Sampling Signals

Overview:

- We use the Fourier transform to understand the discrete sampling and re-sampling of signals.
- One key question is when does sampling or re-sampling provide an adequate representation of the original signal?

Terminology:

- sampling – creating a discrete signal from a continuous process.
- downsampling (decimation) – subsampling a discrete signal
- upsampling – introducing zeros between samples to create a longer signal
- aliasing – when sampling or downsampling, two signals have same sampled representation but differ between sample locations.

Matlab Tutorials: samplingTutorial.m, upSample.m

Sampling Signals

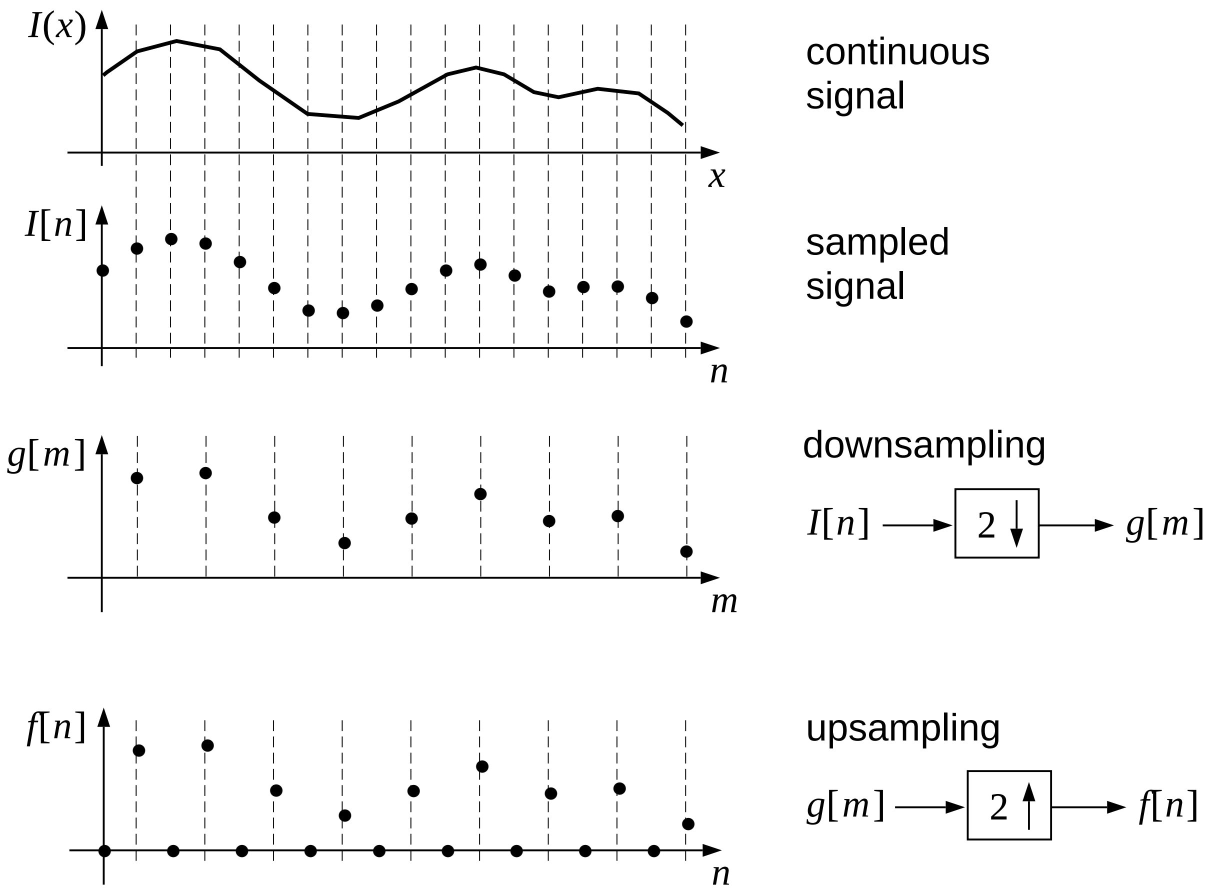

Up-Sampling

Consider up-sampling a signal $I[n]$ of length N:

- Increase number of samples N by a factor of n_s.
- **Step 1.** Place $n_s - 1$ zeros after every sample of $I[n]$, to form $f_0[n]$ of length $n_s N$, namely

$$f_0[n] = \begin{cases} I[n/n_s] & \text{for} \quad n \bmod n_s = 0, \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

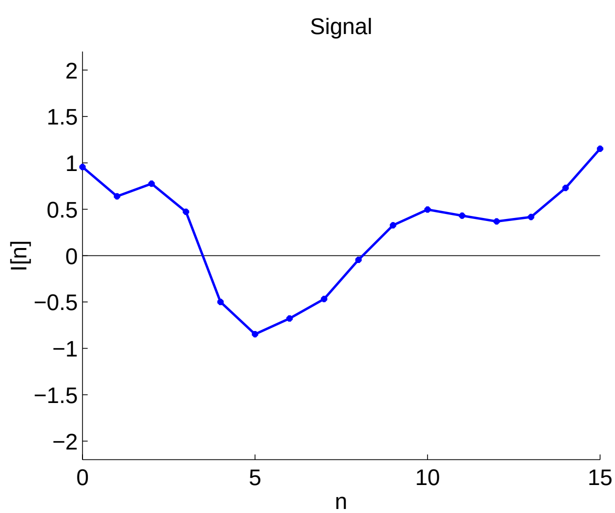

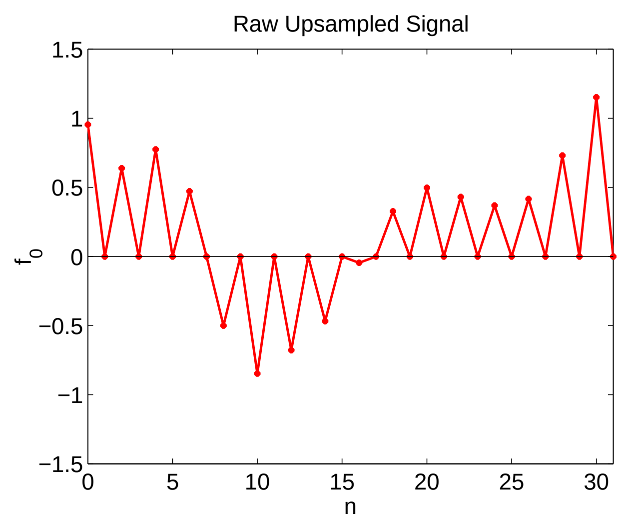

Up-Sampling (Cont.)

- **Step 2.** Interpolate signal f_0:

$$f = S * f_0, \quad \text{for some smoothing kernel } S[n]. \tag{2}$$

  Here, in order for $f[n]$ to interpolate $f_0[n]$, we require that $S[0] = 1$ and $S[jn_s] = 0$ for nonzero integers j.

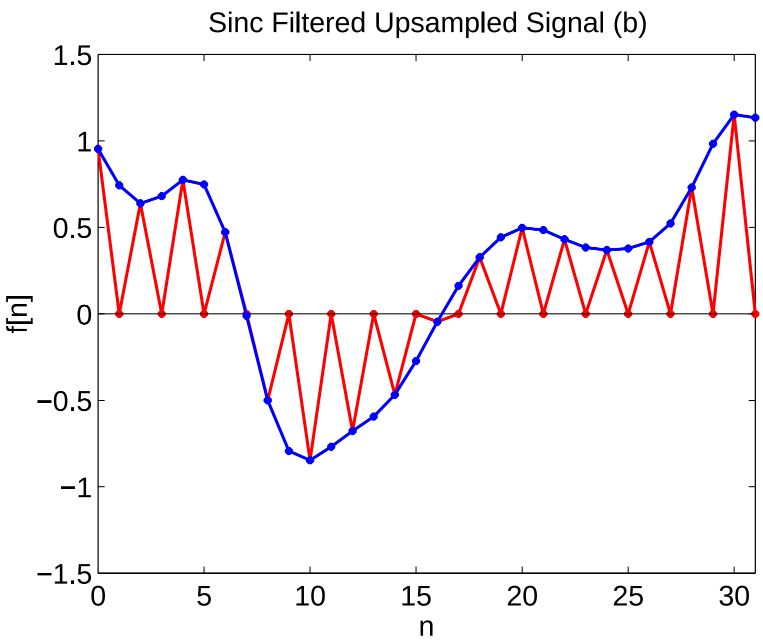

- Many smoothing kernels can be used (see upSample.m).

Frequency Analysis of Up-Sampling

Step 1. Fourier transform of the raw up-sampled signal $f_0[n]$:

$$\begin{aligned}\mathcal{F}(f_0[n])[k] &\equiv \sum_{n=0}^{Nn_s-1} f_0[n]e^{i\omega_k^s n} = \sum_{j=0}^{N-1} f_0[jn_s]e^{i\omega_k^s jn_s} \\ &= \sum_{j=0}^{N-1} I[j]e^{i\omega_k^s jn_s}\end{aligned}$$

Here $\omega_k^s = \frac{2\pi}{Nn_s}k$, so the last line above becomes

$$\mathcal{F}(f_0)[k] = \sum_{j=0}^{N-1} I[j]e^{i\frac{2\pi}{N}kj} \equiv \mathcal{F}(I)[k], \quad \text{for } -\frac{Nn_s}{2} \leq k < \frac{Nn_s}{2}. \quad (3)$$

Since $\mathcal{F}(I)$ is N-periodic, equation (3) implies that $\mathcal{F}(f_0)$ consists of n_s copies of $\mathcal{F}(I)$ concatenated together.

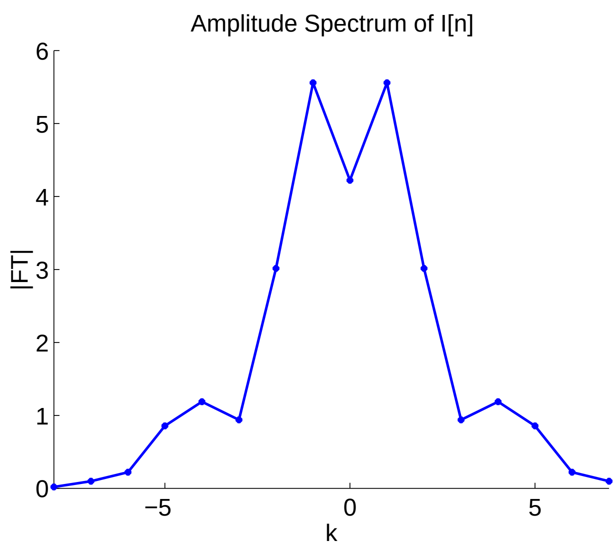

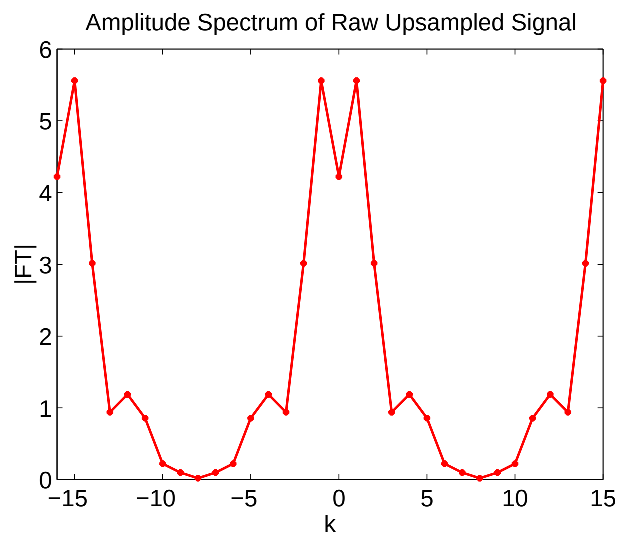

Fourier Analysis of Up-Sampling Step 2

Recall Step 2 is to form $f[n] = S * f_0$, for some interpolation filter S. However, notice from the inverse Fourier transform that (for N even)

$$
\begin{aligned}
I[j] &= \frac{1}{N} \sum_{k=-N/2}^{N/2-1} \hat{I}[k] e^{i\frac{2\pi}{N}kj} \\
&= \frac{n_s}{N n_s} \sum_{k=-N/2}^{N/2-1} \hat{f}_0[k] e^{i\frac{2\pi}{N n_s}kjn_s} = f[jn_s]. \qquad (4)
\end{aligned}
$$

Here we used $f(jn_s) = I(j)$ in the last line. Notice the left term in the last line above is n_s times the inverse Fourier transform of $B[k]\hat{f}_0[k]$ where B is the box function, $B[k] = 1$ for $-N/2 \le k < N/2$ and $B[k] = 0$ otherwise. We can therefore evaluate this inverse Fourier transform at every pixel n, and not just at the interpolation values jn_s, to construct a possible interpolating function $f[n]$

$$f[n] = n_s \mathcal{F}^{-1}(B(k)\hat{f}_0[k]). \qquad (5)$$

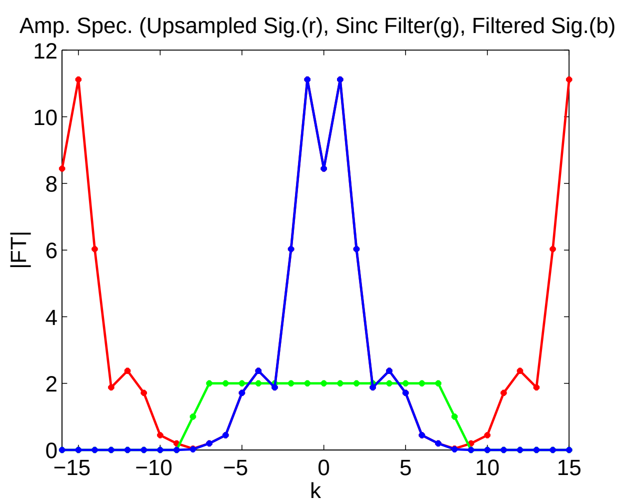

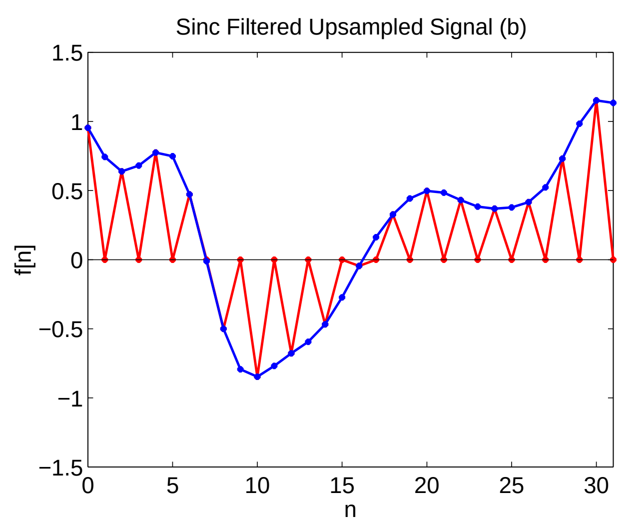

Down-Sampling

Consider down-sampling a signal $I[n]$ of length N:

- Reduce number of samples N by a factor of n_s, where n_s is a divisor of N.
- Define the comb function:

$$C(n; n_s) = \sum_{m=0}^{(N/n_s)-1} \delta_{n,mn_s}$$

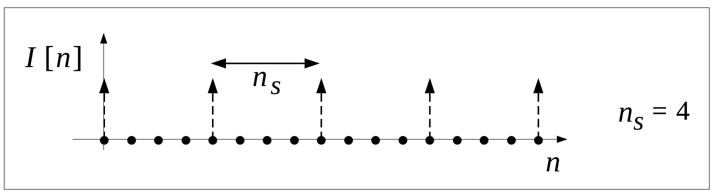

- **Step 1.** Introduce zeros in $I[n]$ at unwanted samples.

$$g_0[n] = C[n; n_s]I[n]. \tag{6}$$

- **Step 2.** Downsample signal g_0:

$$g[m] = g_0[mn_s], \quad \text{for } 0 \leq m < N/n_s. \tag{7}$$

Frequency Domain Analysis of Down-Sampling

Proposition 1. The Fourier transform of the comb function is another comb function:

$$\mathcal{F}(C[n; n_s]) = \frac{N}{n_s} C[k; N/n_s]. \tag{8}$$

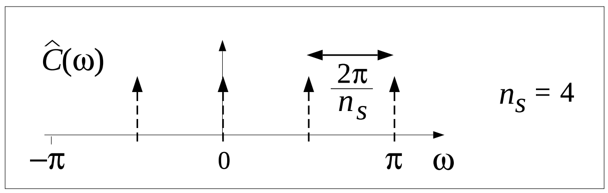

Recall the frequency ω and wave number k are related by $\omega = \frac{2\pi}{N}k$. So the spacing in the plot above is $\omega_s = \frac{2\pi}{N}\frac{N}{n_s} = \frac{2\pi}{n_s}$.

Proposition 2. Pointwise product and convolution of Fourier transforms. Suppose $f[n]$ and $g[n]$ are two signals of length N (extended to be N-periodic). Then

$$\begin{aligned} \mathcal{F}(f[n]g[n]) &= \frac{1}{N}\mathcal{F}(f) * \mathcal{F}(g) \\ &\equiv \frac{1}{N} \sum_{j=-N/2}^{N/2-1} \hat{f}[j]\, \hat{g}[k-j]. \end{aligned} \tag{9}$$

where $\hat{f}$ and $\hat{g}$ denote the Fourier transforms of f and g, respectively.

The proofs of these two propositions are straight forward applications of the definition of the Fourier transform given in the preceeding notes, and are left as exercises.

Fourier Analysis of Down-Sampling Step 1

Recall Step 1 is to form $g_0[n] = C[n; n_s]I[n]$. By Prop. 1 and 2 above, we have

$$\begin{aligned}
\mathcal{F}(g_0) &= \mathcal{F}(C[n; n_s]I[n]) \\
&= \frac{N}{n_s}\frac{1}{N}C[\cdot; N/n_s] * \hat{I}[\cdot] \\
&\equiv \frac{1}{n_s}\sum_{j=-N/2+1}^{N/2} C[j; N/n_s]\hat{I}[k-j] \\
&= \frac{1}{n_s}\sum_{r=-r_0+1}^{n_s-r_0} \hat{I}[k - r\frac{N}{n_s}],
\end{aligned} \tag{10}$$

where, due to the periodicity of $\hat{I}[k]$, we can use any integer r_0 (eg. $r_0 = n_s/2$ for even n_s).

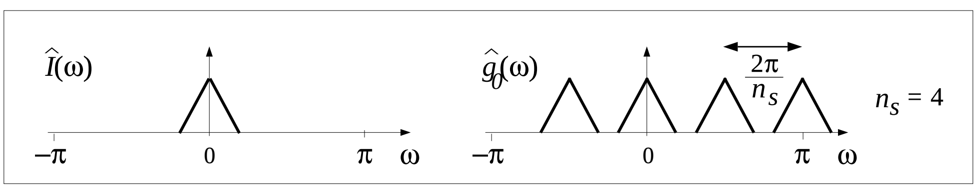

Therefore $\hat{g}_0[k]$ consists of the sum of replicas $\hat{I}[k - rN/n_s]$ of the Fourier transform of the original signal I, spaced by wavenumber N/n_s or, equivalently, by frequency $\omega_s = \frac{2\pi}{n_s}$.

Note the Fourier transform $\hat{g}_0$ has period 2π, so the contribution sketched above for $\omega \geq \pi$ can be shifted 2π to the left.

Fourier Analysis of Down-Sampling Step 2

In Step 2 we simply drop the samples from $g_0[n]$ which were set to zero by the comb function $C[n; n_s]$. That is

$$g[m] = g_0[mn_s], \quad \text{for } 0 \le m < N/n_s.$$

In terms of the Fourier transform, it is easy to show

$$\mathcal{F}(g)[k] = \hat{g}[k] = \hat{g}_0[k], \quad \text{for } -N/(2n_s) \le k < N/(2n_s). \tag{11}$$

Rewriting this in terms of the frequency $\omega_{s,k} = \frac{2\pi}{(N/n_s)}k$ (note $g[m]$ is a signal of length N/n_s), and the corresponding frequency $\omega_k = \frac{2\pi}{N}k = \omega_{s,k}n_s$ of the longer signal $g_0[n]$, we have

$$\hat{g}[\omega_{s,k}] = \hat{g}_0[\omega_k] = \hat{g}_0[\omega_{s,k}/n_s], \quad \text{for } -\pi \le \omega_{s,k} < \pi. \tag{12}$$

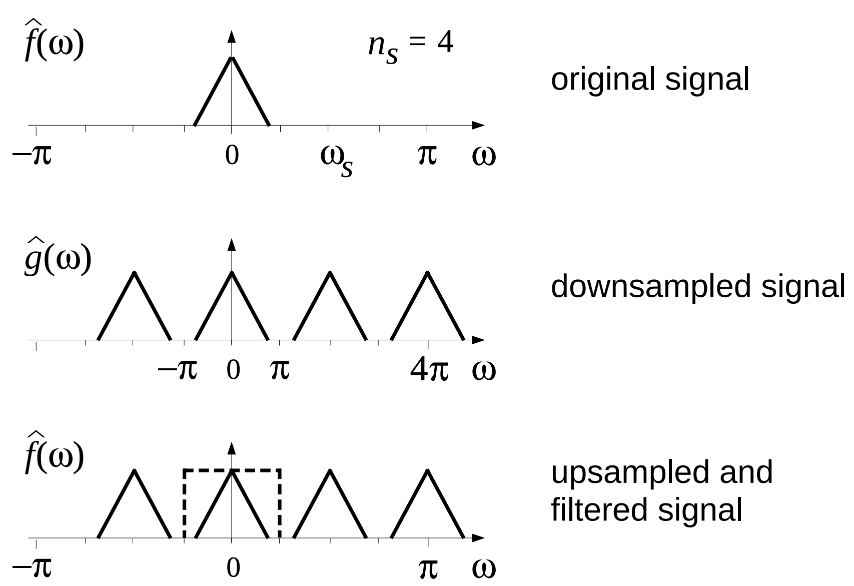

Nyquist Sampling Theorem

Sampling Theorem: Let $f[n]$ be a band-limited signal such that

$$\hat{f}[\omega] \; = \; 0 \quad \text{for all} \;\; |\omega| > \omega_0$$

for some ω_0. Then $f[n]$ is uniquely determined by its samples $g[m] = f[m\, n_s]$ when

$$\omega_s/2 = \frac{\pi}{n_s} > \omega_0 \qquad \text{or equivalently} \quad n_s < \frac{\lambda_0}{2}$$

where $\lambda_0 \; = \; 2\pi/\omega_0$. In words, the distance between samples must be smaller than half a wavelength of the highest frequency in the signal.

In terms of the previous figure, note that the maximum frequency ω_0 must be smaller than one-half of the spacing, ω_s, between the replicas introduced by the sampling. This ensures the replicas **do not overlap**.

When the replicas do not overlap, we can up-sample the signal $g[m]$ and interpolate it to recover the signal $f[n]$, as discussed above.

Otherwise, when the replicas overlap, the Fourier transform $\hat{g}[k]$ contains contributions from more than one replica of $\hat{f}[k]$. Due to these aliased contributions, we cannot then recover the original signal $f[n]$.

Sampling Continuous Signals

A similar theorem holds for sampling signals $f(x)$ for $x \in [0, L)$. We can represent f as the Fourier series

$$f(x) =_{a.e.} \sum_{k=-\infty}^{\infty} \hat{f}[k] e^{i\omega_k x},$$

where $\omega_k = \frac{2\pi}{L} k$ and $=_{a.e.}$ denotes equals almost everywhere. Suppose $f(x)$ is band-limited so that for some $\omega_0 > 0$

$$\hat{f}[k] \ = \ 0 \quad \text{for all } \ |\omega_k| > \omega_0.$$

Then $f(x)$ is uniquely determined by its samples $I[m] = f(m\,\tau)$ when

$$\tau < \frac{\lambda_0}{2} \tag{13}$$

where $\lambda_0 = 2\pi/\omega_0$. In words, the distance between samples must be smaller than half the wavelength of the highest frequency in the signal.

The link with the preceeding analysis is that sampling $f(x)$ with a sample spacing of τ causes replicas in the Fourier transform to appear with spacing $\omega_s = 2\pi/\tau$. As before, the condition that these replicas do not overlap is $\omega_0 < \omega_s/2$, which is equivalent to condition (13).

Aliasing

Aliasing occurs when replicas overlap:

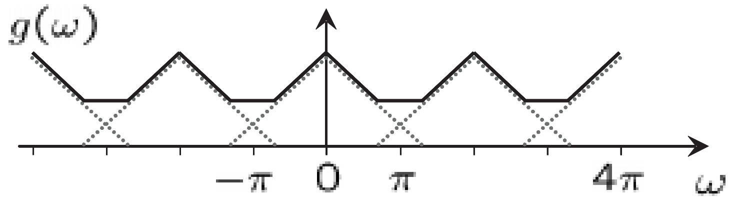

Consider a perspective image of an infinite checkerboard. The signal is dominated by high frequencies in the image near the horizon. Properly designed cameras blur the signal before sampling, using

- the point spread function due to diffraction,
- imperfect focus,
- averaging the signal over each CCD element.

These operations attenuate high frequency components in the signal. Without this (physical) preprocessing, the sampled image can be severely aliased (corrupted):

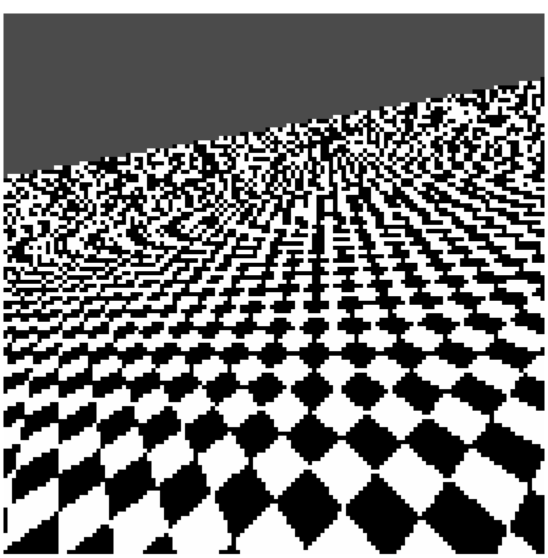

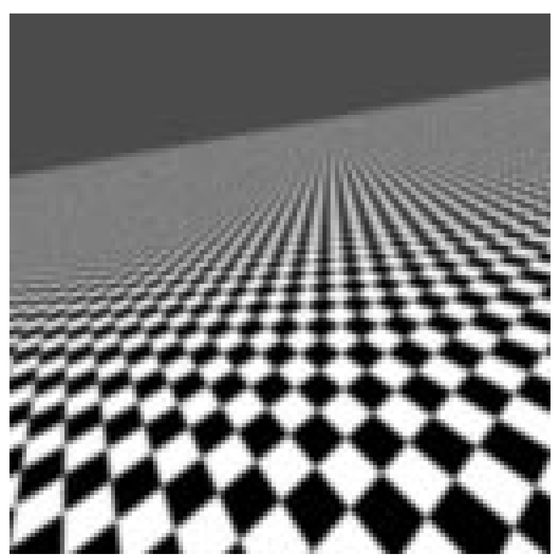

Dimensionality

A guiding principal throughout signal transforms, sampling, and aliasing is the underlying dimension of the signal, that is, the number of linearly independent degress of freedom (dof). This helps clarify many issues that might otherwise appear mysterious.

- Real-valued signals with N samples have N dof. We need a basis of dimension N to represent them uniquely.
- Why did the DFT of a signal of length N use N sinusoids? Because N sinusoids are linearly independent, providing a minimal spanning set for signals of length N. We need no more than N.
- But wait: Fourier coefficients are complex-valued, and therefore have $2N$ dofs. This matches the dof needed for complex signals of length N but not real-valued signals. For real signals the Fourier spectra are symmetric, so we keep half of the coefficients.
- When we down-sample a signal by a factor of two we are moving to a basis with $N/2$ dimensions. The Nyquist theorem says that the original signal should lie in an $N/2$ dimensional space before you down-sample. Otherwise information is corrupted (i.e. signal structure in multiple dimensions of the original N-D space appear the same in the $N/2$-D space).
- The Nyquist theorem is not primarily about highest frequencies and bandwidth. The issue is really one of having a model for the signal; that is, how many non-zero frequency components are in the signal (i.e., the dofs), and which frequencies are they.