ECEU692 Intro to SSI Course Notes

Part 4: Some selected topics in signal and image processing for SSI (1)

Profs. DiMarzio and Brooks
Northeastern University

Spring 2004

Outline of those Selected Topics

- Impulse responses and convolution (review)
- Fourier Transforms and frequency content (review)
- Convolution / Multiplication properties (review)
- Sampling and quantization and their consequences (review)
- Anti-aliasing filtering (review)
- 2-dimensional signals and all of the above
- The Projection-Slice theorem and tomography: basic idea
- What does Fourier phase tell us ?
- Power spectra and autocorrelations
- Resolution and Extent
  - apodization and its dual in time
  - the time-bandwidth product
- Relationships among varieties of Fourier representations
- Hyperspectral Images as signals

Some conceptual stuff to put aside (for now . . .)

- Discrete or continuous signals ?
- 1-D or 2-D or . . . ?

	Discrete	Continuous
1-D	$x(n)$	$x(t)$
2-D	$f(m, n)$	f(x,y)

- Discrete and continous are conceptually very similar in how we treat them
- 1-D and 2-D are also conceptually very similar in how we treat them
- For our purposes, similarities outweigh differences
- We'll start with:
  - Discrete signals but jump back and forth as appropriate
  - 1-D, then move to 2-D

Impulse Response and Convolution

Let's review the 1-D discrete case

- Systems are generally time-dispersive (they can't respond instantaneously)
- We model this as an "impulse response", the response to a very short input
- In discrete time this is the response (denoted $h(n)$ to an input $\delta(n)$ that is non-zero only at $n = 0$
- By superposition we can find the resposne to any other input as a weighted sum of impulse response samples

  $$y(n) = \sum_k h(k)x(n-k)$$

  where the weights come from the inputs. This is the **Convolution Sum**
- We move in opposite directions in input and impulse response inside the sum because **later** parts of response are weighted by **earlier** inputs

Note that no limits on a sum or integral means (here) that it runs from $-\infty$ to ∞

Frequency Content of Signals

Some other examples besides light through a prism . . .

- Why do two instruments playing the same note (say, a violin and a clarinet) sound different ?
- Why can you immediately differentiate between the identity of speakers even if their voices have the same pitch ?
- One major reason: even though the "fundamental frequency" is the same in both cases (the "note" or the "pitch"), there are other (higher harmonic) frequencies present that differ between instruments or people
- Next question: how can you select WBCN when radio waves from WGBH are also hitting your radio receiver at the same time?
- Each station "moves" the frequency content of its music to center around its assigned "broadcast frequency", and so that it does not overlap the frequencies used by other stations
- You select one frequency band to pass and filter out others.
- Bottom line here: The distribution of energy (and phase) with frequency is an extremely useful way to describe many signals

Frequency Response of a System

- The Fourier Transform of an impulse response is called the Frequency Response
- Why? It turns out that it completely describes how a system responds in the steady-state to sinusoidal inputs at each frequency:
  - The magnitude of the frequency response at a given frequency describes the gain or attentuation
  - The angle of the frequency response at a given frequency describes the delay or advance
- For 1-D discrete signals, if $H(\omega)$ is the Discrete Time Fourier Transform (DTFT) of the impulse reponse $h(n)$, then
  - If the input is $A\cos(\omega_0 n + \theta_0)$
  - Then the output is $(|H(\omega_0)|A)\cos(\omega_0 n + (\theta_0 + \angle(H(\omega_0)))$
  - Note that the **frequency** does not change !!!

Convolution and Multiplication

- Applying the Discrete Time Fourier Transform to the Convolution Sum we get that if
$$y(n) = h(n) * x(n)$$
then
$$Y(\omega) = H(\omega)X(\omega).$$
- So the frequency response determines which frequencies in the input are amplified or attenuated, and how much each frequency component is delayed.
- Provides an analysis tool: multiplication is simpler than convolution
- Provides a sythesis tool: describe how you *want* a system to behave in the frequency domain and then design a filter to do it

The Sampling Theorem: a Review

- Shannon sez: To get all the information in a signal and avoid causing distortion (aliasing), you need to sample at least twice as fast as the highest frequency present in the signal
- A simple interpretation: first consider a bar of length one performing uniform circular motion around a circle. How fast do we need to sample it to be able to know for sure where it's been all the time?

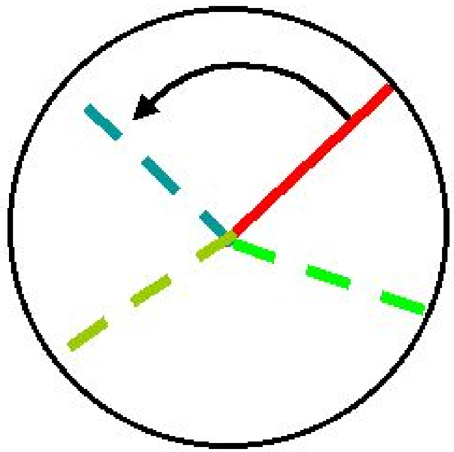

So What's the Answer?

- The first response is generally: once per cycle
- However this cannot distinguish opposite directions: if we sample once per second and at each sample the bar has moved $\pi/4$ counterclockwise
  - it could be moving counterclockwise at $1/8$ cycles per second, or
  - it could be moving **clockwise** at $7/8$ cycles per second.
- Thus we need to sample at least **twice** per cycle.
- If we don't we get ambiguity, called **ALIASING**: for instance, if the bar is going clockwise at $7/8$, or counterclockwise at $9/8$, or clockwise at $15/8$, or counterclockwise at $17/8$, etc., cycles per second, it all looks like $1/8$ cycle per second from the samples

Let's Generalize:

- We can write the bar's movement as $e^{j\Omega t}$
- If we consider a superposition of "bars" of different lengths $X(\Omega)$ moving at different (infinitely close) frequencies we get

$$\int X(\Omega)e^{j\Omega t}d\Omega$$

- This is the Fourier Transform of a signal $x(t)$.
- And this leads directly to Shannon's Theorem.
- Note: if the signal is *band-limited*, to a frequency band $\Omega < B$, this means that $X(\Omega) = 0$ for $\Omega \geq B$.
- In this case, the Shannon rule says we need to sample at least as fast as $2B/2\pi$ samples per second.

A Simple Example

- If we are imaging blood vessels in skin, at 2 alternating wavelengths,
- We image at 30 wavelengths, and each image takes 1/ 30 second,
- The subject's pulse rate is exactly 60 beats / minute (1 Hz), then,
  - one wavelength might capture an image only when the arterial blood is at a maximum, and then
  - another would see it when it is at a minimum.
  - others might see it at various points in-between . . .
  - But you certainly won't see any variation due to the pulse rate.

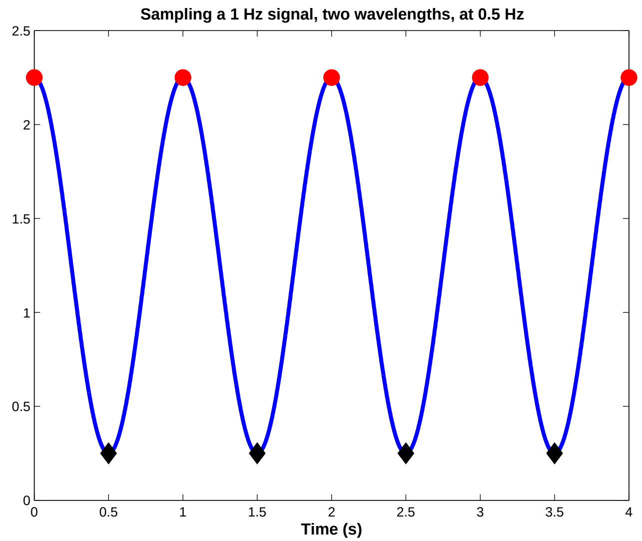

Continuing the Example

- If the subject's pulse rate is 70 beats / minute ($\approx \mathbf{1.2}$ Hz), then
  - This aliases down to about $\mathbf{0.2}$Hz, or every 5 frames, (remember the rotating bar) so
  - we will see a slow variation at each wavelength due to pulse rate

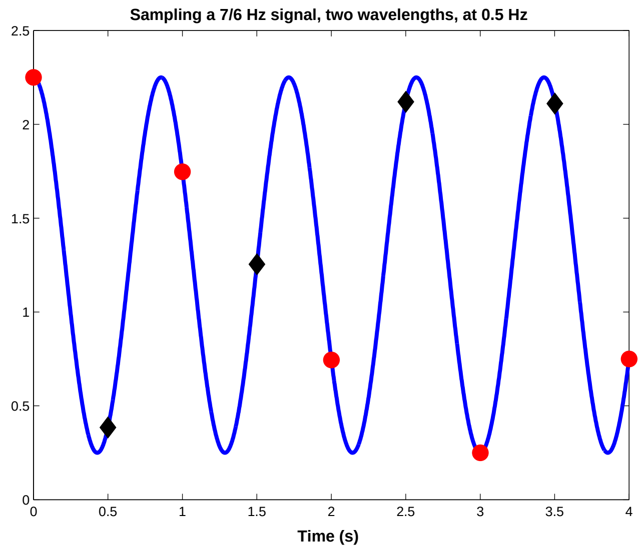

- Bottom line: if you want to see the pulse in each wavelength, you need to sample at least 2 times/pulse/wavelength for the fastest pulse you expect to see, or take fewer wavelengths, or some combination.

Practical Sampling: Anti-Aliasing Filtering

- Some practical considerations
  - No signal is truly band-limited (more on this later)
  - We generally have high-frequency noise
  - Samplers have a finite (non-instantaneous) time aperture (eg they integrate over some time period, like the CCD camera for instance)
- So:
  1. We often want to low-pass filter a signal before sampling it (called anti-aliasing filtering) to get its frequency content below the Shannon limit,
  2. Integration may help, as it's a kind of low-pass filter (details on how this works to be posted on the web site)
  3. The tradeoff is we lose *resolution* in time: everything within the integration interval gets smeared together.
  4. Sometimes we can just *integrate* in time (or space) to do this filtering

How about "Spatial Signals"?

- First, let's think about a signal as a function of two spatial variables, $f(x, y)$, or a discrete version $f(m, n)$.
  - (x, y) (or (m, n) is the location in space.
  - f is the value at that location.
  - usual convention: first variable is horizontal, positive to the right, and second is vertical, positive going up.
- We can now think about filtering, for instance:
  - Low-pass in both directions, or
  - High-pass in both directions, or
  - Low-pass in one direction, High-pass in other, or
  - Even more complicated, even non-separable filters: for example, a circularly-symmetric low-pass filter.
- But what does "low-frequency", for example, mean here ?

Reminder: Sampling on the CCD

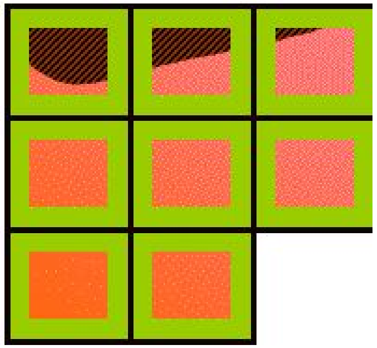

- The CCD camer samples this into pixels, $f(m,n)$ by integrating over a small region on the chip.
- How do we analyze this sampling process?
  The tool we know to analyze how sampling works is "frequency content" — what does this mean for a function of *space*?

Filtering in 2D

- In the general case we have a 2D "impulse response" $h(x, y)$ (or in the discrete case $h(m, n)$, and a 2D definition of convolution (or filtering), so that the output $g(x, y)$relates to the input $f(x, y)$ by

$$g(x, y) = h(x, y) * *f(x, y).$$

- The discrete case looks just the same.
- $h(x, y)$ or $h(m, n)$ are usually called the "point spread function", 2D version of impulse reponse.
- Why? Well, if input is an "impulse", that is concentrated at a single point, then the output will be $h(x, y)$, and wil show how the system "spreads" out the impulse.
- It's a very common and important way to characterize any imaging system — related to its spatial resolution.
- Convolution involves shifting $h(m, n)$ over the image, multiplying, and adding (or integrating for $h(x, y)$.)

Discrete Convolution in 2D

- Let $h(m,n)$ be a discrete psf and $f(m,n)$ be a discrete image.
- Then the output of $f(m,n)$ filtered by $h(m,n)$ can be computed as

$$g(m,n) = \sum_k \sum_j f(k,j)h(m-k,n-j)$$

- Note that for each pixel position in the output, (m,n):
  1. $h(k,j)$ is flipped in both directions and shifted to center at (m,n)
  2. The overlapping pixels of the psf and the image are multiplied and added
  3. And then you move to another position in the output image
- This is the (Double) Convolution Sum
- Usually (almost always) $h(m,n)$ is much smaller than $f(m,n)$ in extent, and it is often called a "mask"

Convolution in 2D: Using a “Mask”