Sampling and the Discrete Fourier Transform

Sampling Methods

- Sampling is most commonly done with two devices, the *sample-and-hold* (S/H) and the *analog-to-digital-converter* (ADC)
- The S/H acquires a CT signal at a point in time and holds it for later use
- The ADC converts CT signal values at discrete points in time into numerical codes which can be stored in a digital system

Sampling Methods

Sample-and-Hold

During the clock, c(*t*), aperture time, the response of the S/H is the same as its excitation. At the end of that time, the response holds that value until the next aperture time.

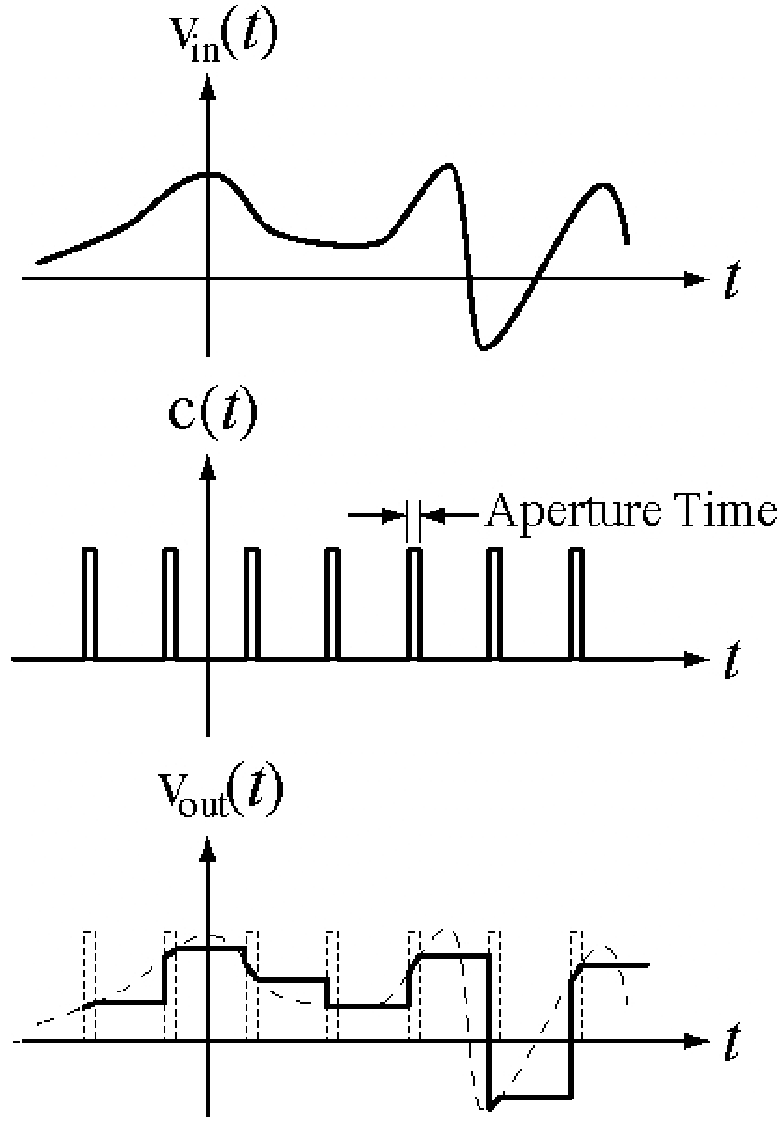

Sampling Methods

An ADC converts its input signal into a code. The code can be output serially or in parallel.

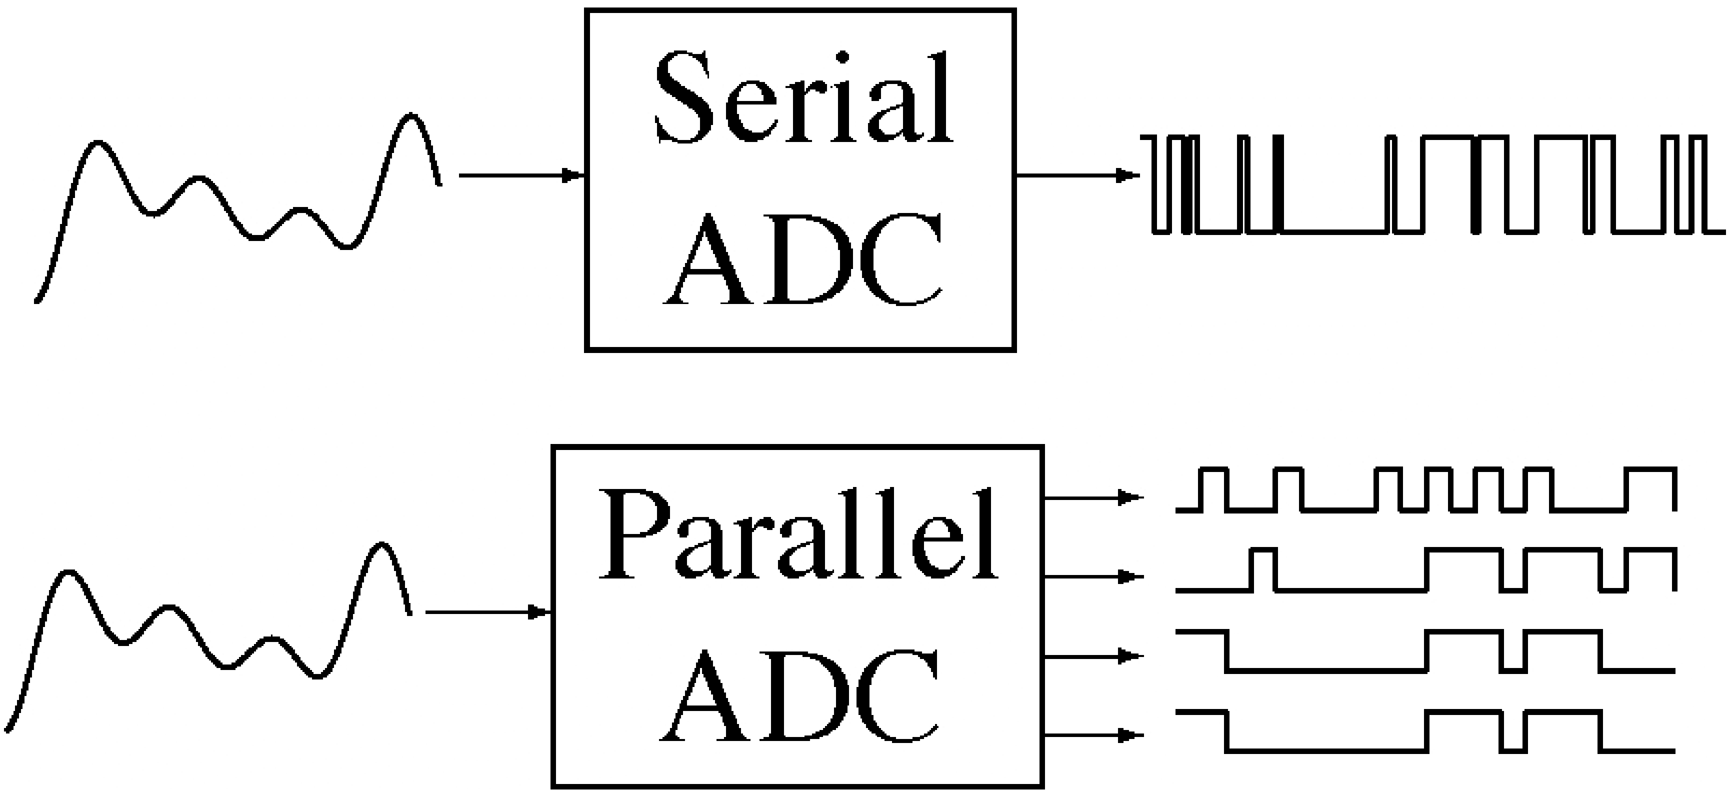

Sampling Methods

Excitation-Response Relationship for an ADC

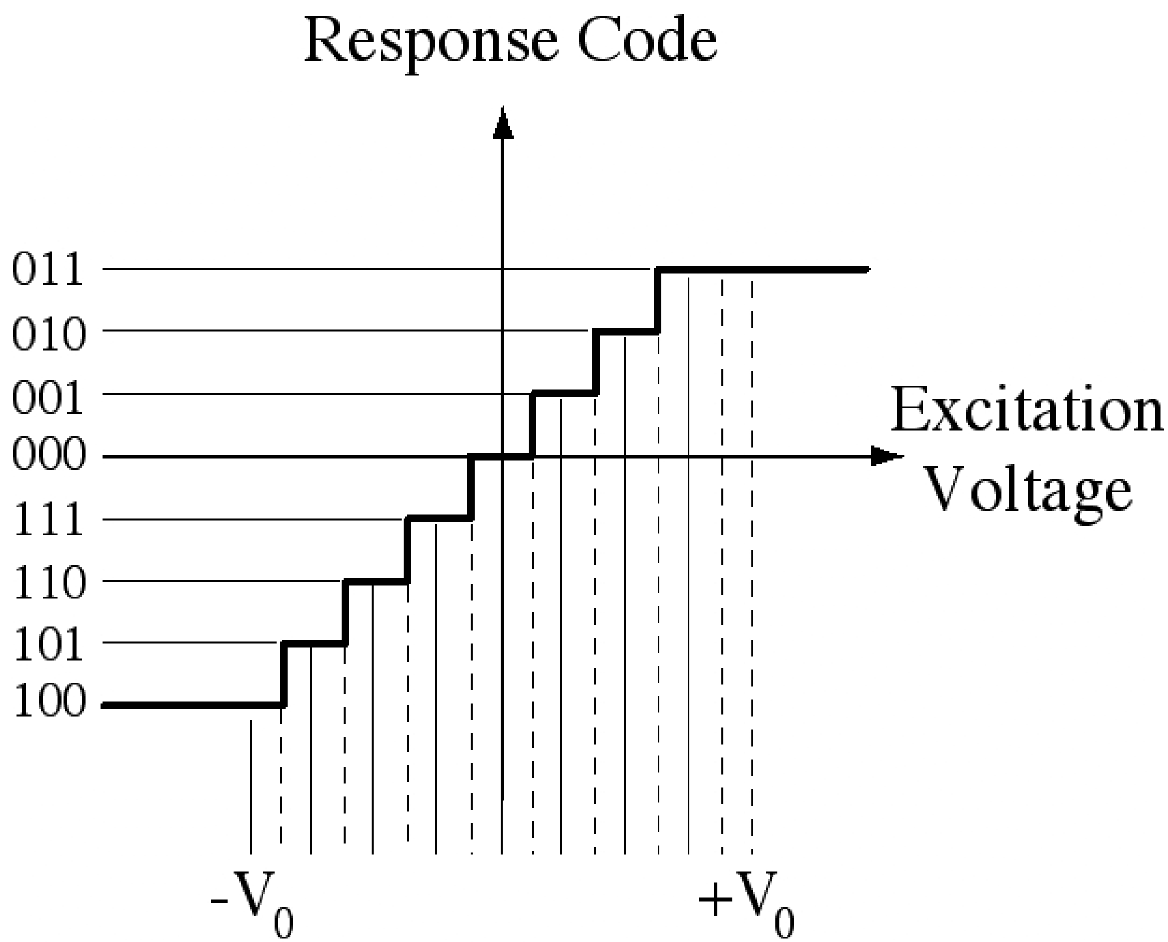

Sampling Methods

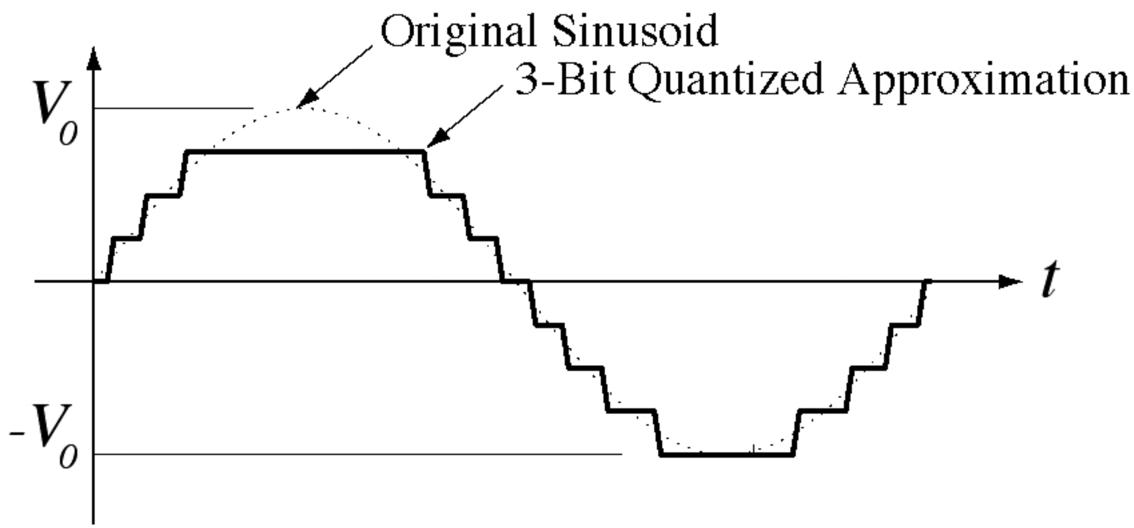

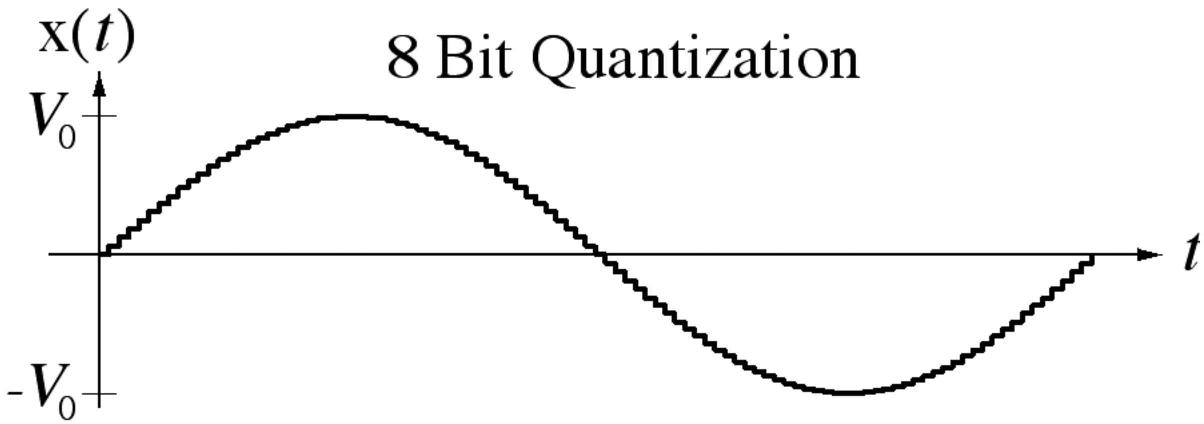

Sampling Methods

Encoded signal samples can be converted back into a CT signal by a *digital-to-analog converter* (DAC).

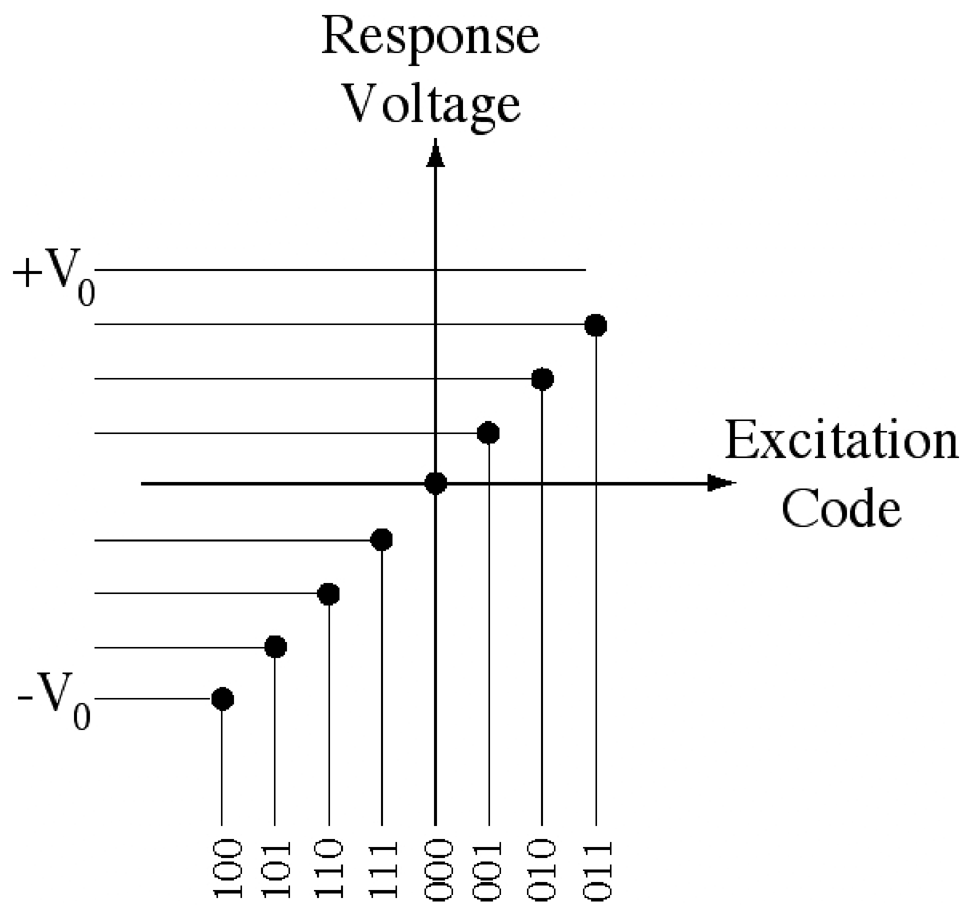

Pulse Amplitude Modulation

Pulse amplitude modulation was introduced in Chapter 6.

Modulator

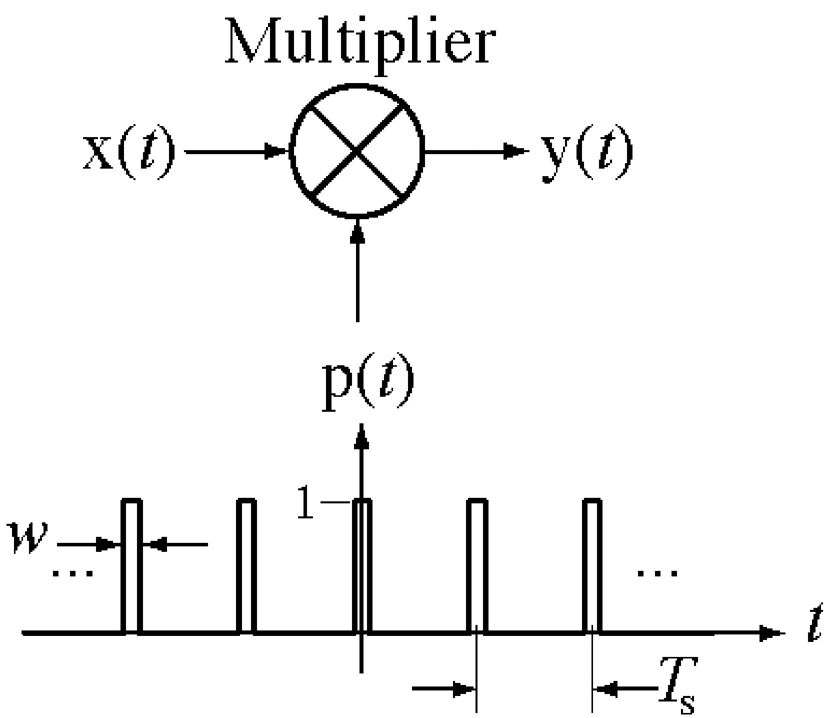

$$\mathrm{p}(t) = \mathrm{rect}\left(\frac{t}{w}\right) * \frac{1}{T_s}\,\mathrm{comb}\left(\frac{t}{T_s}\right)$$

Pulse Amplitude Modulation

The response of the pulse modulator is

$$\mathrm{y}(t) = \mathrm{x}(t)\mathrm{p}(t) = \mathrm{x}(t)\left[\mathrm{rect}\left(\frac{t}{w}\right) * \frac{1}{T_s}\mathrm{comb}\left(\frac{t}{T_s}\right)\right]$$

and its CTFT is

$$\mathrm{Y}(f) = wf_s \sum_{k=-\infty}^{\infty} \mathrm{sinc}(wkf_s)\,\mathrm{X}(f - kf_s)$$

where $f_s = \dfrac{1}{T_s}$

Pulse Amplitude Modulation

The CTFT of the response is basically multiple replicas of the CTFT of the excitation with different amplitudes, spaced apart by the pulse repetition rate.

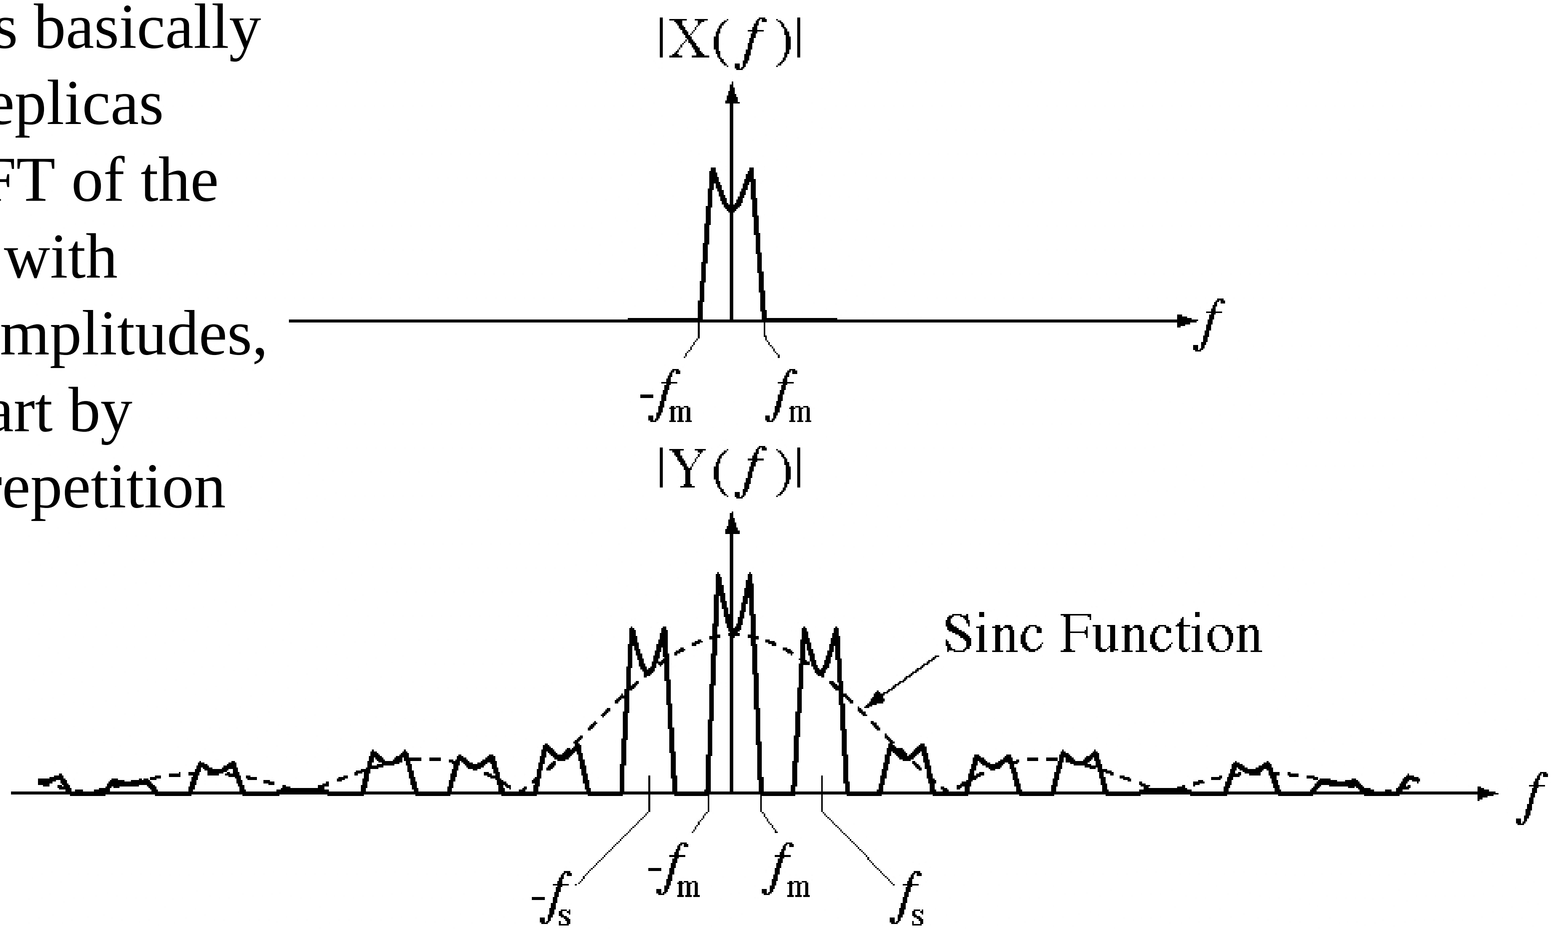

Pulse Amplitude Modulation

If the pulse train is modified to make the pulses have a constant *area* instead of a constant *height*, the pulse train becomes

$$\mathrm{p}(t) = \frac{1}{w}\mathrm{rect}\left(\frac{t}{w}\right) * \frac{1}{T_s}\mathrm{comb}\left(\frac{t}{T_s}\right)$$

and the CTFT of the modulated pulse train becomes

$$\mathrm{Y}(f) = f_s \sum_{k=-\infty}^{\infty} \mathrm{sinc}(wkf_s)\,\mathrm{X}(f - kf_s)$$

Pulse Amplitude Modulation

As the aperture time, *w*, of the pulses approaches zero the pulse train approaches an *impulse* train (a comb function) and the replicas of the original signal's spectrum all approach the same size. This limit is called *impulse sampling*.

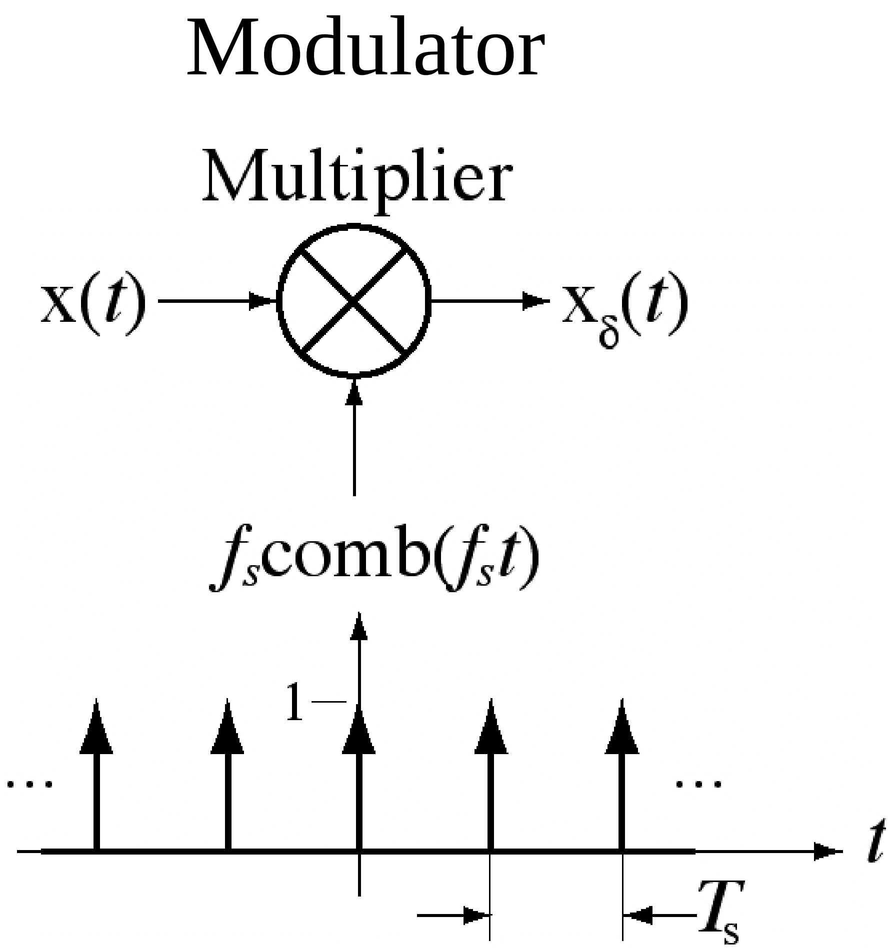

Sampling

The fundamental consideration in sampling theory is how fast to sample a signal to be able to reconstruct the signal from the samples.

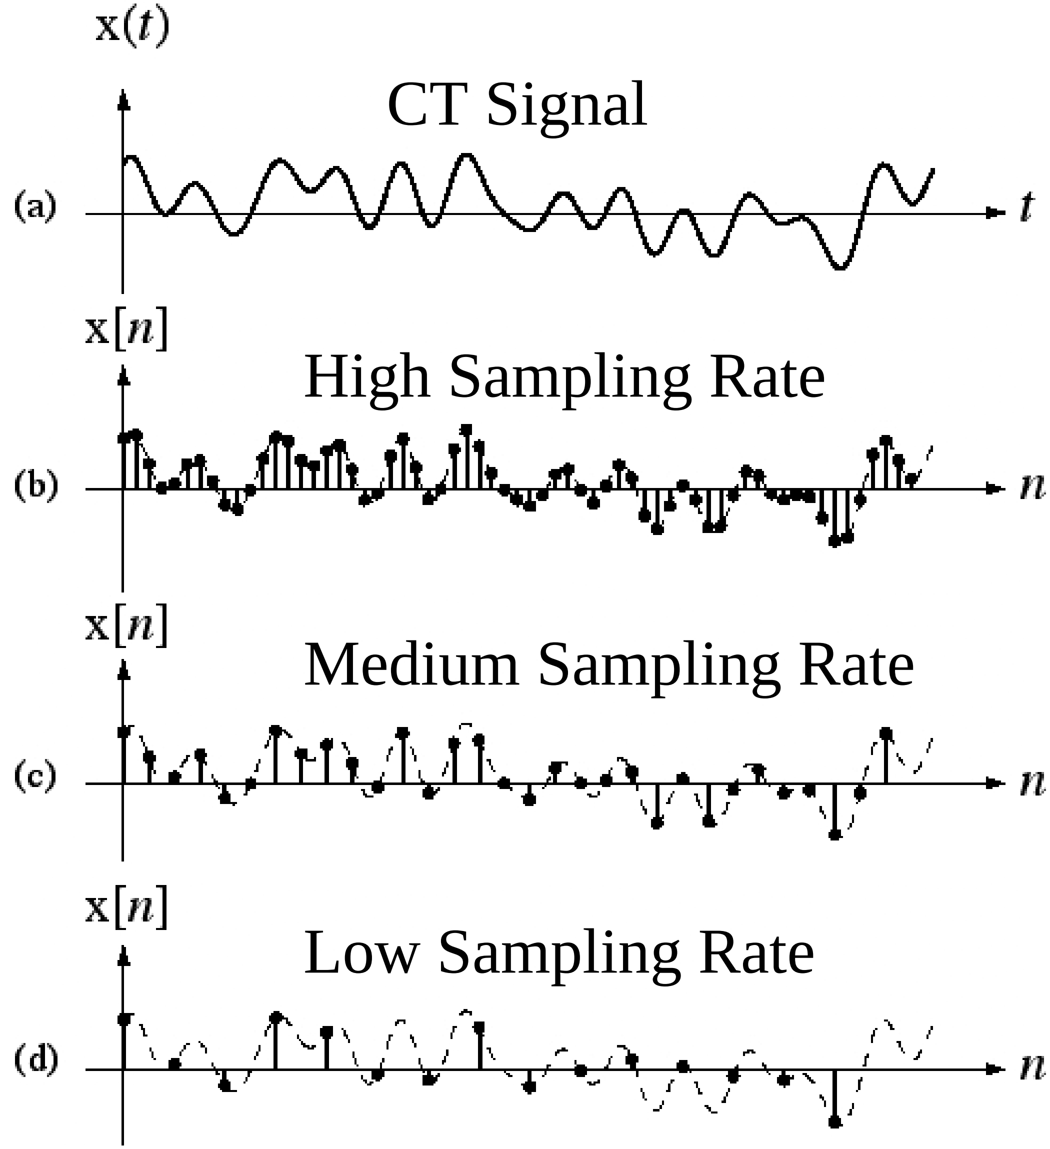

Claude Elwood Shannon

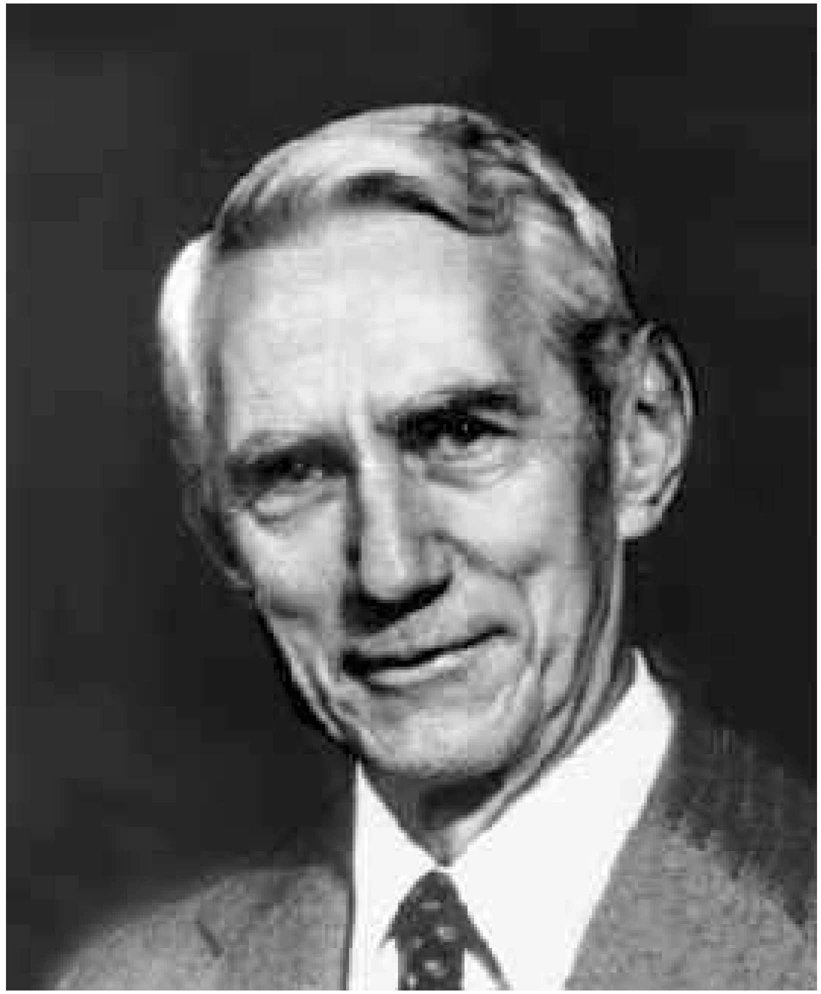

Shannon's Sampling Theorem

As an example,
let the CT signal
to be sampled be

$$x(t) = A\,\mathrm{sinc}\left(\frac{t}{w}\right)$$

Its CTFT is

$$X_{CTFT}(f) = Aw\,\mathrm{rect}(wf)$$

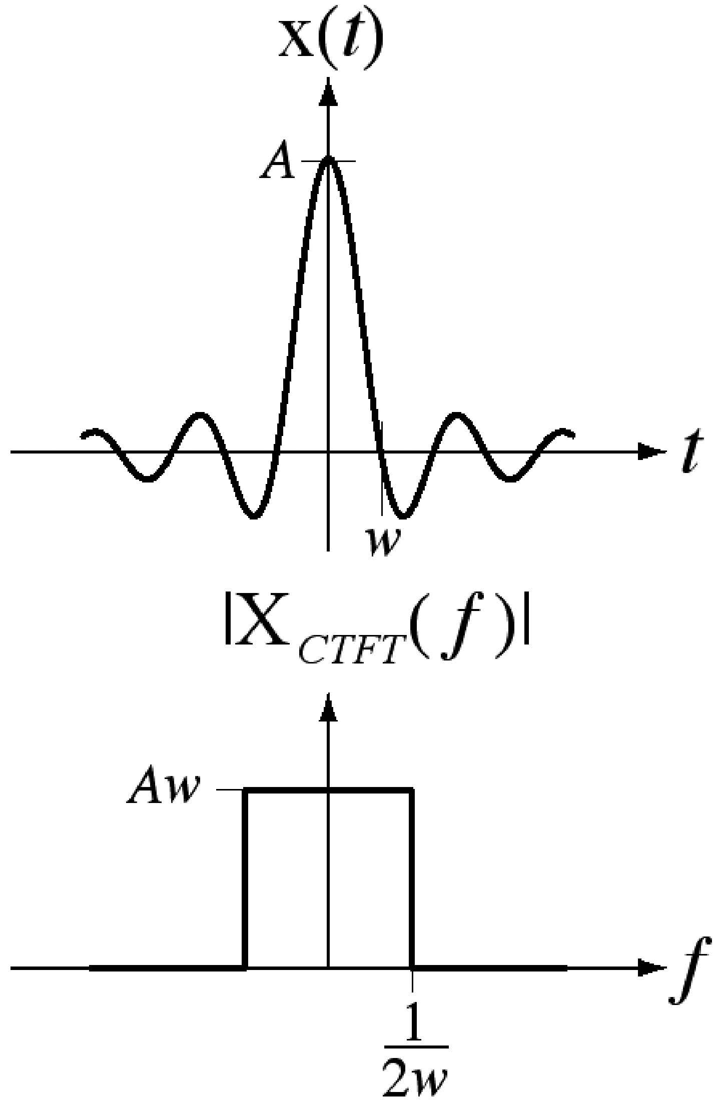

Shannon's Sampling Theorem

Sample the signal to form a DT signal,

$$\mathrm{x}[n] = \mathrm{x}(nT_s) = A\,\mathrm{sinc}\left(\frac{nT_s}{w}\right)$$

and impulse sample the same signal to form the CT impulse signal,

$$\mathrm{x}_\delta(t) = A\,\mathrm{sinc}\left(\frac{t}{w}\right) f_s\,\mathrm{comb}(f_s t) = A\sum_{n=-\infty}^{\infty}\mathrm{sinc}\left(\frac{nT_s}{w}\right)\delta(t - nT_s)$$

The DTFT of the sampled signal is

$$\mathrm{X}_{DTFT}(F) = Awf_s\,\mathrm{rect}(Fwf_s) * \mathrm{comb}(F)$$

Shannon's Sampling Theorem

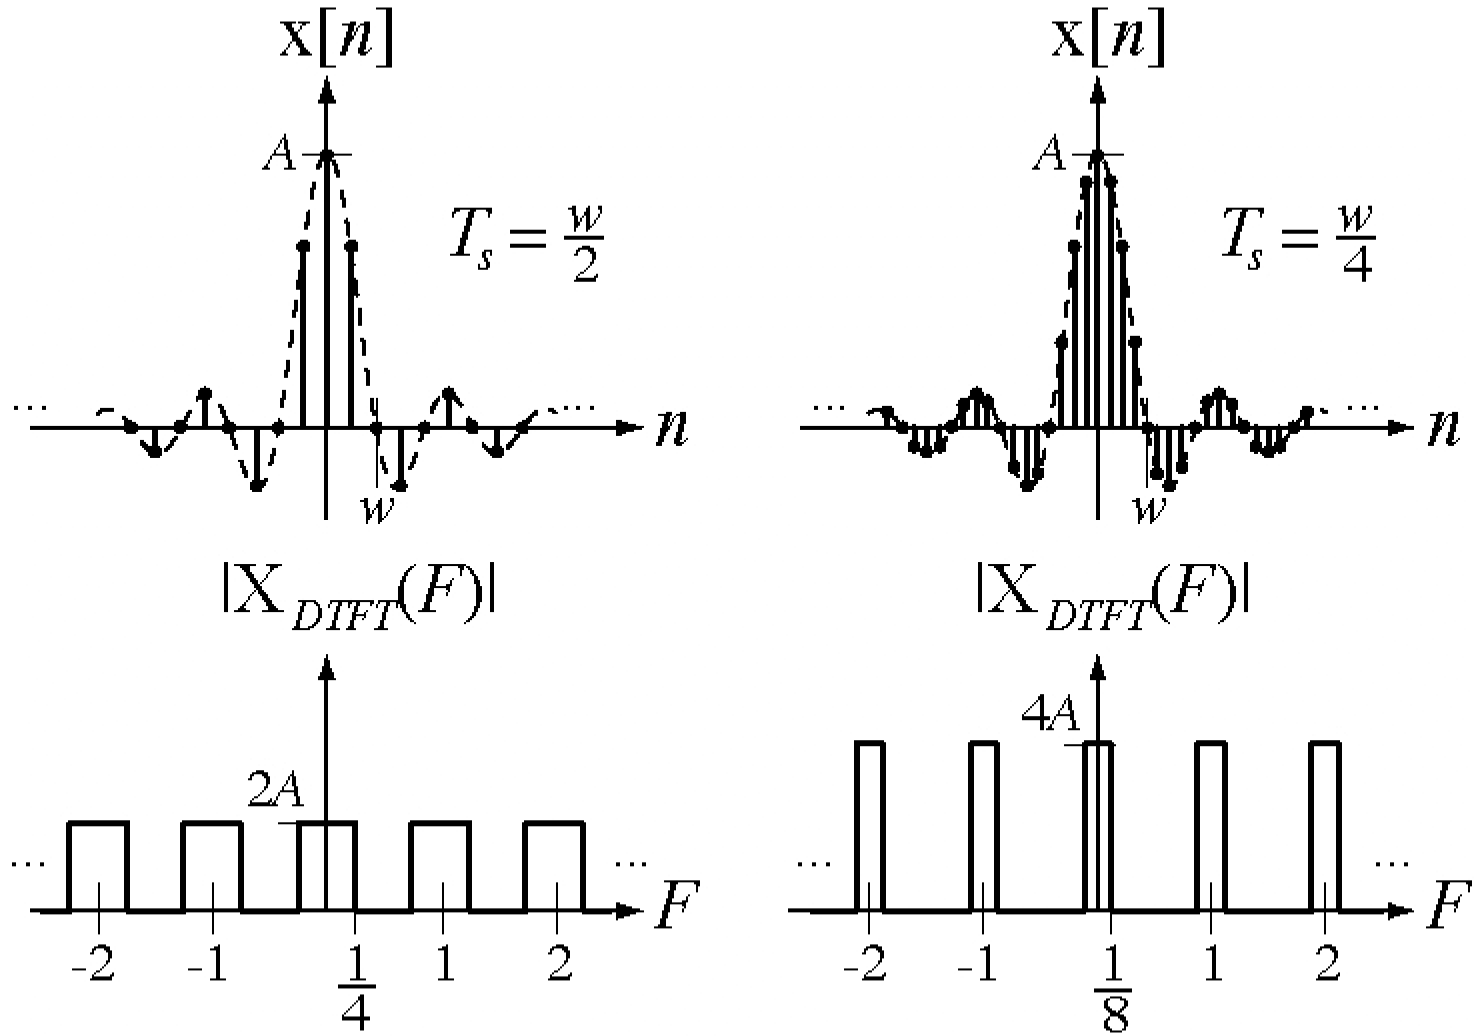

Shannon's Sampling Theorem

The CTFT of the original signal is

$$X_{CTFT}(f) = Aw\,\mathrm{rect}(wf)$$

a rectangle.

The DTFT of the sampled signal is

$$X_{DTFT}(F) = Awf_s\,\mathrm{rect}(Fwf_s) * \mathrm{comb}(F)$$

or

$$X_{DTFT}(F) = Awf_s \sum_{k=-\infty}^{\infty} \mathrm{rect}((F-k)wf_s)$$

a periodic sequence of rectangles.

Shannon's Sampling Theorem

If the "$k = 0$" rectangle from the DTFT is isolated and then the transformation,

$$F \rightarrow \frac{f}{f_s}$$

is made, the transformation is

$$Awf_s \operatorname{rect}(Fwf_s) \rightarrow Awf_s \operatorname{rect}(wf)$$

If this is now multiplied by T_s the result is

$$T_s\left[Awf_s \operatorname{rect}(Fwf_s)\right] = Aw\operatorname{rect}(wf) = \mathrm{X}_{CTFT}(f)$$

which is the CTFT of the original CT signal.

Shannon's Sampling Theorem

In this example (but not for all signals and sampling rates) the original signal can be recovered from the samples by this process:

1. Find the DTFT of the DT signal.
2. Isolate the "k = 0" function from step 1.
3. Make the change of variable, $F \to \frac{f}{f_s}$, in the result of step 2.
4. Multiply the result of step 3 by T_s
5. Find the inverse CTFT of the result of step 4.

The recovery process works in this example because the multiple replicas of the original signal's CTFT do not overlap in the DTFT. They do not overlap because the original signal is *bandlimited* and the sampling rate is high enough to separate them.

Shannon's Sampling Theorem

If the signal were sampled at a lower rate, the signal recovery process would not work because the replicas would overlap and the original CTFT function shape would not be clear.

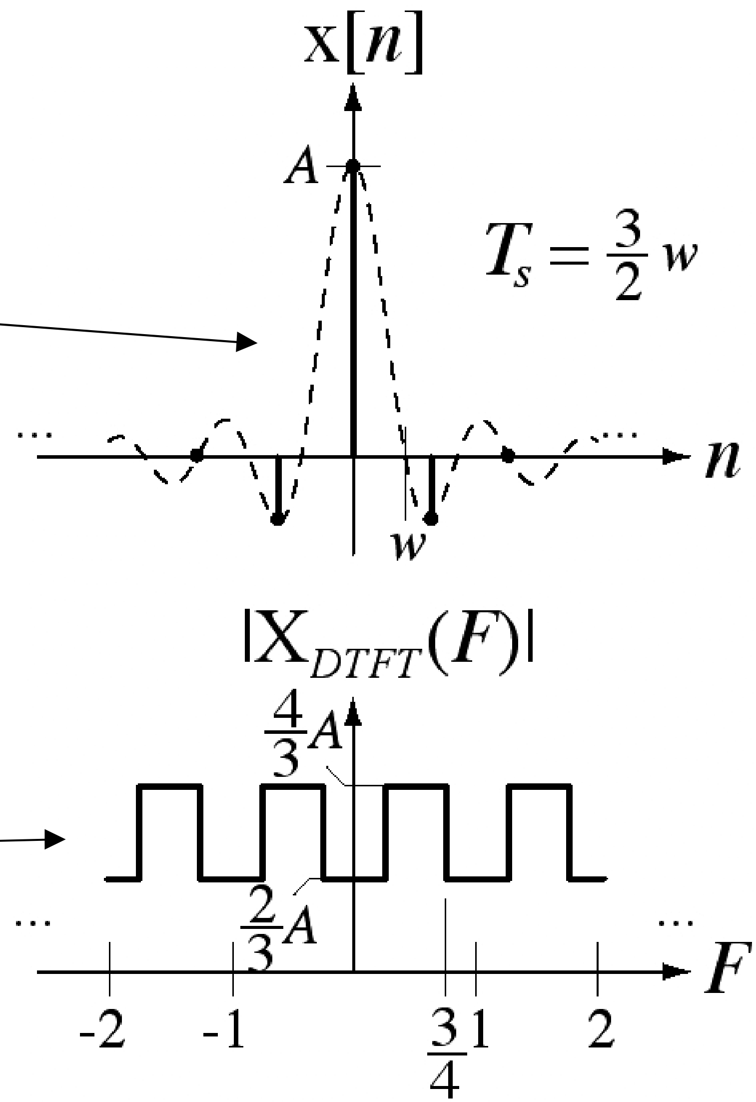

Shannon's Sampling Theorem

If a signal is impulse sampled, the CTFT of the impulse-sampled signal is

$$X_\delta(f) = X_{CTFT}(f) * \text{comb}(T_s f) = f_s \sum_{k=-\infty}^{\infty} X_{CTFT}(f - kf_s)$$

For the example signal (the sinc function),

$$X_\delta(f) = f_s \sum_{k=-\infty}^{\infty} Aw\,\text{rect}(w(f - kf_s))$$

which is the same as

$$X_{DTFT}(F)\Big|_{F \to \frac{f}{f_s}} = Awf_s \sum_{k=-\infty}^{\infty} \text{rect}((f - kf_s)w)$$

Shannon's Sampling Theorem

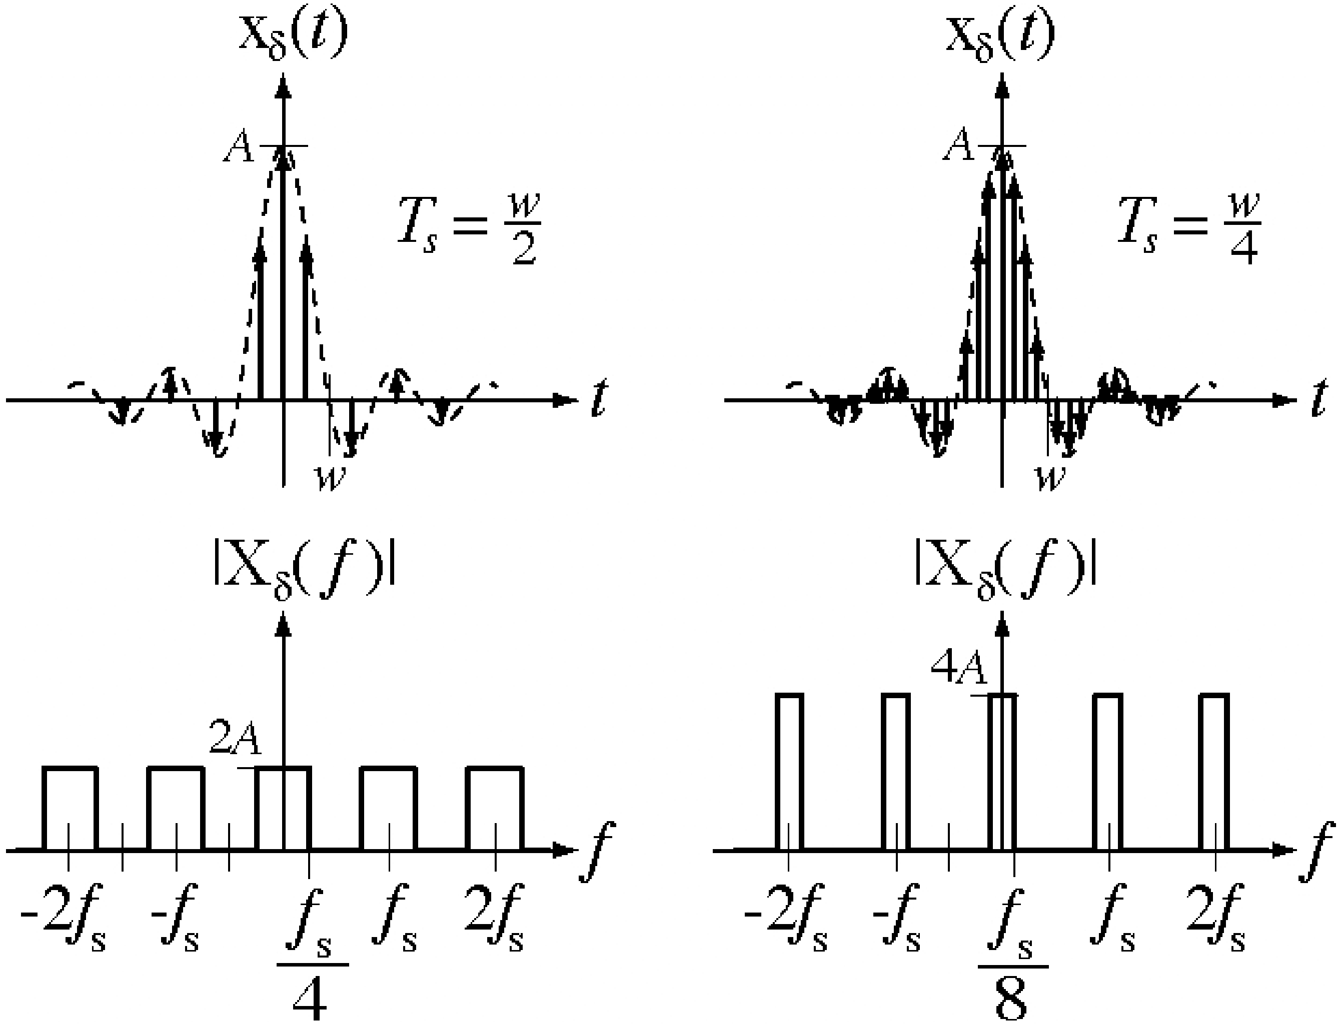

Shannon's Sampling Theorem

If the sampling rate is high enough, in the frequency range,

$$-\frac{f_s}{2} < f < \frac{f_s}{2}$$

the CTFT of the original signal and the CTFT of the impulse-sampled signal are identical except for a scaling factor of f_s . Therefore, if the impulse-sampled signal were filtered by an ideal lowpass filter with the correct corner frequency, the original signal could be recovered from the impulse-sampled signal.

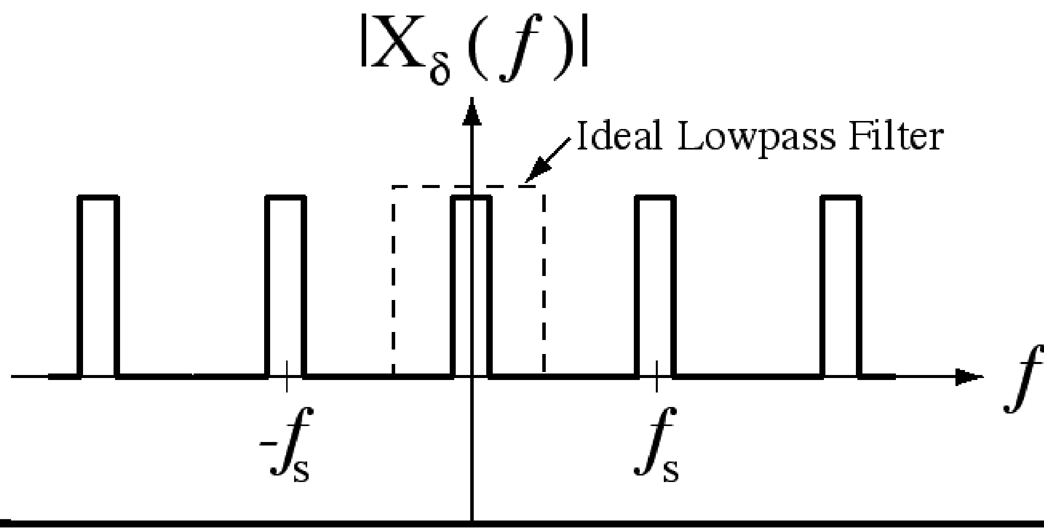

Shannon's Sampling Theorem

Suppose a signal is bandlimited with this CTFT magnitude.

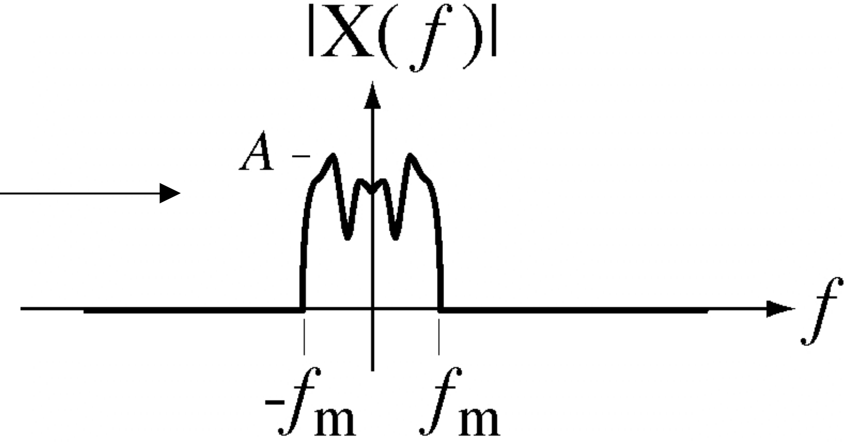

If we impulse sample it at a rate,

$$f_s = 4 f_m$$

the CTFT of the impulse-sampled signal will have this magnitude.

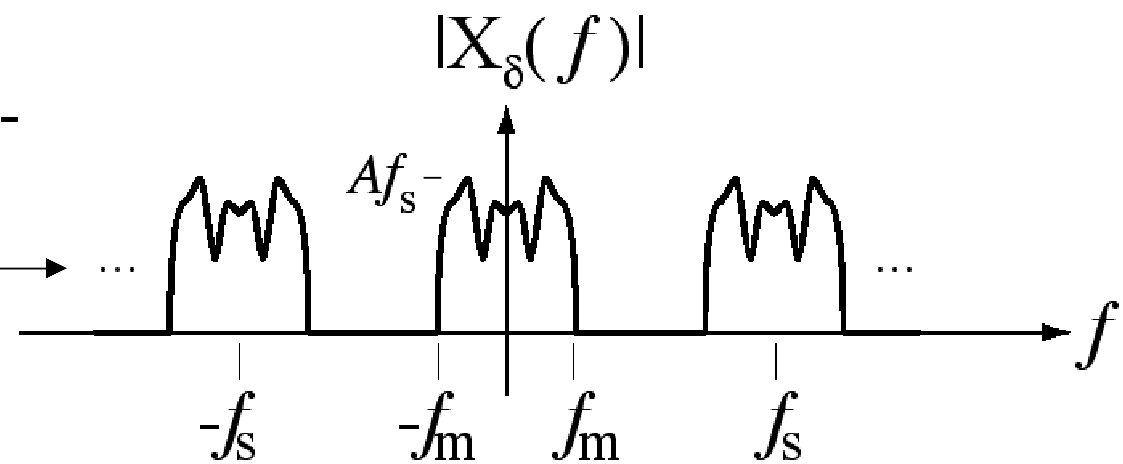

Shannon's Sampling Theorem

Suppose the same signal is now impulse sampled at a rate,

$$f_s = 2 f_m$$

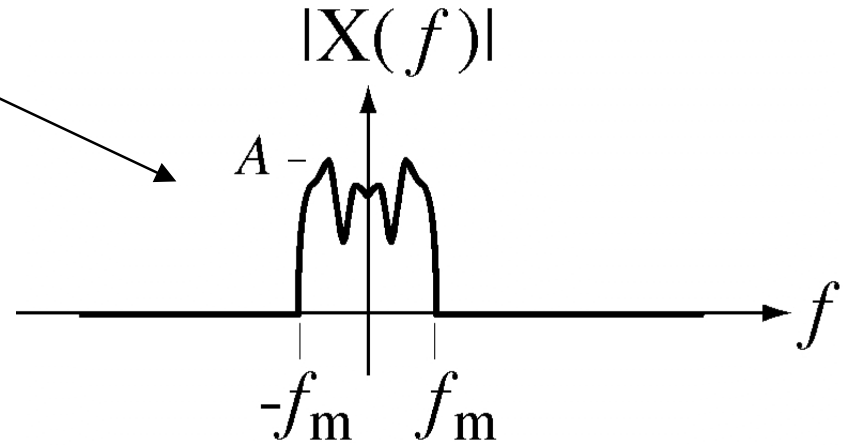

The CTFT of the impulse-sampled signal will have this magnitude.

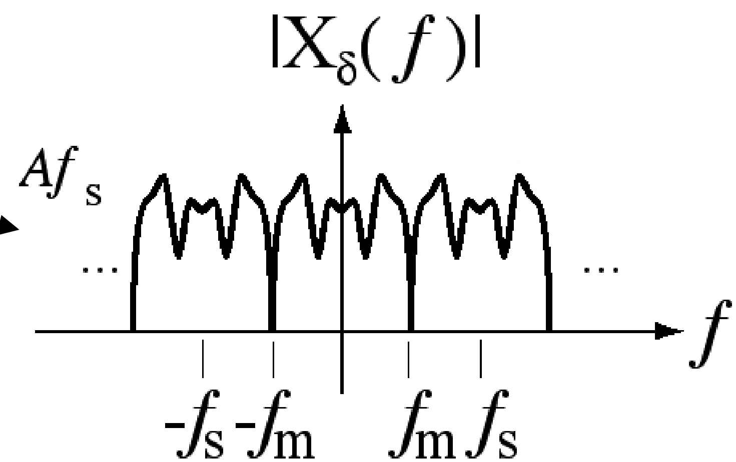

This is the minimum sampling rate at which the original signal could be recovered.

Shannon's Sampling Theorem

Now the most common form of Shannon's sampling theorem can be stated.

If a signal is sampled for all time at a rate more than twice the highest frequency at which its CTFT is non-zero it can be exactly reconstructed from the samples.

The highest frequency present in a signal is called its *Nyquist frequency*. The minimum sampling rate is called the *Nyquist* rate which is twice the Nyquist frequency. A signal sampled above the Nyquist rate is *oversampled* and a signal sampled below the Nyquist rate is *undersampled*.

Harry Nyquist

2/7/1889 - 4/4/1976

Timelimited and Bandlimited Signals

- The sampling theorem says that it is possible to sample a bandlimited signal at a rate sufficient to exactly reconstruct the signal from the samples.
- But it also says that the signal must be sampled for *all time*. This requirement holds even for signals which are *timelimited* (non-zero only for a finite time).

Timelimited and Bandlimited Signals

A signal that is timelimited cannot be bandlimited. Let x(t) be a timelimited signal. Then

$$\mathrm{x}(t) = \mathrm{x}(t)\,\mathrm{rect}\left(\frac{t - t_0}{\Delta t}\right)$$

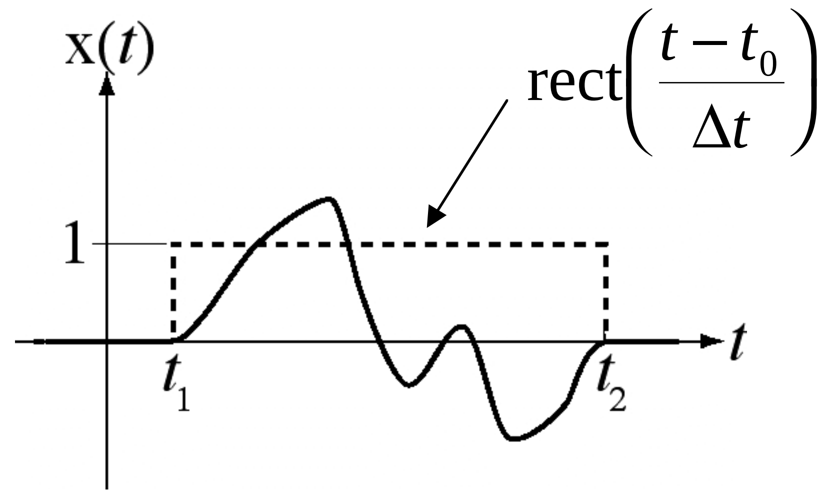

The CTFT of x(t) is

$$\mathrm{X}(f) = \mathrm{X}(f) * \Delta t\,\mathrm{sinc}(\Delta t f)\,e^{-j2\pi f t_0}$$

Since this sinc function of f is not limited in f, anything convolved with it will also not be limited in f and cannot be the CTFT of a bandlimited signal.

Sampling Bandpass Signals

There are cases in which a sampling rate *below* the Nyquist rate can also be sufficient to reconstruct a signal. This applies to so-called *bandpass* signals for which the width of the non-zero part of the CTFT is small compared with its highest frequency. In some cases, sampling below the Nyquist rate will not cause the aliases to overlap and the original signal could be recovered by using a *bandpass* filter instead of a lowpass filter.

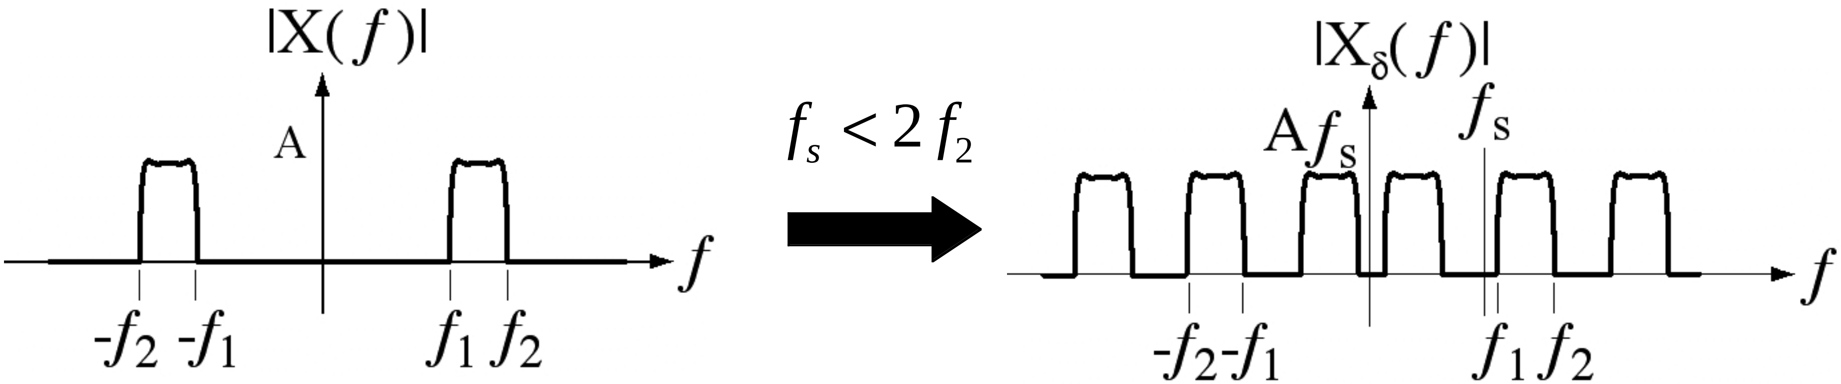

Interpolation

A CT signal can be recovered (theoretically) from an impulse-sampled version by an ideal lowpass filter. If the cutoff frequency of the filter is f_c then

$$X(f) = T_s \operatorname{rect}\left(\frac{f}{2f_c}\right) X_\delta(f) \quad , \quad f_m < f_c < (f_s - f_m)$$

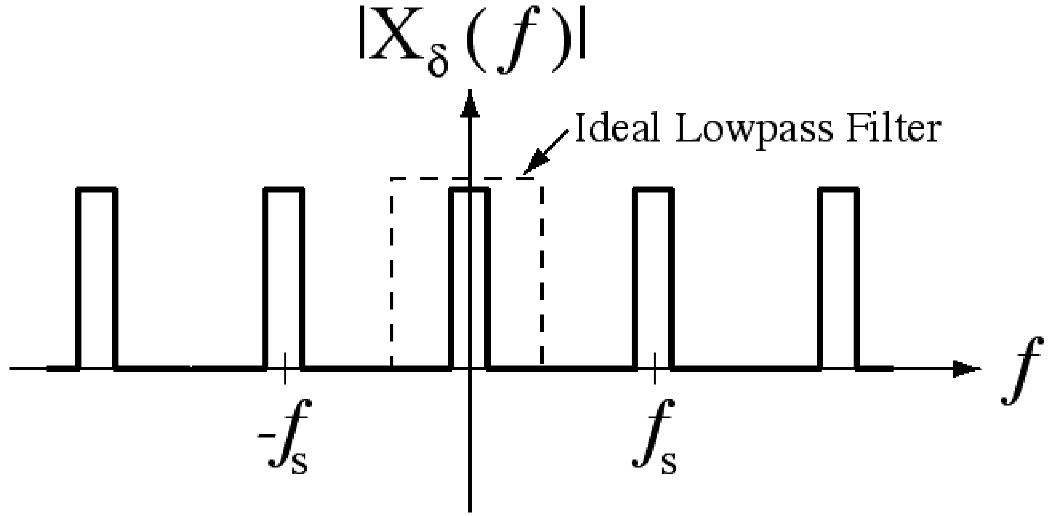

Interpolation

The time-domain operation corresponding to the ideal lowpass filter is convolution with a sinc function, the inverse CTFT of the filter's rectangular frequency response.

$$\mathrm{x}(t) = 2\frac{f_c}{f_s}\mathrm{sinc}(2f_c t) * \mathrm{x}_\delta(t)$$

Since the impulse-sampled signal is of the form,

$$\mathrm{x}_\delta(t) = \sum_{n=-\infty}^{\infty} \mathrm{x}(nT_s)\delta(t - nT_s)$$

the reconstructed original signal is

$$\mathrm{x}(t) = 2\frac{f_c}{f_s}\sum_{n=-\infty}^{\infty} \mathrm{x}(nT_s)\mathrm{sinc}(2f_c(t - nT_s))$$

Interpolation

If the sampling is at exactly the Nyquist rate, then

$$x(t) = \sum_{n=-\infty}^{\infty} x(nT_s)\,\mathrm{sinc}\left(\frac{t - nT_s}{T_s}\right)$$

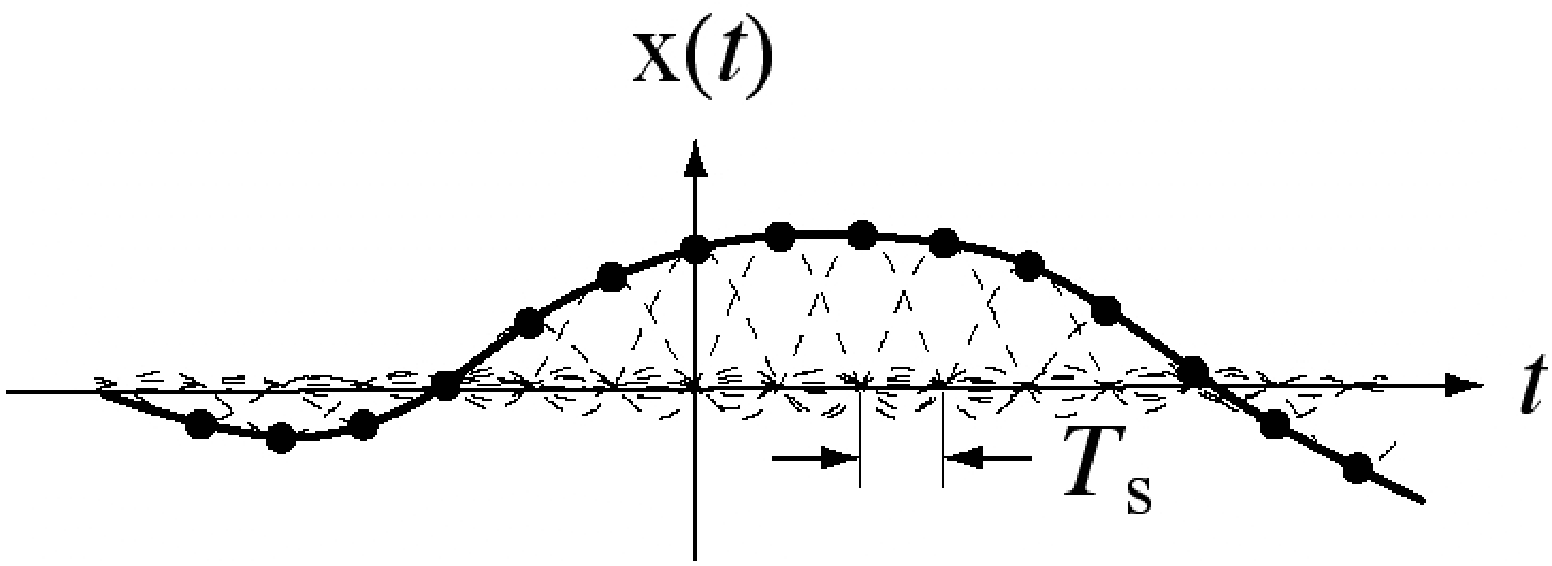

Practical Interpolation

Sinc-function interpolation is theoretically perfect but it can never be done in practice because it requires samples from the signal for all time. Therefore real interpolation must make some compromises. Probably the simplest realizable interpolation technique is what a DAC does.

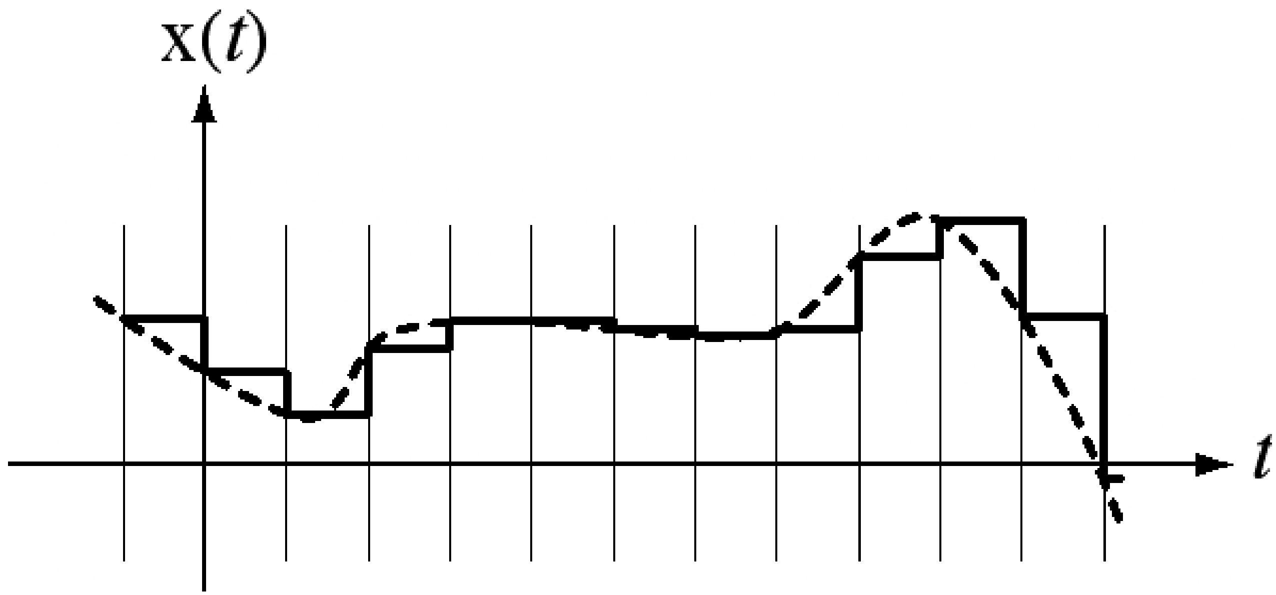

Practical Interpolation

The operation of a DAC can be mathematically modeled by a *zero-order hold (ZOH)*, a device whose impulse response is a rectangular pulse whose width is the same as the time between samples.

$$\mathrm{h}(t) = \begin{cases} 1 \ , & 0 < t < T_s \\ 0 \ , & \text{otherwise} \end{cases} = \mathrm{rect}\left(\frac{t - \frac{T_s}{2}}{T_s} \right)$$

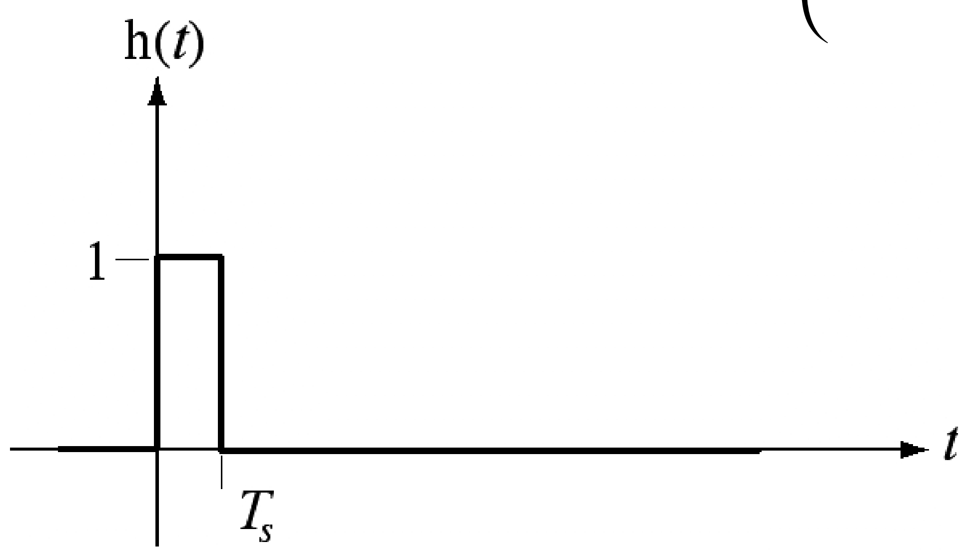

Practical Interpolation

If the signal is impulse sampled and that signal excites a ZOH, the response is the same as that produced by a DAC when it is excited by a stream of encoded sample values. The transfer function of the ZOH is a sinc function with linear phase shift.

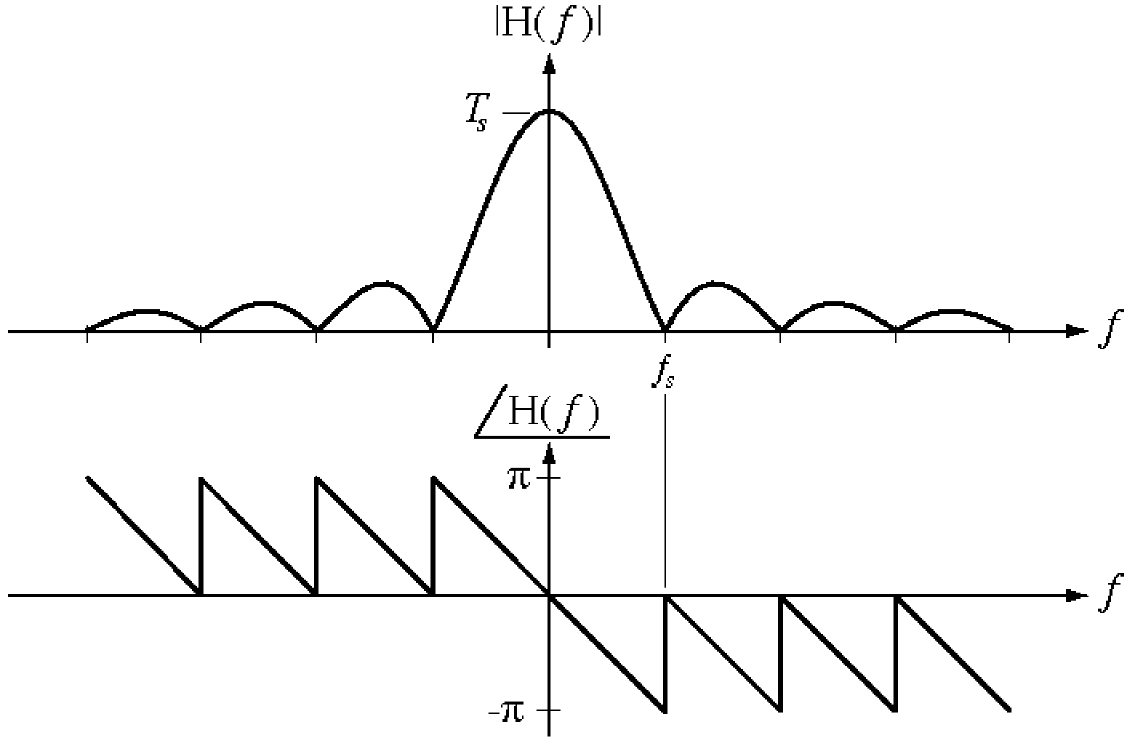

Practical Interpolation

The ZOH suppresses aliases but does not entirely eliminate them.

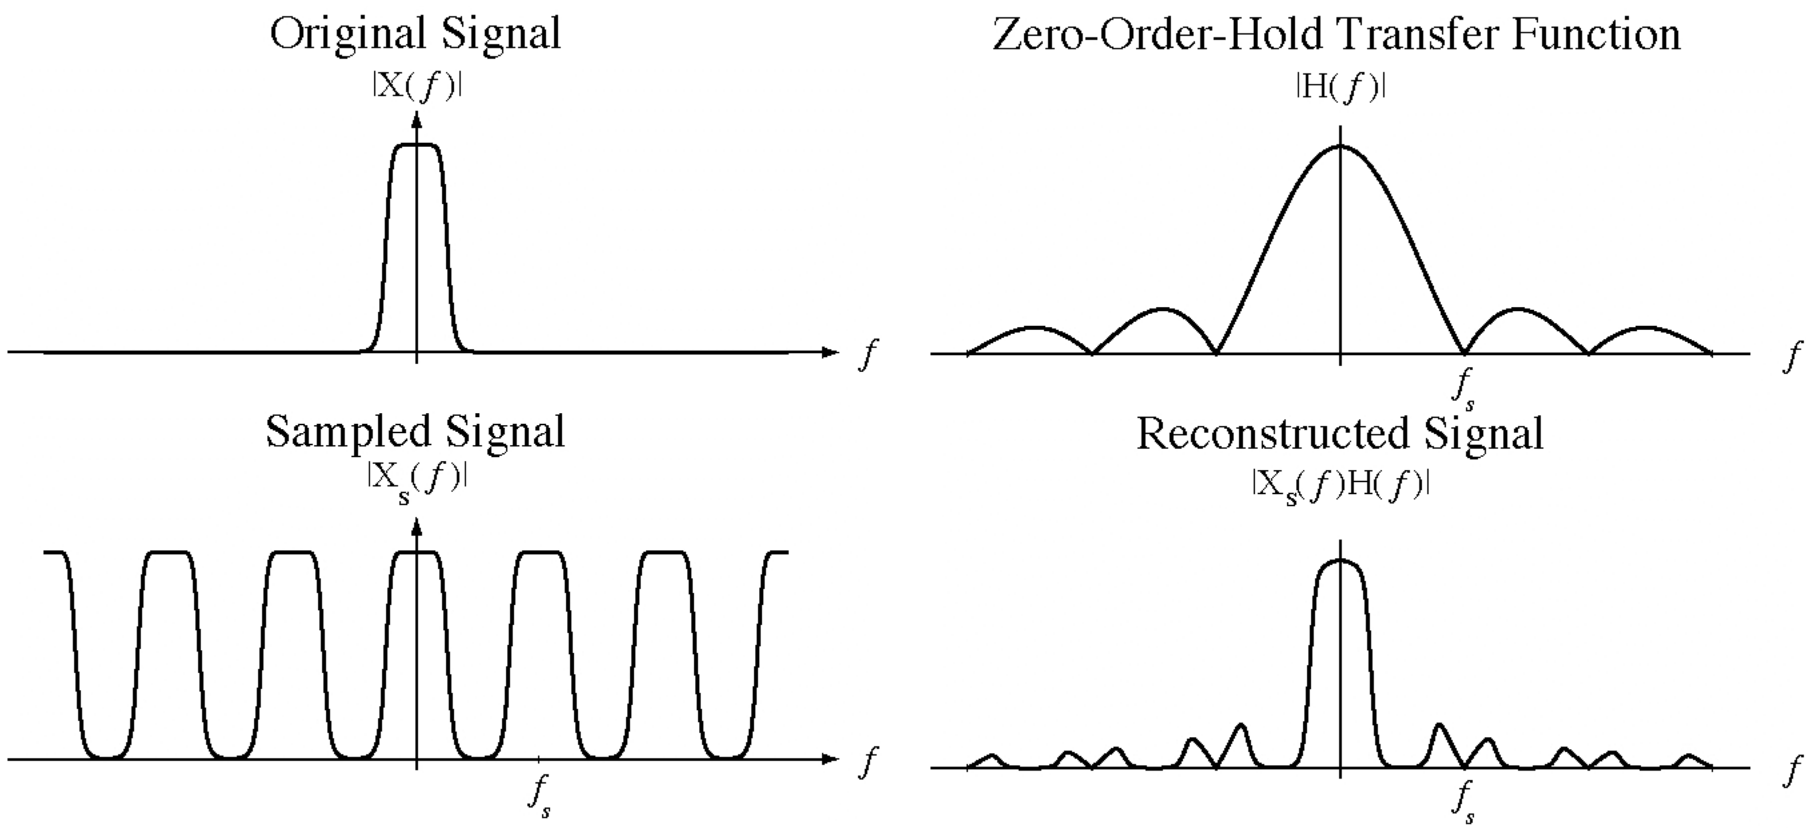

Practical Interpolation

A "natural" idea would be to simply draw straight lines between sample values. This cannot be done in real time because doing so requires knowledge of the "next" sample value before it occurs and that would require a non-causal system. If the reconstruction is delayed by one sample time, then it can be done with a causal system.

Non-Causal First-Order Hold

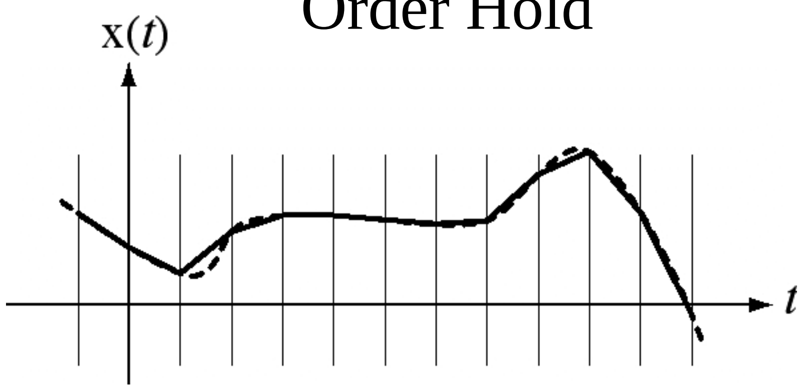

Causal First-Order Hold

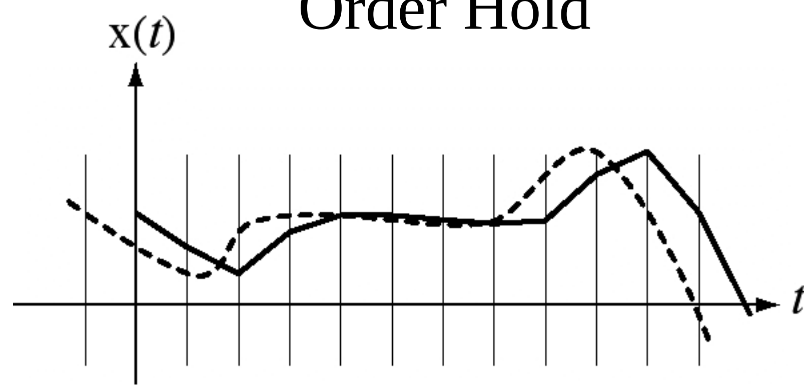

Sampling a Sinusoid

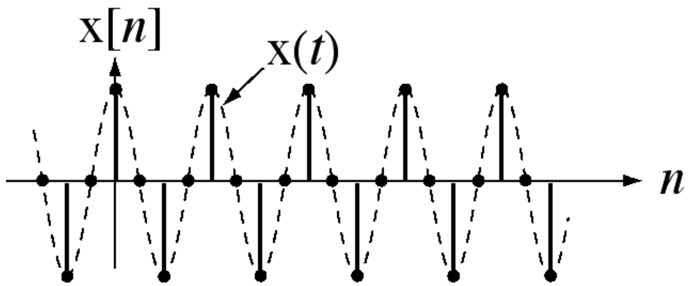

Cosine sampled at twice its Nyquist rate. Samples uniquely determine the signal.

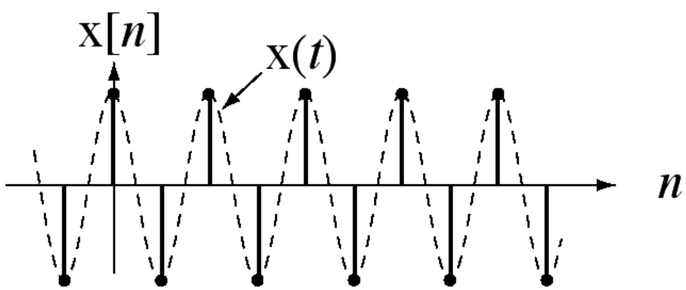

Cosine sampled at exactly its Nyquist rate. Samples *do not* uniquely determine the signal.

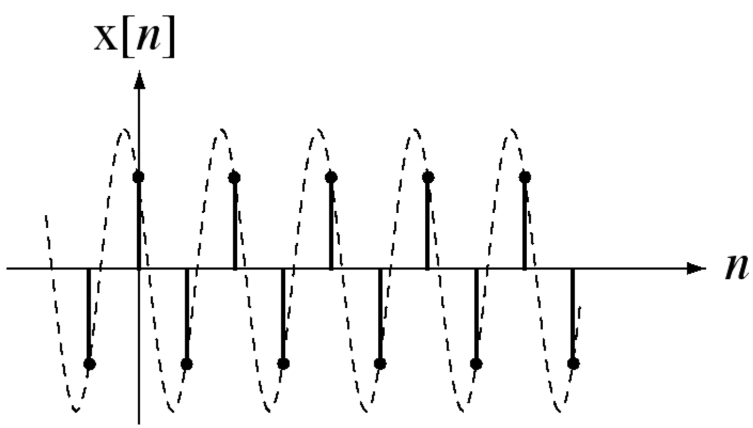

A different sinusoid of the same frequency with exactly the same samples as above.

Sampling a Sinusoid

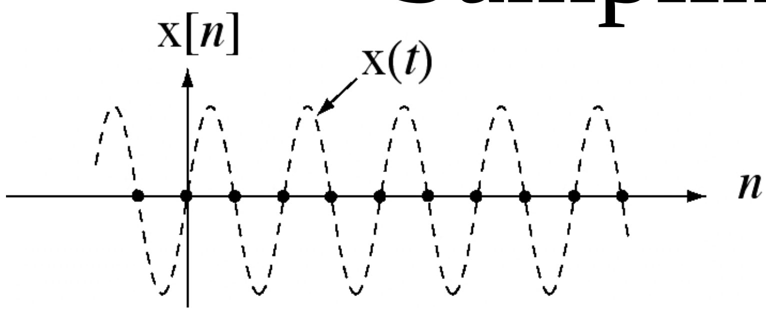

Sine sampled at its Nyquist rate.
All the samples are zero.

Adding a sine at the Nyquist frequency (half the Nyquist rate) to any signal does not change the samples.

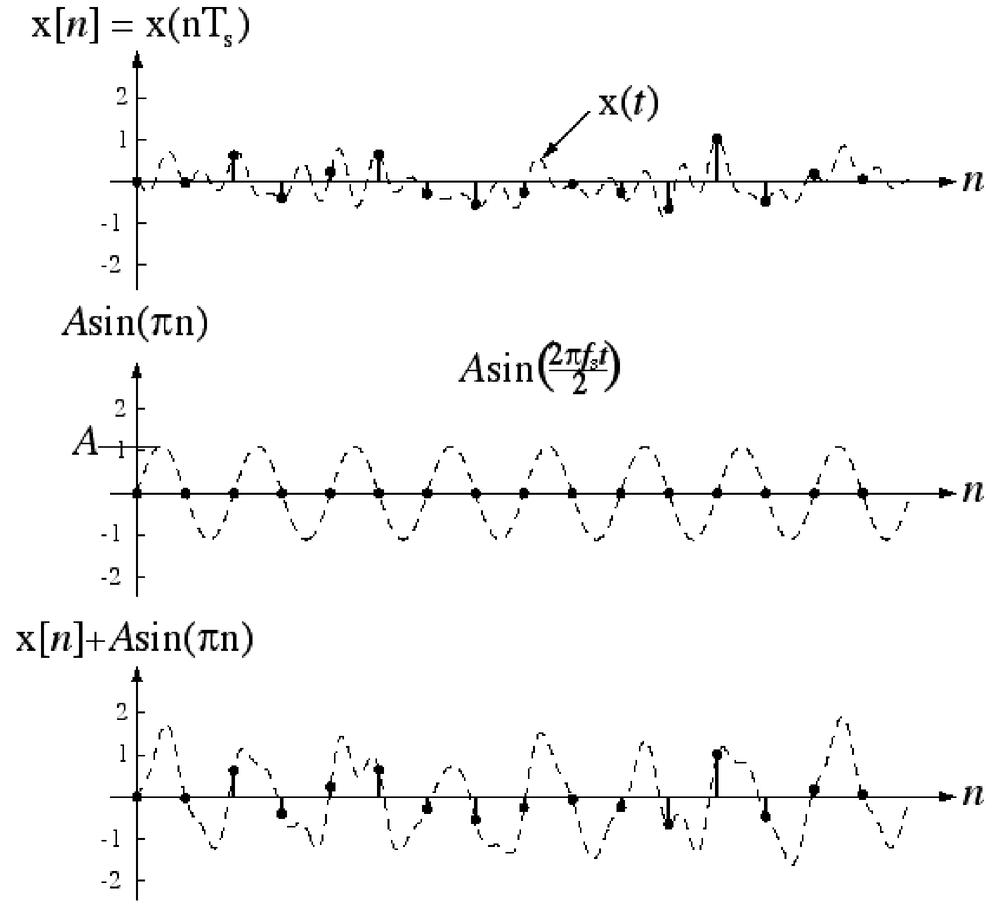

Sampling a Sinusoid

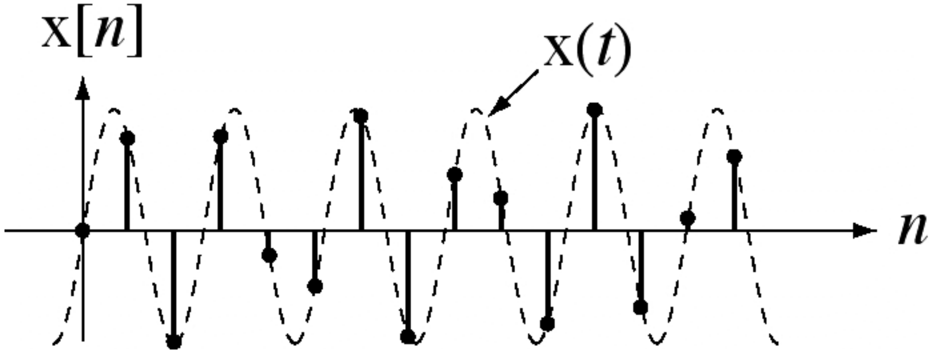

Sine sampled slightly above its Nyquist rate

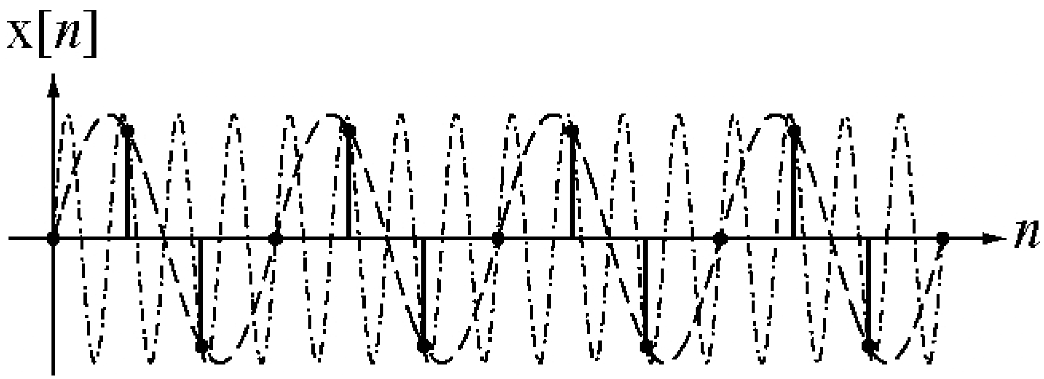

Two different sinusoids sampled at the same rate with the same samples

It can be shown (p. 516) that the samples from two sinusoids,

$$\mathrm{x}_1(t) = A\cos(2\pi f_0 t + \theta) \qquad \mathrm{x}_2(t) = A\cos(2\pi(f_0 + kf_s)t + \theta)$$

taken at the rate, f_s, are the same for any integer value of *k*.

Sampling DT Signals

One way of representing the sampling of CT signals is by impulse sampling, multiplying the signal by an impulse train (a comb). DT signals are sampled in an analogous way. If x[n] is the signal to be sampled, the sampled signal is

$$\mathrm{x}_s[n] = \mathrm{x}[n]\,\mathrm{comb}_{N_s}[n]$$

where N_s is the discrete time between samples and the DT sampling rate is $F_s = \dfrac{1}{N_s}$.

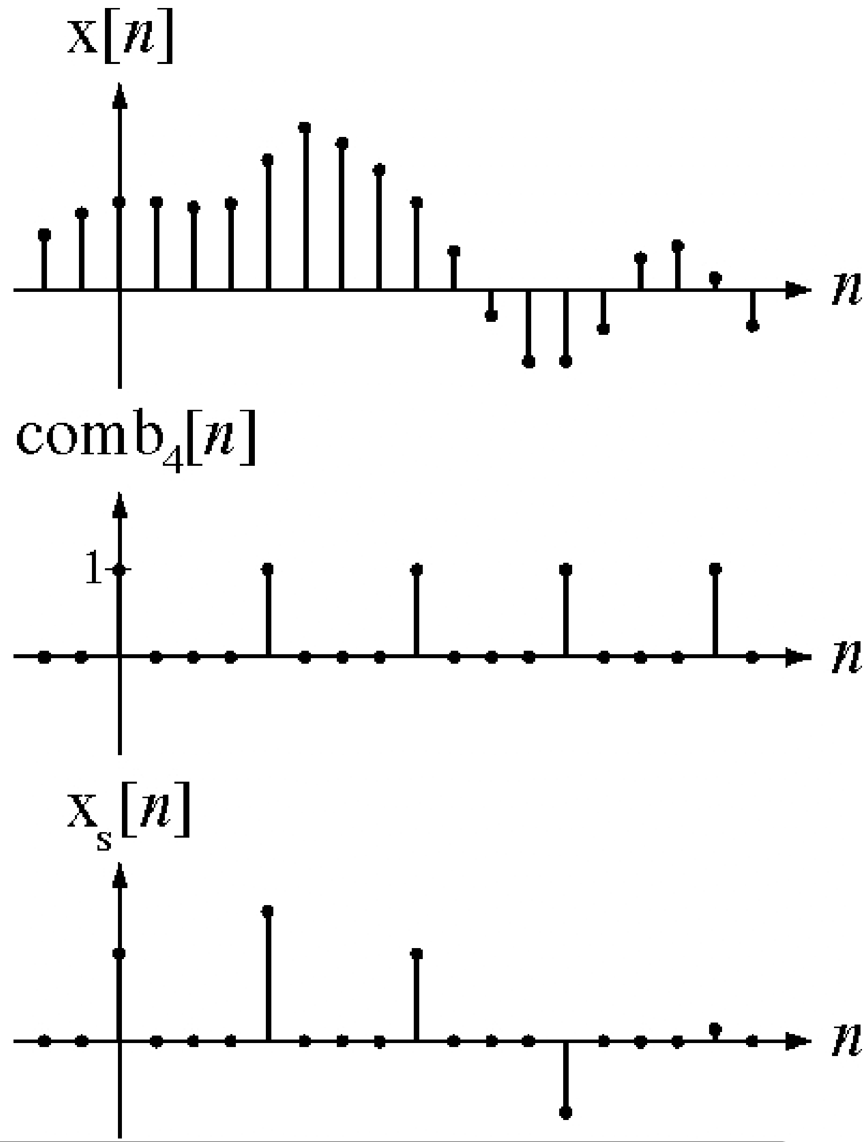

Sampling DT Signals

The DTFT of the sampled DT signal is

$$X_s(F) = X(F) \circledast \mathrm{comb}(N_s F)$$
$$= X(F) \circledast \mathrm{comb}\left(\frac{F}{F_s}\right)$$

In this example the aliases do not overlap and it would be possible to recover the original DT signal from the samples. The general rule is that $F_s > 2F_m$ where F_m is the maximum DT frequency in the signal.

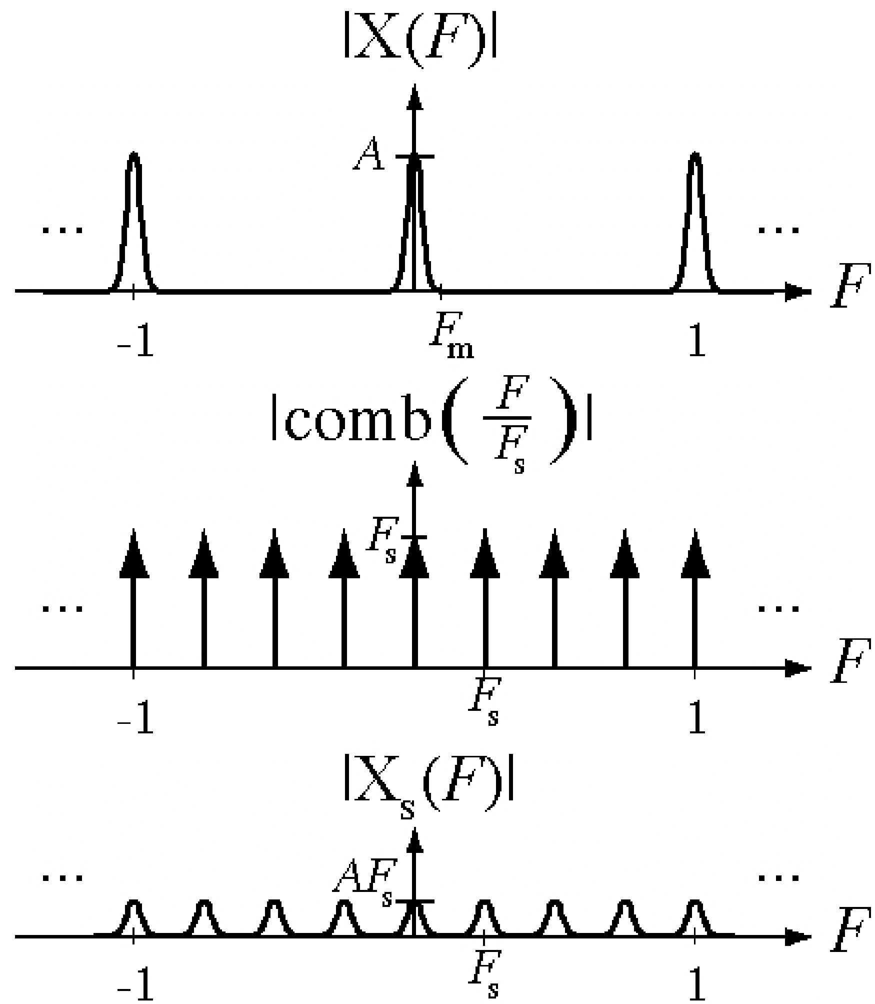

Sampling DT Signals

Interpolation is accomplished by passing the impulse-sampled DT signal through a DT lowpass filter.

$$\mathrm{X}(F) = \mathrm{X}_s(F)\left[\frac{1}{F_s}\mathrm{rect}\left(\frac{F}{2F_c}\right) * \mathrm{comb}(F)\right]$$

The equivalent operation in the discrete-time domain is

$$\mathrm{x}[n] = \mathrm{x}_s[n] * \frac{2F_c}{F_s}\mathrm{sinc}(2F_c n)$$

Sampling DT Signals

Decimation

It is common practice, after sampling a DT signal, to remove all the zero values created by the sampling process, leaving only the non-zero values. This process is *decimation*, first introduced in Chapter 2. The decimated DT signal is

$$\mathrm{x}_d[n] = \mathrm{x}_s[N_s n] = \mathrm{x}[N_s n]$$

and its DTFT is (p. 518)

$$\mathrm{X}_d(F) = \mathrm{X}_s\left(\frac{F}{N_s}\right)$$

Decimation is sometimes called *downsampling*.

Sampling DT Signals

Decimation

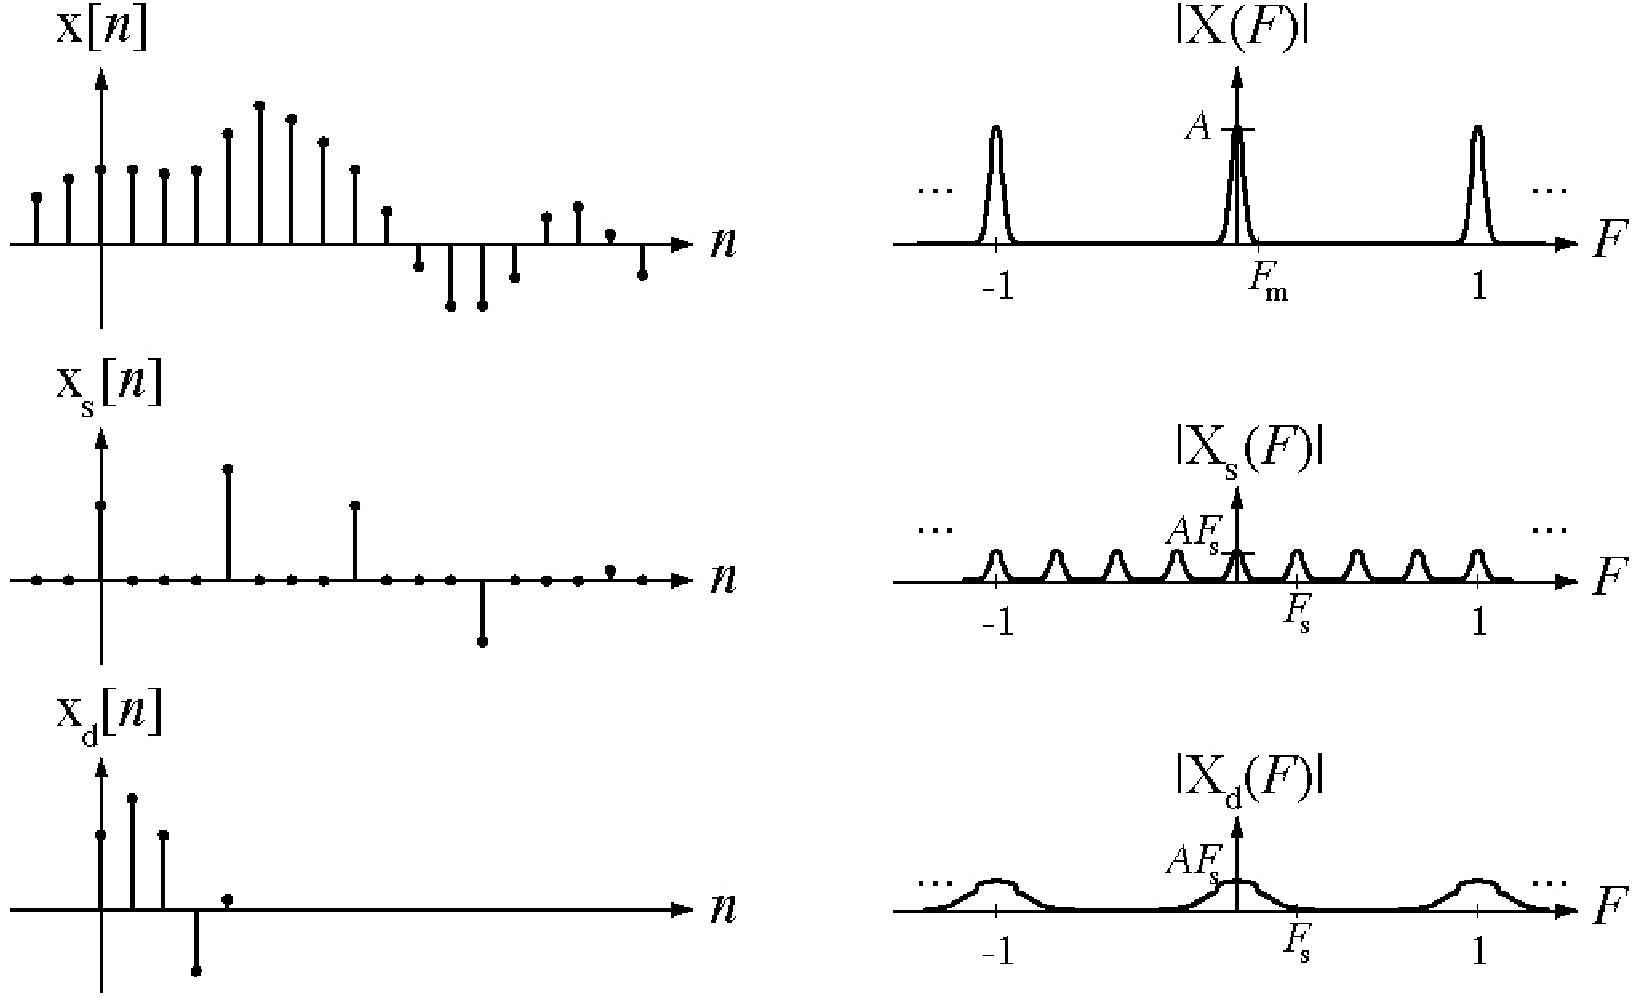

Sampling DT Signals

The opposite of downsampling is *upsampling*. It is simply the reverse of downsampling. If the original signal is x[n], then the upsampled signal is

$$\mathrm{x}_s[n] = \begin{cases} \mathrm{x}\left[\dfrac{n}{N_s}\right] & , \quad \dfrac{n}{N_s} \text{ an integer} \\ 0 & , \quad \text{otherwise} \end{cases}$$

where $N_s - 1$ zeros have been inserted between adjacent values of x[n]. If X(F) is the DTFT of x[n], then

$$\mathrm{X}_s(F) = \mathrm{X}(N_s F)$$

is the DTFT of $\mathrm{x}_s[n]$.

Sampling DT Signals

The signal, $x_s[n]$, can be lowpass filtered to interpolate between the non-zero values and form $x_i[n]$.

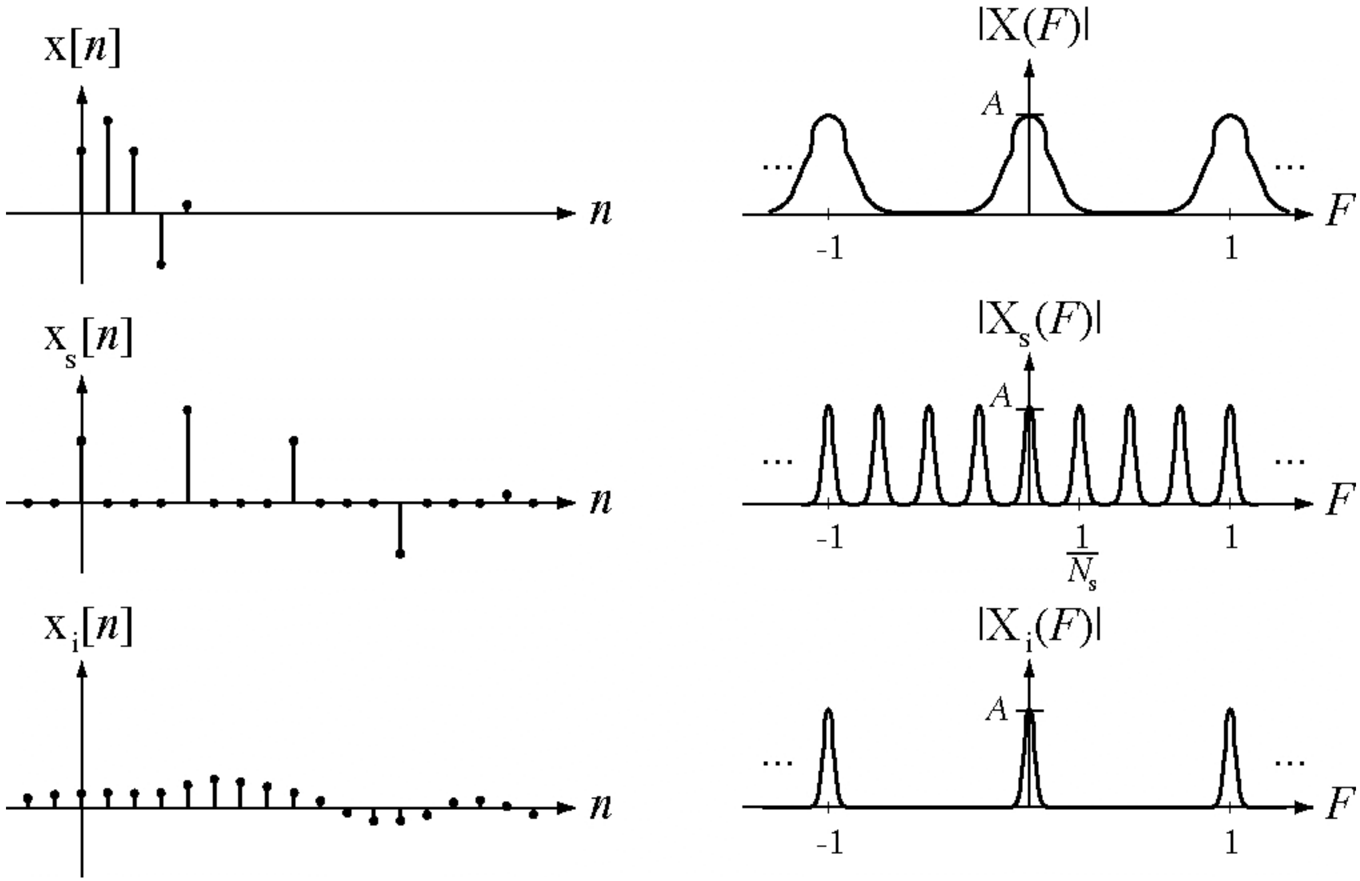

Bandlimited Periodic Signals

- If a signal is bandlimited it can be properly sampled according to the sampling theorem.
- If that signal is also periodic its CTFT consists only of impulses.
- Since it is bandlimited, there is a finite number of (non-zero) impulses.
- Therefore the signal can be exactly represented by a finite set of numbers, the impulse strengths.

Bandlimited Periodic Signals

- If a bandlimited periodic signal is sampled above the Nyquist rate over exactly one fundamental period, that set of numbers is sufficient to completely describe it
- If the sampling continued, these same samples would be repeated in every fundamental period
- So the number of numbers needed to completely describe the signal is finite in both the time and frequency domains

Bandlimited Periodic Signals

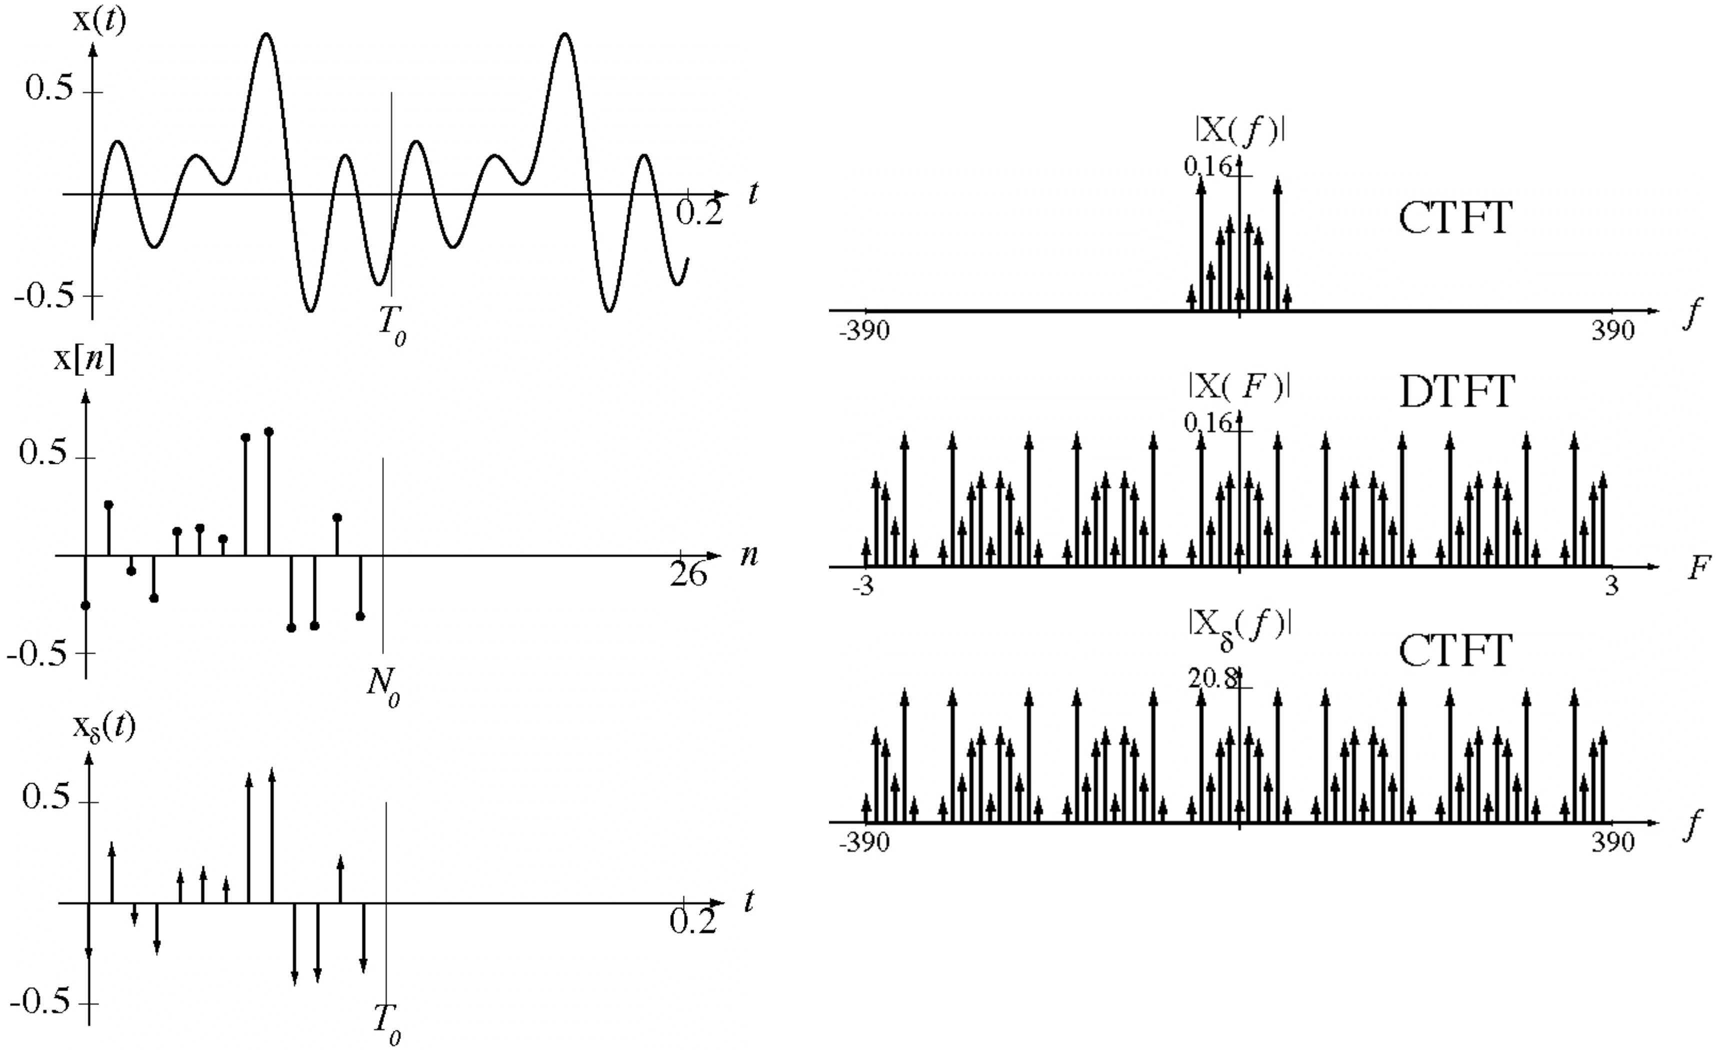

The Discrete Fourier Transform

The most widely used Fourier method in the world is the *Discrete Fourier Transform (DFT)*. It is defined by

$$\mathrm{x}[n] = \frac{1}{N_F}\sum_{k=0}^{N_F-1} \mathrm{X}[k] e^{j2\pi\frac{nk}{N_F}} \xleftrightarrow{\mathcal{DFT}} \mathrm{X}[k] = \sum_{n=0}^{N_F-1} \mathrm{x}[n] e^{-j2\pi\frac{nk}{N_F}}$$

This should look familiar. It is almost identical to the DTFS.

$$\mathrm{x}[n] = \sum_{k=0}^{N_F-1} \mathrm{X}[k] e^{j2\pi\frac{nk}{N_F}} \xleftrightarrow{\mathcal{FS}} \mathrm{X}[k] = \frac{1}{N_F}\sum_{n=0}^{N_F-1} \mathrm{x}[n] e^{-j2\pi\frac{nk}{N_F}}$$

The difference is only a scaling factor. There really should not be two so similar Fourier methods with different names but, for historical reasons, there are.

The Discrete Fourier Transform

The relation between the CTFT of a CT signal and the DFT of samples taken from it will be illustrated in the next few slides. Let an original CT signal, x(t), be sampled N_F times at a rate, f_s .

Original CT Signal

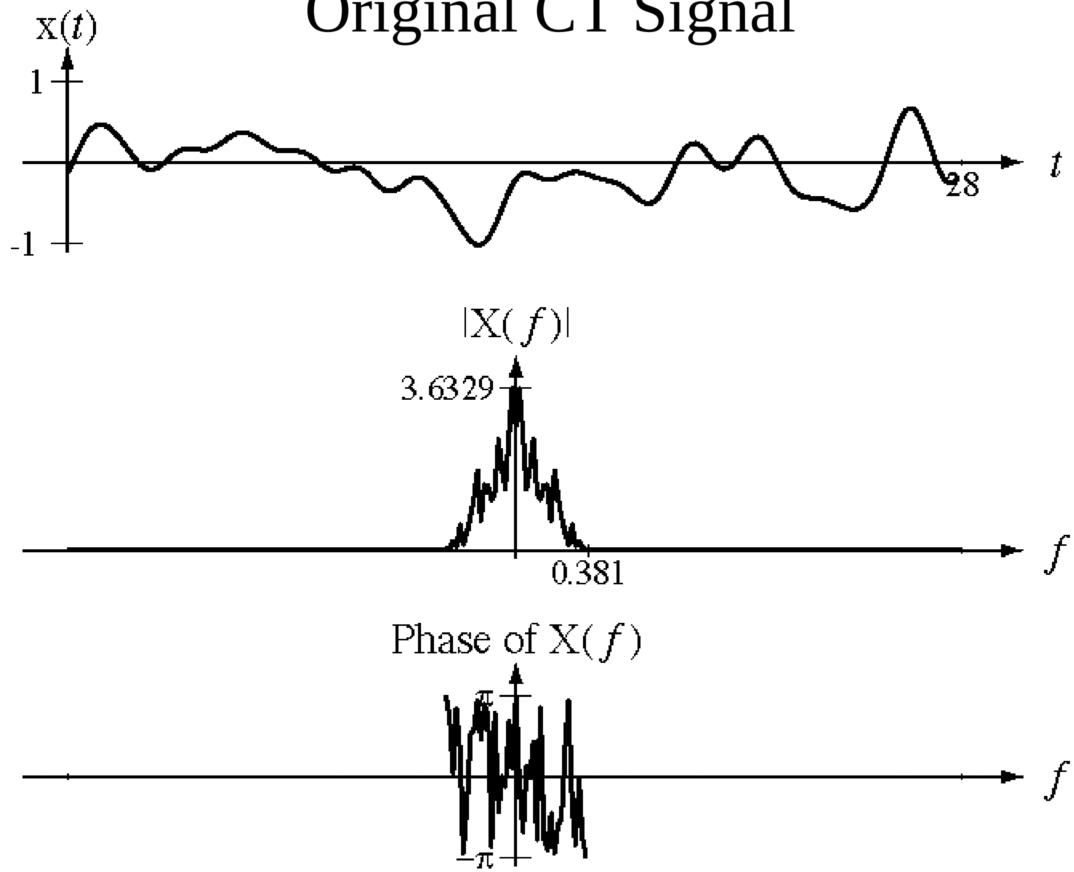

The Discrete Fourier Transform

The sampled signal is

$$x_s[n] = x(nT_s)$$

and its DTFT is

$$X_s(F) = f_s \sum_{n=-\infty}^{\infty} X(f_s(F-n))$$

Samples from Original Signal

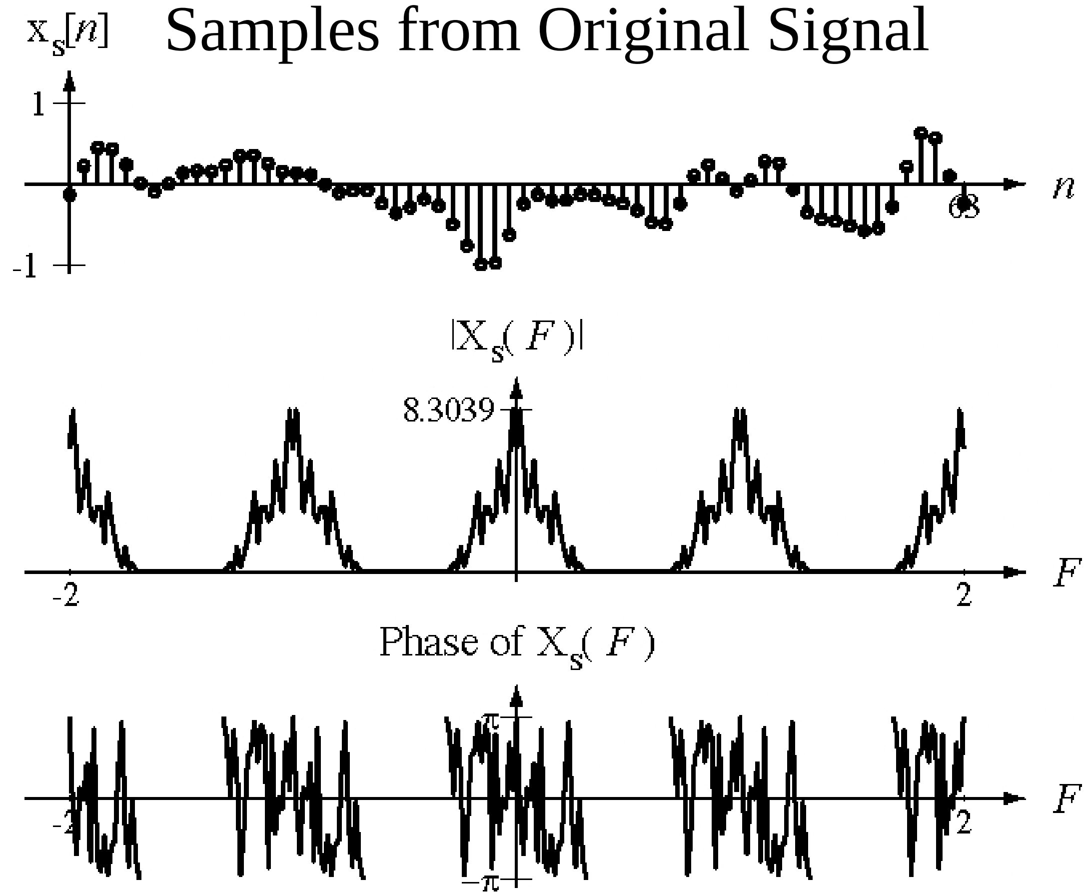

The Discrete Fourier Transform

Only N_F samples are taken. If the first sample is taken at time, $t = 0$ (the usual assumption) that is equivalent to multiplying the sampled signal by the *window* function,

$$\mathrm{w}[n] = \begin{cases} 1 \ , & 0 \le n < N_F \\ 0 \ , & \text{otherwise} \end{cases}$$

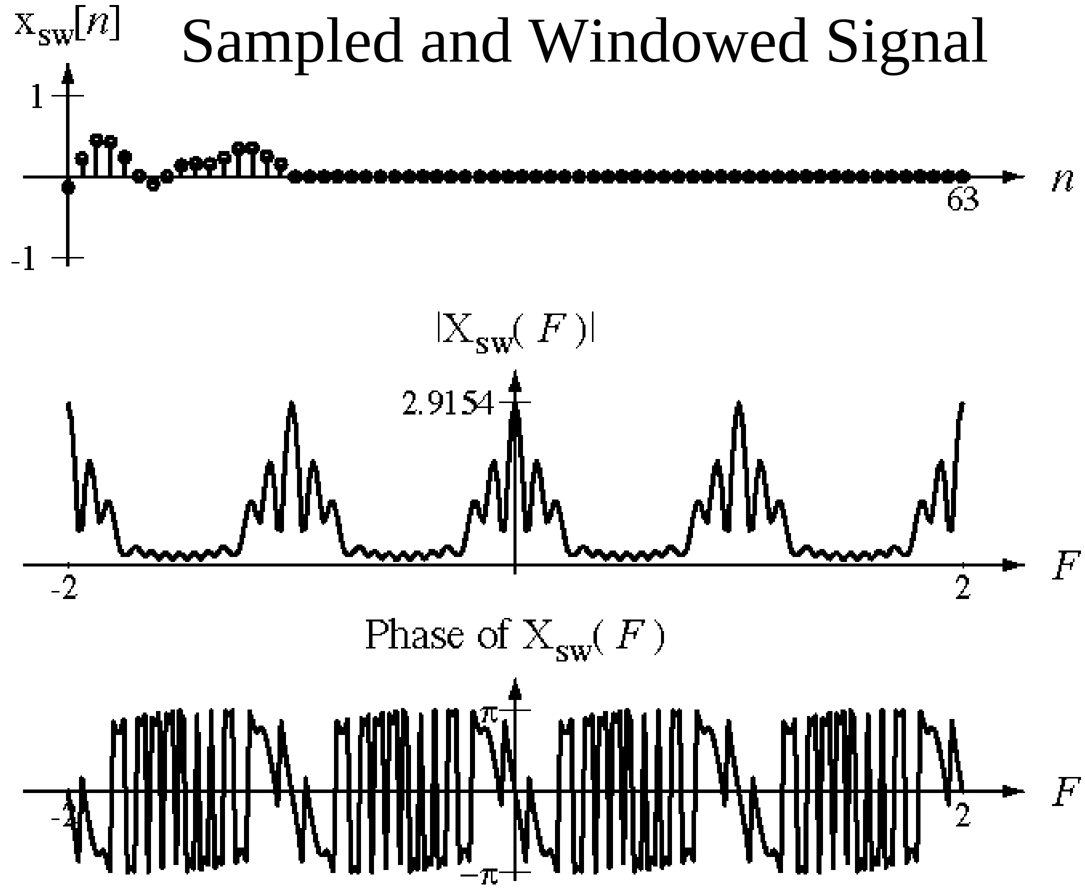

The Discrete Fourier Transform

The last step in the process is to sample the frequency-domain signal which periodically repeats the time-domain signal. Then there are two periodic impulse signals which are related to each other through the DTFS. Multiplication of the DTFS harmonic function by the number of samples in one period yields the DFT.

Sampled, Windowed and Periodically-Repeated Signal

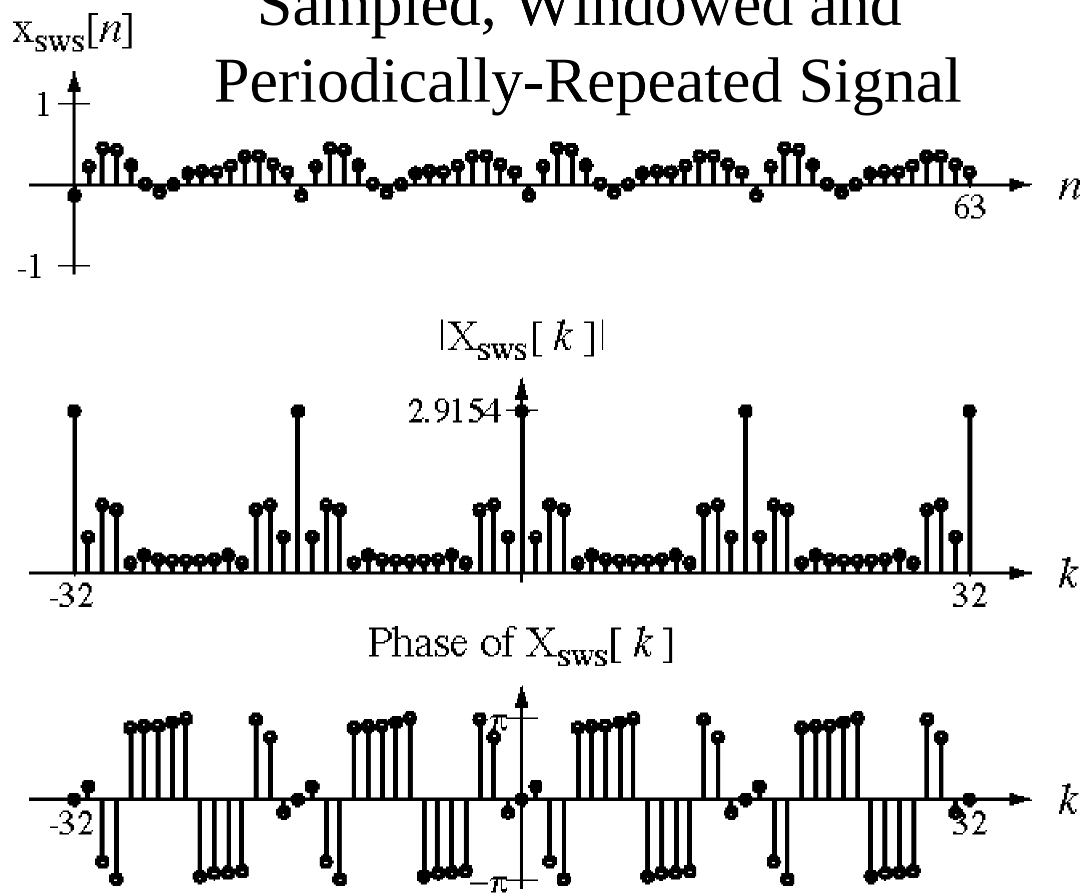

The Discrete Fourier Transform

The original signal and the final signal are related by

$$\mathrm{X}_{sws}[k] = \frac{f_s}{N_F}\Big[\underbrace{e^{-j\pi F(N_F-1)} N_F \operatorname{drcl}(F, N_F)}_{\mathrm{W}(F)} * \mathrm{X}(f_s F)\Big]_{F \to \frac{k}{N_F}}$$

In words, the CTFT of the original signal is transformed by replacing f with $f_s F$. That result is convolved with the DTFT of the window function. Then that result is transformed by replacing F by $\frac{k}{N_F}$. Then that result is multiplied by $\frac{f_s}{N_F}$.

The Discrete Fourier Transform

It can be shown (pp. 530-532) that the DFT can be used to approximate samples from the CTFT. If the signal, x(*t*), is an energy signal and is causal and if N_F samples are taken from it over a finite time beginning at time, *t* = 0, at a rate, f_s, then the relationship between the CTFT of x(*t*) and the DFT of the samples taken from it is

$$\mathrm{X}(kf_F) \cong T_s e^{-j\frac{\pi k}{N_F}} \operatorname{sinc}\left(\frac{k}{N_F}\right) \mathrm{X}_{DFT}[k]$$

where $f_F = \dfrac{f_s}{N_F}$. For those harmonic numbers, *k*, for which $k << N_F$

$$\mathrm{X}(kf_F) \cong T_s\, \mathrm{X}_{DFT}[k]$$

As the sampling rate and number of samples are increased, this approximation is improved.

The Discrete Fourier Transform

If a signal, x(t), is bandlimited and periodic and is sampled above the Nyquist rate over exactly one fundamental period the relationship between the CTFS of the original signal and the DFT of the samples is (pp. 532-535)

$$\mathrm{X}_{DFT}[k] = N_F\,\mathrm{X}_{CTFS}[k] * \mathrm{comb}_{N_F}[k]$$

That is, the DFT is a periodically-repeated version of the CTFS, scaled by the number of samples. So the set of impulse strengths in the base period of the DFT, divided by the number of samples, is the same set of numbers as the strengths of the CTFS impulses.

The Fast Fourier Transform

Probably the most used computer algorithm in signal processing is the *fast Fourier transform (fft)*. It is an efficient algorithm for computing the DFT. Consider a very simple case, a set of four samples from which to compute a DFT. The DFT formula is

$$\mathrm{X}[k] = \sum_{n=0}^{N_F-1} \mathrm{x}[n] e^{-j2\pi \frac{kn}{N_F}}$$

It is convenient to use the notation, $W = e^{-j\frac{2\pi}{N_F}}$, because then the DFT formula can be written as

$$\begin{bmatrix} \mathrm{X}[0] \\ \mathrm{X}[1] \\ \mathrm{X}[2] \\ \mathrm{X}[3] \end{bmatrix} = \begin{bmatrix} W^0 & W^0 & W^0 & W^0 \\ W^0 & W^1 & W^2 & W^3 \\ W^0 & W^2 & W^4 & W^6 \\ W^0 & W^3 & W^6 & W^9 \end{bmatrix} \begin{bmatrix} \mathrm{x}_0[0] \\ \mathrm{x}_0[1] \\ \mathrm{x}_0[2] \\ \mathrm{x}_0[3] \end{bmatrix}$$

The Fast Fourier Transform

The matrix multiplication requires N^2 complex multiplications and $N(N - 1)$ complex additions. The matrix product can be re-written in the form,

$$\begin{bmatrix} X[0] \\ X[1] \\ X[2] \\ X[3] \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & W^1 & W^2 & W^3 \\ 1 & W^2 & W^0 & W^2 \\ 1 & W^3 & W^2 & W^1 \end{bmatrix} \begin{bmatrix} x_0[0] \\ x_0[1] \\ x_0[2] \\ x_0[3] \end{bmatrix}$$

because $W^n = W^{n+mN_F}$, m an integer.

The Fast Fourier Transform

It is possible to factor the matrix into the product of two matrices.

$$\begin{bmatrix} X[0] \\ X[2] \\ X[1] \\ X[3] \end{bmatrix} = \begin{bmatrix} 1 & W^0 & 0 & 0 \\ 1 & W^2 & 0 & 0 \\ 0 & 0 & 1 & W^1 \\ 0 & 0 & 1 & W^3 \end{bmatrix} \begin{bmatrix} 1 & 0 & W^0 & 0 \\ 0 & 1 & 0 & W^0 \\ 1 & 0 & W^2 & 0 \\ 0 & 1 & 0 & W^2 \end{bmatrix} \begin{bmatrix} x_0[0] \\ x_0[1] \\ x_0[2] \\ x_0[3] \end{bmatrix}$$

It can be shown (pp. 552-553) that 4 multiplications and 12 additions are required, compared with 16 multiplications and 12 additions using the original matrix multiplication.

The Fast Fourier Transform

It is helpful to view the fft algorithm in *signal-flow graph* form.

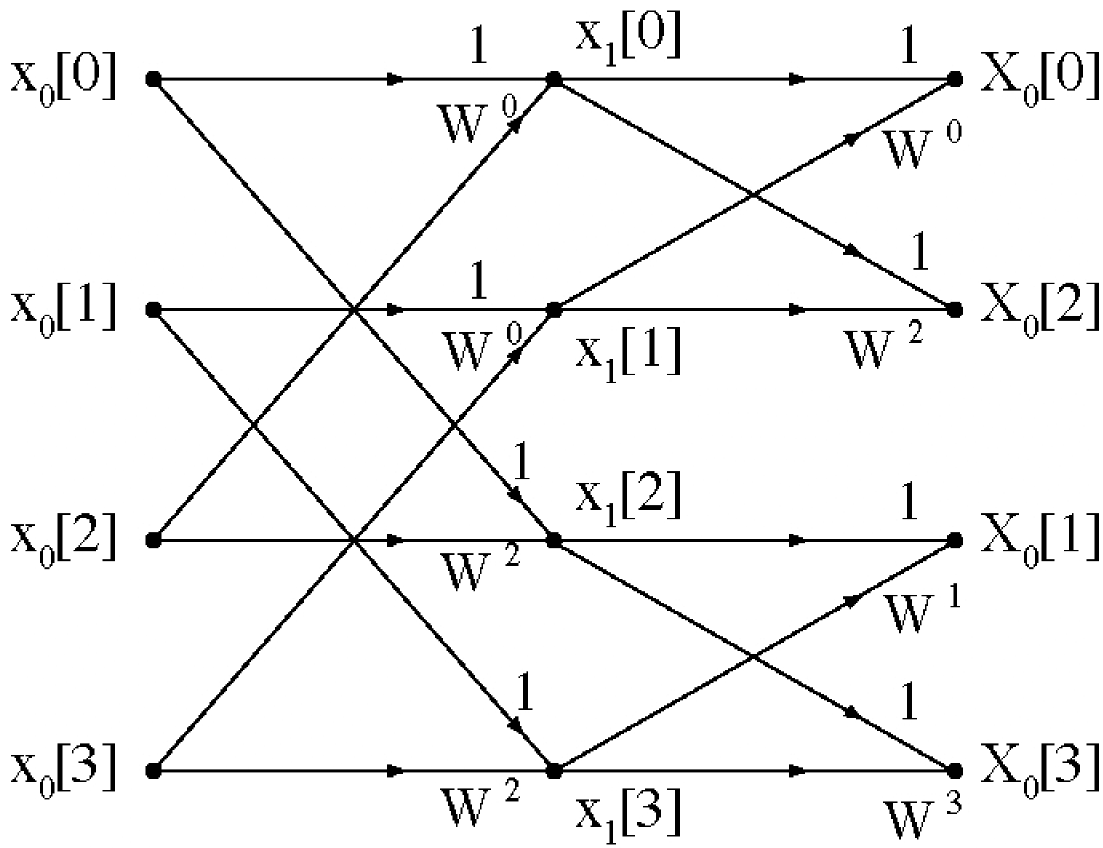

The Fast Fourier Transform

16-Point
Signal-Flow
Graph

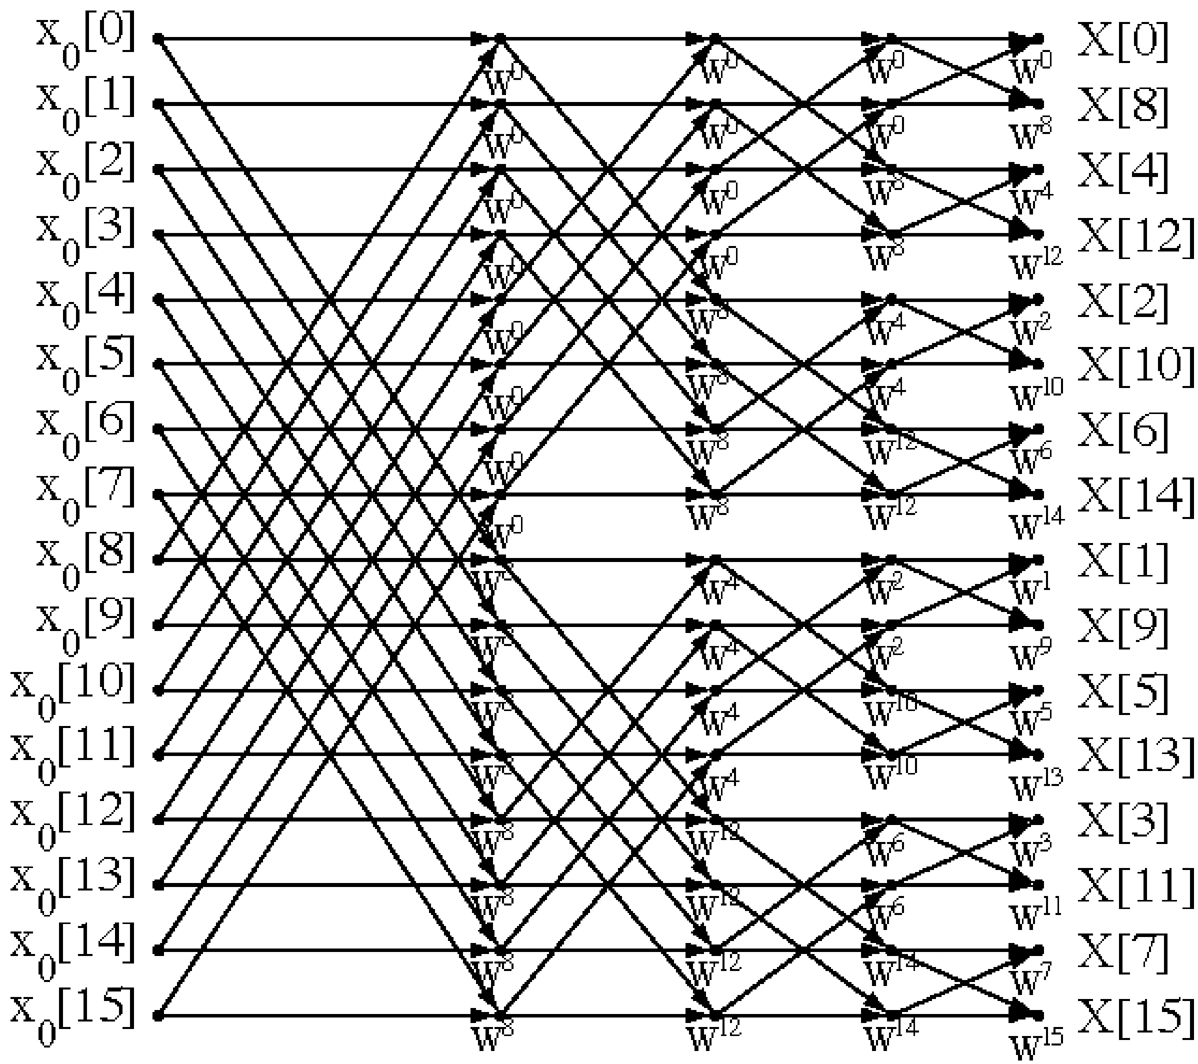

The Fast Fourier Transform

The number of multiplications required for an fft algorithm of length, $N = 2^p$, where p is an integer is $\frac{2N}{p}$. The speed ratio in comparison with the direct DFT algorithm is $\frac{Np}{2}$.

p	N	Speed Ratio FFT/DFT
2	4	4
4	16	8
8	256	64
16	65536	8192