

The Digital Image



Georgios Papaioannou - 2014

The Image as a Digital Signal

- Both computer graphics and photography create a discretized representation of a continuous light signal:
 - Photography: Continuous Incident light \rightarrow Sensor \rightarrow Raster
 - CG: Mathematical geometry representation \rightarrow Raster

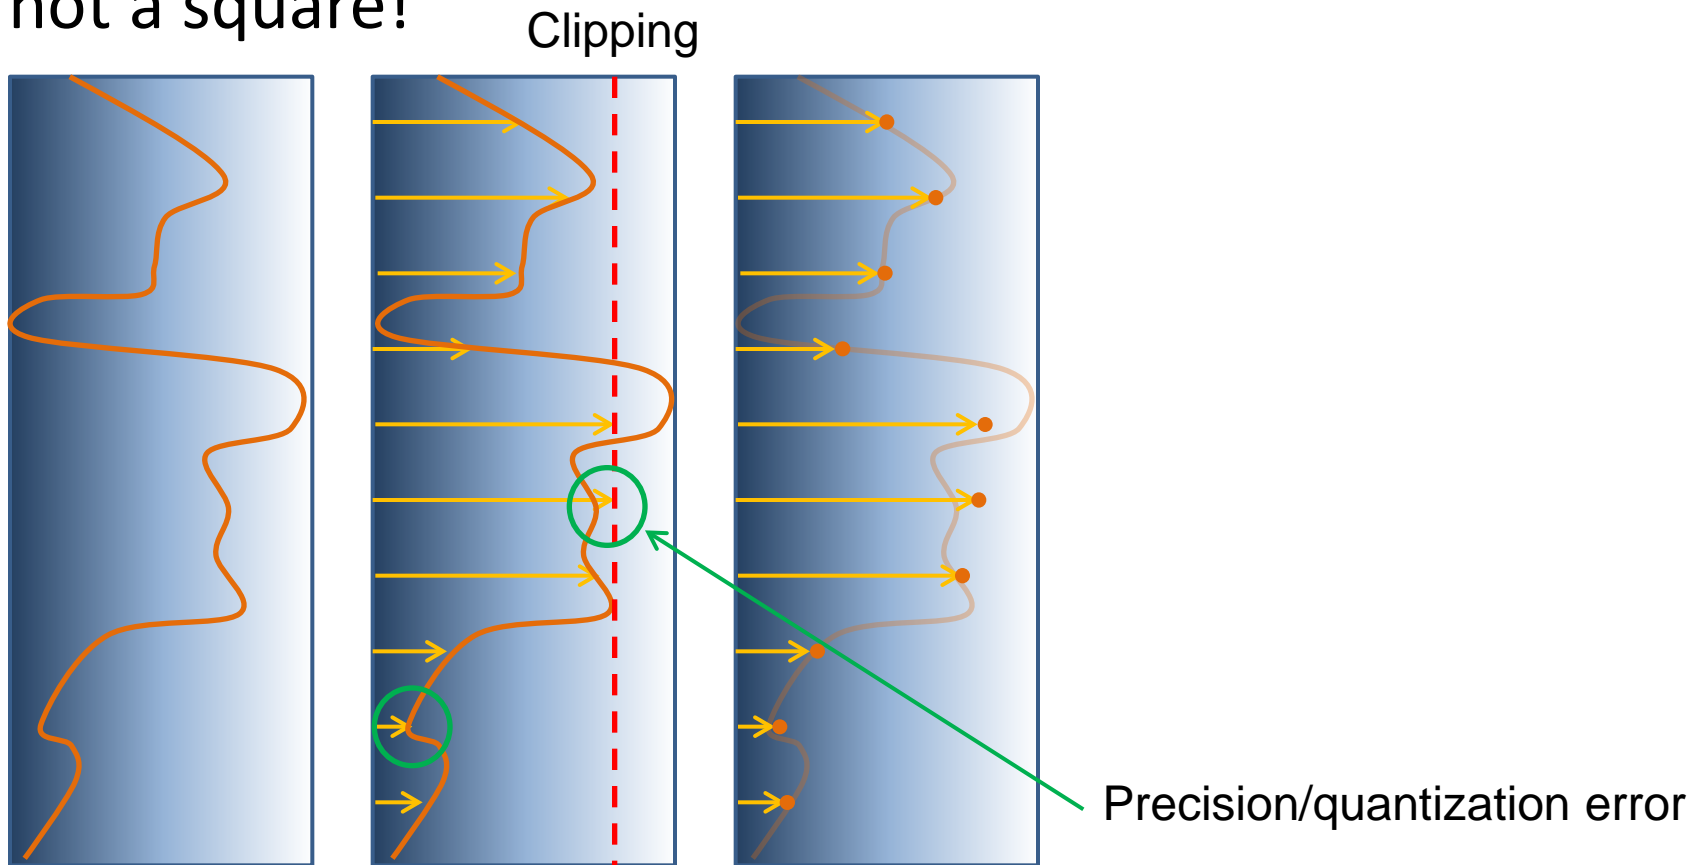


The Image as a Digital Signal

- The digital image signal generation suffers from all digitization problems:
 - Aliasing
 - Quantization errors
- Both digital photography and photorealistic rendering can also suffer from noise
 - Different in nature: statistical vs thermal

The Pixel

- The pixel represents a single sample in the image, not a square!



The Pixel (2)

- The sampling rate in an image is determined by the image resolution vs (physical) image size
 - Denser sampling (higher resolution) can correctly capture higher image detail
 - There is **always** a limit of what a specific resolution can correctly represent (see Nyquist criterion next)
 - E.g. doubling the resolution will only mitigate the problem to a higher image frequency

The Pixel (3)


- Roughly speaking, we can correctly reconstruct (see) a repeating variation in intensity with a (sine) period of 2 pixels or more
- Faster intensity changes cannot be systematically and correctly sampled, leading to **aliasing** and noise
- To alleviate this, we purposefully **limit the frequency** (detail) of the input signal to match our sampling capabilities

The Pixels We See



- In order to perceive the color of an image, we have to go through a reconstruction of an analog intensity from the samples. This involves:
 - A reconstruction filter → obtain a continuous signal
 - A tone mapping stage → adjust intensity to actual displayable range
 - The device's response curve → translates nominal intensities to actual light
 - The device's spatiotemporal impulse response → spreads intensity over screen surface and time

The Pixels We See (2)


Samples





Tone mapping




Reconstruction



Dev. Resp. curve



Device impulse resp.



Reconstruction Filters

- “Spread” isolated discrete samples to form an analog signal
- Try to rectify the original signal



Nearest neighbor
(piecewise constant)




Bilinear



Bicubic

Aliasing

- Aliasing is the miss-interpretation of the samples as a different signal than the original during the reconstruction




Major Aliasing Cases in Graphics

- Image-space
 - Geometry aliasing
 - Erratic and discontinuous sampling of boundaries and thin structures (“juggies”, wholes)
 - Texture aliasing
- Temporal aliasing
 - Unatural apparent motion


Geometric Aliasing

- Erratic and discontinuous sampling of boundaries and thin structures (holes, noise)
- Sampling of smooth structures at regular locations (“juggies”)





Texture Aliasing

- Textures are images themselves or procedural patterns for modifying the appearance of geometry
- These are signals, too
- Sampling them at an inadequate rate in image space, produces significant aliasing
- Manifested as:
 - Noise
 - Irregular patterns
 - Both change erratically with motion

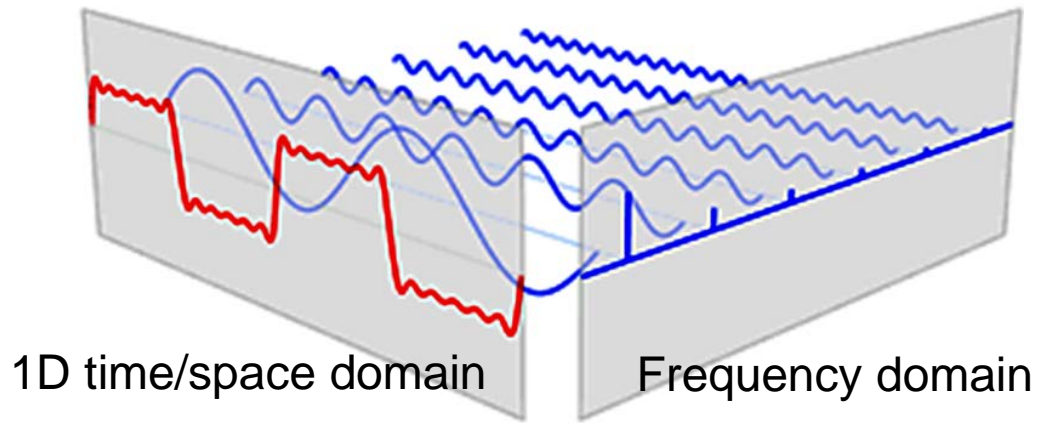


Temporal Aliasing

- Temporal aliasing occurs usually in fast motion
- Frame rate is inadequate to capture the motion frequency (happens to the HVS as well)
- We usually confuse the motion with another



Frequency Domain

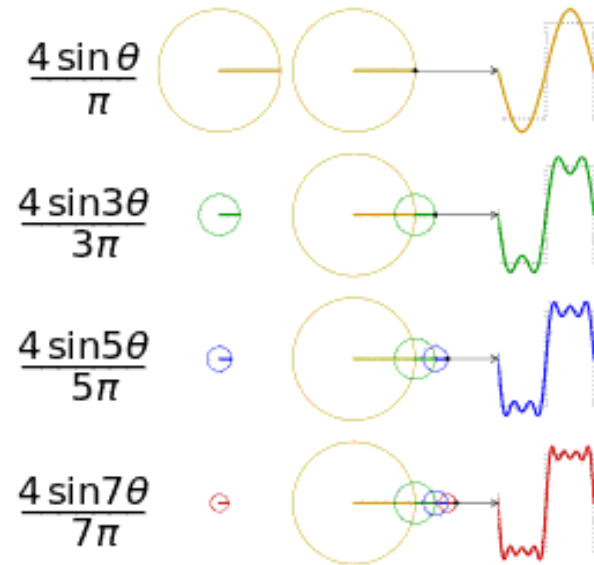



- A periodic signal can be decomposed into a series of overlapping harmonic functions of increasing “frequency”, i.e. shorter period
- The domain for the parameterization of these functions is the *frequency domain*

Frequency Domain (2)

- In the Fourier series expansion of a signal, the signal is analyzed into sinusoids (i.e. projected onto a sinusoid function base):

$$y(t) = a \sin(2\pi \overset{\text{Frequency}}{\xi} t + \varphi) = a \sin(\omega t + \overset{\text{Phase}}{\varphi})$$



The Fourier Transform

- A general transformation to express an analog signal in the frequency domain and back (inverse FT)

$$\hat{f}(\xi) = \int_{-\infty}^{\infty} f(x) e^{-2\pi i x \xi} dx$$

$$f(x) = \int_{-\infty}^{\infty} \hat{f}(\xi) e^{2\pi i x \xi} d\xi$$

Complex Numbers – Euler Formula


$$z = x + iy = |z|(\cos \varphi + i \sin \varphi) = |z|e^{i\varphi}$$

$$\bar{z} = x - iy = |z|(\cos \varphi - i \sin \varphi) = |z|e^{-i\varphi}$$

“conjugate” of z

$$x = \operatorname{Re}\{z\}, y = \operatorname{Im}\{z\}$$


$$|z| = \sqrt{x^2 + y^2}$$



The complex number plane

The Spectrum

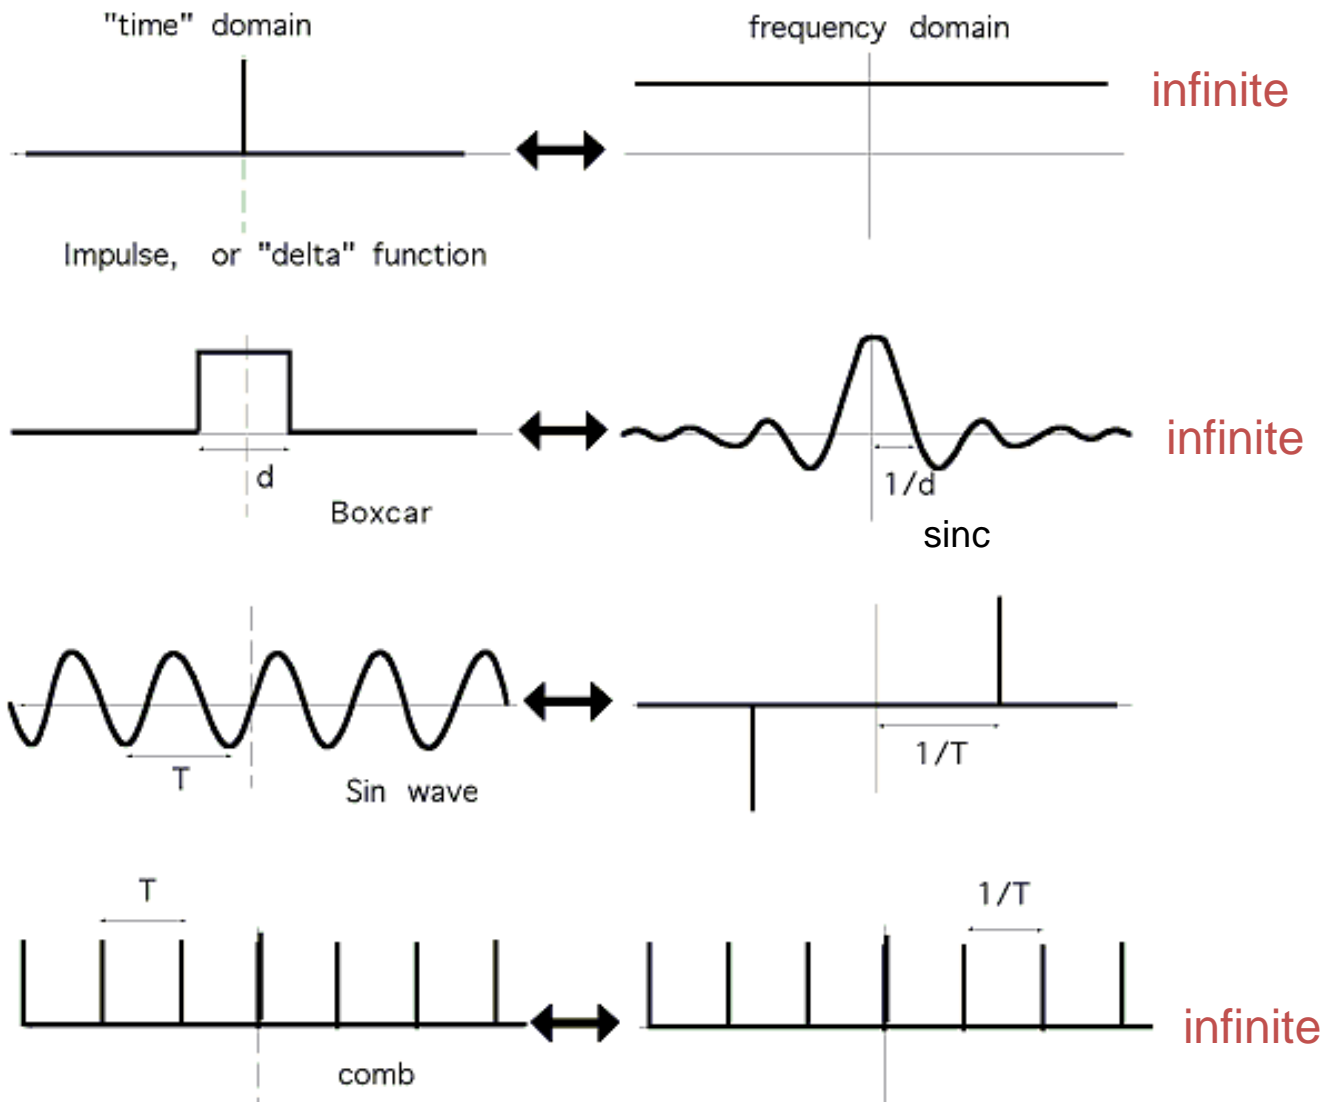
- The Fourier transformation results in an imaginary function
 - Magnitude: the amplitude (or presence) of each frequency
 - Angle: the “phase” (or shift) of each frequency
- They both comprise the *spectrum* of the signal



Types of Spectra

- A signal has an unbounded or infinite spectrum if in order to be completely represented it requires non-zero Fourier coefficients of $\xi \rightarrow \infty$
- All discontinuous signals have an infinite spectrum
 - In reality, there is no ideally abrupt signal, but for all practical purposes, very sharp transitions have a very spread spectrum
 - In graphics we do have discontinuous (mathematical) signals!
- A band-limited signal is one with a finite spectrum (non-zero frequency-domain coefficients)







Typical Spectra



Important signals:
Symmetrical across
the two domains!

$$\text{sinc} = \frac{\sin(x)}{x}$$

Other Common Functions







Fourier Transform in Higher Dimensions

- The FT can be generalized for higher dimensions, even for aperiodic signals:

$$f(\mathbf{x}) = \int_{-\infty}^{+\infty} F(\mathbf{u}) e^{j2\pi\mathbf{x}\cdot\mathbf{u}} du_1 du_2 \dots du_N$$

$$F(\mathbf{u}) = \int_{-\infty}^{+\infty} f(\mathbf{x}) e^{-j2\pi\mathbf{x}\cdot\mathbf{u}} dx_1 dx_2 \dots dx_N$$

Examples in the Image Domain





$$y(x) = (h * s)(x) = \int_{-\infty}^{+\infty} h(x - t)s(t)dt = \int_{-\infty}^{+\infty} s(x - t)h(t)dt$$

- The convolution operation $*$ blends two functions by shifting one over the other and modulating their overlapping values
 - It is like “pushing one function through the other”

Convolution Examples


$$y(x) = (h * s)(x) = \int_{-\infty}^{+\infty} h(x - t)s(t)dt$$



Impulse Response

- A system is characterized by an *impulse response*, i.e. a function $h(x)$ that is the output of the system given a single pulse (impulse) as input
- The impulse in continuous signals is the Dirac function $\delta(x)$, a single pulse at 0 with an integral equal to 1
- The impulse in discrete-time systems is the Kronecker delta function:

$$\delta_{ij} = \begin{cases} 0 & \text{if } i \neq j \\ 1 & \text{if } i = j, \end{cases}$$



Impulse Response (2)

- The impulse response is the result of the convolution of the system with the input singular pulse
- It completely characterizes a time-invariant linear system: it is the (fixed) function that is applied to any input signal:

$$h(x) = \int_{-\infty}^{+\infty} h(x-t)\delta(t)dt$$

Unknown LTI system

- Linear systems are very typically used in image operations to apply filters (i.e. systems) on 2D signals
- Filters are generally characterized as IIR or FIR:
 - Infinite Impulse Response: The domain of support of the impulse response is infinite. Ideal filters are typical IIR ones (more later)
 - Finite Impulse Response: The non-zero values of the impulse response are limited to a finite range
- Filters have their own spectrum, which emphasizes or suppresses certain frequencies

Convolution and Frequency Domain

- Important and useful property:
- If $H(\xi)$ and $S(\xi)$ are the Fourier transforms of two functions $h(x)$ and $s(x)$, then:

$$FT((h * s)(x)) = H(\xi) S(\xi)$$

- I.e: Convolution in the time/space domain becomes multiplication of spectra in the frequency domain
 - Side-effect: Sometimes it is easier to design filters in frequency domain and find their IFT to obtain their impulse response!

Properties of the Fourier Transform

- Linearity: $af_1(t) + bf_2(t) \rightarrow aF_1(\omega) + bF_2(\omega)$
- Input shift: $f(t - t_0) \rightarrow F(\omega)e^{-j\omega t_0}$
- Input scaling: $f(at) \rightarrow \frac{1}{|a|} F\left(\frac{\omega}{a}\right)$
- Frequency shift: $F(\omega - \omega_0) \rightarrow f(t)e^{j\omega_0 t}$
- Convolution:

$$f_1(t) * f_2(t) \rightarrow F_1(\omega)F_2(\omega)$$

$$f_1(t)f_2(t) \rightarrow \frac{1}{2\pi} F_1(\omega) * F_2(\omega)$$

The Sampling Theorem


- In order to ensure that the reconstructed signal is identical to the original, the Nyquist-Shannon sampling theorem states that the original signal has to be *band-limited* and the sampling rate $f_{sampling}$ must be at least twice the highest frequency of the original signal:

$$2 |f_{\max}| \leq f_{sampling}$$


- Let's see why next:

Sampling and Frequency Domain

- What happens to a signal when it gets sampled?
- The samples we take in time or space comprise an (infinite) impulse train:



- Remember the spectrum (*transfer function*) of such a signal? It is also a train of spikes:




Sampling and Frequency Domain (2)

- Sampling a signal $x(t)$ is equal to multiplying it with an impulse train $s(t)$ of sampling period T_s
 - We zero out all non-sample positions and keep the samples
- Now remember the convolution property:


$$s(t)x(t) \rightarrow \frac{1}{2\pi} S(\omega) * X(\omega)$$
- Given the spectrum of $X(\omega)$ of the signal, the spectrum of the sampled signal is an infinitely shifted and scaled version of $X(\omega)$ repeated every $1/T_s$

Sampling and Frequency Domain (3)


Note that the spectrum is “folded” (mirrored) in the “negative” frequency domain



Sampling with $f_s = 1/T_s$



Signal Reconstruction




- To reconstruct the original signal, we need to convolve the sampled one with an appropriate reconstruction filter to isolate the original spectrum


Inseparable Spectra - Aliasing

- If the “ghost” spectra of the original signal overlap with the central spectrum then there is no way to reconstruct it properly!
 - The new (overlapping) spectrum represents now a different signal
 - We call this “aliasing” because the reconstructed signal is the same for many different input ones (aliases)

Inseparable Spectra - Aliasing




This is the spectrum
of a different signal!




The Nyquist Criterion

- According to the above, if B is the maximum frequency of a (band-limited) signal, the frequency f_s of the samples taken must be at least $2XB$ to avoid overlap of spectrum replicas and therefore, aliasing



Antialiasing

- Ok, what about:
 - A fixed sampling rate, that cannot be adjusted according to maximum signal frequency (image case)?
 - Signals with an infinite spectrum (again graphics...)?




Antialiasing (2)


- We must “band-limit” the signal by force to contain its spectrum within the sampling window we have
- So we cut off high frequencies → we smooth out the original signal!
- The original signal cannot ever be correctly reconstructed!
- But we can at least produce a signal free of aliasing, noise and temporal artifacts

Antialiasing Filters

- An antialiasing filter clamps or limits the spectrum of the original signal to the $[-\frac{f_s}{2}, \frac{f_s}{2}]$ frequency range in order to be able to correctly sample it with an f_s rate
- Ideal filters do not emphasize or suppress frequencies




Ideal band-limiter: sinc filter.
Equal response to all usable frequencies



Antialiasing Filters (2)


- Can I have an ideal filter?
 - No... Ideal filters are IIR filters so they cannot be practically constructed
- We use “good” FIR filters, usually approximating truncated or dampened IIR ones



Reconstruction Filters Revisited

- Now that we have limited our (image) signal and sampled the analog domain, we need to create proper reconstruction filters to output the digital signal
- Remember:
 - We have (forcefully) abided the Nyquist criterion:
 - Either by increasing the sampling rate, or by limiting the signal spectrum
 - Side bands are “clearly” (sufficiently) separated
- Now we need a filter that does not overlap the side bands

Reconstruction Filters Revisited (2)



- Ideal Filter
- Too wide frequency response
- Too de-emphasizing (suppresses higher frequencies)
- Good filter

Moving to the Discrete Input Domain

- The Fourier Transform and the convolution have their own versions in the discrete domain
- What about aperiodic discrete signals?
 - The same theory still applies but with some modifications both to what the input domain represents and how the frequency domain is interpreted

The Discrete Fourier Transform (DFT)

- Applies to discrete, even samples of a function (discrete signal):

$$X_k = \sum_{n=0}^{N-1} x_n e^{-i2\pi kn/N}, k \in \mathbb{Z}$$

- It is typically computed via the Fast Fourier Transform (FFT) algorithm
 - Fast, highly parallel algorithm (CPU, GPU, ASIC, FPGA implementations exist)

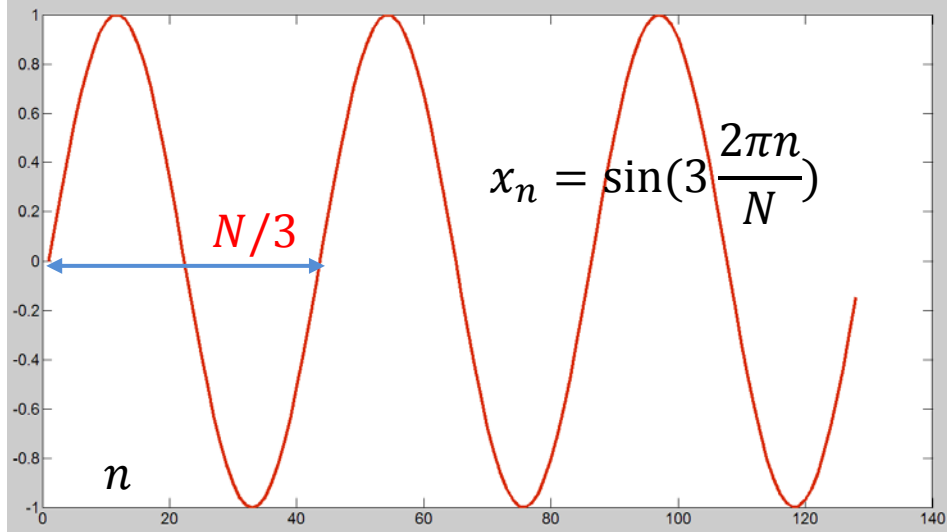
The Inverse Discrete Fourier Transform

$$x_k = \frac{1}{N} \sum_{n=0}^{N-1} X_n e^{i2\pi kn/N}, k \in \mathbb{Z}$$


- It is typically computed via the Fast Fourier Transform (FFT) algorithm

Reading the FFT

Input signal



DFT (FFT): $|X_k|$



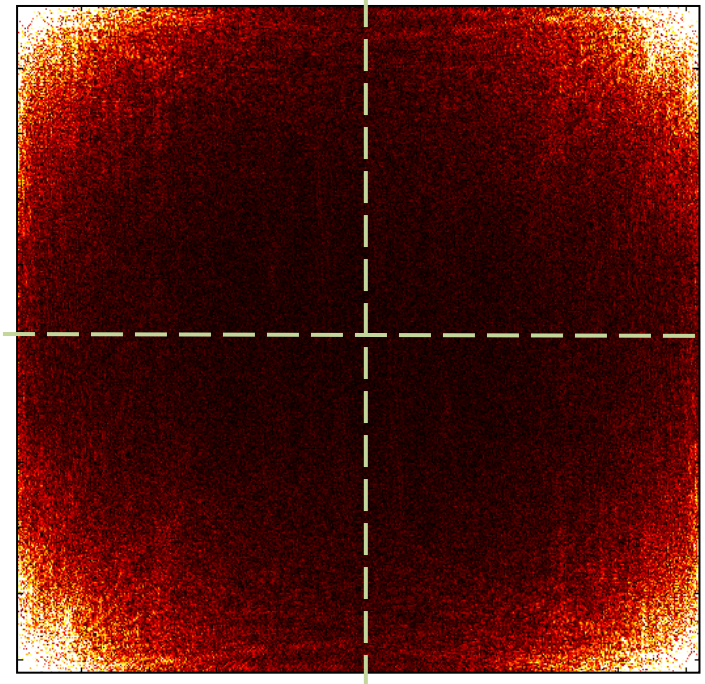
N (represents a full-width "period")

0 $\frac{1}{N}$ $\frac{2}{N}$ $\frac{3}{N}$...

$\frac{N/2-1}{N}$ $-\frac{N/2}{N}$

$-\frac{3}{N}$ $\frac{1}{N}$

DFT in Image Space



Interpreting the DFT in Image Space





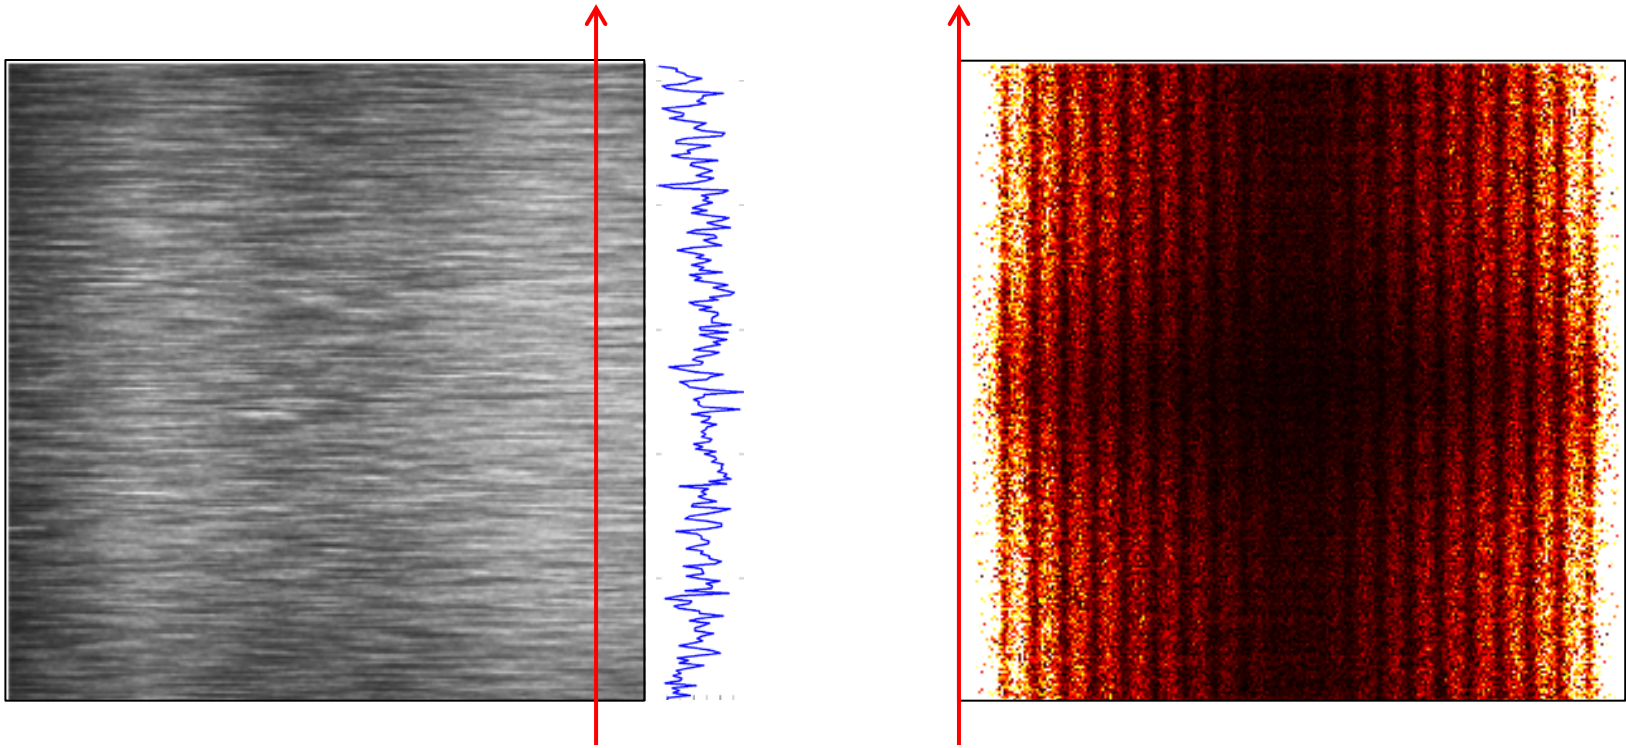


Image Bandwidth Examples


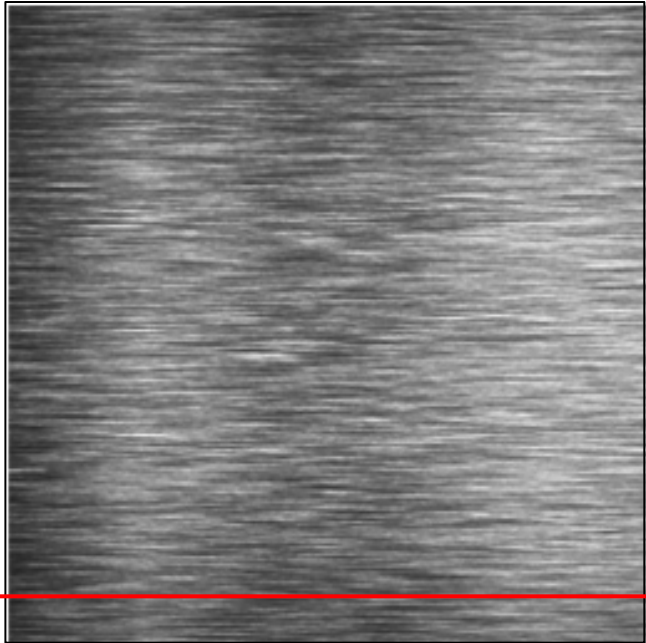


Interpretation of 2D DFT



- The vertical intensity variations in the above image look like noise
- The blue graph shows the intensity along the red section of the image
- The vertical frequency distribution is widely spread, typical of noisy input

Interpretation of 2D DFT (2)



- The horizontal intensity varies more slowly and sinusoids can be detected (brush strokes)
- Brush strokes appear as high concentration of frequencies around certain “prominent” ones

- Georgios Papaioannou
- Sources:
 - Wikipedia
 - T. Theoharis, G. Papaioannou, N. Platis, N. M. Patrikalakis, Graphics & Visualization: Principles and Algorithms, CRC Press