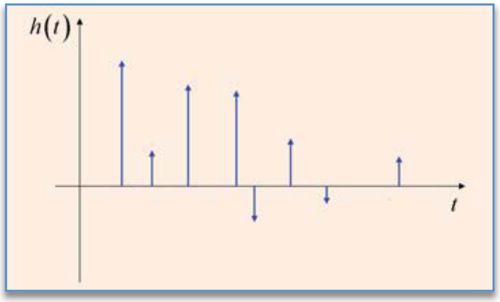

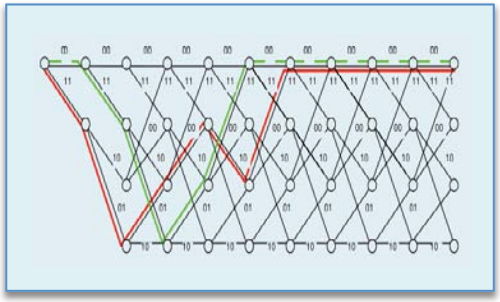

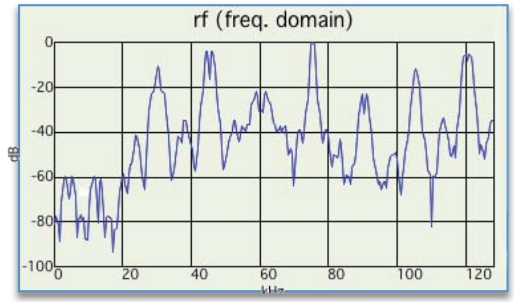

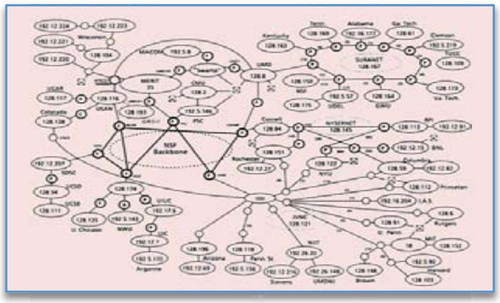

INTRODUCTION TO EECS II

DIGITAL COMMUNICATION SYSTEMS

6.02 Fall 2012
Lecture #8

- Noise: bad things happen to good signals!
- Signal-to-noise ratio and decibel (dB) scale
- PDF's, means, variances, Gaussian noise
- Bit error rate for bipolar signaling

Single Link Communication Model

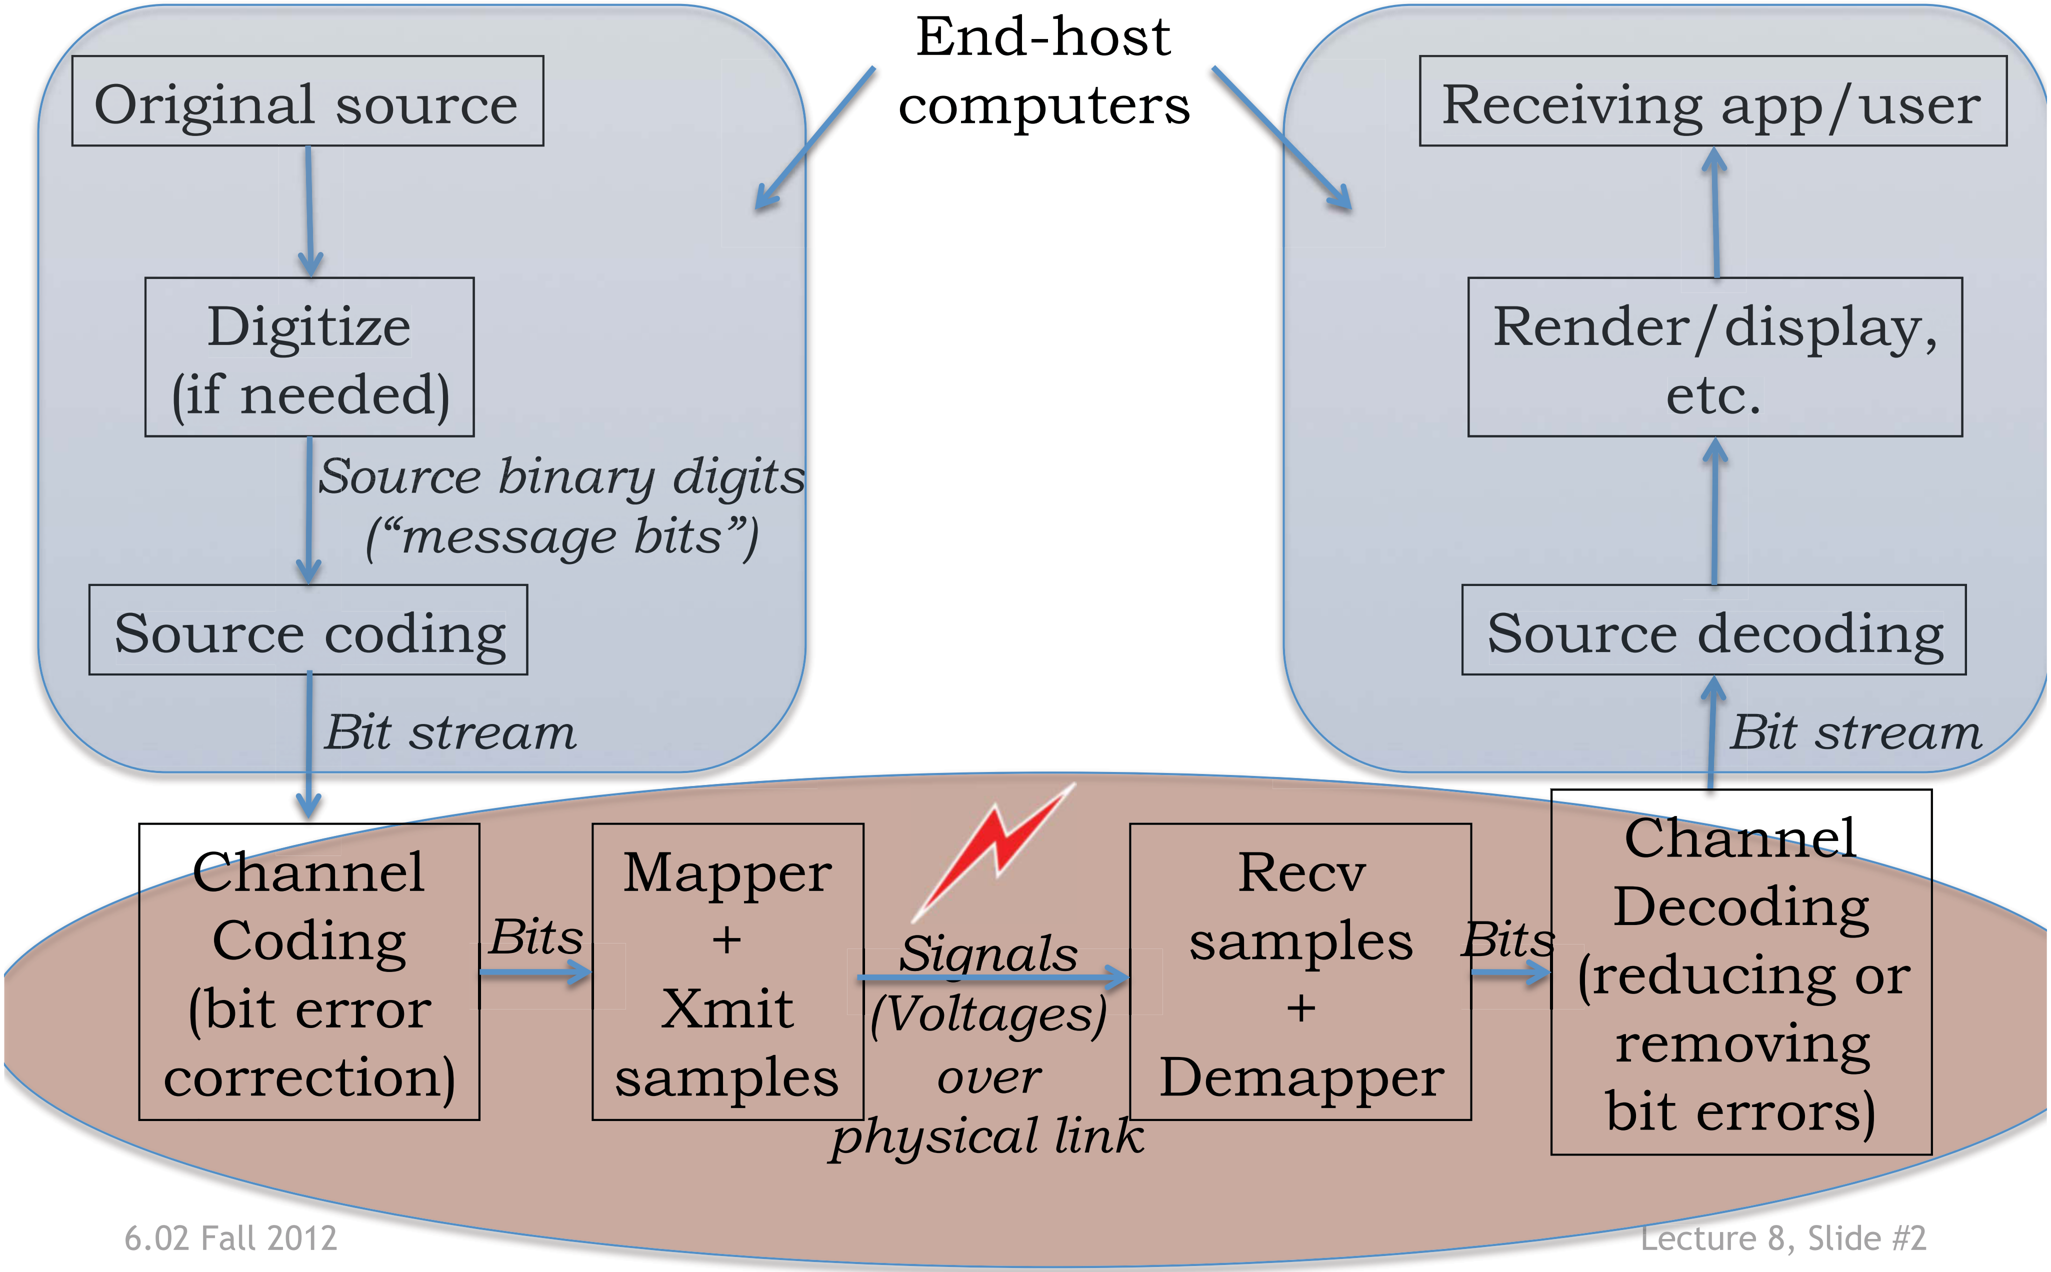

From Baseband to Modulated Signal, and Back

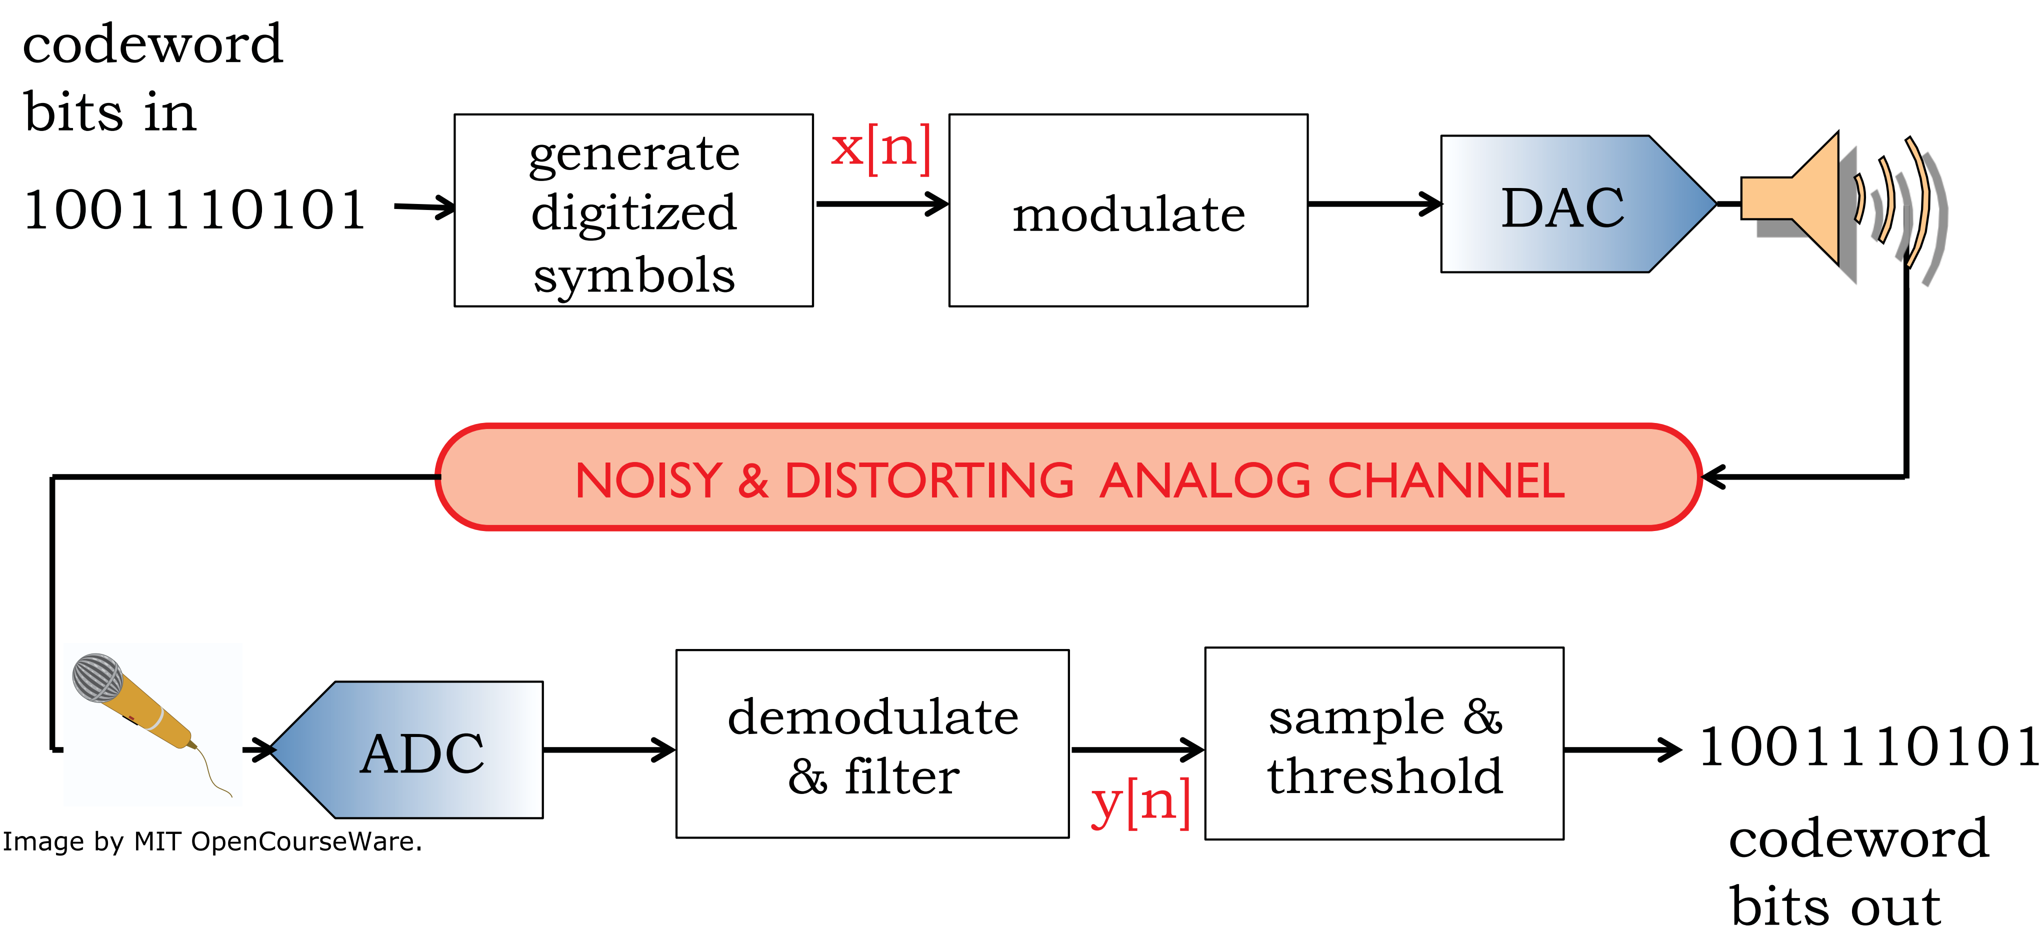

Image by MIT OpenCourseWare.

Mapping Bits to Samples at Transmitter

codeword
bits in

1001110101 → generate digitized symbols → x[n]

16 samples per bit

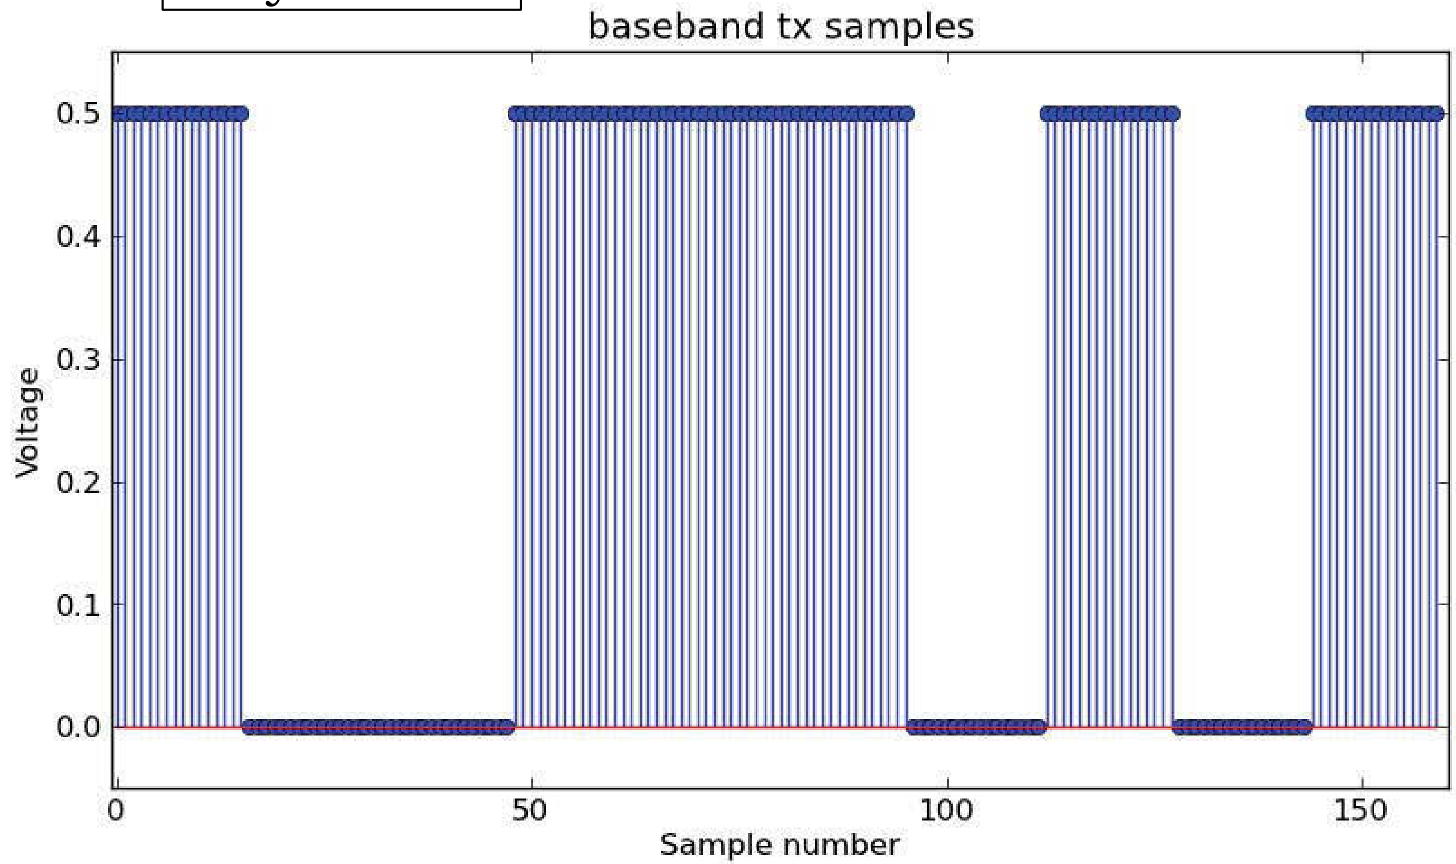

1 0 0 1 1 1 0 1 0 1

Samples after Processing at Receiver

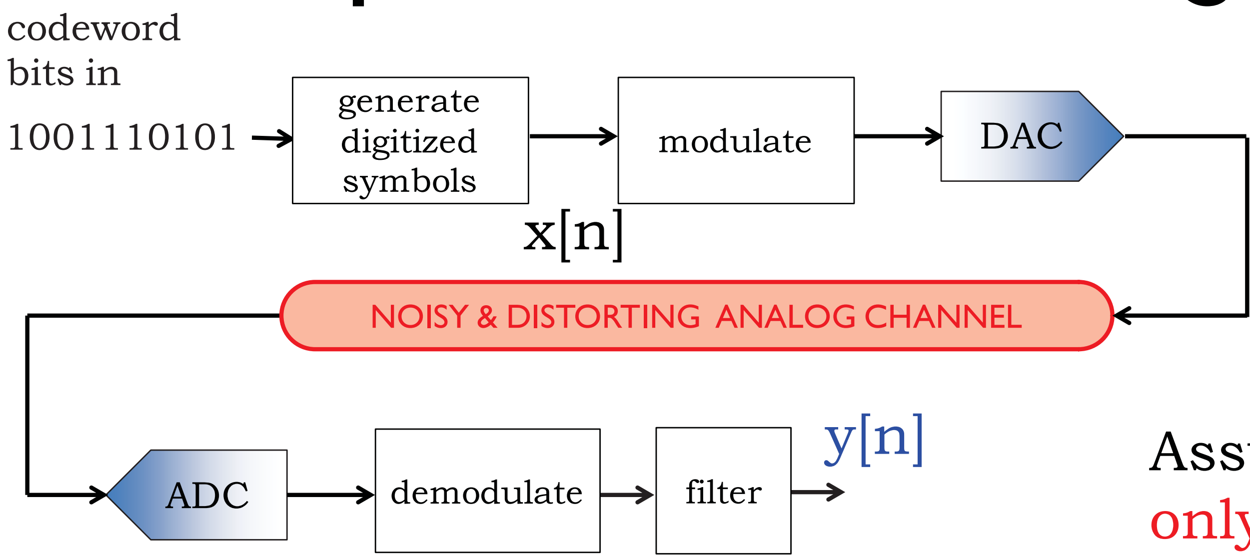

Assuming no noise,
only end-to-end distortion

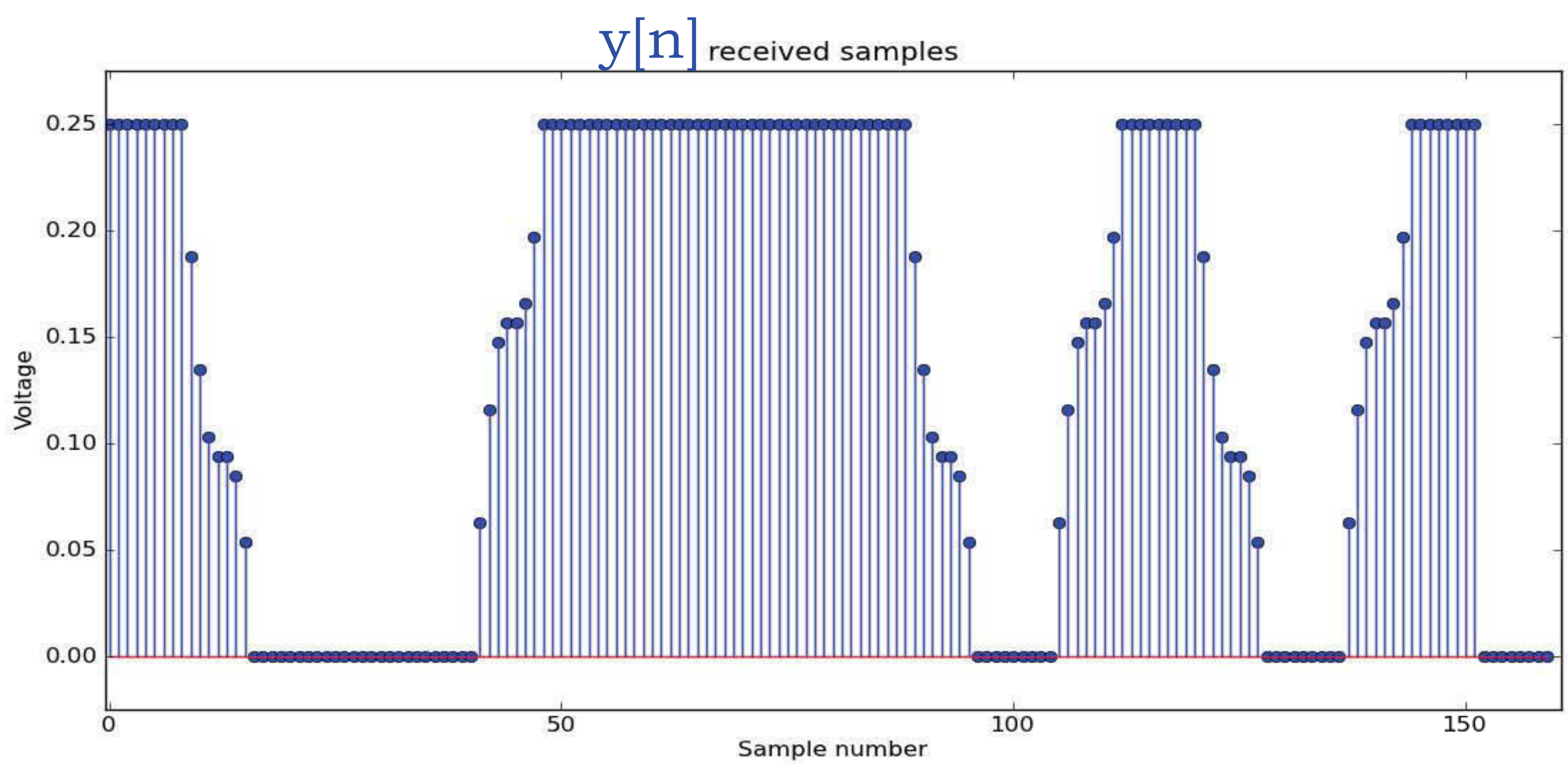

Mapping Samples to Bits at Receiver

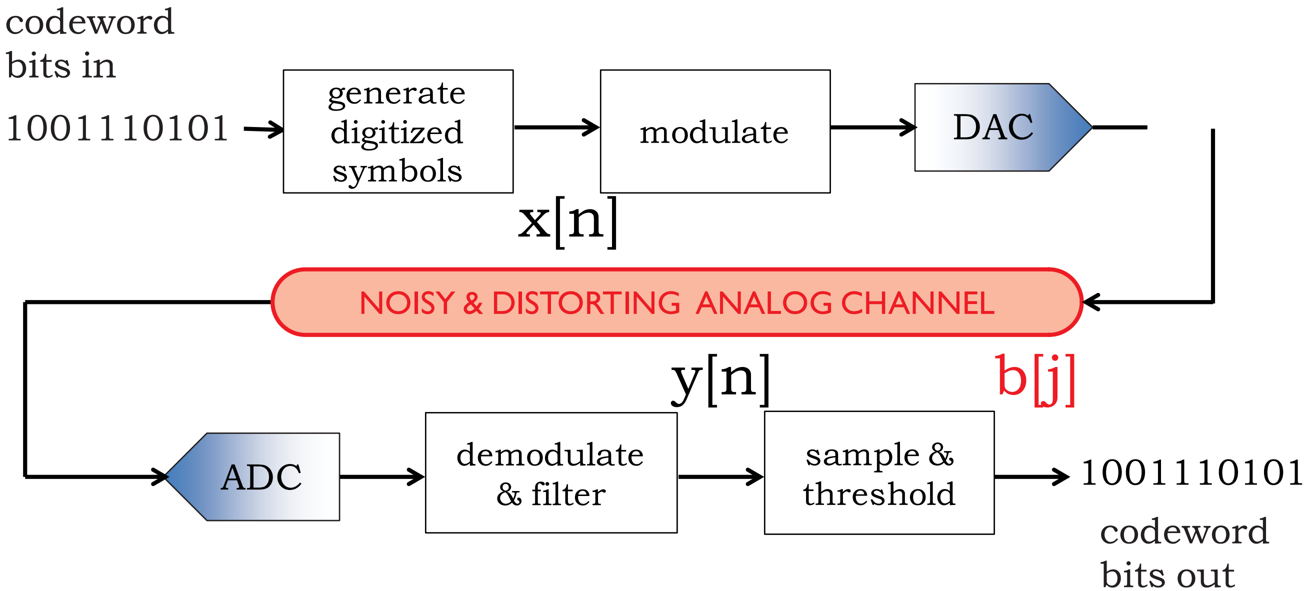

n = sample index

j = bit index

16 samples per bit

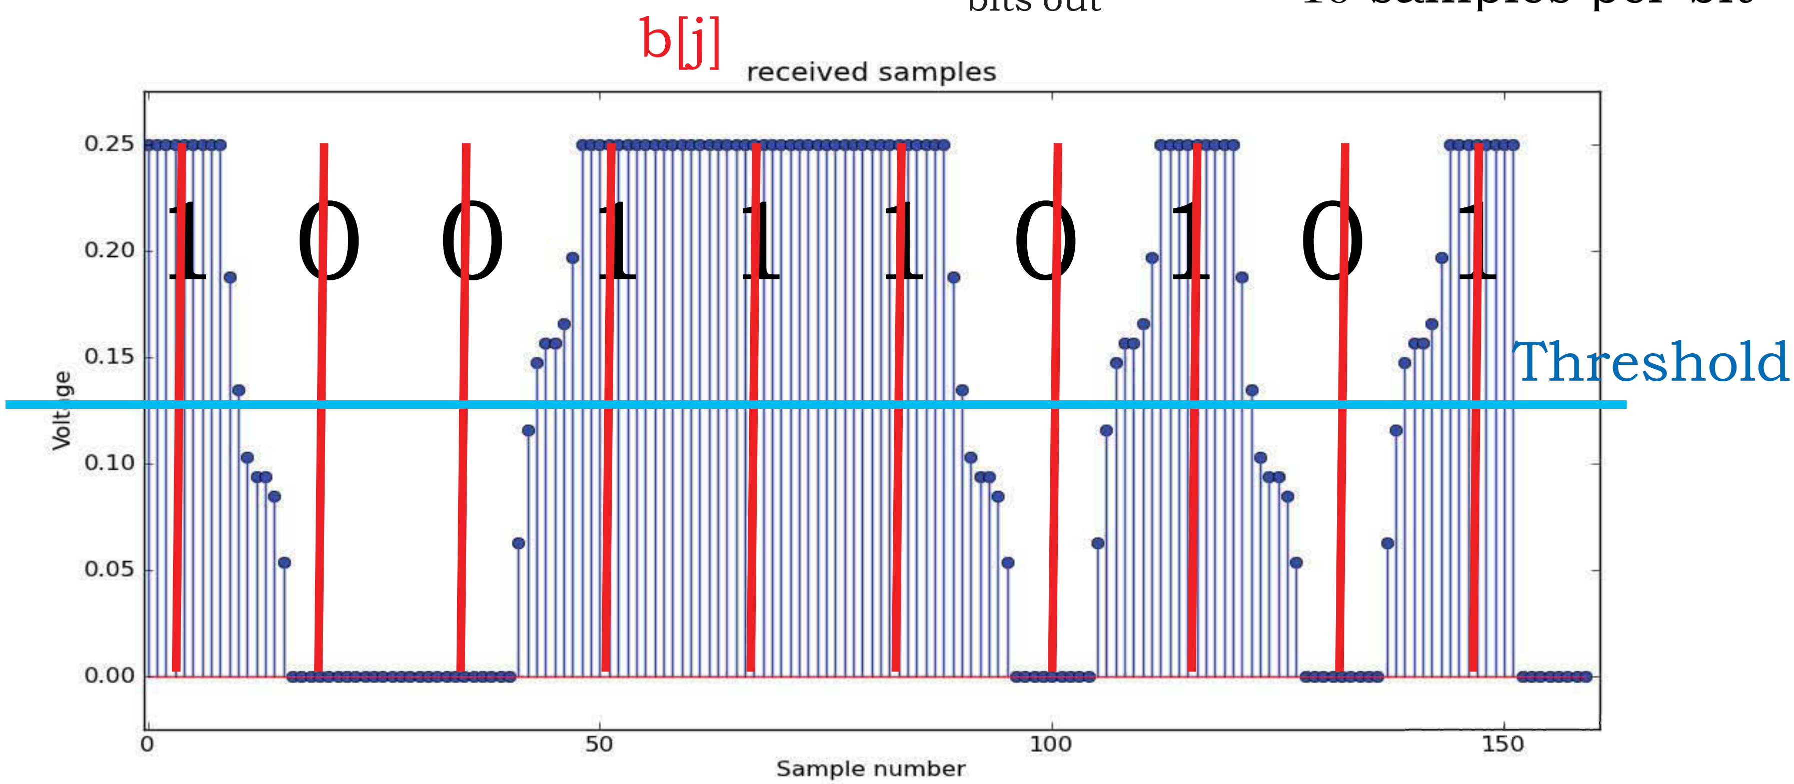

For now, assume no distortion, only Additive Zero-Mean Noise

- Received signal

$$y[n] = x[n] + w[n]$$

i.e., received samples y[n] are

the transmitted samples x[n] +

zero-mean noise w[n] on each sample, assumed iid (independent and identically distributed at each n)

- Signal-to-Noise Ratio (SNR)

– usually denotes the ratio of

(time-averaged or peak) signal power, i.e., *squared amplitude* of x[n]

to

noise variance, i.e., *expected squared amplitude* of w[n]

SNR Example

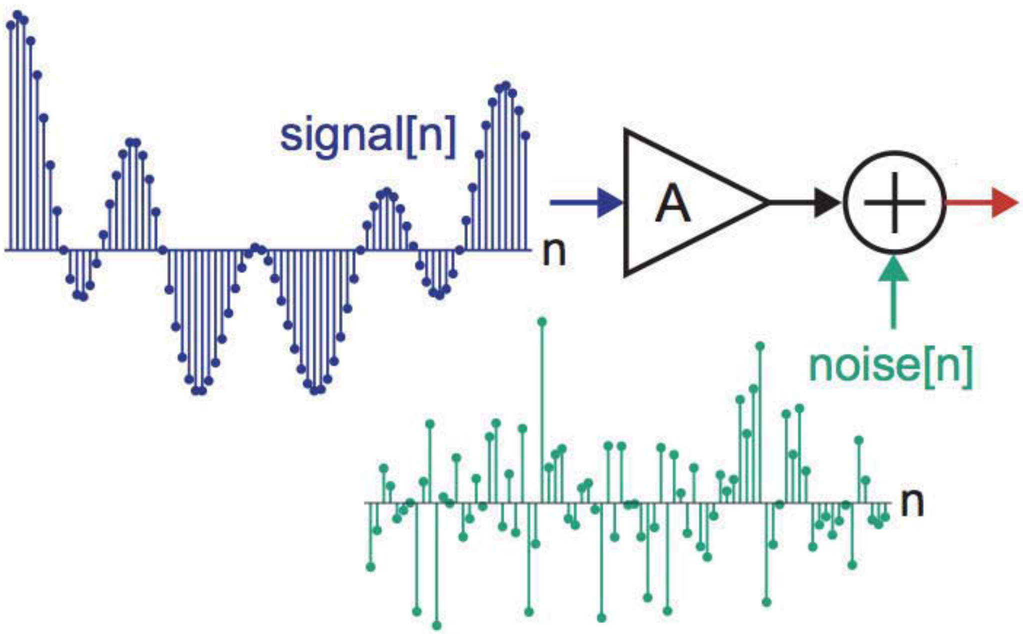

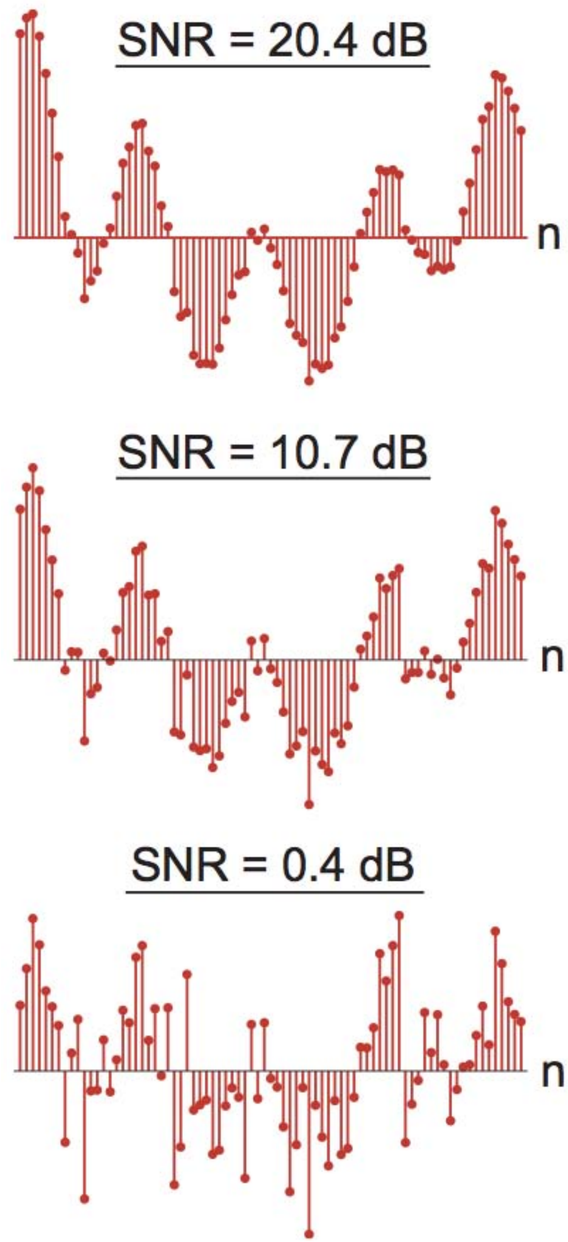

Changing the amplification factor (gain) A leads to different SNR values:

- Lower A → lower SNR
- Signal quality degrades with lower SNR

Signal-to-Noise Ratio (SNR)

The Signal-to-Noise ratio (SNR) is useful in judging the impact of noise on system performance:

$$\text{SNR} = \frac{\tilde{P}_{signal}}{\tilde{P}_{noise}}$$

SNR for power is often measured in decibels (dB):

$$\text{SNR (db)} = 10\log_{10}\left(\frac{\tilde{P}_{signal}}{\tilde{P}_{noise}}\right)$$

Caution: For measuring ratios of *amplitudes* rather than powers, take 20 $\log_{10}$ (ratio).

10logX	X
100	10000000000
90	1000000000
80	100000000
70	10000000
60	1000000
50	100000
40	10000
30	1000
20	100
10	10
0	1
-10	0.1
-20	0.01
-30	0.001
-40	0.0001
-50	0.000001
-60	0.0000001
-70	0.00000001
-80	0.000000001
-90	0.0000000001
-100	0.00000000001

3.01db is a factor of 2 in power ratio

Noise Characterization: From Histogram to PDF

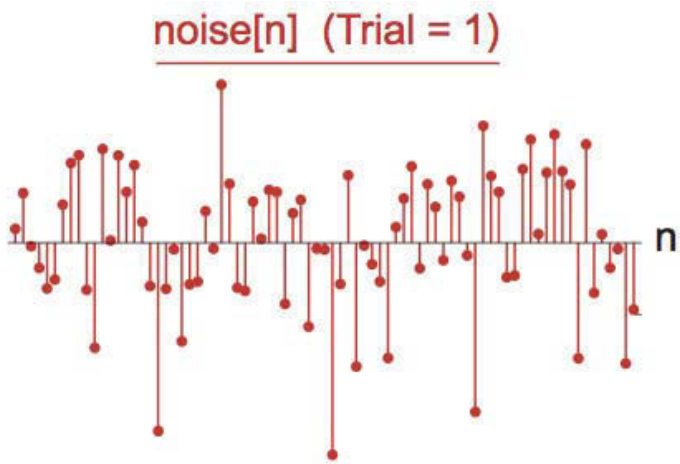

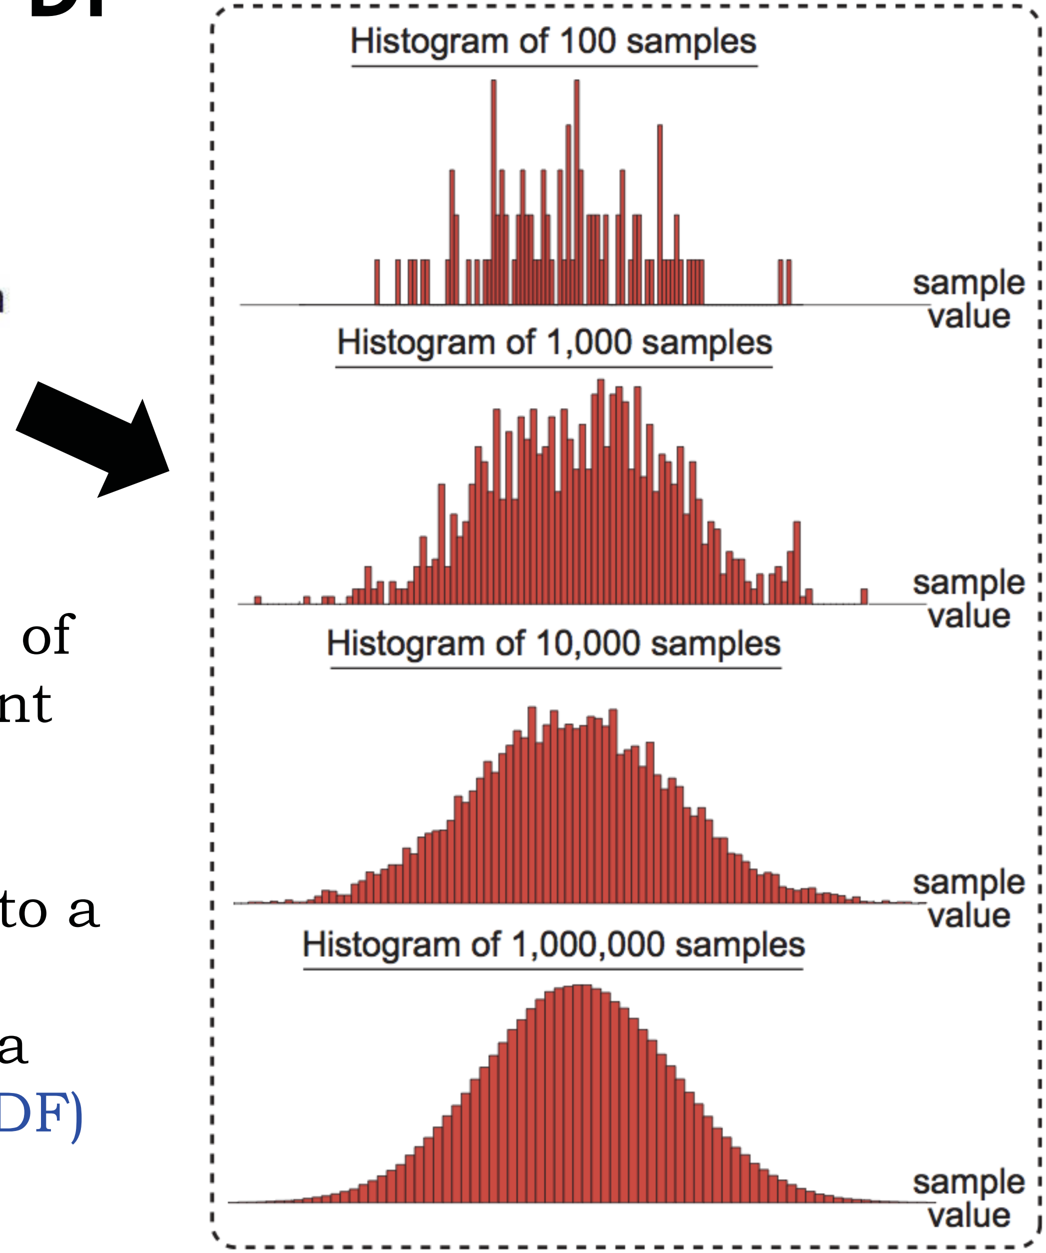

Experiment: create histograms of sample values from independent trials of increasing lengths.

Histogram typically converges to a shape that is known – after *normalization to unit area* – as a probability density function (PDF)

Using the PDF in Probability Calculations

We say that X is a random variable governed by the PDF $f_X(x)$ if X takes on a numerical value in the range of x_1 to x_2 with a probability calculated from the PDF of X as:

$$p(x_1 < X < x_2) = \int_{x_1}^{x_2} f_X(x)dx$$

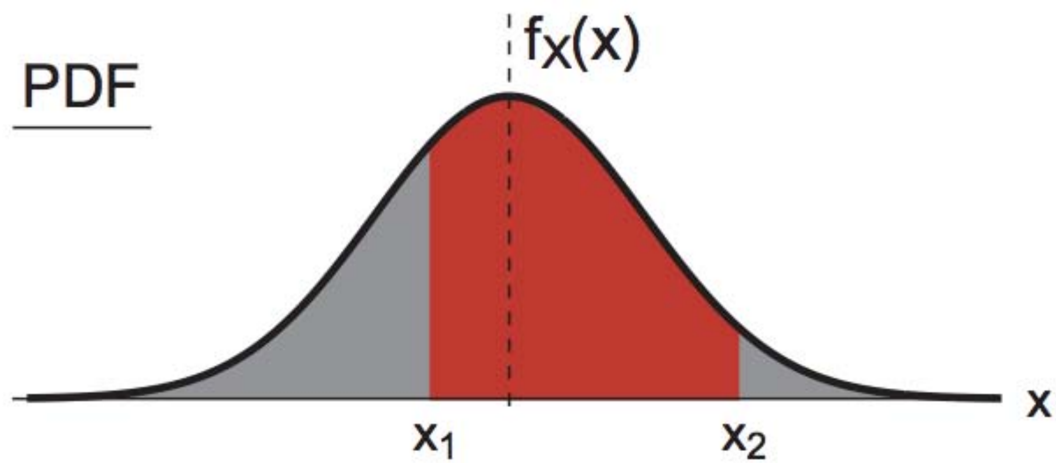

A PDF is **not** a probability – its associated *integrals* are.
Note that probability values are always in the range of 0 to 1.

The Ubiquity of Gaussian Noise

The net noise observed at the receiver is often the sum of many small, independent random contributions from many factors. If these independent random variables have finite mean and variance, the Central Limit Theorem says their sum will be a *Gaussian*.

The figure below shows the histograms of the results of 10,000 trials of summing 100 random samples drawn from [-1,1] using two different distributions.

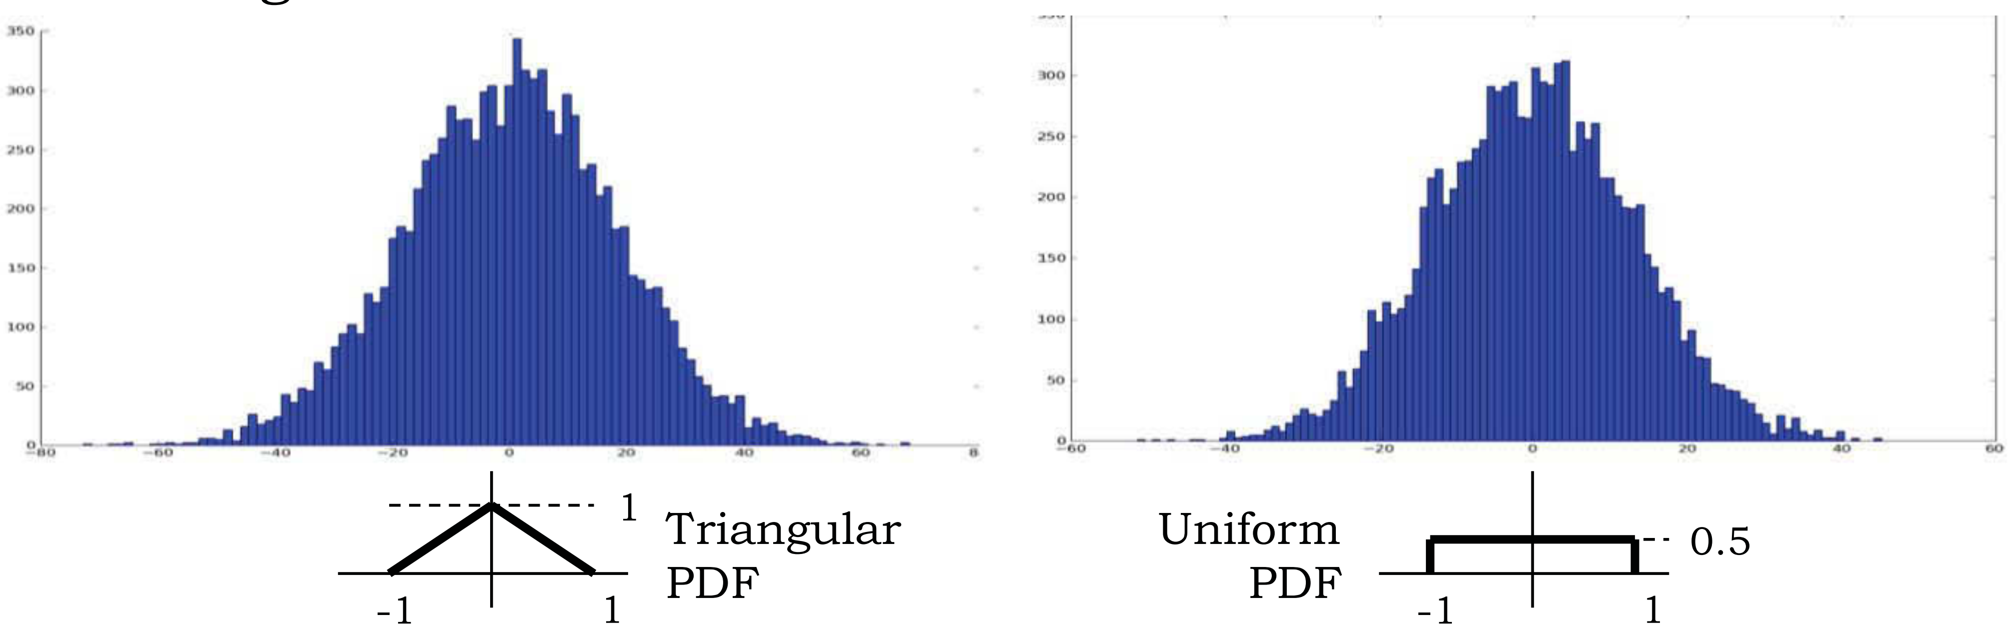

Mean and Variance of a Random Variable *X*

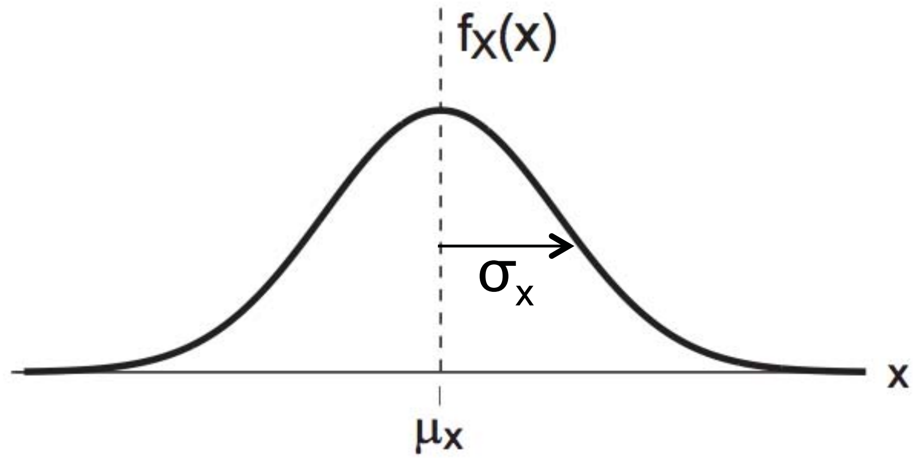

The *mean* or *expected value* μ_X is defined and computed as:

$$\mu_X = \int_{-\infty}^{\infty} x\, f_X(x)\,dx$$

The *variance* σ_X^2 is the expected squared variation or deviation of the random variable around the mean, and is thus computed as:

$$\sigma_X^2 = \int_{-\infty}^{\infty} (x-\mu_X)^2 f_X(x)\,dx$$

The square root of the variance is the *standard deviation*, σ_X

Visualizing Mean and Variance

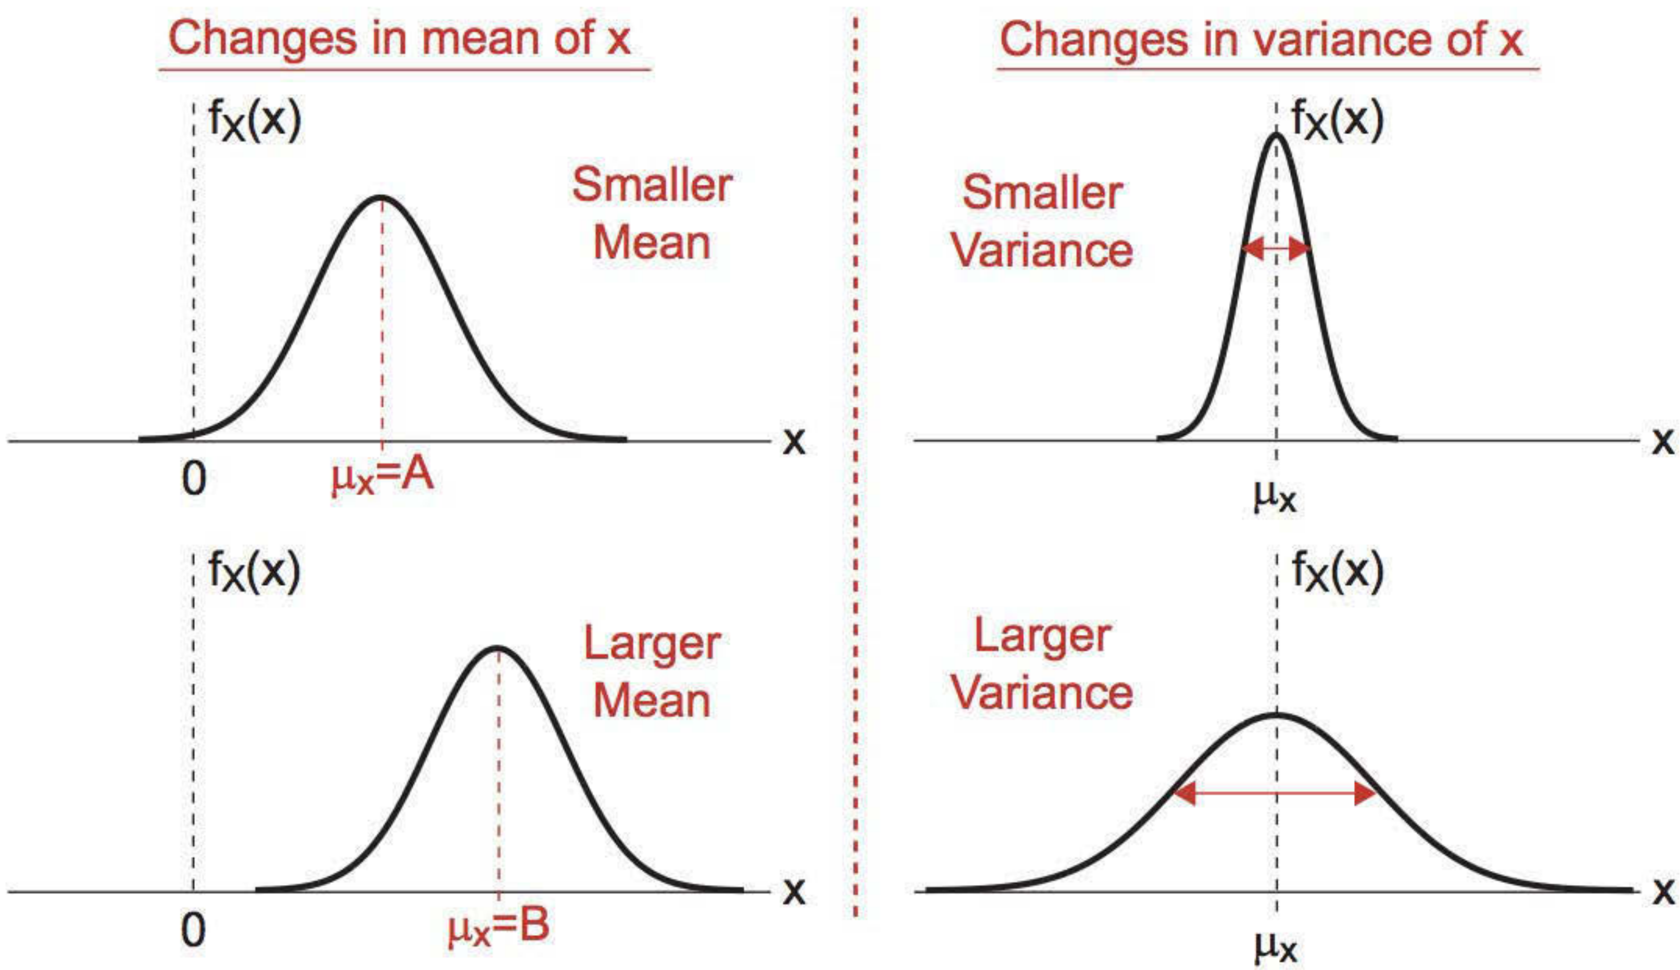

Changes in mean shift the center of mass of PDF

Changes in variance narrow or broaden the PDF (but area is always equal to 1)

The Gaussian Distribution

A Gaussian random variable W with mean μ and variance σ^2 has a PDF described by

$$f_W(w) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{\frac{-(w-\mu)^2}{2\sigma^2}}$$

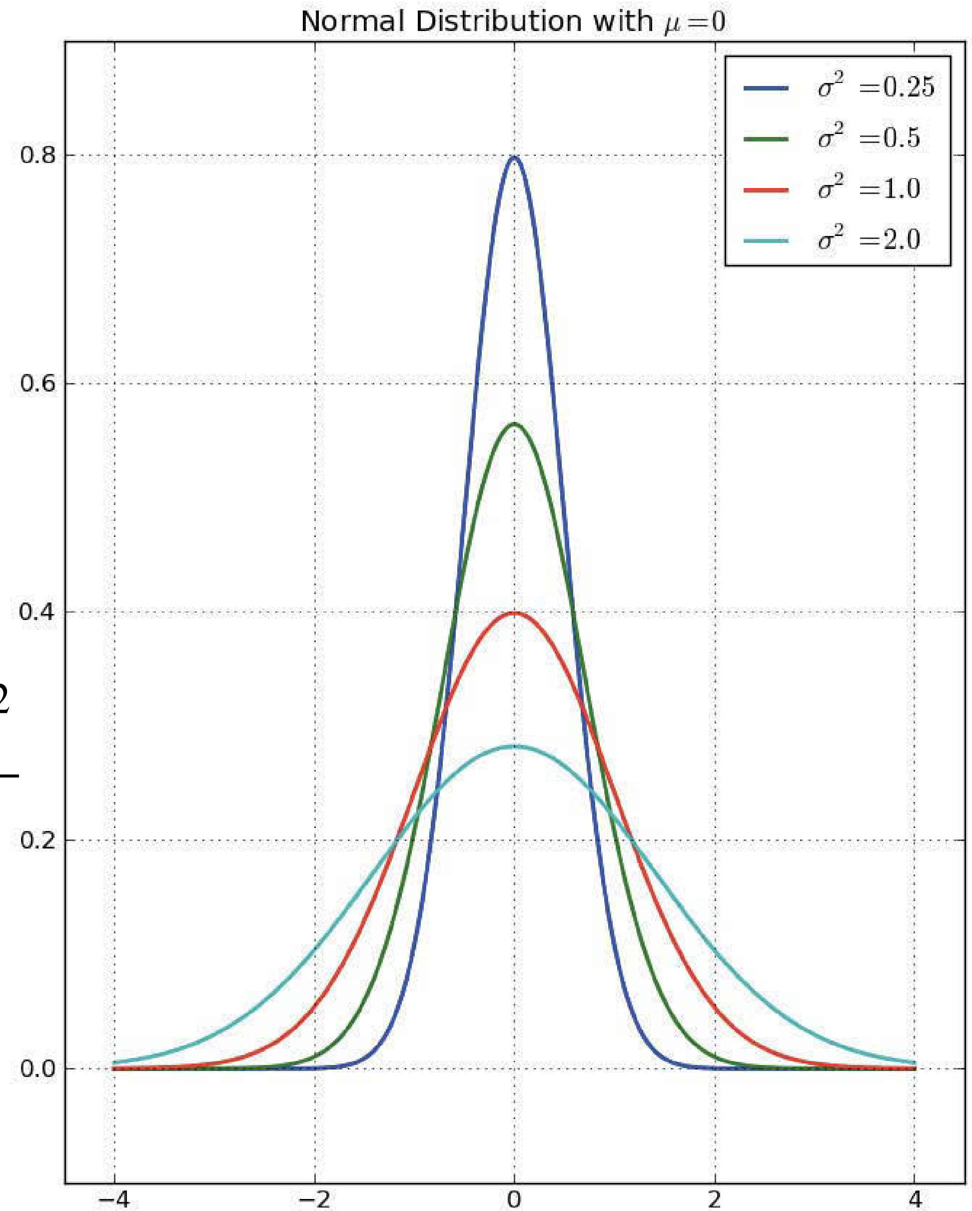

Noise Model for iid Process w[n]

- Assume each w[n] is distributed as the Gaussian random variable W on the preceding slide, but with mean 0, and independently of w[.] at all other times.

Estimating noise parameters

- Transmit a sequence of "0" bits, i.e., hold the voltage V_0 at the transmitter
- Observe received samples y[n], n = 0, 1, . . . , K − 1
  - Process these samples to obtain the statistics of the noise process for additive noise, under the assumption of iid (independent, identically distributed) noise samples (or, more generally, an "ergodic" process – beyond our scope!).
- Noise samples $w[n] = y[n] - V_0$
- For large K, can use the sample mean *m* to estimate μ, and sample standard deviation *s* to estimate σ:

$$m = \frac{1}{K}\sum_{k=0}^{K-1} w[k] \qquad s^2 = \frac{1}{K}\sum_{k=0}^{K-1}(w[k]-m)^2$$

Back to distinguishing "1" from "0":

- Assume bipolar signaling:

  Transmit L samples x[.] at $+V_p$ ($=V_1$) to signal a "1"

  Transmit L samples x[.] at $-V_p$ ($=V_0$) to signal a "0"

- Simple-minded receiver: take a single sample value $y[n_j]$ at an appropriately chosen instant n_j in the
- j-th bit interval. Decide between the following two hypotheses:

$$y[n_j] = +V_p + w[n_j] \quad (==> \text{"1"})$$

or

$$y[n_j] = -V_p + w[n_j] \quad (==> \text{"0"})$$

where $w[n_j]$ is Gaussian, zero-mean, variance σ^2

Connecting the SNR and BER

$$V_p = \sqrt{E_S}$$

$$2\sigma^2 = N_0$$

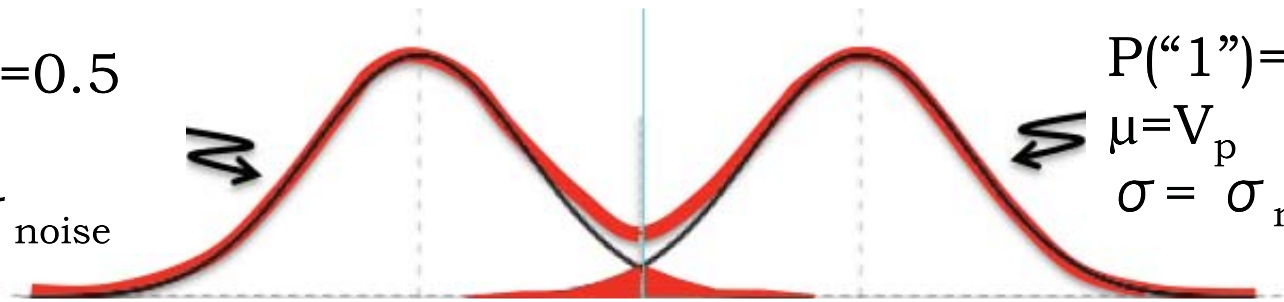

$$\text{BER} = \mathbb{P}(\text{error}) = \frac{1}{\sqrt{2\pi\sigma^2}} \int_{\sqrt{E_s}}^{\infty} e^{-w^2/(2\sigma^2)}\, dw$$

$$\text{erf}(z) = \frac{2}{\sqrt{\pi}} \cdot \int_{-\infty}^{z} e^{-v^2}\, dv\ ,$$

$$\text{BER} = \mathbb{P}(\text{error}) = \frac{1}{\sqrt{\pi}} \cdot \int_{\sqrt{E_s/N_0}}^{\infty} e^{-v^2}\, dv$$

$$\text{erfc}(z) = 1 - \text{erf}(z) = \frac{2}{\sqrt{\pi}} \cdot \int_{z}^{\infty} e^{-v^2}\, dv$$

$$BER = P(error) = \frac{1}{2} erfc\left(\sqrt{\frac{E_S}{N_0}}\right) = \frac{1}{2} erfc\left(\frac{V_p}{\sigma\sqrt{2}}\right)$$

Bit Error Rate for Bipolar Signaling Scheme with Single-Sample Decision

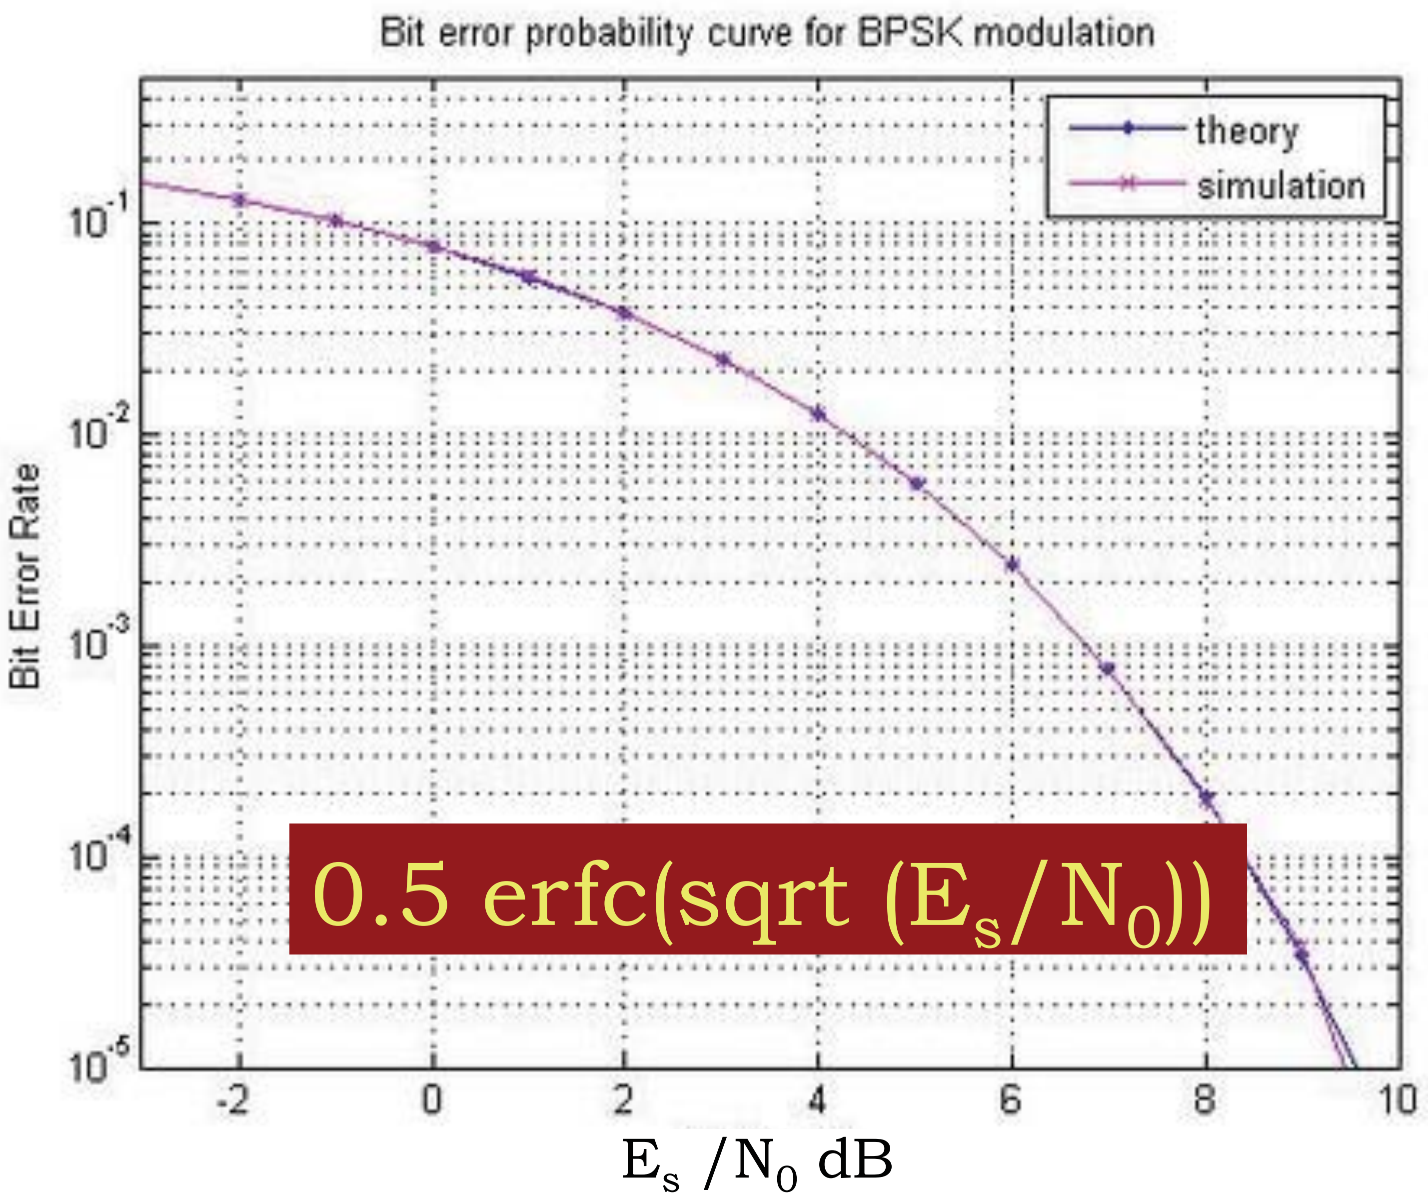

Source: http://www.dsplog.com/2007/08/05/bit-error-probability-for-bpsk-modulation/.
Courtesy of Krishna Sankar Madhavan Pillai. Used with permission.

But we can do better!

- Why just take a single sample from a bit interval?
- Instead, average M (≤L) samples:

$y[n] = +V_p + w[n]$ so $avg\{y[n]\} = +V_p + avg\{w[n]\}$

- $avg\{w[n]\}$ is still Gaussian, still has mean 0, but its variance is now σ^2/M instead of σ^2 ➔ SNR is increased by a factor of M

- Same analysis as before, but now bit energy E_b = M.E_s instead of sample energy E_s:

$$BER = P(error) = \frac{1}{2}erfc(\sqrt{\frac{E_b}{N_0}}) = \frac{1}{2}erfc(\frac{V_p\sqrt{M}}{\sigma\sqrt{2}})$$

Implications for Signaling Rate

- As the noise intensity increases, we need to slow down the signaling rate, i.e., increase the number of samples per bit (K), to get higher energy in the (M≤K) samples extracted from a bit interval, if we wish to maintain the same error performance.
  - e.g. Voyager 2 was transmitting at 115 **kbits**/s when it was near Jupiter in 1979. Last month it was over 9 billion miles away, 13 light hours away from the sun, twice as far away from the sun as Pluto. And now transmitting at only 160 **bits**/s. The received power at the Deep Space Network antennas on earth when Voyager was near Neptune was on the order of 10^(-16) watts!! --- 20 billion times smaller than an ordinary digital watch consumes. The power now is estimated at 10^(-19) watts.

Flipped bits can have serious consequences!

- "On **November 30, 2006**, a telemetered command to *Voyager 2* was incorrectly decoded by its on-board computer—in a random error—as a command to turn on the electrical heaters of the spacecraft's magnetometer. These heaters remained turned on until December 4, 2006, and during that time, there was a resulting high temperature above 130 °C (266 °F), significantly higher than the magnetometers were designed to endure, and a sensor rotated away from the correct orientation. It has not been possible to fully diagnose and correct for the damage caused to the *Voyager 2's* magnetometer, although efforts to do so are proceeding."

- "On **April 22, 2010**, *Voyager 2* encountered scientific data format problems as reported by the Associated Press on May 6, 2010. On **May 17, 2010**, JPL engineers revealed that a flipped bit in an on-board computer had caused the issue, and scheduled a bit reset for May 19. On **May 23, 2010**, *Voyager 2* has resumed sending science data from deep space after engineers fixed the flipped bit."

http://en.wikipedia.org/wiki/Voyager_2

What if the received signal is distorted?

- Suppose $y[n] = \pm x_0[n] + w[n]$ in a given bit slot (L samples), where $x_0[n]$ is known, and $w[n]$ is still iid Gaussian, zero mean, variance σ^2.
- Compute a weighted linear combination of the y[.] in that bit slot:

$$\sum a_n y[n] = \pm \sum a_n x_0[n] + \sum a_n w[n]$$

This is still Gaussian, mean $\pm \sum a_n x_0[n]$, but now the variance is $\sigma^2 \sum (a_n)^2$

- So what choice of the $\{a_n\}$ will maximize SNR? Simple answer: $a_n = x_0[n]$ ➔ **"matched filtering"**
- Resulting SNR for receiver detection and error performance is $\sum (x_0[n])^2 / \sigma^2$, i.e., again the ratio of bit energy E_b to noise power.

The moral of the story is …

… if you're doing appropriate/optimal processing at the receiver, your signal-to-noise ratio (and therefore your error performance) in the case of iid Gaussian noise is the ratio of bit energy (not sample energy) to noise variance.