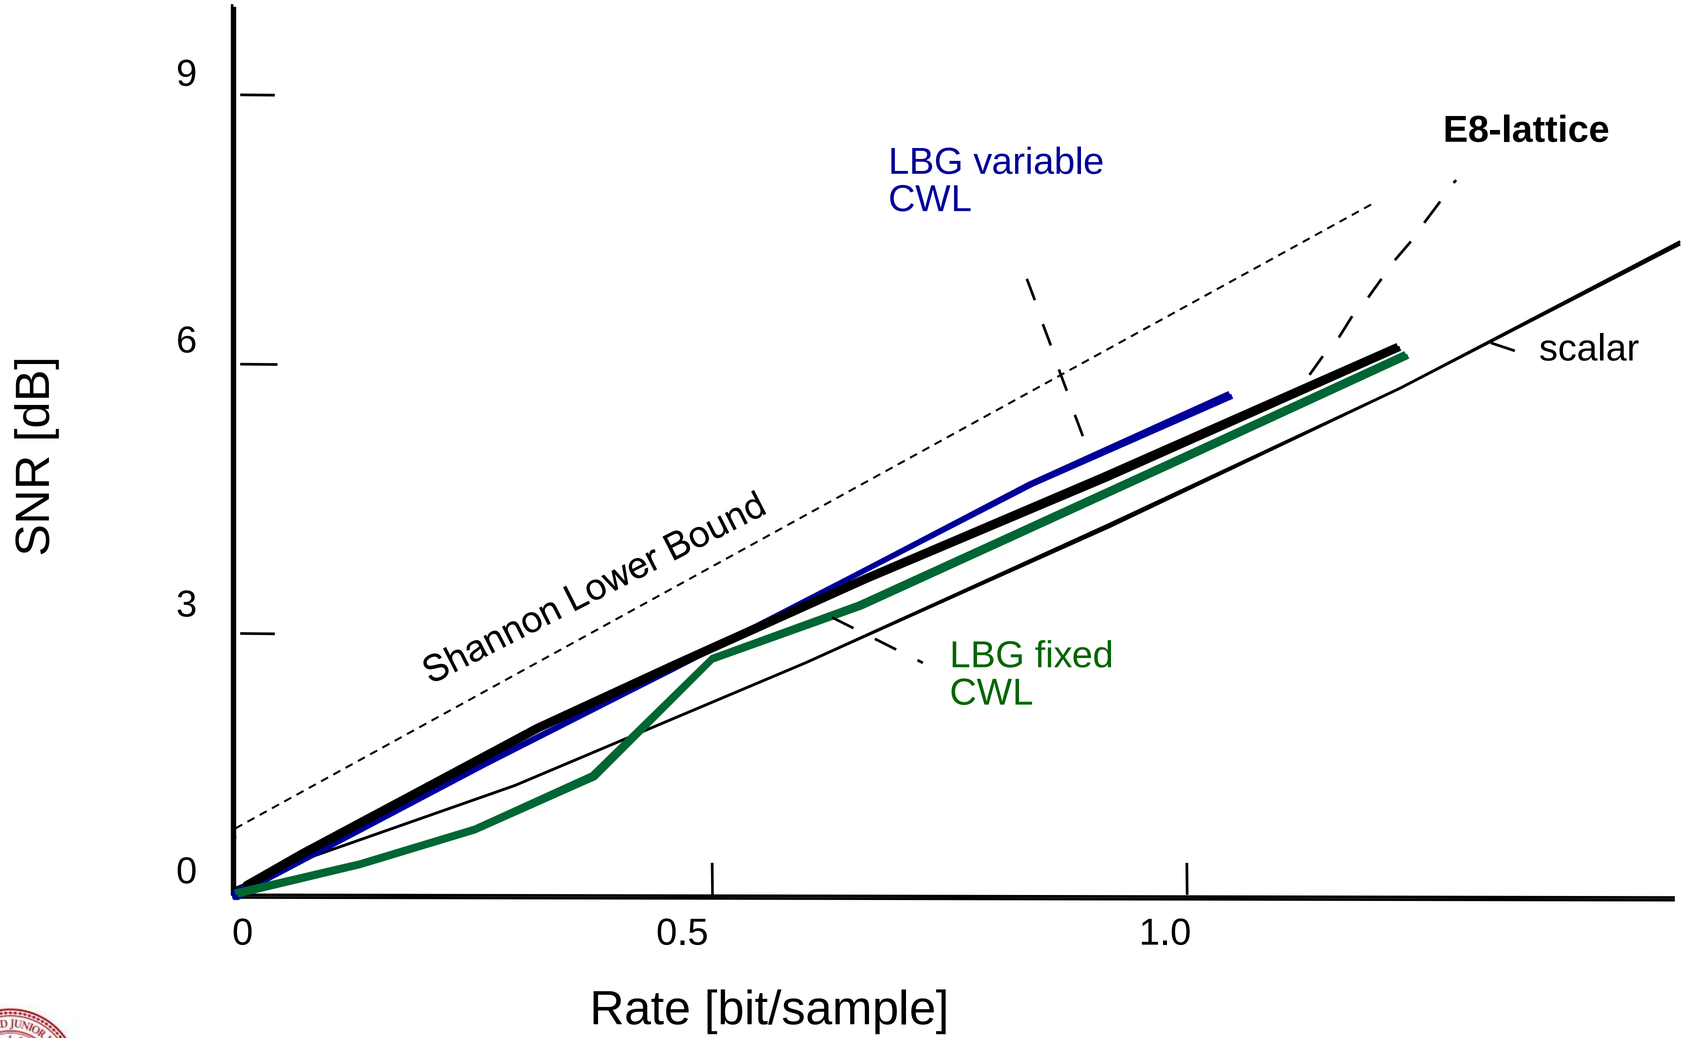

Reading

- Taubman, Marcellin, Sections 3.2, 3.4
- J. Max, "Quantizing for Minimum Distortion," IEEE Trans. Information Theory, vol. 6, no. 1, pp. 7-12, March 1960.
- S. P. Lloyd, "Least Squares Quantization in PCM," IEEE Trans. Information Theory, vol. 28, no. 2, pp. 129-137, March 1982.
- P. A. Chou, T. Lookabaugh, R. M. Gray, "Entropy-constrained vector quantization," IEEE Trans. Signal Processing, vol. 37, no. 1, pp. 31-42, January 1989.