

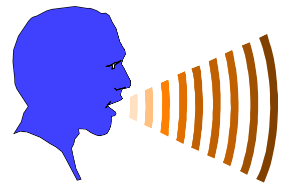


EE3414
Multimedia Communication Systems - I

Quantization

Yao Wang
Polytechnic University, Brooklyn, NY11201
http://eeweb.poly.edu/~yao

Outline

- Review the three process of A to D conversion
- Quantization
  - Uniform
  - Non-uniform
    - Mu-law
  - Demo on quantization of audio signals
  - Sample Matlab codes
- Binary encoding
  - Bit rate of digital signals
- Advantage of digital representation

Three Processes in A/D Conversion

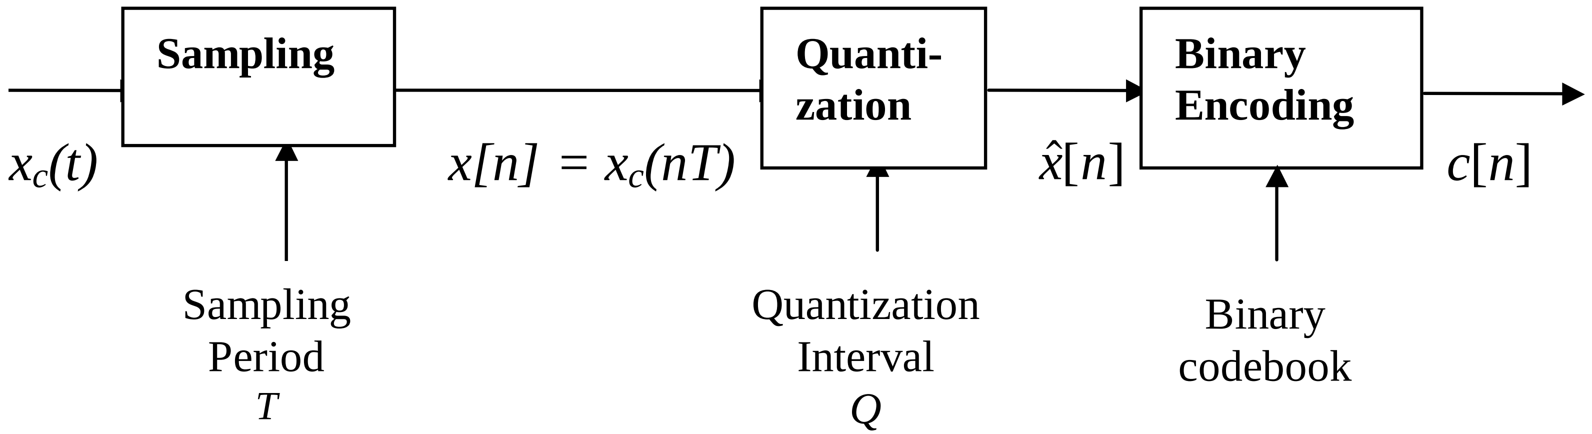

- Sampling: take samples at time *nT*
  - T: sampling period;
  - f_s = 1/T: sampling frequency
- Quantization: map amplitude values into a set of discrete values *kQ*
  - Q: quantization interval or stepsize
- Binary Encoding
  - Convert each quantized value into a binary codeword

Analog to Digital Conversion

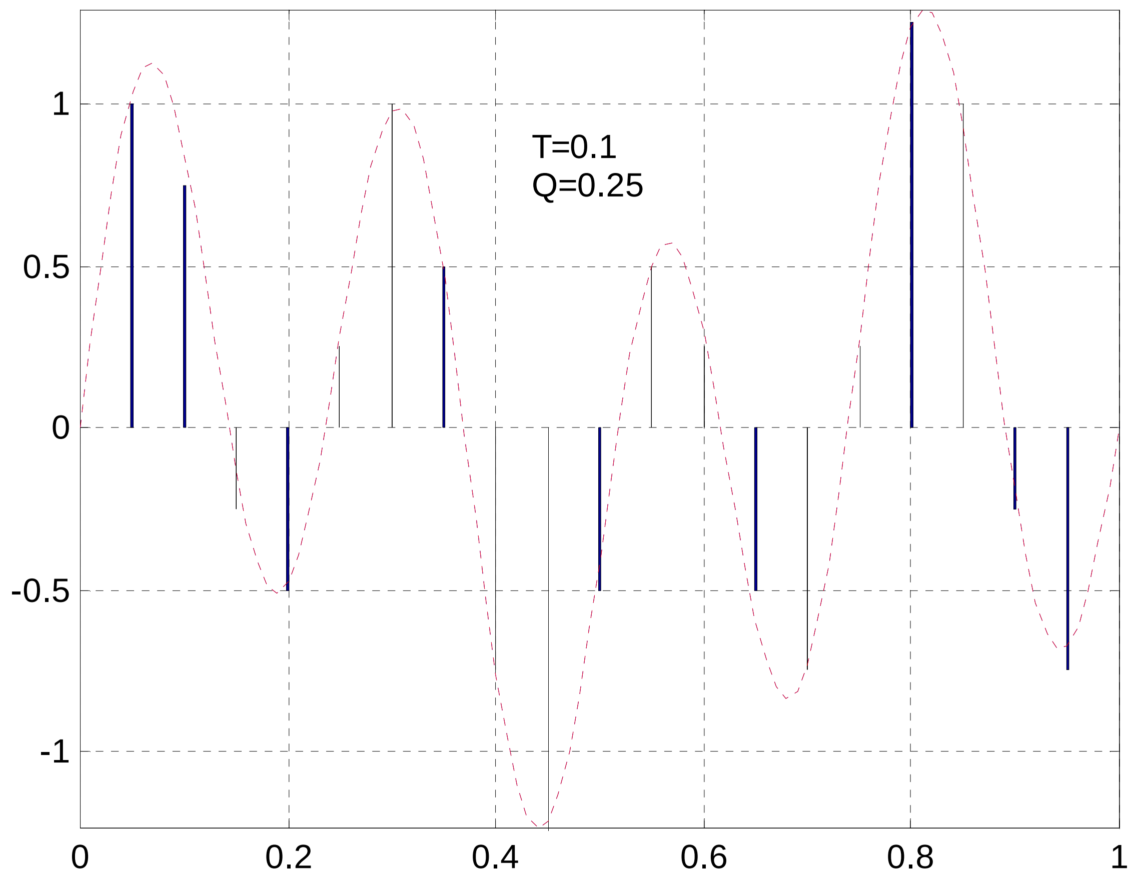

A2D_plot.m

How to determine T and Q?

- T (or f_s) depends on the signal frequency range
  - A fast varying signal should be sampled more frequently!
  - Theoretically governed by the Nyquist sampling theorem
    - $f_s > 2 f_m$ (f_m is the maximum signal frequency)
    - For speech: $f_s >= 8$ *KHz;* For music: $f_s >= 44$ *KHz;*
- Q depends on the dynamic range of the signal amplitude and perceptual sensitivity
  - Q and the signal range D determine bits/sample R
    - $2^R = D/Q$
    - For speech: $R = 8$ *bits;* For music: $R = 16$ *bits;*
- One can trade off T (or f_s) and Q (or R)
  - lower R -> higher f_s; higher R -> lower f_s
- We considered sampling in last lecture, we discuss quantization in this lecture

Uniform Quantization

- Applicable when the signal is in a finite range (f_{min}, f_{max})

- The entire data range is divided into L equal intervals of length Q *(known as quantization interval or quantization step-size)*

$$Q=(f_{max}-f_{min})/L$$

- Interval i is mapped to the middle value of this interval

- We store/send only the index of quantized value

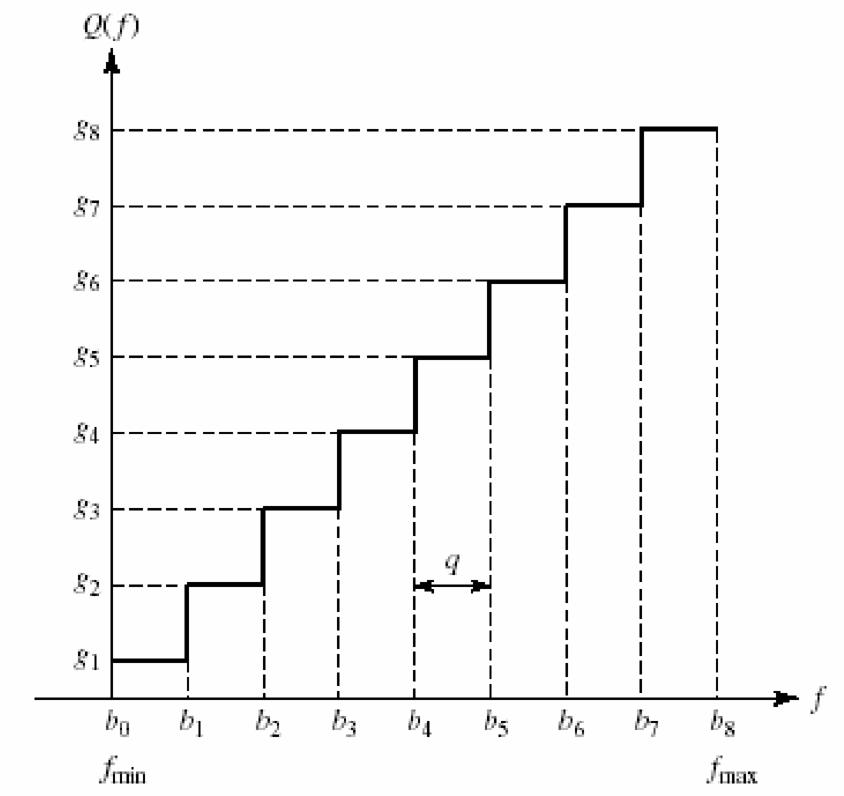

$$\text{Index of quantized value} = Q_i(f) = \left\lfloor \frac{f - f_{\min}}{Q} \right\rfloor$$

$$\text{Quantized value} = Q(f) = Q_i(f)Q + Q/2 + f_{\min}$$

Special Case I: Signal range is symmetric

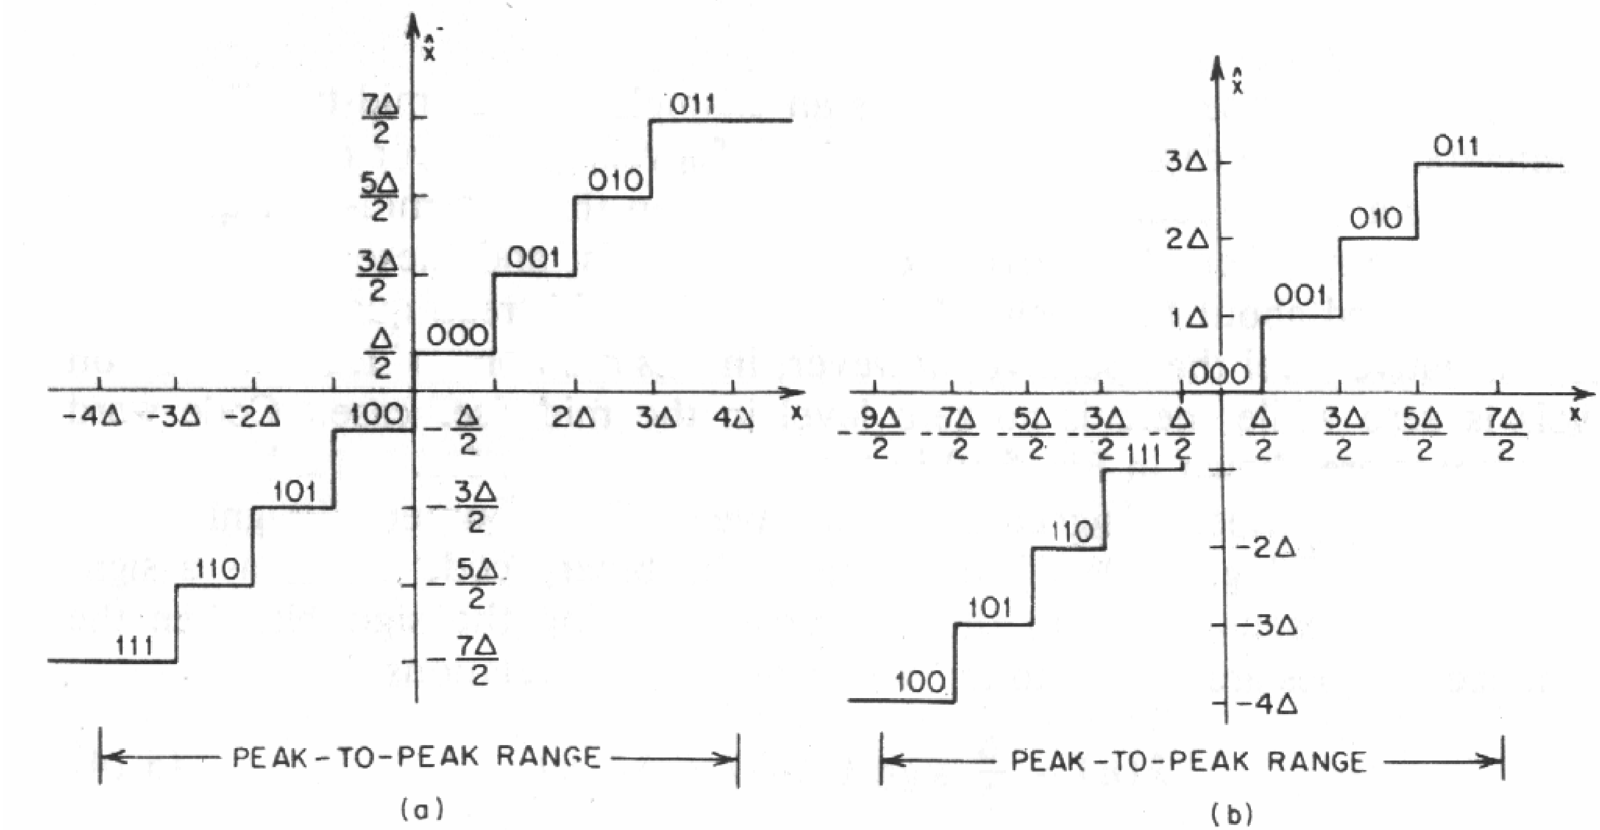

L = even, Mid - Riser

$$Q_i(f) = floor(\frac{f}{Q}), \quad Q(f) = Q_i(f) * Q + \frac{Q}{2}$$

L = odd, Mid - Tread

$$Q_i(f) = round(\frac{f}{Q}), \quad Q(f) = Q_i(f) * Q$$

Special Case II: Signal range starts at 0

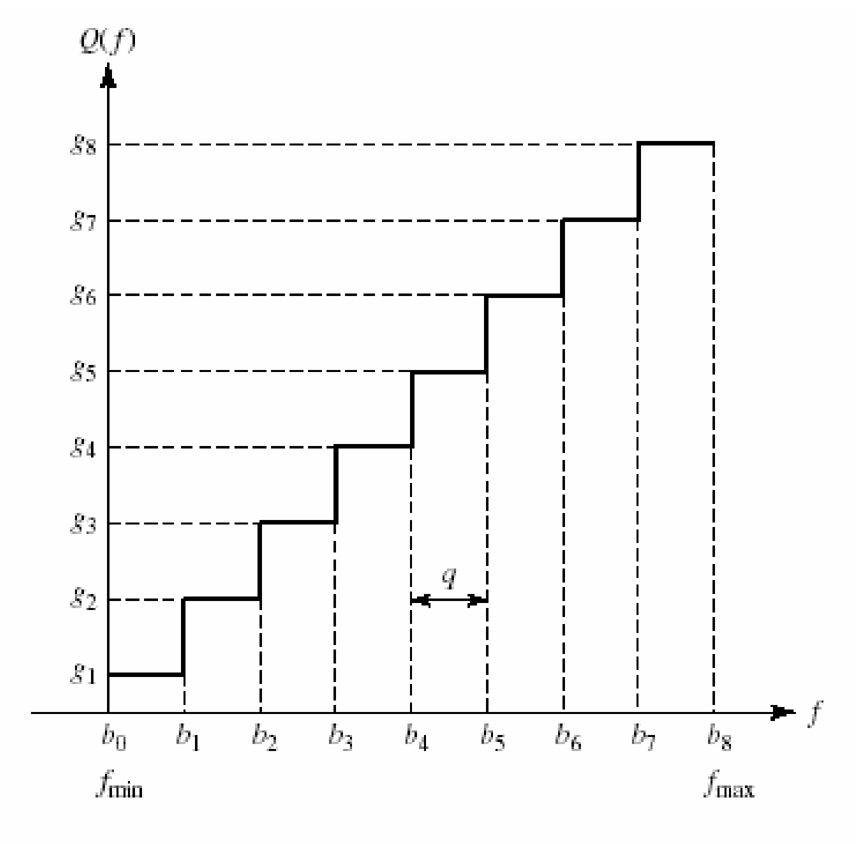

$$f_{\min} = 0, B = f_{\max}, q = f_{\max} / L$$

$$Q_i(f) = floor(\frac{f}{Q})$$

$$Q(f) = Q_i(f) * Q + \frac{Q}{2}$$

Example

- For the following sequence {1.2,-0.2,-0.5,0.4,0.89,1.3…}, Quantize it using a uniform quantizer in the range of (-1.5,1.5) with 4 levels, and write the quantized sequence.
- Solution: Q=3/4=0.75. Quantizer is illustrated below.

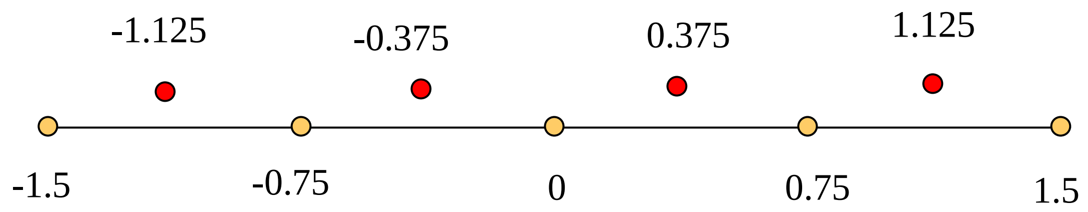

Yellow dots indicate the partition levels (boundaries between separate quantization intervals)
Red dots indicate the reconstruction levels (middle of each interval)

1.2 fall between 0.75 and 1.5, and hence is quantized to 1.125

- Quantized sequence:

{1.125,-0.375,-0.375,0.375,1.125,1.125}

Effect of Quantization Stepsize

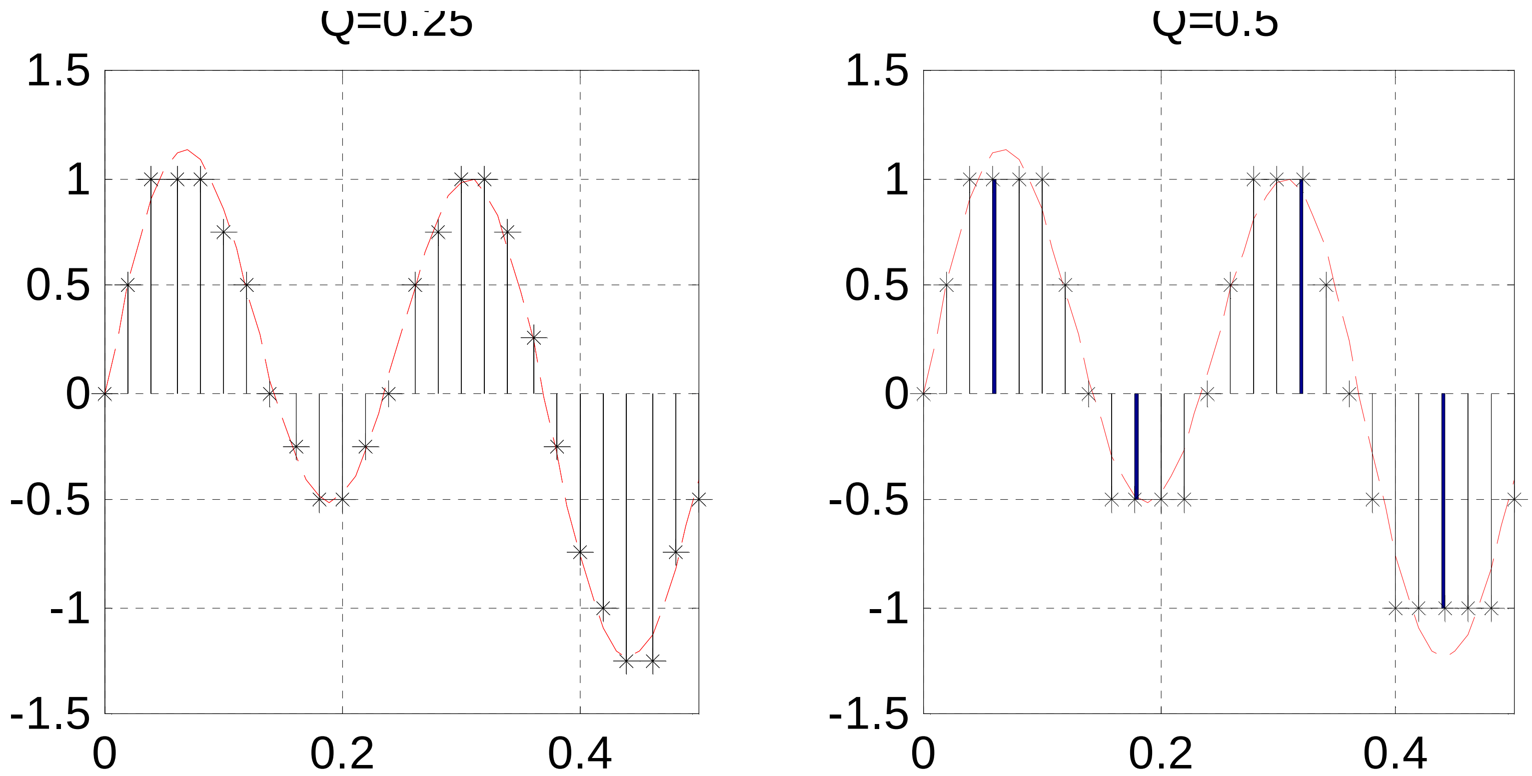

demo_sampling_quant.m

Demo: Audio Quantization

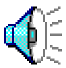

Original
Mozart.wav

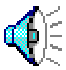

Quantized
Mozart_q16.wav

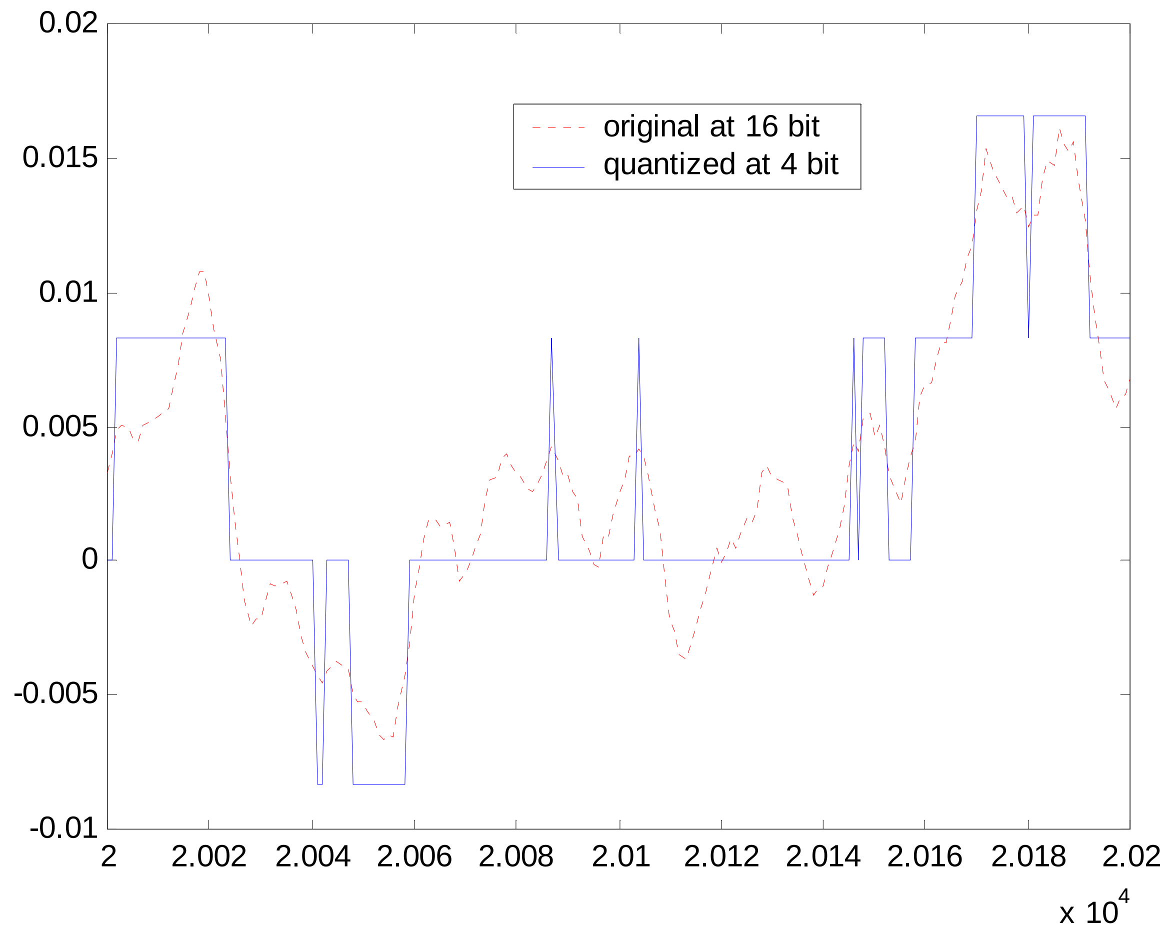

Demo: Audio Quantization (II)

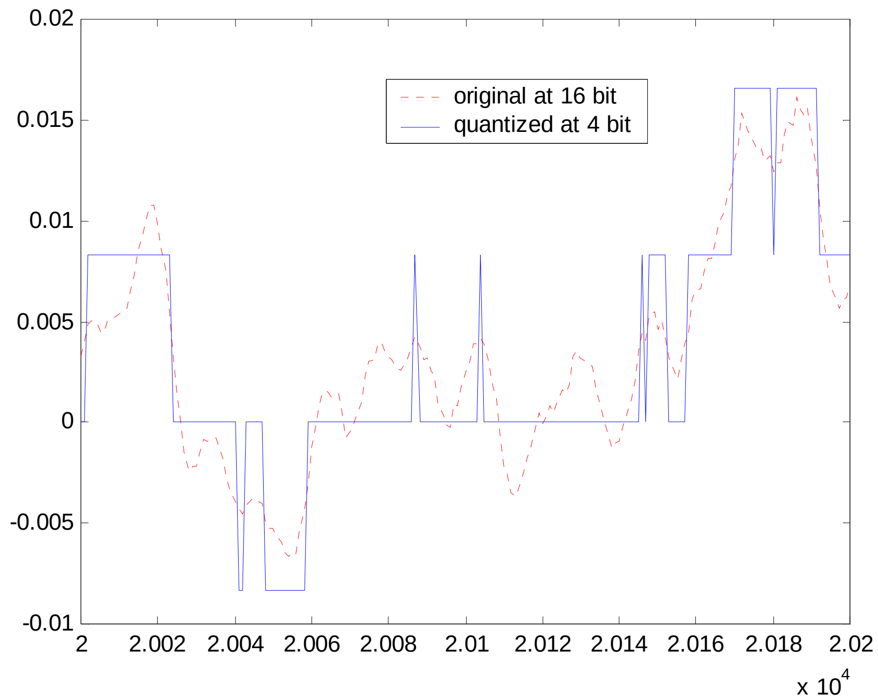

Quantized
Mozart_q16.wav

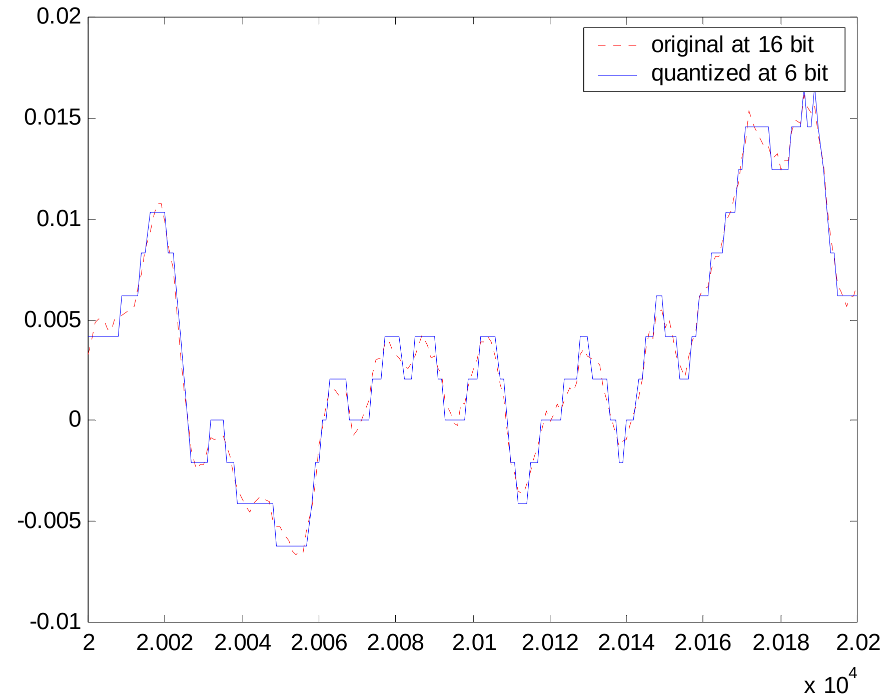

Quantized
Mozart_q64.wav

Non-Uniform Quantization

- Problems with uniform quantization
  - Only optimal for uniformly distributed signal
  - Real audio signals (speech and music) are more concentrated near zeros
  - Human ear is more sensitive to quantization errors at small values
- Solution
  - Using non-uniform quantization
    - quantization interval is smaller near zero

Quantization: General Description

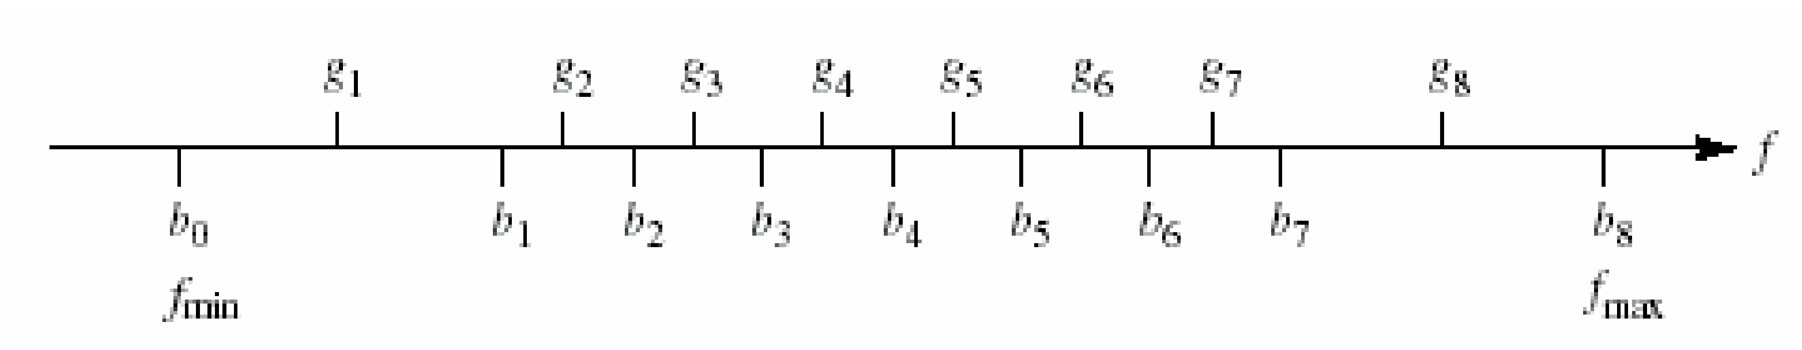

Quantization levels: L

Partition values: b_l

Partition regions: $B_l = [b_{l-1}, b_l)$

Reconstruction values: g_l

Quantized Index: $Q_i(f) = l$, if $f \in B_l$

Quantizer value: $Q(f) = g_l$, if $f \in B_l$

Function Representation

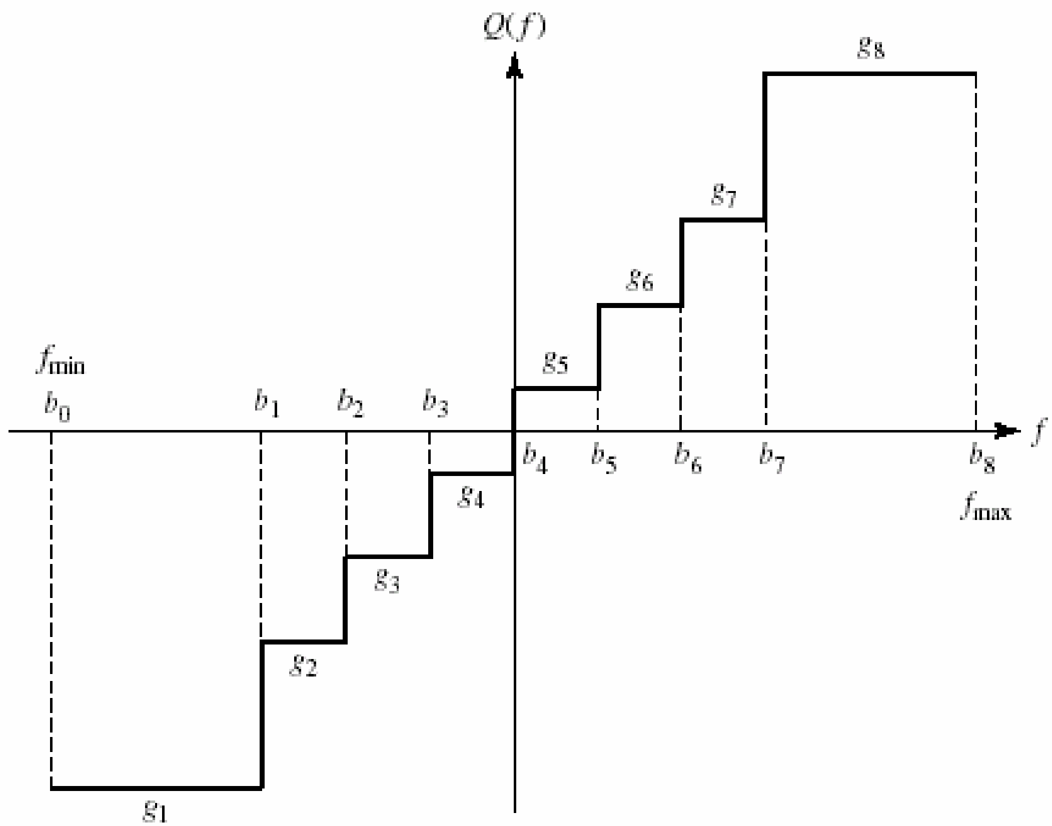

$$Q(f) = g_l, \text{ if } \quad f \in B_l$$

Design of Non-Uniform Quantizer

- Directly design the partition and reconstruction levels
- Non-linear mapping+uniform quantization

  μ-law quantization

μ-Law Quantization

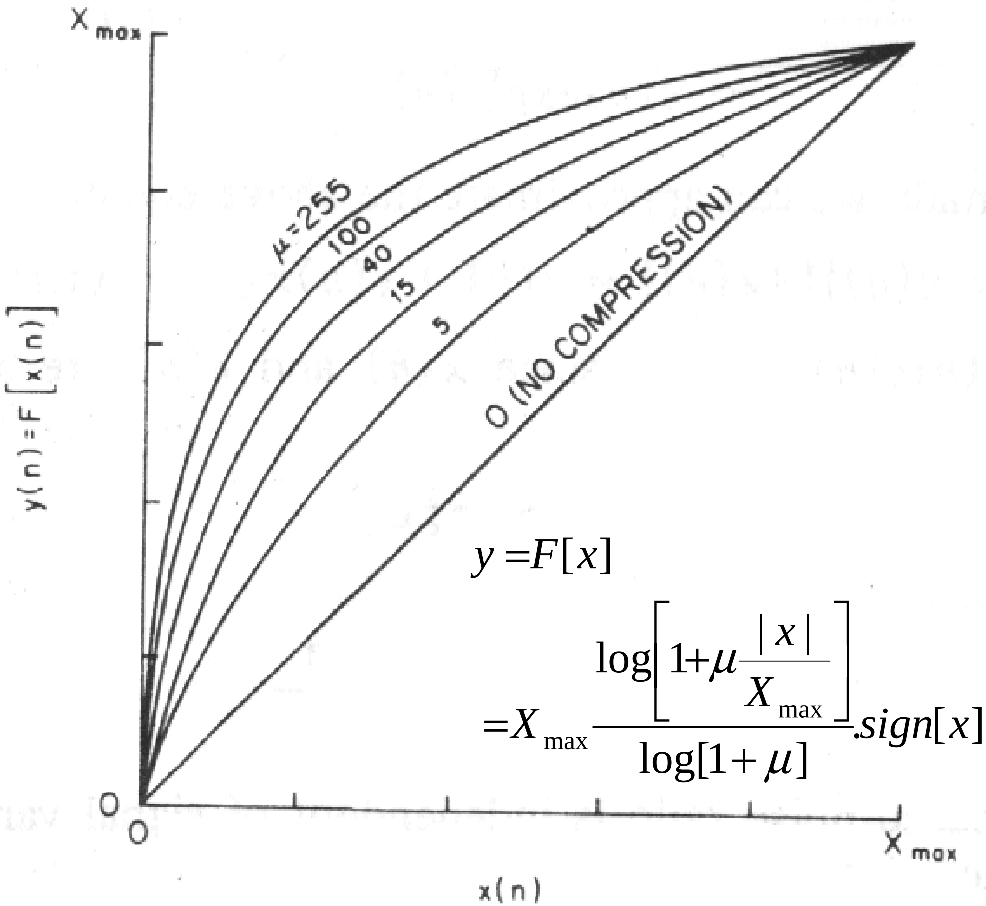

$$y = F[x]$$

$$= X_{\max} \frac{\log\left[1+\mu \frac{|x|}{X_{\max}}\right]}{\log[1+\mu]} . sign[x]$$

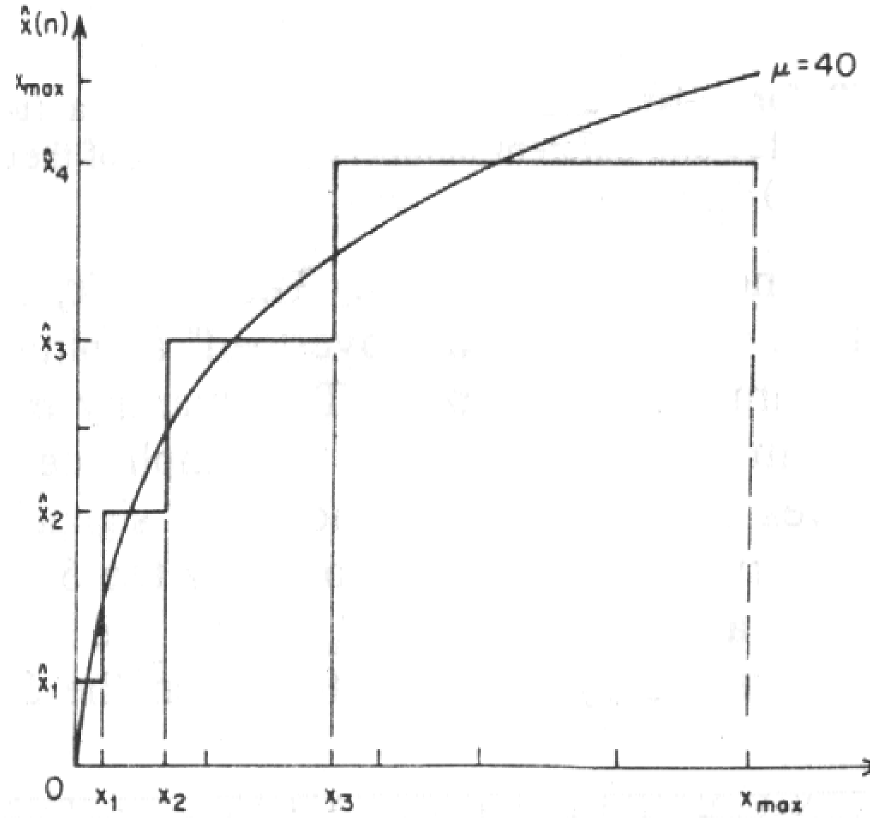

Implementation of μ-Law Quantization (Direct Method)

– Transform the signal using μ-law: x->y

$$y = F[x] = X_{\max} \frac{\log\left[1 + \mu \frac{|x|}{X_{\max}}\right]}{\log[1+\mu]} . sign[x]$$

– Quantize the transformed value using a uniform quantizer: y->y^

– Transform the quantized value back using inverse μ-law: y^->x^

$$x = F^{-1}[y] = \frac{X_{\max}}{\mu} \left(10^{\frac{\log(1+\mu)}{X_{\max}}|y|} - 1 \right) \mathrm{sign(y)}$$

Implementation of μ-Law Quantization (Indirect Method)

- Indirect Method:
  - Instead of applying the above computation to each sample, one can pre-design a quantization table (storing the partition and reconstruction levels) using the above procedure. The actual quantization process can then be done by a simple table look-up.
  - Applicable both for uniform and non-uniform quantizers
  - How to find the partition and reconstruction levels for mu-law quantizer
    - Apply inverse mu-law mapping to the partition and reconstruction levels of the uniform quantizer for y.
    - Note that the mu-law formula is designed so that if x ranges from (-x_max, x_max), then y also has the same range.

Example

- For the following sequence {1.2,-0.2,-0.5,0.4,0.89,1.3…}, Quantize it using a mu-law quantizer in the range of (-1.5,1.5) with 4 levels, and write the quantized sequence.

- Solution (indirect method):
  - apply the inverse formula to the partition and reconstruction levels found for the previous uniform quantizer example. Because the mu-law mapping is symmetric, we only need to find the inverse values for y=0.375,0.75,1.125

    μ=9, x_max=1.5, 0.375->0.1297, 0.75->0.3604, 1.125->0.7706
  - Then quantize each sample using the above partition and reconstruction levels.

Example (cntd)

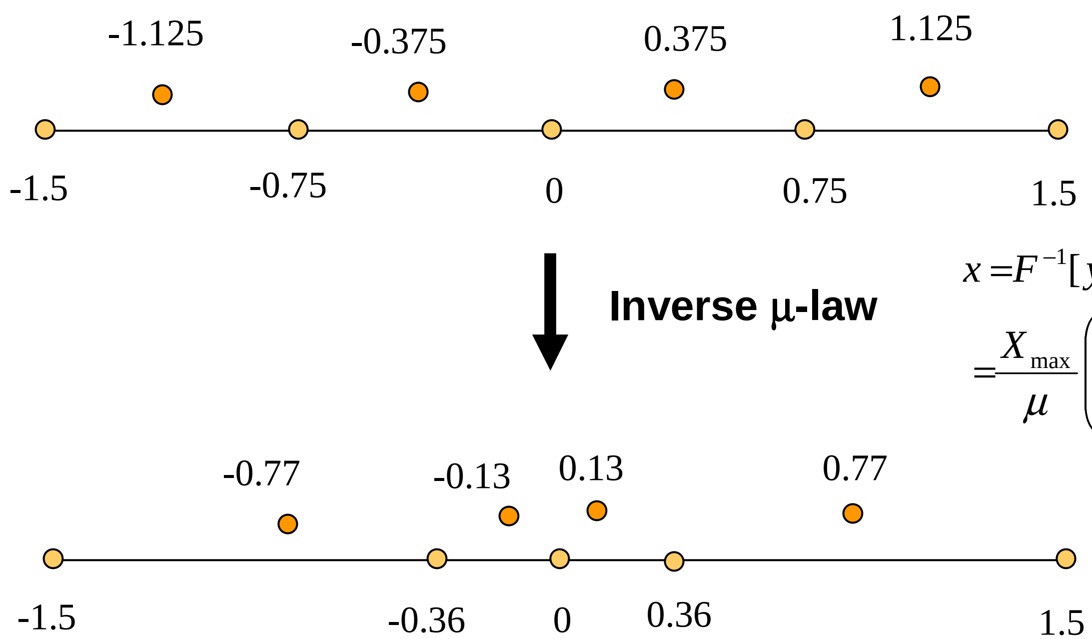

$$x = F^{-1}[y]$$

$$= \frac{X_{\max}}{\mu}\left(10^{\frac{\log(1+\mu)}{X_{\max}}|y|} - 1\right)\operatorname{sign}(y)$$

- Original sequence: {1.2,-0.2,-0.5,0.4,0.89,1.3…}
- Quantized sequence
  - {0.77,-0.13,-0.77,0.77,0.77,0.77}

Uniform vs. μ-Law Quantization

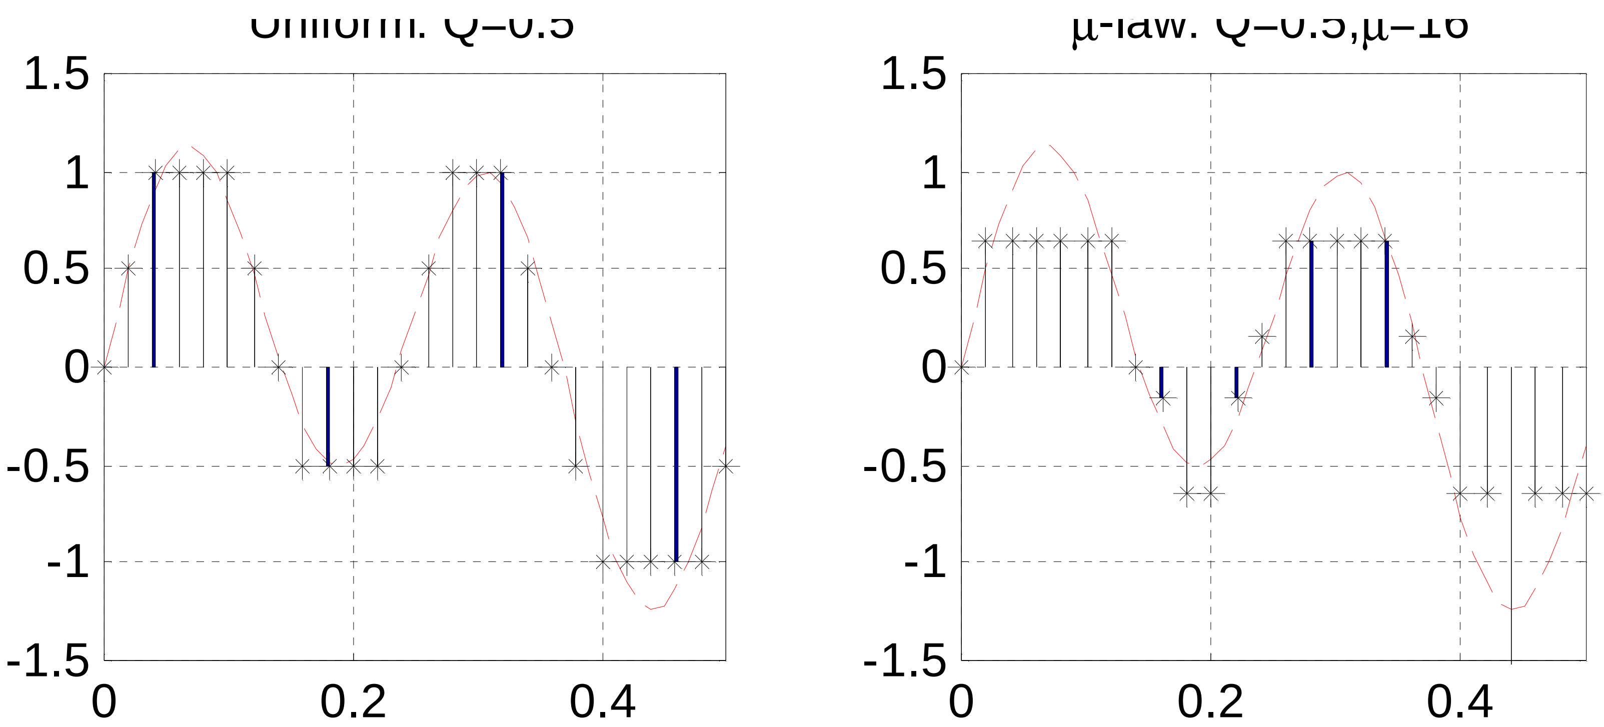

With μ-law, small values are represented more accurately, but large values are represented more coarsely.

Uniform vs. μ-Law for Audio

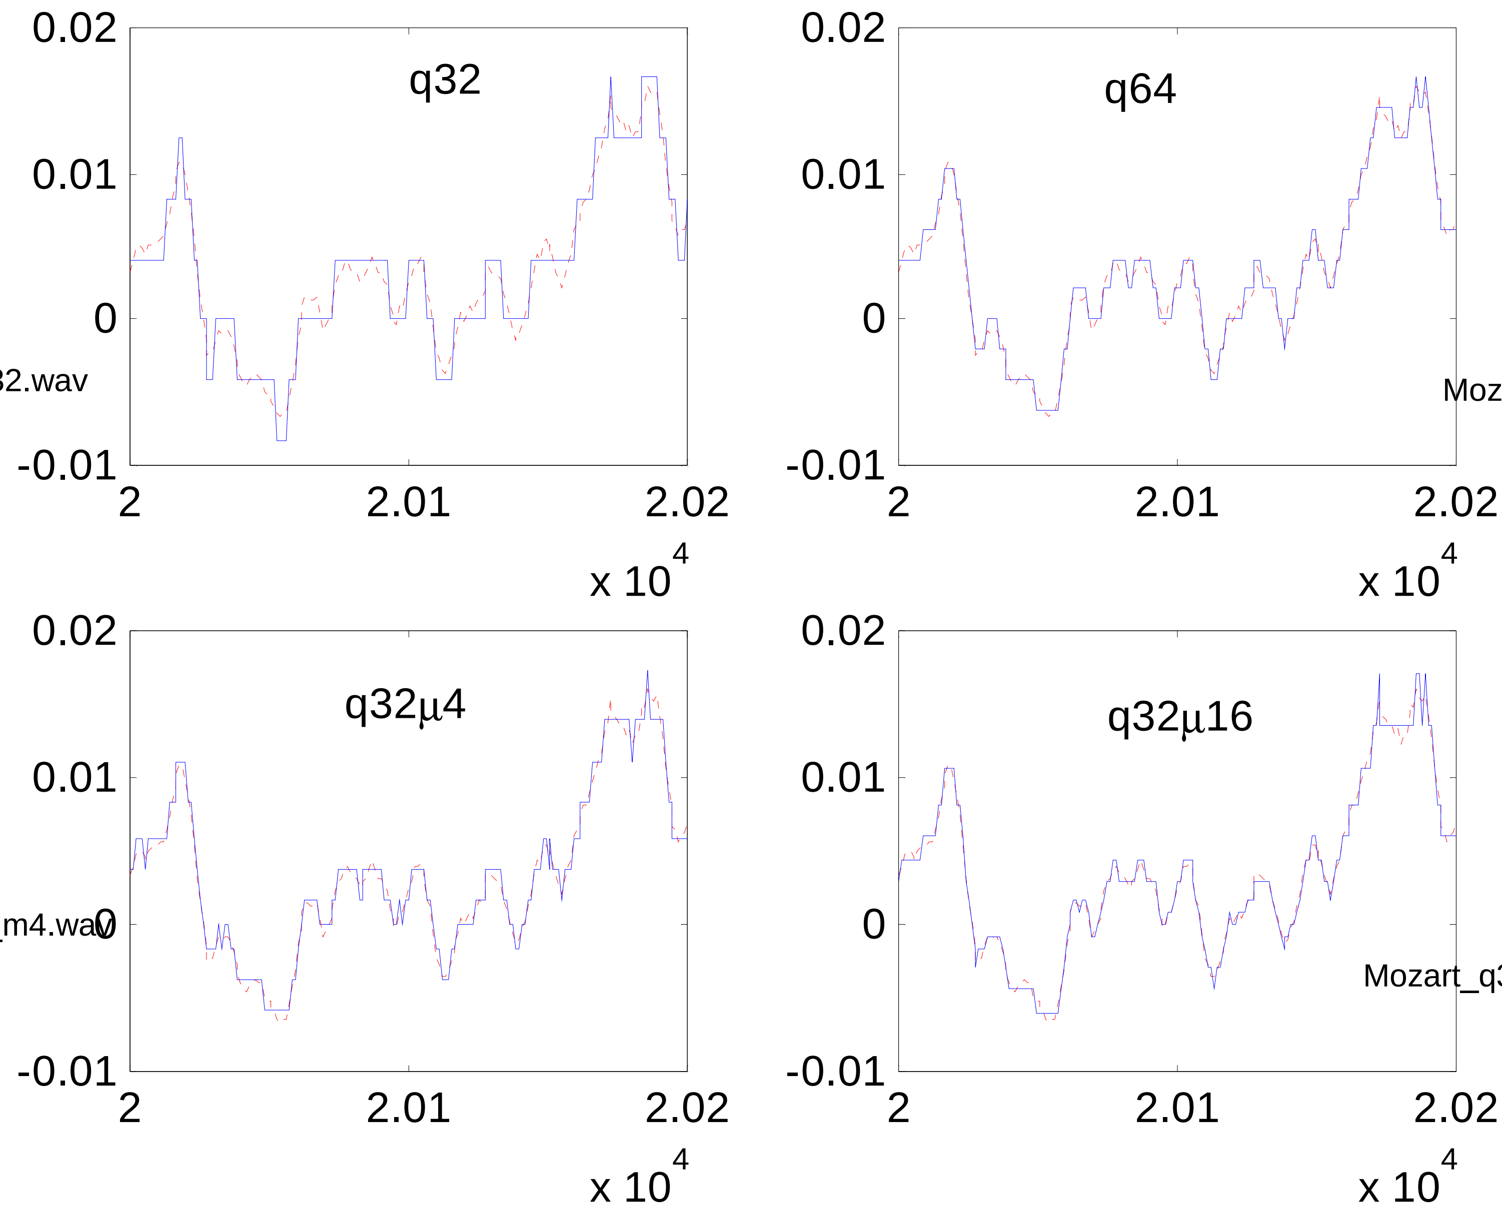

Mozart_q32.wav

Mozart_q64.wav

Mozart_q32_m4.wav

Mozart_q32_m16.wav

Evaluation of Quantizer Performance

- Ideally we want to measure the performance by how close is the quantized sound to the original sound to our ears -- Perceptual Quality
- But it is very hard to come up with a objective measure that correlates very well with the perceptual quality
- Frequently used objective measure – mean square error (MSE) between original and quantized samples or signal to noise ratio (SNR)

$$\text{MSE}: \quad \sigma_q^2 = \frac{1}{N}\sum_n (x(n) - \hat{x}(n))^2$$

$$\text{SNR(dB)}: \quad \text{SNR} = 10\log_{10}\left(\sigma_x^2 / \sigma_q^2\right)$$

where N is the number of samples in the sequence.

σ_x^2 is the variance of the original signal, $\sigma_z^2 = \frac{1}{N}\sum_n (x(n))^2$

Sample Matlab Code

Go through "quant_uniform.m", "quant_mulaw.m"

Binary Encoding

- Convert each quantized level index into a codeword consisting of binary bits
- Ex: natural binary encoding for 8 levels:
  - 000,001,010,011,100,101,110,111
- More sophisticated encoding (variable length coding)
  - Assign a short codeword to a more frequent symbol to reduce average bit rate
  - To be covered later

Example 1: uniform quantizer

- For the following sequence {1.2,-0.2,-0.5,0.4,0.89,1.3…}, Quantize it using a uniform quantizer in the range of (-1.5,1.5) with 4 levels, and write the quantized sequence and the corresponding binary bitstream.
- Solution: Q=3/4=0.75. Quantizer is illustrated below.
- Codewords: 4 levels can be represented by 2 bits, 00, 01, 10, 11

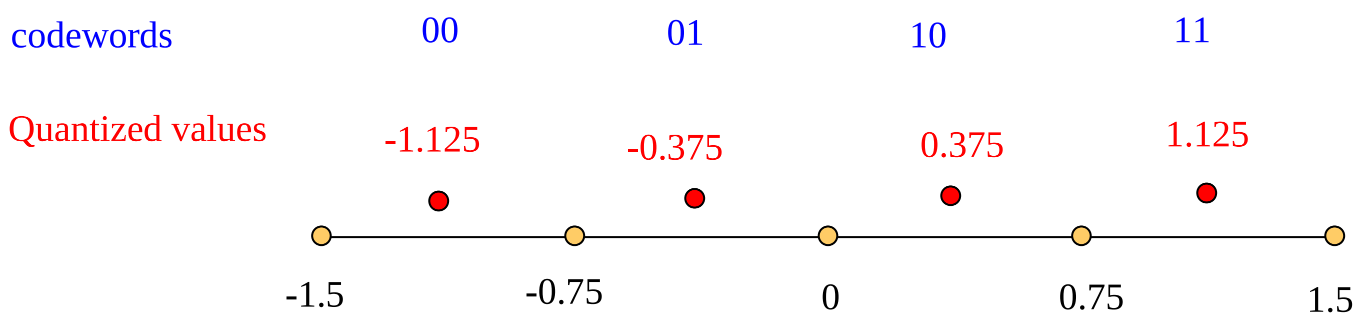

- Quantized value sequence:
  {1.125,-0.375,-0.375,0.375,1.125,1.125}
- Bitstream representing quantized sequence:
  {11, 01, 01, 10, 11, 11}

Example 2: mu-law quantizer

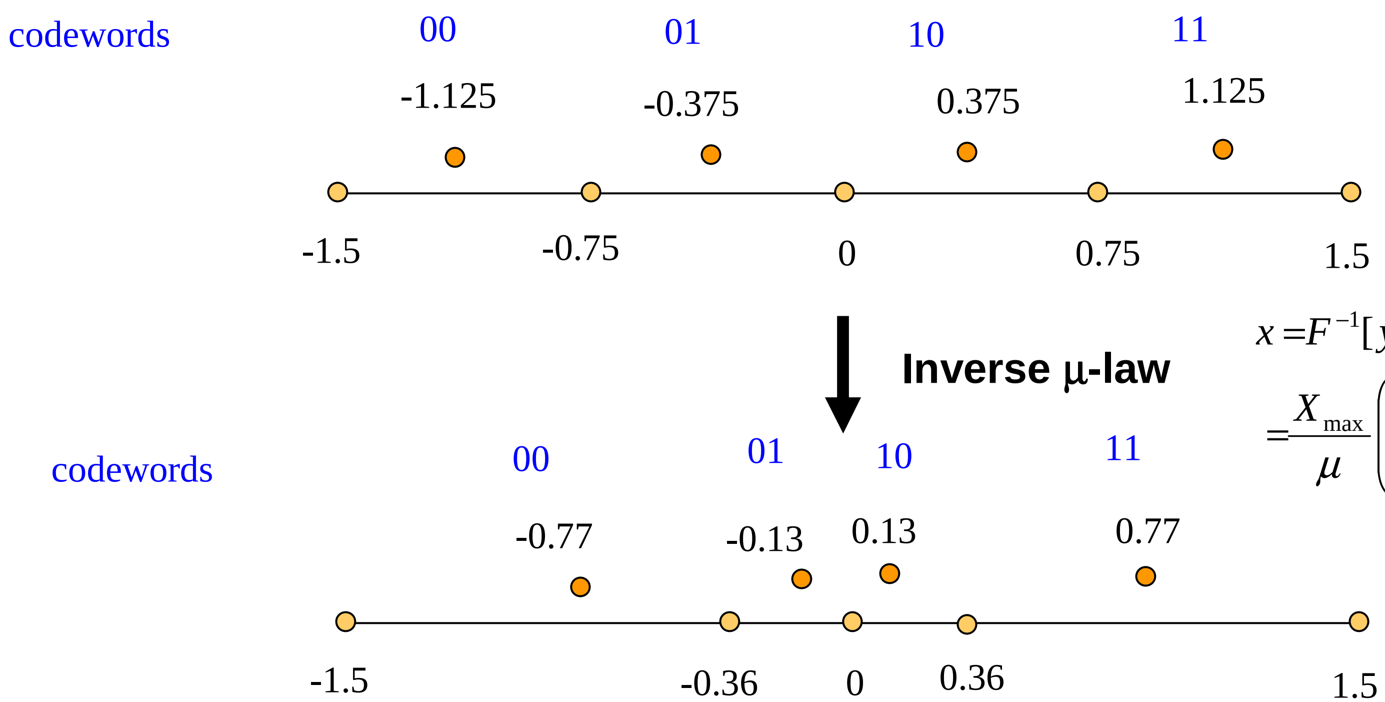

$$x = F^{-1}[y]$$

$$= \frac{X_{\max}}{\mu}\left(10^{\frac{\log(1+\mu)}{X_{\max}}|y|} - 1\right)\operatorname{sign}(y)$$

- Original sequence: {1.2,-0.2,-0.5,0.4,0.89,1.3…}
- Quantized sequence: {0.77,-0.13,-0.77,0.77,0.77,0.77}
- Bitstream: {11,01,00,11,11,11}

Bit Rate of a Digital Sequence

- Sampling rate: f_s sample/sec
- Quantization resolution: B bit/sample, B=[log2(L)]
- Bit rate: R=f_s B bit/sec
- Ex: speech signal sampled at 8 KHz, quantized to 8 bit/sample, R=8*8 = 64 Kbps
- Ex: music signal sampled at 44 KHz, quantized to 16 bit/sample, R=44*16=704 Kbps
- Ex: stereo music with each channel at 704 Kbps: R=2*704=1.4 Mbps
- Required bandwidth for transmitting a digital signal depends on the modulation technique.
  - To be covered later.
- Data rate of a multimedia signal can be reduced significantly through lossy compression w/o affecting the perceptual quality.
  - To be covered later.

Advantages of Digital Representation (I)

More immune to noise added in channel and/or storage

The receiver applies a threshold to the received signal:

$$\hat{x} = \begin{cases} 0 & \text{if} \quad x < 0.5 \\ 1 & \text{if} \quad x \geq 0.5 \end{cases}$$

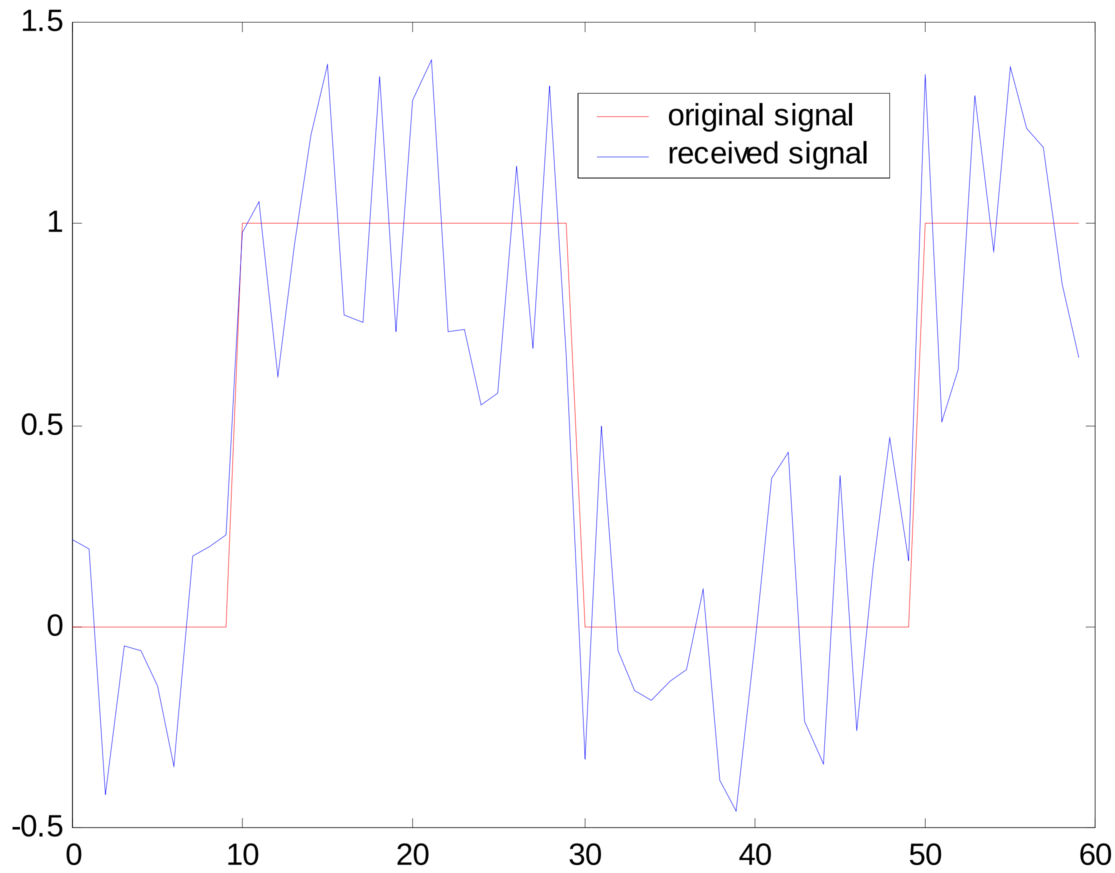

Advantages of Digital Representation (II)

- Can correct erroneous bits and/or recover missing bits using “forward error correction” (FEC) technique
  - By adding “parity bits” after information bits, corrupted bits can be detected and corrected
  - Ex: adding a “check-sum” to the end of a digital sequence (“0” if sum=even, “1” if sum=odd). By computing check-sum after receiving the signal, one can detect single errors (in fact, any odd number of bit errors).
  - Used in CDs, DVDs, Internet, wireless phones, etc.

What Should You Know

- Understand the general concept of quantization
- Can perform uniform quantization on a given signal
- Understand the principle of non-uniform quantization, and can perform mu-law quantization
- Can perform uniform and mu-law quantization on a given sequence, generate the resulting quantized sequence and its binary representation
- Can calculate bit rate given sampling rate and quantization levels
- Know advantages of digital representation
- Understand sample matlab codes for performing quantization (uniform and mu-law)

References

- Y. Wang, Lab Manual for Multimedia Lab, Experiment on Speech and Audio Compression. Sec. 1-2.1. (copies provided).