Some notes on sampling theory for the Part II Information Theory & Coding course

Neil Dodgson*
based on Computer Lab. Tech. Report No. 261

February 18, 2003

These notes are based on a Technical Report which discusses the resampling of images. A lot of the text therefore refers to images but it can easily be generalised to all forms of data which need to be sampled. In this text, the term “image” will usually mean “digital image” and the the concept of a real image is captured by the phrase “intensity surface”, i.e. a function from the 2D real plane to intensity.

1 Sampling

Sampling is the process of converting a continuous function into a discrete representation. Conventional sampling theory deals with regularly spaced point samples, with each sample being produced by taking the value of the continuous function at a single point.

1.1 Converting the continuous into the discrete

Sampling takes a continuous signal, one defined over all space, and produces a discrete signal, one defined over only a discrete set of points. In practice the two signals are only defined within a region of interest. For example, when sampling images, the continuous signal is an intensity surface; that is: a three-dimensional function with two spatial dimensions and one intensity dimension where each point in the spatial plane has a single intensity. When dealing with colour, the number of dimensions increases, typically to three colour dimensions (but still only two spatial dimensions). When dealing with one spatial dimension (e.g. sound) or three spatial dimensions (e.g. a CAT or MRI scan of the human body) the situation also changes but the underlying principle is the same: at every spatial point there is either a scalar or vector value, which is the value of the function at that point, be it intensity, colour, pressure, density or whatever.

The discrete signal has the same number of spatial and function value dimensions as the continuous signal but is discrete in the spatial domain. Furthermore, a digital computer cannot represent a continuous quantity and so the sampled signal will be discrete in value as well. The process of converting a continuous value to a discrete one is called quantisation.

1.2 Quantisation

Ideally, infinite precision in recorded value is desirable because no errors would then arise from slightly imprecise values. In practice, the sampling device has only finite precision and the values are recorded to only finite precision.

The important point is that sufficient quantisation levels must be used that any errors caused by imprecision are unimportant. For example, with sound and images, the number of quantisation levels must be sufficient to fool the human observer into believing that the signal is continuous. We considered this in the Part IB Computer Graphics course, and the conclusions we drew there are reproduced below as an example.

1.2.1 The minium necessary number of intensity levels required to represent images

For images, quantisation errors are most noticeable when only a few discrete intensity levels are available, and most image processing texts include an example showing the effect on a displayed image of allowing a larger or smaller number of intensity levels [see, for example, Rosenfeld and Kak, 1982, pp.107–8 figs. 14 and 15; Gonzalez and Wintz, 1977, fig. 2.8]. The extreme case is where there are only two levels: black and white. Such an image takes comparatively little storage (one bit per pixel) but has a limited usefulness. With digital half-toning techniques it can be made to simulate a wide range of grey shades, but, for resampling, it is preferable to have true shades of grey rather than simulated ones.

The general case, then, is one where a significant number of discrete intensity levels are used to represent a continuous range of intensities. But how many intensity levels are sufficient?

The human visual system is limited in how small an intensity change it can detect. Research suggests that, for two large areas of constant intensity, a two percent difference can be just detected [Crow, 1978, p.4]. The minimum difference that can be detected rises to a higher value for very dark or very bright areas [Pratt, 1978, pp.17–18]. It also rises when comparing small areas of constant intensity [*ibid.*, pp.18–19, fig. 2.5]. When dealing with colour images the minimum noticeable differences for pure colour information (that is with no intensity component) are much larger than those for intensity information, hence broadcast television has an intensity channel with high spatial resolution and two channels carrying the colour information at low spatial resolution [NMFPT, 1992]. Here, we shall consider only intensity. The number of intensity levels required to produce a faithful representation of an original intensity surface depends on the image itself [Gonzalez and Wintz, pp.27-28] and on the display device.

For example, we concluded, in Part IB, that, for display on a CRT, ten bits of intensity data are sufficient. For any image without large areas of slowly varying dark intensities, ten bits of intensity information is too much, and so, for many images, if the intensity values are stored to a precision of eight bits, few visible errors will result from the quantisation of intensity. If they are stored to a precision of ten bits or more, practically no visible errors will arise.

1.3 The sampling process

Let us now assume that our function values are quantised to sufficient precision that the errors introduced by quantisation are neglibible. We therefore go on to consider the process of sampling.

The aim of sampling is to generate sample values so that they "best represent" the original continuous function. This is a good concept but what is meant by 'best represent'? To some it may mean that the samples obey classical sampling theory. To others it may mean that the resulting samples, when reconstructed on a particular device (e.g. a monitor or a loudspeaker), is indistinguishable from the original as far as is possible.

1.3.1 Classical sampling theory

The roots of sampling theory go back to 1915, with Whittaker's work on interpolation of a set of equispaced samples. However, most people attribute the sampling theorem to Shannon [1949], who acknowledges a debt to many others in its development. Attribution of the theorem has been jointly given to Whittaker and Shannon [Gonzalez and Wintz, 1977, p.72] Shannon and Nyquist [Turkowski, 1986], and to Shannon, Kotel'nikof and Whittaker [Petersen and Middleton, 1962, p.279]. Shannon's statement of the sampling theorem is[1]:

> If a function $f(t)$ contains no frequencies higher that W cps it is completely determined by giving its ordinates at a series of points spaced $\frac{1}{2W}$ seconds apart.
> [Shannon, 1949, Theorem 1] [2].

Mathematically, the sampling process is described as the product of $f(t)$ with the comb function:

$$\text{comb}(t) = \sum_{n=-\infty}^{\infty} \delta(t - nT)$$

where $\delta(t)$ is the Dirac delta-function[3]. This gives the sampled function:

$$\hat{f}(t) \quad = \quad \left(\sum_{n=-\infty}^{\infty} \delta(t - nT) \right) f(t)$$

[1]The theory laid out by Shannon [1949] and others is for one dimensional sampling only. Petersen and Middleton [1962] extended Shannon's work to many dimensions.

[2]cps = cycles per seconds ≡ Hertz

[3]The Dirac delta function is zero everywhere except at $t = 0$. The area under the function is unity, that is $\int_{-\infty}^{\infty} \delta(t)\, dt = 1$. One definition of the Dirac delta function is:

$$\delta(t) = \lim_{a \to \infty} \sqrt{a}\, e^{-a\pi x^2}$$

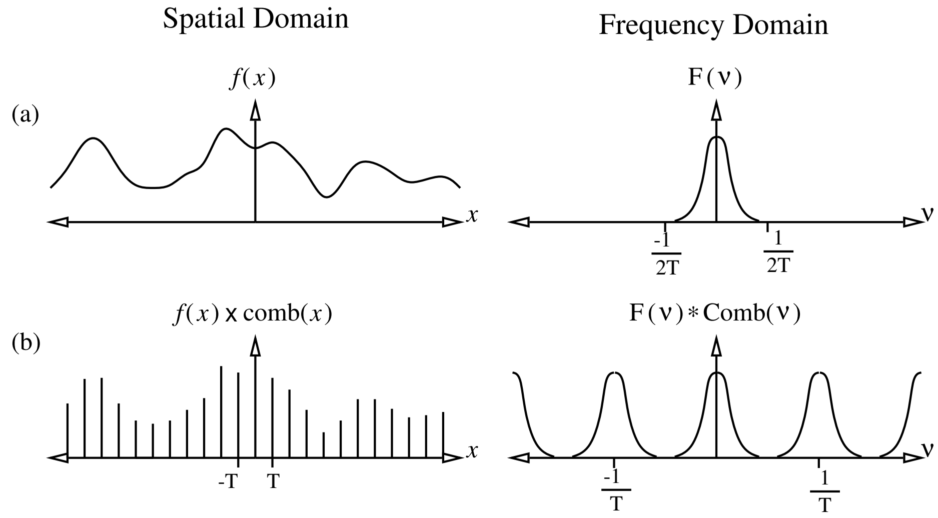

Figure 1: The sampling process: (a) the continuous spatial domain function, $f(x)$, has a Fourier transform, $F(\nu)$, bandlimited to below half the sampling frequency, $\frac{1}{2T}$; (b) when it is sampled ($f(x) \times \mathrm{comb}(x)$) its Fourier transform is convolved with $\mathrm{Comb}(\nu)$ producing replicas of the original Fourier transform at a spacing of $\frac{1}{T}$.

$$= \sum_{n=-\infty}^{\infty} (\delta(t - nT) f(t - nT)) \tag{1}$$

This is not the same as:

$$\hat{f}_n = f(nT) \tag{2}$$

Equation 1 is a continuous function which consists of an infinite sum of weighted, shifted Dirac delta functions and is zero everywhere except at $t = nT, n \in \mathcal{Z}$. Equation 2 is a discrete function which is defined on the set $n \in \mathcal{Z}$. The values of the discrete function are the weights on the delta functions that make up the continuous function, that is:

$$\hat{f}(t) = \sum_{n=-\infty}^{\infty} \hat{f}_n \, \delta(t - nT)$$

In a computer we, of course, store the discrete version; but mathematically and theoretically we deal with the continuous one. The sampling theorem can be justified by considering the function in the frequency domain.

The continuous spatial domain function, $f(t)$, is bandlimited. That is it contains no frequencies higher than ν_b (W in the statement of Shannon's theorem above). Its Fourier transform, $F(\nu)$ is thus zero outside the range $(-\nu_b, \nu_b)$. Sampling involves multiplying $f(t)$ by $\mathrm{comb}(t)$. The equivalent operation in the frequency domain is to convolve $F(\nu)$ by $\mathrm{Comb}(\nu)$, the Fourier transform of $\mathrm{comb}(t)$. $\mathrm{Comb}(\nu)$ is composed of Dirac delta-functions at a spacing of $\frac{1}{T}$ (for a proof of this see Marion [1991, pp.31–32]). Convolving this with $F(\nu)$ produces replicas of $F(\nu)$ at a spacing of $\frac{1}{T}$. Figure 1 illustrates this process.

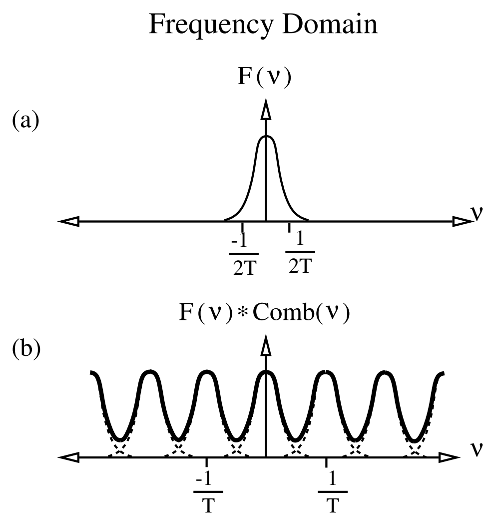

Figure 2: An example of sampling at too low a frequency. Here we see the frequency domain representations of the functions. The original function, (a), is not bandlimited to within half the sampling frequency and so when it is sampled, (b), the copies overlap and add up producing the function shown by the dark line. It is impossible to recover the original function from this aliased version.

If $T < \frac{1}{2\nu_b}$ then the copies of $F(\nu)$ will not overlap and the original $F(\nu)$ can be retrieved by multiplying $\hat{F}(\nu)$ by a box function:

$$\text{Box}(\nu) = \begin{cases} 1, & |\nu| < \nu_b \\ 0, & \text{otherwise} \end{cases}$$

This removes all the copies of the original except for the copy centred at $\nu = 0$. As this is the original frequency domain function, $F(\nu)$, the spatial domain function will also be perfectly reconstructed.

Multiplying by a box function in the frequency domain is equivalent to convolving in the spatial domain by the box's inverse Fourier transform, $s(x)$. This can be shown to be $s(x) = 2\nu_b \text{sinc}(2\nu_b x)$, where $\text{sinc}(x) = \frac{sin(\pi x)}{\pi x}$ (a proof of this can be found in appendix A).

If $T \geq \frac{1}{2\nu_b}$ then the copies of $F(\nu)$ will overlap (figure 2). The overlapping parts sum to produce $\hat{F}(\nu)$. There is no way that this process can be reversed to retrieve $F(\nu)$ and so $f(t)$ cannot be perfectly reconstructed. If $\hat{F}(\nu)$ is multiplied by the box function (the perfect reconstructor) then the resulting function is as if $F(\nu)$ had been folded about the frequency $\frac{1}{2T}$ and summed (figure 3). For this reason $\nu = \frac{1}{2T}$ is known as the folding frequency. The effect of this folding is that the high frequencies in $F(\nu)$ alias into low frequencies. This causes artifacts in the reconstructed image which are collectively known as 'aliasing'. In computer graphics the term aliasing is usually incorrectly used to refer to both aliasing and rastering artifacts [Pavlidis, 1990]. This point is picked up in section 1.5. Figure 4 gives an example of aliasing. It normally manifests as unsightly ripples, especially near sharp changes in intensity.

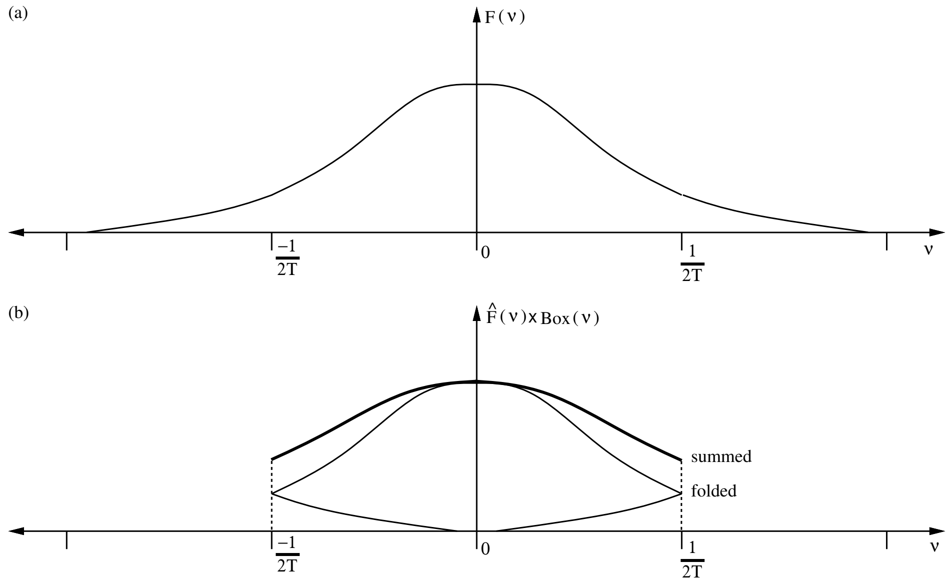

Figure 3: An example of aliasing. (a) is the frequency spectrum of the original image. If the original image is sampled and then reconstructed by multiplying its frequency spectrum by a box function then (b) is the result. The copies of the spectrum overlap and add up, as if $F(\nu)$ in (a) had been folded back about $\frac{1}{2T}$ and summed.

To avoid aliasing we must ensure that the sampled intensity surface is bandlimited to below the folding frequency. If it is not then it can be prefiltered to remove all information above this frequency. The procedure here is to multiply the intensity surface's Fourier transform by a box function, a process known as *bandlimiting* the intensity surface. Foley *et al* [1990, fig. 14.29] give an example of this procedure. This prefiltering followed by point sampling is equivalent to an area sampling process, as is explained later.

These procedures are mathematically correct and produce an image free of aliasing. However they have their drawbacks. Firstly, if an intensity surface has to be prefiltered then the image does not represent the original intensity surface but rather the filtered one. For certain applications this may be undesirable. To represent an intensity surface with flat or linearly sloped areas, infinite frequencies *are* required. Bandlimiting prevents such surfaces from being perfectly represented, and so any 'flat' part of a bandlimited image will be ripply. An example of this effect is that any sharp edge in a band-unlimited intensity surface will have a 9% overshoot either side if it is bandlimited, no matter what the bandlimit is [Lynn and Fuerst, 1989, p.145]; ripples will also propagate out from the discontinuity, reducing in magnitude as they get farther away. Thus a bandlimited intensity surface is, by its very nature, ripply. This can be seen in figure 8(a); there is no aliasing in this figure, the intensity surface is inherently ripply due to being bandlimited. The human visual system abhors ripples, they seriously degrade an image. If however, the ripples in the reconstructed intensity surface are undetectable by the human eye then the surface looks very good.

More importantly, it is practically impossible to achieve perfect prefiltering and so some aliasing will creep in. In an image capture device, some (imperfect) prefiltering will

Figure 4: An example of a reconstructed intensity surface which contains aliases. These manifest themselves as ripples around the sharp edges in the image. The small image of the letter 'P' is the original. The large is an intensity surface reconstructed from this small original using an approximation to sinc reconstruction.

occur; that is: the image capture device does not perform perfect prefiltering. Fraser [1987] asserts that most people depend on this prefiltering to bandlimit the intensity surface enough that little aliasing occurs. In the digital operations of rendering and resampling, perfect prefiltering is simply impossible, ensuring that some aliasing always occurs. This is because perfect prefiltering involves either continuous convolution by an *infinite* sinc function or multiplication of the *continuous* Fourier transform by a bandlimiting function. Both of these operations are impossible in a discrete computer.

Finally, perfect reconstruction is also practically impossible because it requires an infinite image. Thus, whilst it would be theoretically possible to exactly recreate a bandlimited intensity surface from its samples, in practice it is impossible. In fact, most display devices reconstruct so badly that correct sampling can produce visually worse results than the optimal incorrect sampling method.

Before discussing this, however, we need to examine the various types of sampling. These fall into two broad categories: area samplers and point samplers.

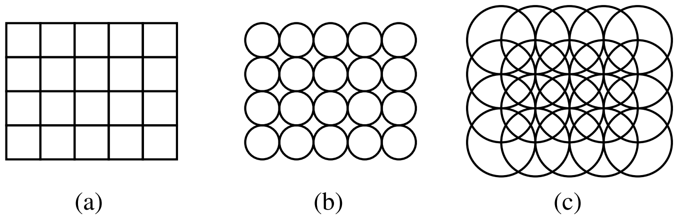

Figure 5: Three common assumptions about pixel shape: (a) abutting squares, (b) abutting circles, and (c) overlapping circles.

1.4 Area *vs* point sampling

Any sampling method will fall into one of these two categories. Area sampling produces a sample value by taking into account the values of the intensity surface over an area. These values are weighted somehow to produce a single sample value. Point sampling takes into account the values of the intensity surface only at a finite set of distinct points[4]. If the set contains more than one point then these values must be combined in some way to produce a single sample value. This is not the usual description of point sampling, because 'point sampling' is usually used in its narrow sense: to refer to single point sampling. Fiume [1989, section 3.2.4] studied the various sampling techniques, the following sections draw partly on his work.

1.4.1 Area sampling

Exact-area sampling The assumptions behind exact-area sampling are that each pixel has a certain area, that every part of the intensity surface within that area should contribute equally to the pixel's sample value, and that any part of the intensity surface outside that area should contribute nothing. Hence, if we make the common assumption that pixels are abutting squares then for a given pixel the sample value produced by the exact area sampler will be the average intensity of the intensity surface within that pixel's square (figure 5(a)) Alternately, we could assume that pixels are abutting circles [Durand, 1989] or overlapping circles [Crow, 1978] and produce the average value of the intensity surface over the relevant circular area (figure 5(b) and (c)). Obviously other assumptions about pixel shape are possible.

General area sampling The general case for area sampling is that some area is chosen, inside which intensity surface values will contribute to the pixel's sample value. This area is usually centred on the pixel's centre. Some weighting function is applied to this area and the pixel's value is the weighted average over the area. Exact-area sampling is obviously a special case of this.

Equivalence to prefiltering Area sampling is equivalent to prefiltering followed by single point sampling. In area sampling, a weighted average is taken over an area of

[4]Fiume [1989, p.82] states that the set of point samples should be countable, or more generally, a set of measure zero.

the intensity surface. With prefiltering a filter is convolved with the intensity surface and a single point sample is taken for each pixel off this filtered intensity surface. To be equivalent the prefilter will be the same function as the area sample weighting function. The two concepts are different ways of thinking about the same process. Heckbert [1989] proposed that sampling could be decomposed into prefiltering and single point sampling. Whilst this is entirely appropriate for all area samplers, not all point samplers can be represented in this way, because some point samplers cannot be represented by filters. Specific examples are adaptive super-samplers and stochastic point samplers, discussed below.

1.4.2 Point sampling

Unlike an area-sampling process, a point-sampler ignores all but a countable set of discrete points to define the intensity value of a pixel. Point sampling comes in two flavours: regular and irregular (or stochastic) point sampling. Wolberg [1990, sections 6.2 and 6.3] summarises the various types of point sampling; here we only outline them.

Single point sampling In single point sampling, the samples are regularly spaced and each pixel's sample value is produced by a single point. In fact this is exactly the method of sampling described in classical sampling theory (section 1.3.1). All other point sampling methods are attempts to approximate area sampling methods or attempts to reduce artifacts in the reconstructed intensity surface (often they attempt to be both).

Super-sampling As in single point sampling, the samples are arranged in a regular fashion, but more that one sample is used per pixel. The pixel's intensity value is produced by taking a weighted average of the sample values taken for that pixel. This can be seen as a direct approximation to area sampling. Fiume [1989, p.96, theorem 6] proves that, given a weighting function, a super-sampling method using this weighting function for its samples converges, as the number of samples increases, toward an area sampling method using the same weighting function.

Such a super-sampling technique can be represented as a prefilter followed by single point sampling; on the other hand, the adaptive super-sampling method cannot.

Adaptive super-sampling Adaptive super-sampling is an attempt to reduce the amount of work required to produce the samples. In order to produce a good quality image, many super-samples need to be taken in areas of high intensity variance but a few samples would give good results in areas of low intensity variance. Ordinary super-sampling would need to take many samples in all areas to produce a good quality image, because it applies the same sampling process for every pixel. Adaptive super-sampling takes a small number of point samples for a pixel and, from these sample values, ascertains whether more samples need to be taken to produce a good intensity value for that pixel. In this way fewer samples need to be taken in areas of lower intensity variance and so less work needs to be done.

Stochastic sampling Stochastic, or irregular, sampling produces a sample value for each pixel based on one or more randomly placed samples. There are obviously some bounds within which each randomly placed sample can lie. For example it may be constrained to lie somewhere within the pixel's area, or within a third of a pixel width from the pixel's centre. The possible locations may also be constrained by some probability distribution so that, say, a sample point has a greater chance of lying near the pixel centre than near its edge. Stochastic methods cannot be represented as a prefilter followed by single point sampling, because of this random nature of the sample location.

Stochastic sampling is in favour amongst the rendering community because it replaces the regular artifacts which result from regular sampling with irregular artifacts. The human visual system finds these irregular artifacts far less objectionable than the regular ones and so an image of equal quality can be achieved with fewer stochastic samples than with regularly spaced samples (see Cook [1986] but also see Pavlidis' [1990] comments on Cook's work).

Fiume [1989, p.98, theorem 7] proves that the choice of point-samples stochastically distributed according to a given probability distribution will converge to the analagous area-sampling process as the number of points taken increases, a result similar to that for super-sampling.

1.4.3 Summary

Area sampling yields mathematically precise results, if we are attempting to implement some prefilter (section 1.3.1). However, in digital computations it may be impossible, or difficult to perform area sampling because it involves continuous convolution. The point sampling techniques have been shown to be capable of approximating the area sampling methods. Such approximations are necessary in cases where area sampling is impracticable [Fiume, 1989, p.102].

1.5 Anti-aliasing

The purpose of all sampling, other than single-point sampling, is ostensibly to prevent any artifacts from occuring in the image. In fact it is usually used to prevent any artifacts from occuring in the intensity surface reconstructed from the image by the display device. This prevention is generally known in computer graphics as anti-aliasing. This is something of a misnomer as many of the artifacts do not arise from aliases. The perpetuation of the incorrect terminology is possibly due to the fact that the common solution to cure aliasing also alleviates rastering, the other main source of artifacts in images [Pavlidis, 1990, p.234].

Crow [1977] was the first to show that aliasing (sampling a signal at too low a rate) was the cause of some of the artifacts in digital imagery. He also notes [*ibid.*, p. 800] that some artifacts are due to failing to reconstruct the signal properly. This latter problem he termed 'rastering'. 'Aliasing', however, became the common term for both of these effects, with some subsequent confusion, and it is only occassionaly that one sees the term 'rastering' [Foley *et al*, 1990, p.641]. Mitchell and Netravali [1988], for example,

Figure 6: Two very similar images: (a) on the left, was generated by summing the first twenty-nine terms of the Fourier series representation of the appropriate square wave; (b) on the right, was generated by super-sampling a perfect representation of the stripes.

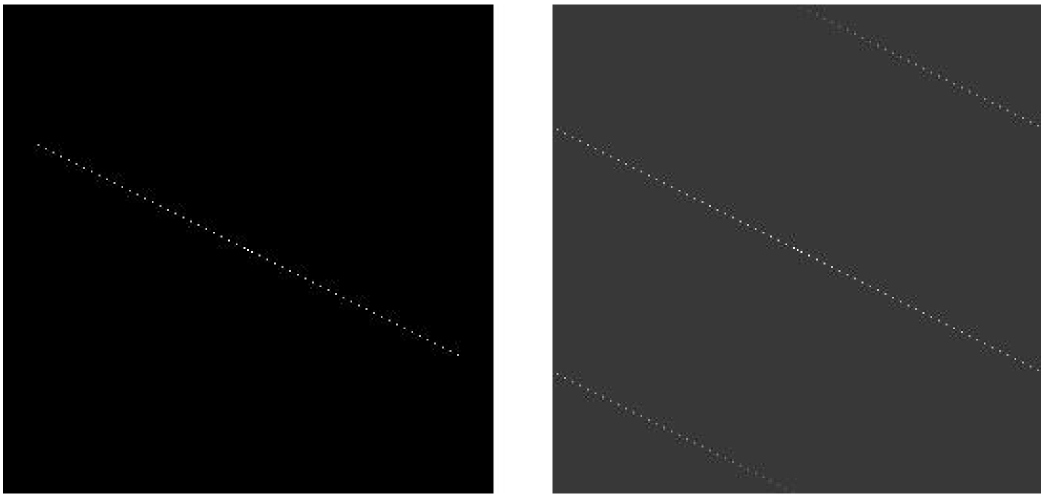

Figure 7: The Fourier transforms of the images in figure 6. (a) on the left and (b) on the right. The Fourier transforms are shown on a logarithmic scale.

Figure 8: The two images of figure 6 reconstructed using as near perfect reconstruction as possible. (a) on the left and (b) on the right. This figure shows part of the reconstructed intensity surface. Aliasing artifacts occur in (b) only

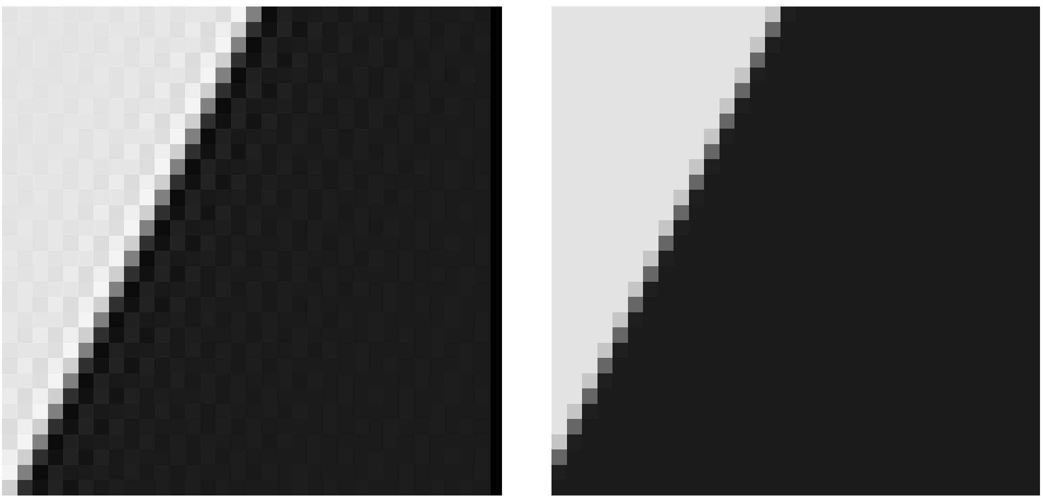

Figure 9: The two images of figure 6 reconstructed using as nearest-neighbour reconstruction. (a) on the left and (b) on the right. This figure shows part of the reconstructed intensity surface. (a) contains only rastering atifacts, while (b) contains a combination of rastering and aliasing artifacts.

discuss the two types of artifact as distinct effects but perpetuate the terminology by naming aliasing and rastering, pre-aliasing and post-aliasing respectively.

Sampling has thus been used to alleviate both types of artifact. This has an inbuilt problem that, to correct for reconstruction artifacts (rastering), one needs to know the type of reconstruction that will be performed on the image. Most images are displayed on a CRT, and all CRTs have a similar reconstruction method, so this is not too big a problem. However, when displayed on a different device (for example a film recorder) artifacts may appear which were not visible on the CRT because the reconstruction method is different.

Figures 6 through 9 illustrate the distinction between aliasing artifacts and reconstruction artifacts. Two images are shown of alternate dark and bright stripes. These images were specially designed so that no spurious information due to edge effects would appear in their discrete Fourier transforms (that is: the Fourier transforms shown here are those of these simple patterns copied off to infinity so as to fill the whole plane).

Figure 6(a) was generated by summing the first twenty-nine terms of the Fourier series representation of the appropriate square wave as can be seen by its Fourier transform (figure 7(a)). Figure 6(b) was generated by one of the common anti-aliasing techniques. It was rendered using a 16×16 super-sampling grid on each pixel with the average value of all 256 super-samples being assigned to the pixel. Figure 7(b) shows its Fourier transform. It is similar in form to figure 7(a) but the aliasing can be clearly seen in the wrap around effect of the line of major components, and also in that the other components are not zero, as they are in the unaliased case.

When these are reconstructed using as near-perfect reconstruction as possible we get the intensity surfaces shown in figure 8. The ripples in figure 8(a) are not an aliasing artifact but the correct reconstruction of the function; it is, after all, a sum of sine waves. The ripply effect in figure 8(b)*is* due to aliasing. The intensity surface that was sampled had constant shaded stripes with infinitely sharp transitions between the dark and light stripes. The perfectly reconstructed intensity surface shown here perfectly reconstructs all of the aliases caused by sampling the original intensity surface.

By contrast, figure 9 shows the intensity surfaces which result from an imperfect reconstructor: the nearest-neighbour interpolant. The artifacts in figure 9(a) are entirely due to *rastering* (the image contains no aliasing). This is the familiar blocky artifact so often attributed, incorrectly, to aliasing. The artifacts in figure 9(b) are due to a combination of aliasing and rastering. Oddly, it is this latter intensity surface which we tend to find intuitively preferable. This is probably due to the areas perceived as having constant intensity in figure 6(b) retaining this constant intensity in figure 9(b).

How these intensity surfaces are perceived does depend on the scale to which they are reconstructed. The 'images' in figure 6 are of course intensity surfaces but are reconstructed to one eighth the scale of those in figure 8 and figure 9. If one stands far enough back from these larger-scale figures then they all look identical.

It is fascinating that the intensity surface with both types of artifact in it appears to be the prefered one. This is possibly due to the facts that (a) the human visual system is good at detecting intensity changes, hence rippling is extremely obvious; and (b) most scenes consist of areas of constant intensity or slowly varying intensity separated by fairly sharp edges, hence representing them as a bandlimited sum of sinusoids will produce what we perceive as an incorrect result: most visual scenes simply do not contain areas of sinsoidally varying intensity.

Thus there is a tension between the sampling theory and the visual effect of the physical reconstruction on the display device. Indeed many researchers, instead of turning to the perfect sampling theory method of sampling, turn instead to methods which take the human observer into account. The main point to take home here is that, in order to make truly successful compression algorithms, we need an understanding of both sampling theory and of how human beings perceive the world. Thus we get to Markus Kuhn's part of the course, where he discusses practical algorithms which are used for compressing data intended for human perception, i.e. images and sound.

A The inverse Fourier transform of a box function

$$\begin{aligned}
s(x) &= \int_{-\infty}^{\infty} \mathrm{Box}(\nu) e^{i2\pi\nu x}\, d\nu \\
&= \int_{-\nu_b}^{\nu_b} e^{i2\pi\nu x}\, d\nu \\
&= \frac{1}{i2\pi\ x} e^{i2\pi\nu x}\Big|_{-\nu_b}^{\nu_b} \\
&= \frac{1}{i2\pi\ x}\left(e^{i2\pi\nu_b x} - e^{-i2\pi\nu_b x}\right) \\
&= \frac{1}{i2\pi\ x}\left(\cos(2\pi\nu_b x) + i\sin(2\pi\nu_b x) - (\cos(2\pi\nu_b x) - i\sin(2\pi\nu_b x))\right) \\
&= \frac{1}{i2\pi\ x}\left(2i\sin(2\pi\nu_b x)\right) \\
&= \frac{\sin(2\pi\nu_b x)}{\pi x} \\
&= 2\nu_b \mathrm{sinc}(2\nu_b x)
\end{aligned}$$

$\mathrm{sinc}(x)$ is defined in the literature as either $\mathrm{sinc}(x) = \frac{sin(x)}{x}$ or as $\mathrm{sinc}(x) = \frac{sin(\pi x)}{\pi x}$. In this thesis we will stick with the latter.

Bibliography

Blinn, J.F. [1990] "Wonderful world of video", *IEEE Computer Graphics and Applications*, Vol. 10, No. 3, May, 1990, pp.83–87.

Cook, R.L. [1986] "Stochastic sampling in computer graphics", *ACM Transactions on Graphics*, Vol. 5, No. 1, Jan, 1986, pp.51–72.

Crow, F.C [1978] "The use of grayscale for improved raster display of vectors and characters", *Computer Graphics*, Vol. 12, No. 3, Aug, 1978, pp.1–5.

Durand, C.X. [1989] "Bit map transformations in computer 2-D animation", *Computers and Graphics*, Vol. 13, No. 4, 1989, pp.433–440.

Dwight, H.B. [1934] *Tables of Integrals and other Mathematical Data*, The MacMillan Company, New York, 1934.

Fiume, E.L. [1986] "A mathematical semantics and theory of raster graphics", PhD thesis, technical report CSRI–185, Computer Systems Research Group, University of Toronto, Toronto, Canada M5S 1A1. Published, with minor additions, as Fiume [1989].

Fiume, E.L. [1989] *The Mathematical Structure of Raster Graphics*, Academic Press, San Diego, 1989, ISBN 0–12–257960–7. Based on Fiume [1986].

Foley, J.D. and van Dam. A. [1982] *Fundamentals of Interactive Computer Graphics*, Addison-Wesley, Reading, Massachusetts, USA, 1982, ISBN 0–201–14468–9.

Foley, J.D., van Dam, A., Feiner, S.K. and Hughes, J.F. [1990] *Computer Graphics: Principles and Practice* ('second edition' of Foley and van Dam [1982]), Addison-Wesley, Reading, Massachusetts, USA, 1990, ISBN 0–201–12110–7.

Fraser, D. [1987] "A conceptual image intensity surface and the sampling theorem", *Australian Computer Journal*, Vol. 19, No. 3, Aug, 1987, pp.119–125.

Fraser, D. [1989a] "Interpolation by the FFT revisited — an experimental investigation", *IEEE Transactions on Acoustics, Speech and Signal Processing*, Vol. 37, No. 5, May, 1989, pp.665–675.

Fraser, D. [1989b] "Comparison at high spatial frequencies of two-pass and one-pass geometric transformation algorithms", *Computer Vision, Graphics and Image Processing*, Vol. 46, No. 3, Jun, 1989, pp.267–283.

Gonzalez, R.C. and Wintz, P. [1977] *Digital Image Processing*, Addison-Wesley, 1977, ISBN 0–201–02596–5.

Heckbert, P.S. [1989] "Fundamentals of texture mapping and image warping", Master's thesis, report number UCB/CSD 89/516, Computer Science Division (EECS), University of California, Berkeley, California 94720, USA; June, 1989.

Lynn, P.A. and Fuerst, W. [1989] *Introductory digital signal processing with computer applications*, John Wiley and Sons, Ltd, Chichester, 1989, ISBN 0–471–91564–5.

Marion, A. [1991] *An introduction to image processing*, Chapman and Hall, 2–6 Boundary Row, London, 1991, ISBN 0–412–37890–6.

Mitchell, D.P. and Netravali, A.N. [1988] "Reconstruction filters in computer graphics", *Computer Graphics*, Vol. 22, No. 4, Aug, 1988, pp.221–288.

Murch, G.M. and Weiman, N. [1991] "Gray-scale sensitivity of the human eye for different spacial frequencies and display luminance levels", *Proceedings of the Society for Information Display*, Vol. 32, No. 1, 1991, pp.13–17.

NMFPT, [1992] National Museum of Film, Photography, and Television, Bradford, West Yorkshire.

Pavlidis, T. [1982] *Algorithms for Graphics and Image Processing*, Springer-Verlag, 1982.

Pavlidis, T.[1990] "Re: Comments on 'Stochastic Sampling in Computer Graphics' [Cook, 1986]", Letter to the Editor, *ACM Transactions on Graphics*, Vol. 9, No. 2, Apr, 1990, pp.233–326.

Petersen, D.P. and Middleton, D. [1962] "Sampling and reconstruction of wave-number-limited functions in N-dimensional Euclidean spaces", *Information and Control*, Vol. 5, No. 4, Dec, 1962, pp.279–323.

Prasad, K.P. and Satyanarayana, P. [1986] "Fast interpolation algorithm using FFT", *Electronics Letters*, Vol. 22, No. 4, 13 Feb, 1986, pp.185–187.

Pratt, W.K [1978] *Digital Image Processing*, John Wiley and Sons, New York, USA, 1978, ISBN 0–471–01888–0.

Rosenfeld, A. and Kak, A.C. [1982] *Digital Picture Processing* (second edition, two volumes), Academic Press, Orlando, Florida, USA, 1982, ISBN 0–12–597301–2.

Schafer, R.W. and Rabiner, L.R. [1973] "A digital signal processing approach to interpolation", *Proceedings of the IEEE*, Vol. 61, No. 6, Jun, 1973, pp.692–702.

Shannon, C.E. [1949] "Communication in the presence of noise", *Proceedings of the IRE*, Vol. 37, No. 1, Jan, 1949, pp.10–21.

Turkowski, K. [1986] "Anti-aliasing in topological color spaces", *Computer Graphics*, Vol. 20, No. 4, Aug, 1986, pp.307–314.

Watson, A.B. [1986] "Ideal shrinking and expansion of discrete sequences", NASA Technical Memorandum 88202, Jan, 1986, NASA, Ames Research Center, Moffett Field, California 94035, USA.

Whittaker, E.T. [1915] "On the functions which are represented by the expansions of the interpolation theory", *Proceedings of the Royal Society of Edinburgh*, Vol. 35, Part 2, 1914–1915 session (issued separately July 13, 1915), pp.181–194.