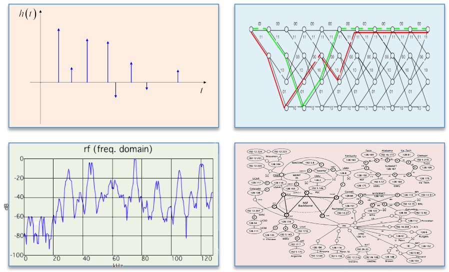

INTRODUCTION TO EECS II

DIGITAL COMMUNICATION SYSTEMS

6.02 Spring 2011
Lecture #6

- Mean, power, energy, SNR
- Metrics for random processes
- Normal PDF, CDF
- Calculating p(error), BER vs. SNR

Bad Things Happen to Good Signals

Noise, broadly construed, is any change to the signal from its expected value, x[n]*h[n], when it arrives at the receiver.

We'll look at *additive noise* and assume the noise in our systems is independent in value and timing from the nominal signal, $y_{nf}[n]$, and that the noise can be described by a random variable with a known probability distribution.

We'll model the received signal as $y_{nf}[n]$ + noise[n].

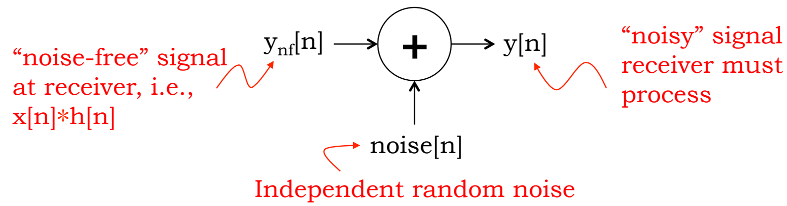

Definition of Mean, Power, Energy

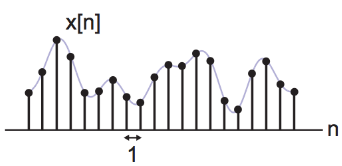

Slides 3-16 derived from 6.02 slides by Mike Perrott

Some interesting statistical metrics for x[n]:

Mean: $\mu_x = \frac{1}{N}\sum_{n=1}^{N} x[n]$

Power: $P_x = \frac{1}{N}\sum_{n=1}^{N} x[n]^2$ $\quad \tilde{P}_x = \frac{1}{N}\sum_{n=1}^{N}\left(x[n]-\mu_x\right)^2$

Energy: $E_x = \sum_{n=1}^{N} x[n]^2$ $\quad \tilde{E}_x = \sum_{n=1}^{N}\left(x[n]-\mu_x\right)^2$

In analyzing our systems, we often use metrics where the mean has been factored out.

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 is often measured in decibels (dB):

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

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

3db is a factor of 2

SNR Example

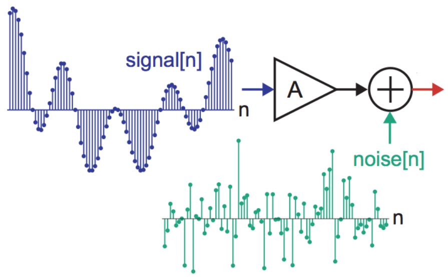

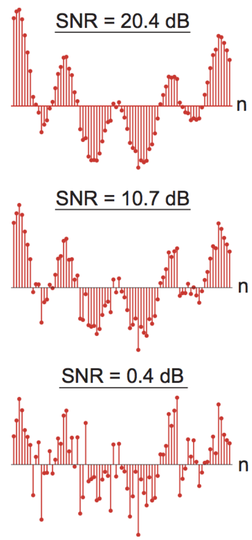

Changing the amplification factor (gain) A leads to different SNR values:

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

Analysis of Random Processes

- Random processes, such as noise, take on different sequences for different trials
  - Think of trials as different measurement intervals from the same experimental setup (as in lab)
- For a *given* trial, we can apply our standard analysis tools and metrics
  - mean and power calculations, etc…
- When trying to analyze the *ensemble* (i.e., all trials) of possible outcomes, we find ourselves in need of new tools and metrics

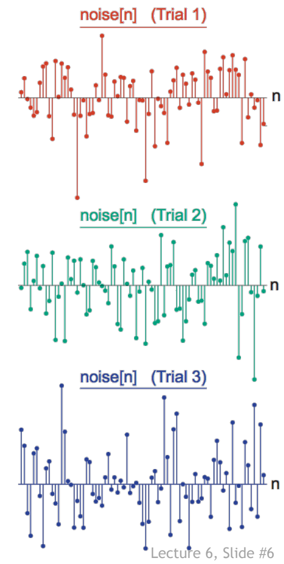

Stationary and Ergodic Random Processes

<u>Stationary</u>

statistical behavior is independent of shifts in time in a given trial. Implies noise[k] is statistically indistinguishable from noise[k+N]

<u>Ergodic</u>

statistical sampling can be performed at one sample time (i.e., n=k) across different trials, or across different sample times of the same trial with no change in the measured result

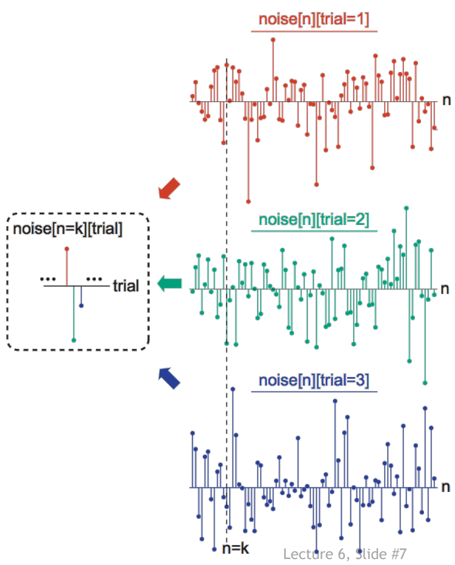

Experiment to See Statistical Distribution

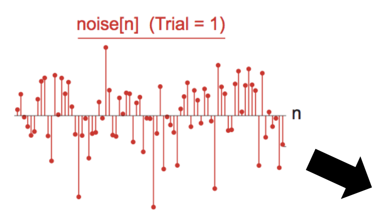

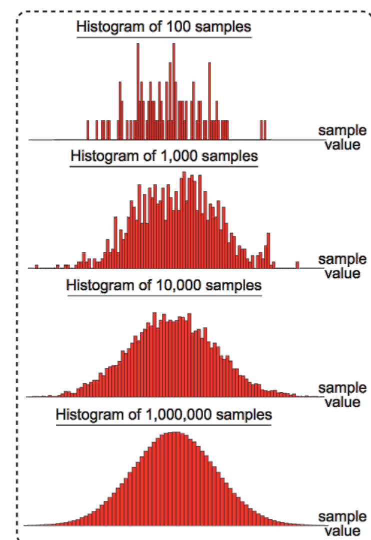

Experiment: create histograms of sample values from trials of increasing lengths.

Assumption of stationarity implies histogram should converge to a shape known as a probability density function (PDF)

Formalizing the PDF Concept

Define x as a random variable whose PDF has the same shape as the histogram we just obtained.

Denote the PDF of x as $f_x(x)$ and scale $f_x(x)$ such that its overall area is 1:

$$\int_{-\infty}^{\infty} f_x(x) = 1$$

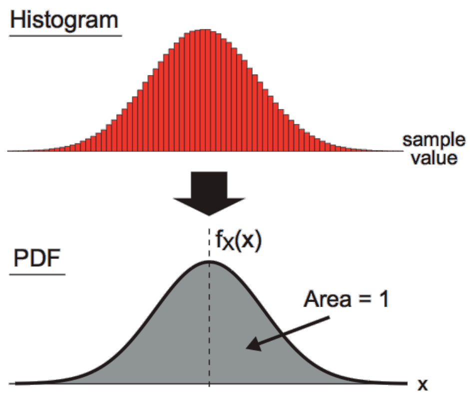

Formalizing Probability

The probability that random variable x takes on a value in the range of x_1 to x_2 is calculated from the PDF of x as:

$$p(x_1 \le x \le x_2) = \int_{x_1}^{x_2} f_x(x)dx$$

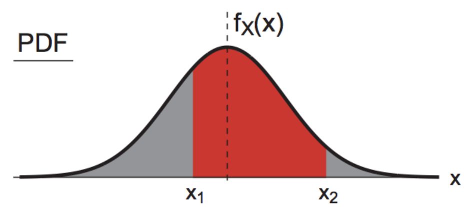

Note that probability values are always in the range of 0 to 1.

Example Probability Calculation

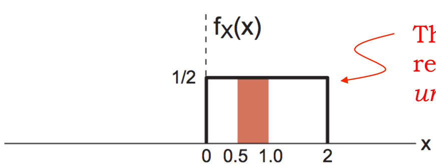

This shape is referred to as a *uniform* PDF.

Verify that overall area is 1:

$$\int_{-\infty}^{\infty} f_x(x)dx = \int_{0}^{2} 0.5\,dx = 1$$

Probability that x takes on a value between 0.5 and 1:

$$p(0.5 \le x \le 1.0) = \int_{0.5}^{1} 0.5\,dx = 0.25$$

Examination of Sample Value Distribution

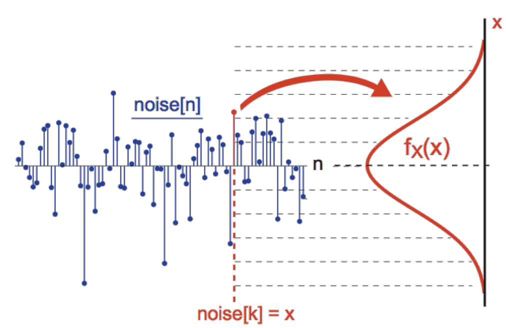

Assumption of ergodicity implies the value occurring at a given time sample, noise[k], across many different trials has the same PDF as estimated in our previous experiment of many time samples and one trial.

Thus we can model noise[k] using the random variable x.

Probability Calculation

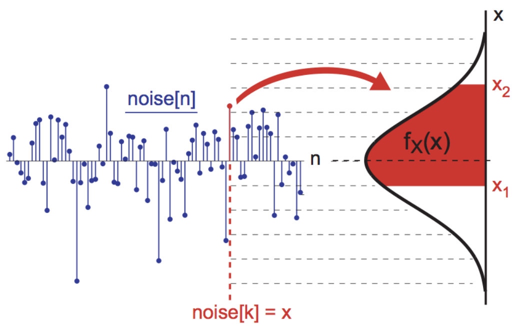

In a given trial, the probability that noise[k] takes on a value in the range of x_1 to x_2 is computed as

$$p(x_1 \le x \le x_2) = \int_{x_1}^{x_2} f_x(x)\,dx$$

Mean and Variance

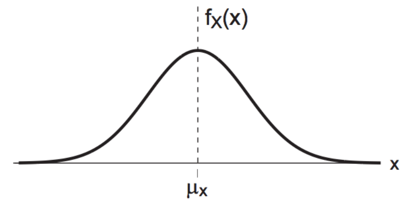

The *mean* of a random variable x, μ_x, corresponds to its average value and computed as:

$$\mu_x = \int_{-\infty}^{\infty} x\, f_x(x)\,dx$$

The *variance* of a random variable x, σ_x^2, gives an indication of its variability and is computed as:

$$\sigma_x^2 = \int_{-\infty}^{\infty} (x-\mu_x)^2 f_x(x)\,dx$$

Compare with power calculation

Visualizing Mean and Variance

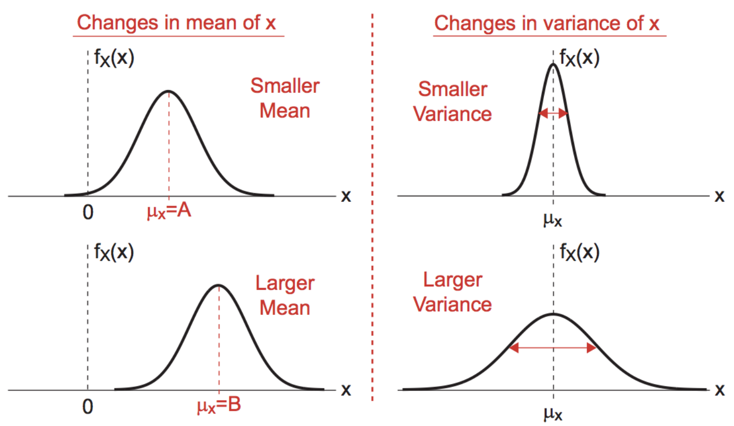

Changes in mean shift the center of mass of PDF

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

Example Mean and Variance Calculation

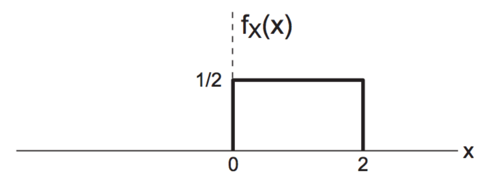

Mean:

$$\mu_x = \int_{-\infty}^{\infty} x\, f_x(x)\,dx = \int_0^2 x\frac{1}{2}dx = \frac{1}{4}x^2\Big|_0^2 = 1$$

Variance:

$$\sigma_x^2 = \int_{-\infty}^{\infty} (x-\mu_x)^2 f_x(x)\,dx = \int_0^2 (x-1)^2\frac{1}{2}dx = \frac{1}{6}(x-1)^3\Big|_0^2 = \frac{1}{3}$$

Noise on a Communication Channel

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

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

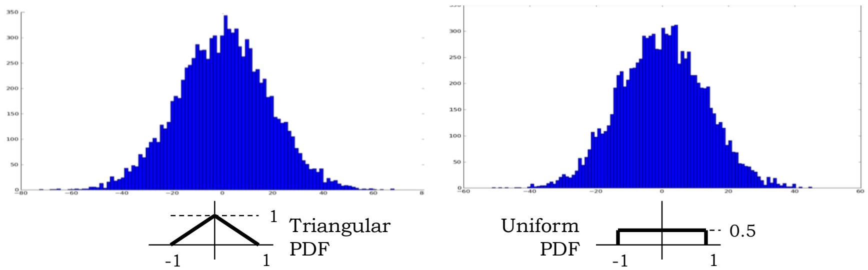

The Normal Distribution

A normal or Gaussian distribution with mean μ and variance σ^2 has a PDF described by

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

The normal distribution with $\mu=0$ and $\sigma^2=1$ is called the "standard" or "unit" normal.

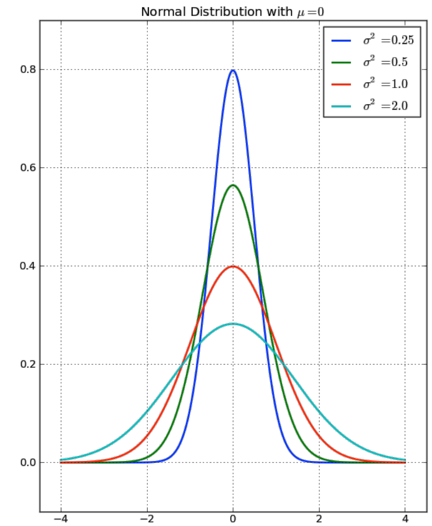

Cumulative Distribution Function

When analyzing the effects of Gaussian noise, we'll often want to determine the probability that the noise is larger or smaller than a given value x_0. From slide #10:

$$p(x \le x_0) = \int_{-\infty}^{x_0} \frac{1}{\sqrt{2\pi\sigma^2}} e^{\frac{-(x-\mu)^2}{2\sigma^2}} dx \equiv \Phi_{\mu,\sigma}(x_0)$$

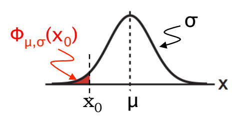

$$p(x \ge x_0) = \int_{x_0}^{\infty} \frac{1}{\sqrt{2\pi\sigma^2}} e^{\frac{-(x-\mu)^2}{2\sigma^2}} dx = 1 - \Phi_{\mu,\sigma}(x_0)$$

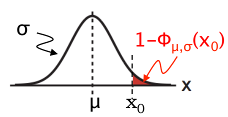

Where $\Phi_{\mu,\sigma}(x)$ is the cumulative distribution function (CDF) for the normal distribution with mean μ and variance σ^2. The CDF for the unit normal is usually written as just $\Phi(x)$.

$$\boxed{\Phi_{\mu,\sigma}(x) = \Phi\left(\frac{x-\mu}{\sigma}\right)}$$

Φ(x) = CDF for Unit Normal PDF

Most math libraries don't provide $\Phi(x)$ but they do have a related function, erf(x), the *error function*:

$$\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$$

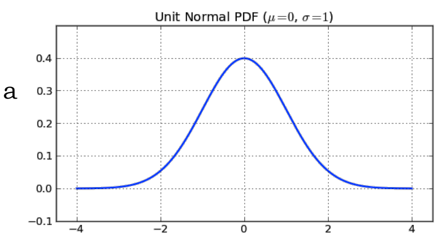

For Python hackers:

```
from math import sqrt
from scipy.special import erf

# CDF for Normal PDF
def Phi(x,mu=0,sigma=1):
  t = erf((x-mu)/(sigma*sqrt(2)))
  return 0.5 + 0.5*t
```

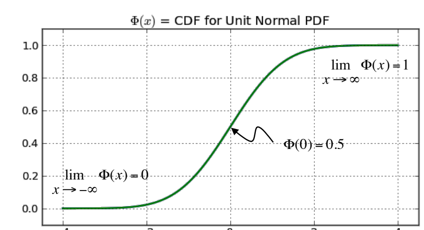

Bit Error Rate

The *bit error rate* (BER), or perhaps more appropriately the *bit error ratio*, is the number of bits received in error divided by the total number of bits transferred. We can estimate the BER by calculating the probability that a bit will be incorrectly received due to noise.

Using our normal signaling strategy (0V for "0", 1V for "1"), on a noise-free channel with no ISI, the samples at the receiver are either 0V or 1V. Assuming that 0's and 1's are equally probable in the transmit stream, the number of 0V samples is approximately the same as the number of 1V samples. So the mean and power of the noise-free received signal are

$$\mu_{y_{nf}} = \frac{1}{N}\sum_{n=1}^{N} y_{nf}[n] = \frac{1}{N}\frac{N}{2} = \frac{1}{2}$$

$$\tilde{P}_{y_{nf}} = \frac{1}{N}\sum_{n=1}^{N}\left(y_{nf}[n]-\frac{1}{2}\right)^2 = \frac{1}{N}\sum_{n=1}^{N}\left(\frac{1}{2}\right)^2 = \frac{1}{N}\frac{N}{4} = \frac{1}{4}$$

p(bit error)

Now assume the channel has Gaussian noise with $\mu=0$ and variance σ^2. And we'll assume a digitization threshold of 0.5V. We can calculate the probability that noise[k] is large enough that $y[k] = y_{nf}[k] + \text{noise}[k]$ is received incorrectly:

p(error | transmitted "0"):

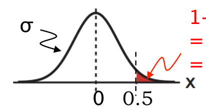

$1-\Phi_{\mu,\sigma}(0.5) = \Phi_{\mu,\sigma}(-0.5)$
$= \Phi((-0.5-0)/\sigma)$
$= \Phi(-0.5/\sigma)$

Plots of noise-free voltage + Gaussian noise

p(error | transmitted "1"):

$\Phi_{\mu,\sigma}(0.5)$
$= \Phi((0.5-1)/\sigma)$
$= \Phi(-0.5/\sigma)$

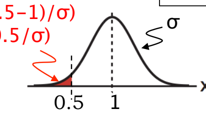

p(bit error) = p(transmit "0")*p(error | transmitted "0") + p(transmit "1")*p(error | transmitted "1")

$= 0.5*\Phi(-0.5/\sigma) + 0.5*\Phi(-0.5/\sigma)$

$= \Phi(-0.5/\sigma)$

BER (no ISI) vs. SNR

We calculated the power of the noise-free signal to be 0.25 and the power of the Gaussian noise is its variance, so

$$\text{SNR (db)} = 10\log\left(\frac{\tilde{P}_{signal}}{\tilde{P}_{noise}}\right) = 10\log\left(\frac{0.25}{\sigma^2}\right)$$

Given an SNR, we can use the formula above to compute σ^2 and then plug that into the formula on the previous slide to compute p(bit error) = BER.

The BER result is plotted to the right for various SNR values.