Illusions
by Gerald Fitton

It's life's illusions I recall,
I really don't know life at all.

From "Both Sides Now"
A hit for Judy Collins (1939 -) in 1967
Music and lyrics by Joni Mitchell (1943 -)

Joni Mitchell

Joni Mitchell describes herself as, "Primarily a Painter; Secondarily as a Musician". It was she who coined the words for the song, *Woodstock*:

We are stardust
Billion year old carbon
We are golden
Caught in the devil's bargain
And we've got to get ourselves
Back to the garden

Of course you will know that we are all made of the dust strewn in space when a star or two exploded billions of years ago. Indeed we are made from "Billion year old carbon". Also, you will appreciate that the "garden" is the Garden of Eden; this is an allusion to changing our life style to one which welcomes simpler, more virtuous, fundamental, even illusory, values. I would suggest that Joni, like so many singer song writers, is also a poet.

The Discrete Fourier Transform (DFT)

In the last couple of issues of Archive I have described how to use a PipeDream spreadsheet to find the frequency spectrum of a time dependent, periodic wave. The DFT finds the amplitude and phase of a set of harmonically related, sinusoidal frequencies which can be used to reconstruct a wave passing through the original sampled points.

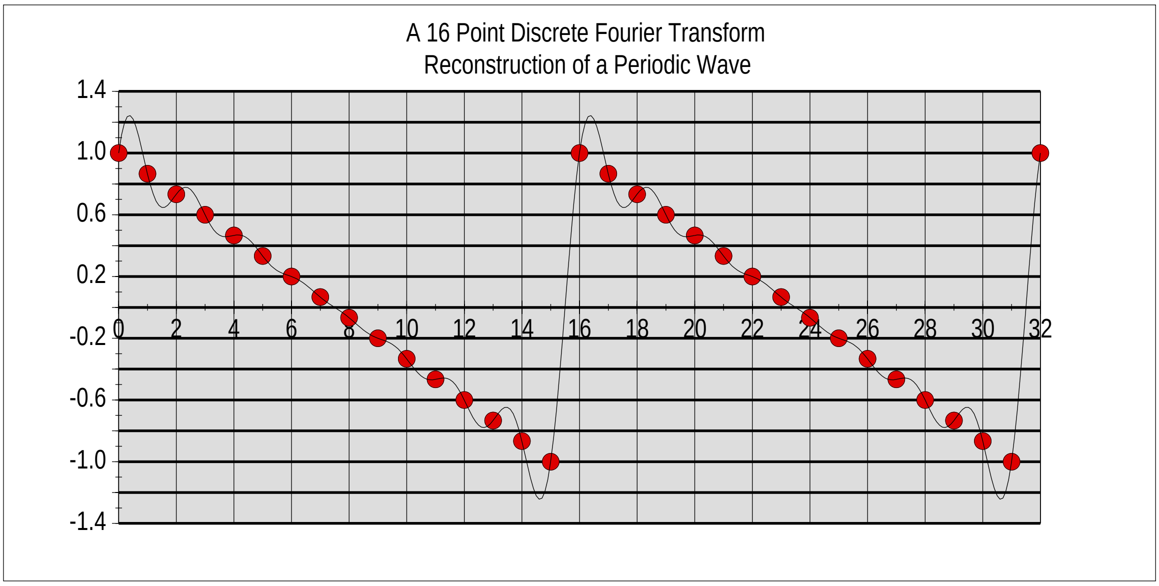

The graphic above shows two cycles of such a reconstructed wave. You will see that the curve which passes through the original 16 points, 0 to 15, is repeated exactly for the 16 points points numbered 16 to 31. The graph is repeated exactly every cycle.

The Discrete Cosine Transform (DCT)

Contrast the graph of the 16 point Fourier Transform with the graph of the 9 point Cosine Transform shown below. The original 9 points are those numbered 0 to 8.

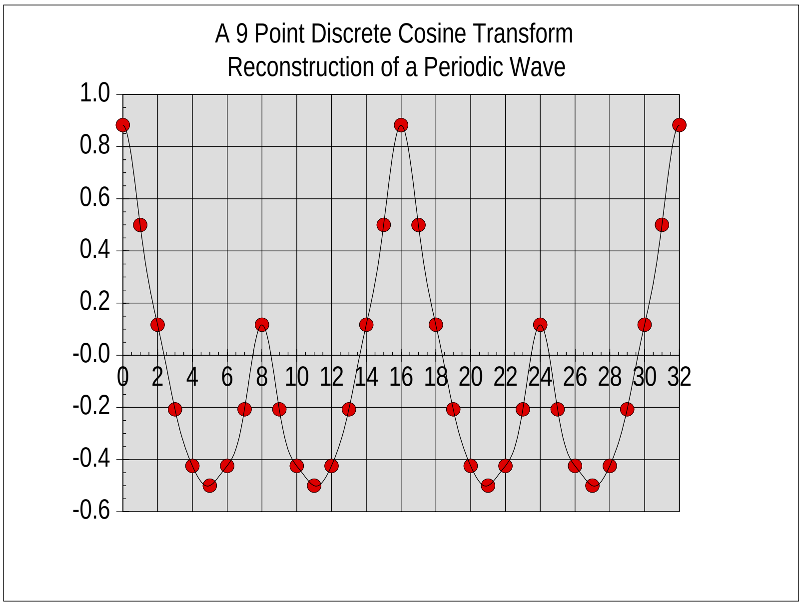

What I would like you to notice is that, outside these original 9 points (numbered 0 to 8), the graph for the points 8 to 16 does not repeat in an identical way. The graph for the points 8 to 16 are a mirror image reflected in the vertical line, x = 8. Similarly, for points 16 to 32, the graph is a mirror image of the graph for the points 0 to 16 reflected in the line x = 16.

Interpolation and Extrapolation

For interpolation, both the DFT and the DCT return an identical reconstruction. It is only when the sample is used to extrapolate outside the sampled region that the differences shown above appear. The reconstruction from the frequencies of the DFT repeats exactly the same waveform every cycle; the frequencies of the DCT returns a series of mirror images of the original waveform.

Why Choose DCT rather than DFT?

Why am I bothering to discuss DCT at all when a DFT provides us with both Amplitude and Phase information in full? The answer is that the DCT is a much more efficient transform, than is the DFT, when extrapolation outside the original points is not required. The DFT requires both cosine and sine functions whereas the DCT uses only the even cosine functions. Examples for which extrapolation is not required are the JPEG compression of a digitally encoded photograph and the MP3 compression of audio.

Compression is used extensively to reduce file sizes by discarding some of the fine detail which our ears, eyes and our brains replace with what we expect to be there. The DCT carries out the data processing necessary to enact this illusion with much greater efficiency than does the DFT because, unlike the DFT, all the sine componets are discarded.

The Amplitude Spectrum

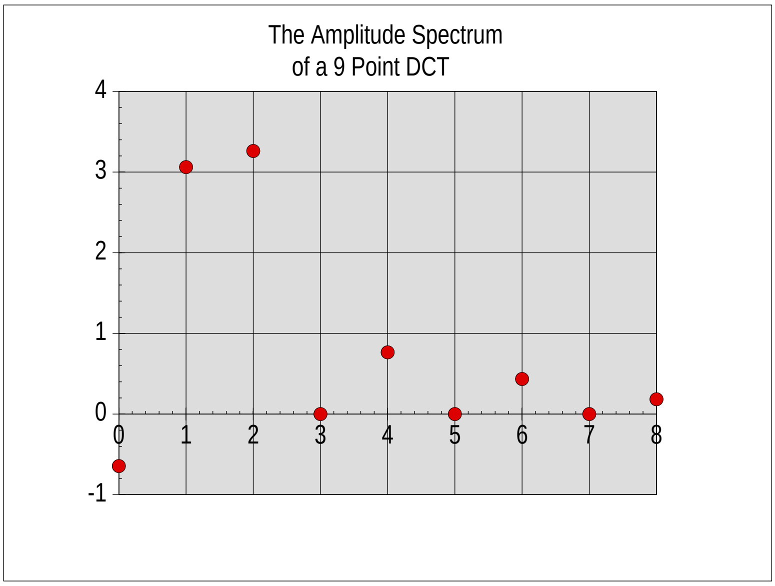

The amplitude spectrum of the 9 point DCT is shown in the above screenshot. The amplitude of the lower frequencies is greater than the amplitude of the higher frequencies. This is typical of DCTs, including multi dimensional DCTs, but is more unusual for a DFT; it is this fact which makes the DCT a more efficient means of compression than the DFT.

Reconstruction

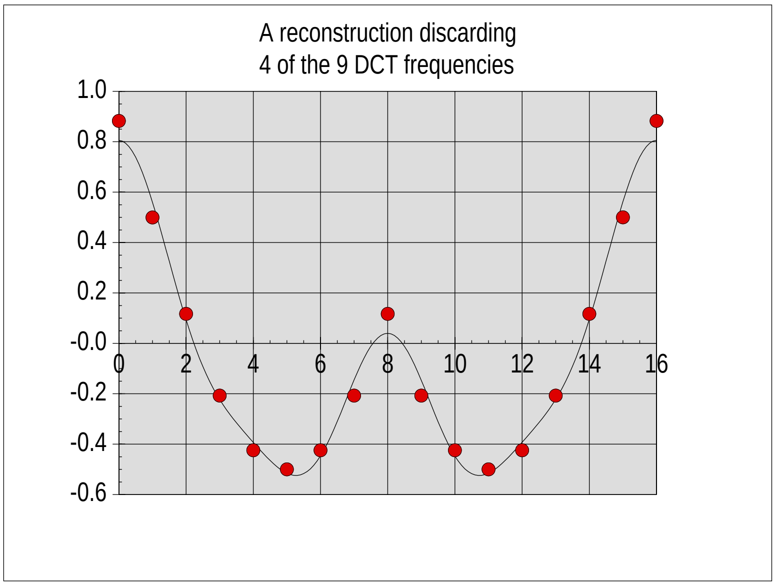

In the reconstruction above, the original 9 points are those numbered 0 to 8. In use, the original is reconstructed only between the points 0 to 8; I have included points 9 to 16 solely for reference. In this reconstruction, rather than use all 9 frequencies, I have discarded the higher frequencies. The resulting waveform is a close approximation to the original. The drawfile shows in red the original 9 points (plus the reflected points); the black line shows the reconstruction from only 5 of the 9 original amplitude components.

The error between the original and the reconstructed values is about 10%. As we shall see in the pictures below, this 10% error is not easily discernible to the ear nor the eye because we 'fill in the gaps' with illusory, smooth lines. It is the illusion which we hear or see.

Two Dimensional Discrete Cosine Transform

I shall not continue discussing one dimensional DCTs even though they are used extensively for creating the ubiquitous MP3 files which now fill countless portable devices. Instead I shall consider what is usually called a two dimensional DCT.

My 2D DCT example is a picture of a lady called Lenna. This picture of Lenna is famous because it has been, and still is, used extensively to test image processing algorithms.

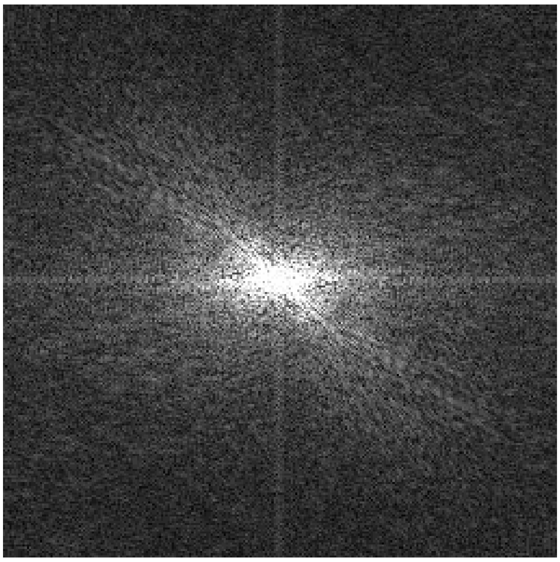

The graphic which looks a little like a galaxy of stars, is a pictorial representation of the Discrete Cosine Transform of this picture of Lenna. The brightness of the 'stars' represents the amplitude of different frequencies with the lower frequencies being near the centre.

The classic image of Lenna (which I have used here) is 512 x 512 pixels. The fundamental frequency of the DCT is such that half a wavelength spans 512 pixels both vertically and horizontally; the amplitude of these 'signals' are placed nearly at the centre of the graphic. The highest frequency is such that half a wavelength is just 1 pixel wide; the amplitude of these components are placed on the perimeter of a circle surrounding this centre point.

Generally, if the photograph consists of vertical lines, then the 'stars' will fall in a horizontal line. The picture of Lenna has lines running from the bottom left to the top right, mainly her hat. This is why there is a band of 'stars' running from top left to bottom right.

Reconstruction

My second picture of Lenna is 256 x 256 pixels. It is a reconstruction from the DCT but with the highest frequency components omitted. The graphic on the right is a 3D representation of the amplitudes of the frequency components making up the 2D DCT.

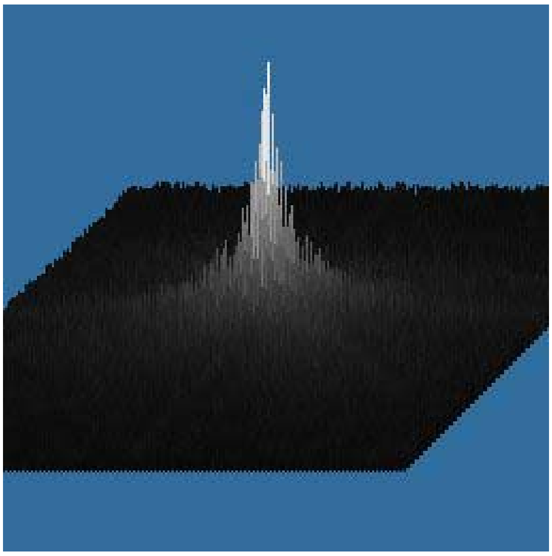

Although you might not realise it when you look at the full 256 x 256 reconstruction, the picture is very slightly, almost unnoticeably, blurred when compared with the original 512 x 512 pixels. However, if you look at the blown up version of a detail taken from each of the two pictures, you will be able to see this difference. The picture on the left is clipped from the original (512 x 512 pixels); the one on the right is clipped from the reconstructed picture (256 x 256 pixels). The reason that we don't notice these genuine differences when we look at the full pictures is that, whatever part of our brains we use to process the pictures, fills in the jagged bits with the smooth lines we expect; we experience this as 'slight blurring'. The reduction in resolution and file size is no illusion; the illusion, that the two pictures are virtually the same, takes place within our brains.

Box-Out Lenna

The November 1972 issue of Playboy featured a full length picture of Lenna Söderberg. A group of engineers at the University of California Signal and Image Processing Institute needed a digitised image for a conference paper. The standard test images which were around at that time were television test pages with the usual sharp edged lines and circles; they wanted something eye catching with graduated shading for their presentation.

They found the November 1972 issue of Playboy with a full length Lenna.

Their scanner was a Muirhead 'wirephoto' scanner. In 1973 this machine was in common use by newspapers to send pictures around the world by 'wire'. It worked rather like a cylindrical drum type, fax machine. The 'standard' was 100 lines per inch; they wanted a 512 x 512 image so they used the top 5.12 inches of Lenna. They wrapped the cropped picture showing just her head and shoulders around the drum of the scanner and created the first digitised picture of the now famous Lenna.

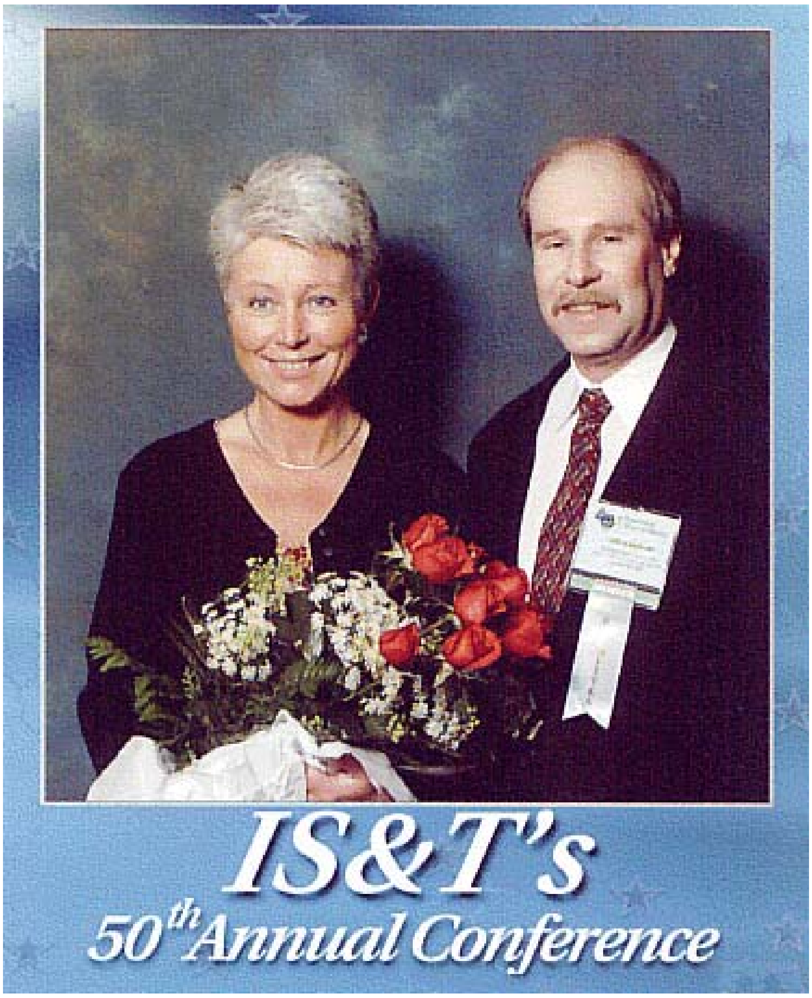

In 1997 Lenna was a guest speaker at the 50th annual conference of the Society for Imaging and Technology; she gave a presentation about herself. Currently she lives in Sweden with her husband and 3 children. Today her charity work includes helping handicapped people use computers. If you google "First Lady of the Internet" you will find that this is Lenna Söderberg.

Of course the original full length image of Lenna taken from the November 1972 edition of Playboy (their best selling issue - over 7 million copies) can be found easily on the internet but, rather than include it here in the Archive magazine, for my second picture of Lenna, I have chosen one taken in 1997 at the 50th annual conference of the IS&T.

Bio-Bit

Gerald is aware that we are all stardust - stardust which is under the illusion that it is alive. However, as Joni says, "It's life's illusions" that he experiences daily. Gerald is under no illusion that "Life's Illusions" is the only reality we have!