A Guide to the Fourier Transform

The first of two parts by George Brown, BSc, PhD, CEng, FIEE, G1VCY*

IGNORED BY MANY, revered by some, the techniques pioneered by Fourier form the basis of much of the theory of communications. We all use them (perhaps unwittingly) when we discuss the bandwidth of a signal. Here, just enough of the specifics of the subject are presented in order that those interested may pursue the concept in more detail.

HISTORY

JEAN BAPTISTE JOSEPH FOURIER, born in 1768, fought for Napoleon I in Egypt, and was an expert on Egyptian antiquities. He was made a Baron by Napoleon in 1808 and was elected to the Academy of Sciences in 1816. It was in 1822 that he published *The Analytical Theory of Heat* and, from this arguably obscure treatise, came the mathematics which were to change the evolution of science. In it, he used a trigonometrical series (later to become known as the *Fourier Series*) by which a complicated mathematical function can be expressed as a sum of an infinite series of sine and cosine functions. He was, justifiably, elected to the French Academy in 1827, three years before his death.

INTRODUCTION

FROM THESE UNLIKELY roots have evolved the techniques which are so fundamental to our appreciation and understanding. With the common availability of the computer, non-mathematicians can now enjoy realising the potential of Fourier's work without needing to understand the intricacies of the equations involved. Just as electronics is becoming more biased towards systems as being assemblies of 'black boxes', with no knowledge being necessary of what goes on inside them, so the worlds of mathematics and communications are being opened-up by the use of software packages which can be used as tools to solve specific problems. Such problems can now be addressed (and solved) by those who, ordinarily, could never understand the underlying complexities.

This article was prompted by the description

1 Langford Crescent, Benfleet, Essex SS7 3JP.

of the use of digital signal processing (DSP) by Peter Martinez, G3PLX, for the study of long-delay echoes, reported in *TT* in the January and February 1998 editions of *RadCom*. Because the Fourier transform and DSP are very closely linked, a simple discussion of the principles involved was thought to be timely. AF9Y has also reported using the FFT to identify moonbounce signals, and to detect the 70cm beacon from the Mars Global Surveyor (complete with 30Hz earth-rotation Doppler shift) at a distance of 5 x 10⁶km! [1].

Unfortunately, it is still true that Fourier Series are quickly dismissed in A-level maths

Fig 1: A trace of power spectral density, from *one line* of Photo 1. This can be compared directly with the graph in Photo 2, which averages all 370 lines of the display.

courses, with little or no mention of their application in *any* area of science, thus consigning them at the first opportunity to the student's waste paper basket. Fourier's work, complex though it may be, is readily understandable in a general sense, provided it is acknowledged at the outset that its basic premises are true. We can paraphrase these for our own purposes as:

● Any complex waveform can be *synthesised* by the addition of a number of sine waves of different frequencies, amplitudes and phases.

And its converse statement:

● Any complex waveform can be *analysed* into a number of sine waves of different frequencies, amplitudes and phases.

Strictly speaking, the words 'a number' should be replaced by 'an infinite number', but since we are restricting our descriptions to practical cases the word 'infinite' is not in our vocabulary, and we are restricted to using finite numbers. This statement, as will be realised later, is simply recognising the very practical restriction that the bandwidths of our

signals (at audio and RF) must be limited! The two statements also introduce the concepts of *analysing* a signal *into* its components, and *synthesising* a signal *from* its components; thus we talk about Fourier Analysis or Fourier Synthesis, depending upon the *direction* in which we are interested. In general, the expression 'Fourier transform' is used, and is applicable both to the synthesis and analysis processes. The transform enables processes which occur in the *time domain* (e.g. changes of any parameter with *time*) to be examined in the *frequency domain* (e.g. changes of an equivalent parameter with *frequency*) and *vice versa*.
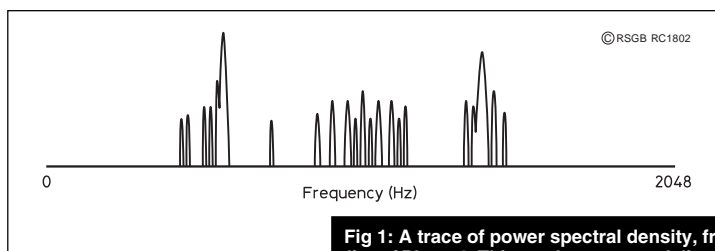
ANALYSING A (RTTY) SIGNAL

BY FAR THE MOST common application of Fourier's work to amateur radio lies in the *analysis* of an existing signal. This can be a live signal, extracted from the audio stages of a transceiver, or a recorded version of the same signal. The two should be the same, provided the recording process does not introduce artefacts of its own (bandwidth reduction, wow and flutter, etc). Computer programs exist which will perform the analysis process very rapidly, depending upon the speed of the particular processor. My personal favourite, which is extremely easy to use, is *SbFFT* by Kevin McWilliams, KW5Q [2]. It uses a digital technique known as the *Fast Fourier Transform* (FFT), and displays on the screen an intensity (or colour) plot of frequency horizontally and time vertically. It requires a Soundblaster SB16 (or 100% compatible) sound card, into which the audio waveform is fed from the signal source. The display is in real-time, and can be saved as a PCX file for future study.

An example of the display of the *SbFFT* programme is shown in **Photo 1**. The legend on the top five lines need not concern us; it relates to the parameters that can be used to customise the display. The sixth line is the frequency axis calibration, in this case extending from 0 to 2048Hz. The horizontal axis is frequency, and each horizontal line of the display is a single

Fourier transform. The power (ie a number proportional to the voltage squared) at any frequency is represented by the brightness of the display along a single line. The process is repeated, and the next transform is displayed along the next lowest line of the screen, and so on. This enables 370 lines of FFT to be displayed on a low-resolution screen (640 x 480 pixels). On a high-resolution display (1024 x 768 pixels), 658 lines (transforms) are available. The vertical time axis, which starts at the top and progresses downwards, is not calibrated, as it is seldom used for measurement purposes. In this particular case, the line rate was 4 lines per second, thus completing the 370 active-line display in 370 / 4 = 92.5 seconds. The signal being analysed is a normal, two-tone commercial RTTY transmission. Purists should note that the signal is not tuned-in as it would be for normal reception, but to centre it in the display. In this way no information is lost. The two tones are the two prominent vertical lines. The 'graininess' of the picture indicates the presence of noise (of which there was a significant amount), and the effects of frequency-switching transients. The cursor can be seen at the centre of the display. This can be moved about to any frequency of interest, and the value of that frequency is given beside the word 'Freq' on line two of the legend. The display in question has a horizontal frequency resolution of 4Hz per pixel, meaning that any frequency up to 2048Hz can be identified within 4Hz. The two tones can be measured in this way to be (592 ± 4)Hz and (1436 ± 4)Hz, showing the frequency shift of the RTTY signal to be (844 ± 8)Hz, the standard value being 850Hz.

The second feature of interest is the series of vertical stripes between (and outside) the two RTTY tones. These are caused by the frequency switching which is used to produce the binary codes corresponding to the characters being transmitted. All modulation, of whatever sort, causes extra frequencies to be added to the basic spectrum. The frequency spacing of these subsidiary lines will tell us the frequency of the switching. Moving the cursor to a peak near the upper tone gives a frequency of 1300Hz; repeating the measurement 11 lines lower (near the lower tone) gives 750Hz, showing a frequency separation of (550 ± 8)Hz. Because of the 11 subsidiary

lines between these measurements, the spacing of adjacent lines is 550/11 = (50.0 ± 0.7)Hz. This particular station uses RTTY at 50 baud (50 bits per second), hence the appearance of the harmonics of 50Hz!

The third feature of this photograph was not intended at the outset, but appeared as the data capture proceeded - the diagonal 'empty' areas between the two tones. This is a classic example of frequency-selective fading. Anyone who has listened to HF RTTY signals knows that, from time to time, one or other of the two tones disappears, and the decoder (depending upon its design) will either carry on working or produce gibberish until the missing tone returns. This behaviour is very interesting to observe on the photograph. It begins at high fre-

quency, almost removing the upper tone, and then reduces its frequency quite steadily, taking out the 50Hz harmonics as it does so, then removes the lower tone. By this time, another round of fading has begun at the upper tone, and so it continues. As time passes, the time interval between successive rounds of fading increases.

This photograph shows the intriguing nature of the FFT, especially when displayed as a time-dependent feature, as with the *SbFFT* software. It is easy to spend hours tuning through various signals and studying their spectra.

Few signals can be received noise-free, as the photograph shows. If a single line of this FFT is displayed as a simple graph of power against frequency (**Fig 1**), it means surprisingly little, because of the noise. However, for signals like this one where the noise is random, and where the spectral information of the wanted signal is not radically changing, there is a way in which science can come to the

rescue of noisy FFTs. It uses the fact that noise is random in nature and will average (given infinite time) to zero. The 'wanted' signal spectrum remains essentially unchanged. Consequently, if the FFT is averaged over all 370 lines, the noise will reduce and the signal will be seen more clearly. My own programme takes the FFT display from *SbFFT* and averages it, line by line, over the 370 lines of the original FFT. The result is shown in **Photo 2**.
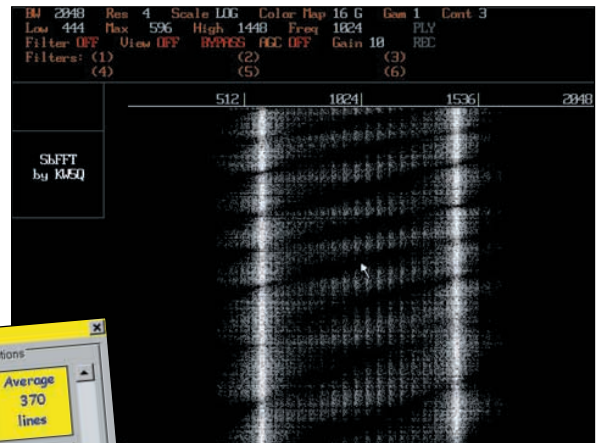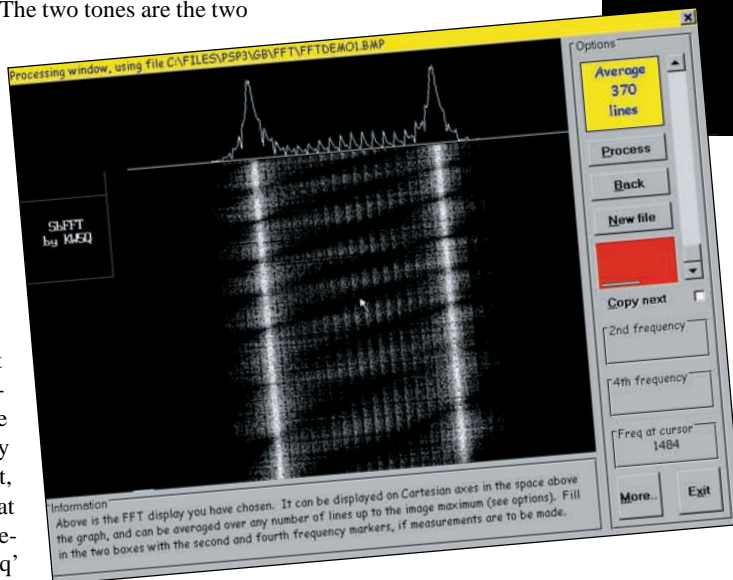
Photo 1 (above): SbFFT display of a RTTY signal, showing the frequency shift of the keying, and fading of the signal.

Photo 2 (left): RTTY signal, with the averaged spectrum displayed above it.

The averaged spectrum is displayed in the area above the original data, on the same frequency axis, and can be compared directly with the single-line graph of Fig 1. Averaged signals can usually be measured more accurately than unprocessed data, although time-dependent displays do show artefacts that averaged data can remove. Frequency-selective fading is a good example of this.

Next month we will take a look at another interesting application, namely Doppler shifts. In view of the article by Peter Martinez, G3PLX, in the May *RadCom*, it is a topic which is extremely timely. Afterwards, an explanation of the techniques involved will be given.

REFERENCES

[1] 'A Conversation with Mike Cook, AF9Y', *QST*, January 1998, p9 and pp 56, 57. The FFT software used here is available from http://www.webcom.com/af9y, and costs $35 to register.

[2] The DOS freeware programme *SbFFT*, by Kevin McWilliams, KW5Q, is available by anonymous ftp from ftp.funet.fi/pub/ham/. Alternative sources have been mentioned by Roger Cooke, G3LDI, in his *Data* column (*RadCom* July 1997 and February 1998).

♦

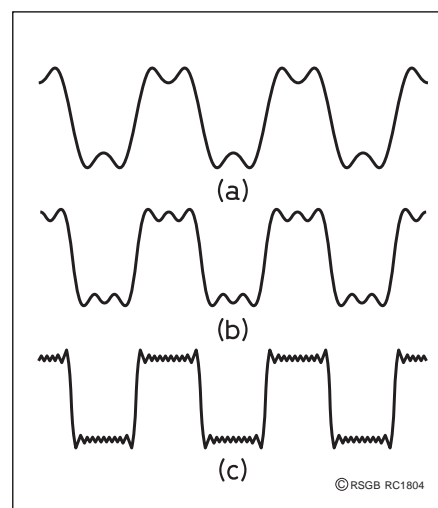***. . . to be continued***

A Guide to the Fourier Transform

The second of two parts by George Brown, BSc, PhD, CEng, FIEE, G1VCY*

I N THE FIRST part of *A Guide to the Fourier Transform* we looked at what the Fourier transform is, and a sample application of *SbFFT* software. This month we look at another application, then go on to explain the techniques involved.
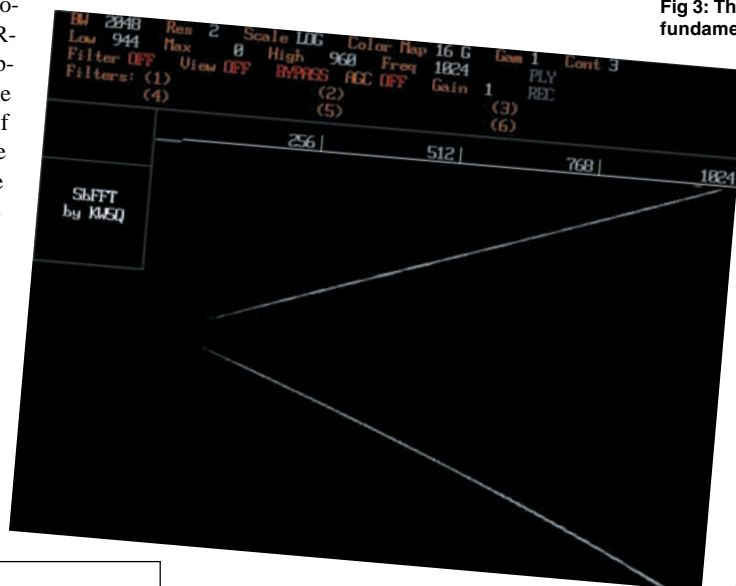
FFTS AND DOPPLER SHIFTS

ONE OF THE MOST eye-catching and interesting uses of programmes like *SbFFT*, is to display the Doppler-shifted signal received from a polar-orbiting satellite. First, a few words about the equipment and the receiver settings are in order. A Kenwood R-5000 receiver with a VHF adapter was used, the aerial being a standard 2m colinear, mounted in the house loft. The 136.770MHz beacon on the NOAA-12 weather satellite provided the signal source. The R-5000 was set for USB reception, the idea being to use the carrier insertion oscillator of the receiver to mix with the incoming carrier, and produce an audible beat note in the speaker. The beat note will be the absolute difference between the frequencies of the carrier insertion oscillator and the incoming signal. With the receiver tuned to the source frequency of 136.770MHz, the beat note

*1 Langford Crescent, Benfleet, Essex SS7 3JP.

is exactly equal to the magnitude of the Doppler shift! What is heard is a high-pitched whistle which slowly decreases in frequency, until it is finally removed (at around 200Hz) by the receiver filtering. The satellite is at its closest approach when the two frequencies are at zero-beat. Unfortunately, this crucial moment is missed. However, the R-5000 has the facility to remove the normal sideband filtering so that heterodynes are heard from frequencies in what would normally be the 'other sideband'. Switching the filters out in this way allows the Doppler shift to be heard as it comes down in frequency, is lost around 200Hz due to the RC-coupled audio stages, and then reappears at 200Hz again, this time as a rising frequency. **Photo 3** shows just what happens. An equation exists [3] which allows the mini-
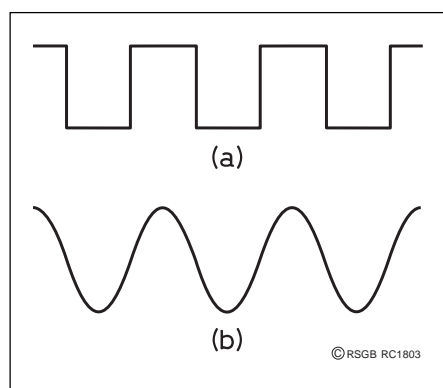
mum slant range of the satellite to be calculated from the maximum rate of change of the Doppler shift, and I have written a programme which automatically calculates this range from the display shown in this photograph. The results agree very well with theory. It is also useful, and instructive, to compare the FFT displays for near-overhead satellite passes, and other more distant passes. The characteristic 'V' shape has a very narrow apex angle for an overhead pass, the angle increasing for passes at greater distances, as the rate of change of received frequency decreases.

Fig 3: The waveform with (a) two components (the fundamental and the 3rd harmonic), (b) three components (the fundamental, 3rd and 5th harmonics), (c) the first 10 terms (ie with harmonics up to the 19th).

Photo 3 (left): The Doppler shift of a passing weather satellite.

A final (practical) note about the *SbFFT* programme. It also allows six independent digital filters to be simultaneously applied to the signal, and the result fed to the speakers attached to the sound card. Each filter can be toggled in or out of the signal chain at will. These filters all have -3dB to -50dB skirts of 25Hz, and passband filters down to 1Hz bandwidth are possible. CW signals only a few Hertz apart can be resolved in this way, making it a truly remarkable package!

FOURIER SYNTHESIS AND ANALYSIS

SOME EXPLANATION OF the processes underlying such applications of the work of Fourier now helps to build an appreciation of the techniques involved. Let us start with the process which is the reverse of that which we have just considered.

Fourier synthesis concerns the creation of a complex signal from simple sine waves. Consider the rectangular wave of **Fig 2(a)**. There

Fig 2: (a) The desired waveform. (b) The fundamental frequency sine wave (the first term of the series).

is little comparison between its shape and that of a pure sine wave of the same frequency, **Fig 2(b)**, except that each has an equal positive and negative excursion, once per cycle. Despite this, Fourier says that Fig 2(a) *can* be produced by the sum of sine waves alone. This can be proved very simply on a computer [4], using the standard Fourier equations for the components of a rectangular wave [5]:

$$V = \frac{4}{\pi} E \left[\cos(f) - \frac{1}{3}\cos(3f) + \frac{1}{5}\cos(5f) - \frac{1}{7}\cos(7f) + \ldots \right],$$

where:

$\frac{4E}{\pi}$ allows a particular amplitude of the resultant waveform to be specified, and f is the frequency of the synthesised wave.

Notice the following aspects of the equation:

● the signs of the terms alternate, positive and negative;

● the terms represent *only odd harmonics* (3, 5, 7, etc) of the fundamental frequency, f;

● each term is preceded by a decreasing factor ($\frac{1}{3}$, $\frac{1}{5}$, $\frac{1}{7}$, etc) showing that the amplitudes of the harmonics decrease quite rapidly;

● the dots at the end show that the equation can be extended as far as is meaningful in any particular context.

Using this equation, the shape of the resultant waveform can be derived for any number of terms in this expression. **Fig 3** shows the resultant waveform for 2, 3 and 10 terms. It is easily seen that the first term gives rise to the basic frequency of the final waveform, while the harmonics serve to add the fine detail to the wave - the rapid rise and fall and the sharp corners - both being attributes which we intuitively associate with a high frequency response. The more terms we use in the evaluation, the better the final wave approximates to the shape of Fig 2(a).

The same can be done with a sawtooth wave, shown in **Fig 4(a)**. This wave can be synthesised using the Fourier Series:

$$V = \frac{2}{\pi} E \left[\sin(f) - \frac{1}{2}\sin(2f) + \frac{1}{3}\sin(3f) - \frac{1}{4}\sin(4f) + \ldots \right]$$

This series has *all* harmonics, both even and odd; the terms still alternate their signs, and the denominator of each fraction matches the harmonic number. **Fig 4(b)** shows the fundamental frequency used as a basis for the

synthesised wave. **Fig 4(c)** shows the wave synthesised with four terms only, ie using the expression as it stands, with no additional terms. Notice that the peak of the fundamental does not coincide with the peak of the sawtooth; its amplitude is less than that of the sawtooth; as before, the smoothness of the rise and the sharpness of the fall are both due to the harmonic content of the wave.

All the diagrams shown so far are illustrations of processes taking place in the *time domain* (characterised by the use of *time* as the horizontal axis).

The frequencies needed to synthesise the two waveforms are shown in **Fig 5**. Note that, as the equations suggest, the higher harmonic

Fig 4:(a) The sawtooth wave to be synthesised. (b) The fundamental frequency sine wave. (c) The fundamental frequency sine wave with the first four terms (ie with harmonics up to the 4th).

Fig 5: (a) The component frequencies and intensities for the rectangular wave up to the 9th harmonic. (b) The component frequencies and intensities of the sawtooth wave up to the 10th harmonic.

terms become smaller and smaller in intensity.

BASEBAND FREQUENCIES AND BANDWIDTH

FOR OUR PURPOSES, we can define *baseband frequencies* as those present in the modulating signal *before* they are used to modulate the carrier. **Fig 5(a)** shows the

baseband frequencies of a rectangular wave, and **Fig 5(b)** those of a sawtooth wave. The range of baseband frequencies, from the fundamental to the highest harmonic present, defines the *bandwidth* of the modulating signal. These diagrams show that, for practical purposes, there comes a point beyond which adding higher harmonics becomes inefficient, as their presence changes the waveform only slightly at the expense of greatly increased bandwidth. This illustrates the well-known compromise between acceptable intelligibility at narrow bandwidth and acceptable fidelity at wide bandwidth.
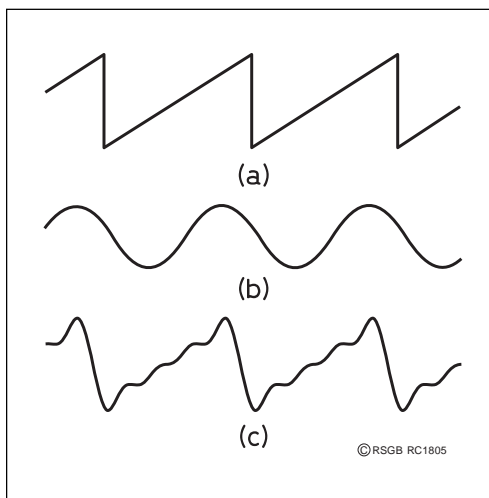
Fig 5 shows two examples of processes occurring in the *frequency domain*. They can be transformed into their rectangular-wave and sawtooth-wave time-domain equivalents (and back again) by repeated use of the Fourier transform.

Fourier synthesis performed in this way requires the availability of a mathematical equation which specifies the frequencies and amplitudes to be used. In the case of amateur radio, this is not very common. However, the way in which a well-known waveform is built up from its component parts, as in Fig 2 and Fig 3, gives an intuitive 'feel' for the way in which other waveforms will be affected.

By now, you should have no real problem in accepting that any wave can be synthesised from a set of sine-wave components. Performing the reve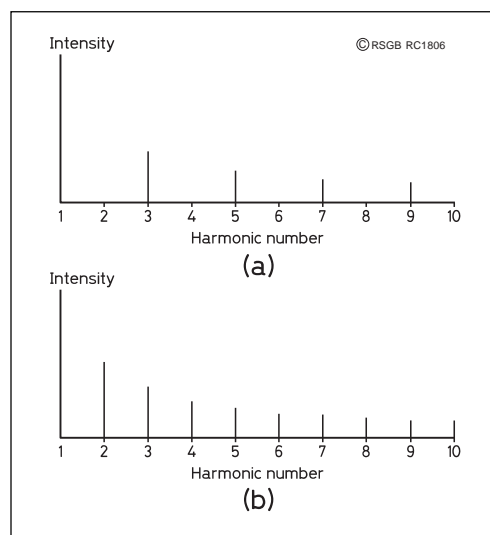rse process is computationally more complex, but should not represent a conceptual problem. Any 3-year old will know that he can build a house with wooden blocks, and then demonstrate with great glee the reverse process! The reverse process in our case is called *Fourier analysis*, analysing a wave into its building blocks - a number of sine waves.

We have looked at the use of the *SbFFT* [2] freeware programme for the analysis of RTTY and Doppler signals in Photo 1 and Photo 3. Mathematical software packages, such as *Mathcad* [6], can also perform the Fourier transform on digital functions, allowing theory and practice to be compared. Some of the graphs used later on were derived using this software.

The interpretation of the output of a standard FFT requires a little mental agility. Take
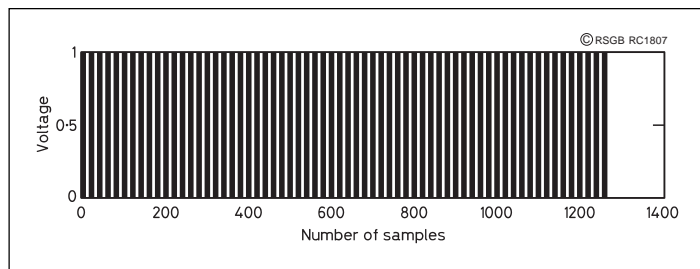
Fig 6: A digital input signal comprising a rectangular wave of 10 samples of 1V and 10 samples of 0V, repeated 64 times.

an example, shown by **Fig 6**, of an input signal consisting of a rectangular wave. This wave is represented digitally in the form of 1280 samples, which means that the voltage of the wave has been examined at 1280 equal increments of time, and the results plotted. The wave is a repetition (64 times) of 10 samples of 1V and 10 samples of 0V. The 1V areas are blocked in, so that it is visually obvious that the wave is rectangular, and has a 1:1 mark-space ratio.

The unique part of the FFT output (now in the frequency domain) is shown in **Fig 7**, and is taken from the *Mathcad* programme. There is little doubt that it is identical with Fig 5(a), but the labelling of the horizontal axis is not at all clear. I shall try to keep the explanation brief! The maximum frequency that can be represented by our 1280 samples is represented by one in which the sample values are '10101010 . . .'; all other combinations of zeroes and ones give lower frequencies (which is a first-principles statement of the famous theorem attributed to Shannon and Nyquist, that the lowest possible sampling frequency must be at least twice that of the highest frequency present in the signal). It is not necessary to know the time period of the 1280 samples; this is quite arbitrary. To make things simple, suppose this time period is one second. Our maximum frequency has a wavelength of two samples, since it is the signal '10' repeated 640 times to give the 1280 samples. 640 cycles in the one second signal interval correspond to a frequency of 640Hz. Our hypothetical input signal (Fig 6), occupies 20 samples, therefore it must have a frequency which is $\frac{1}{10}$ that of the highest frequency, making it $\frac{640}{10}$ = 64Hz.

The graph in Fig 7 has 640 entries, corresponding to frequencies from DC to 640Hz, in 1Hz increments. If the graph is examined closely, there are entries *only* at 64Hz, 192Hz, 320Hz, 448Hz and 576Hz, corresponding to the fundamental frequency together with the 3rd, 5th, 7th and 9th harmonics! The amplitudes of these signals are in the ratio 1: $\frac{1}{3}$: $\frac{1}{5}$: $\frac{1}{7}$: $\frac{1}{9}$, as we would expect from our original synthesis of the rectangular wave. We need not concern ourselves about the absolute magnitudes of these individual entries, as it is the relative magnitudes in which we are interested. Our allocation of one second to the 1280 samples is also quite arbitrary, as the fundamental frequency will always be $\frac{1}{10}$ of the highest frequency for this particular waveform; all the other frequencies will maintain the same relationship to each other.

We have now demonstrated the two converse activities of breaking a rectangular wave into its sine-wave components, and reconstituting it by the algebraic addition of these components. This Jekyll and Hyde existence of waves permits great flexibility in their processing, because this can occur in the time domain or in the frequency domain. Although digital processing in the frequency domain is much more complex, the results *can* be quite remarkable, as shown by the examples given here. Perhaps this article may stimulate anyone interested in waves and their significance to experiment further, using packages such as *SbFFT* and *Mathcad*.
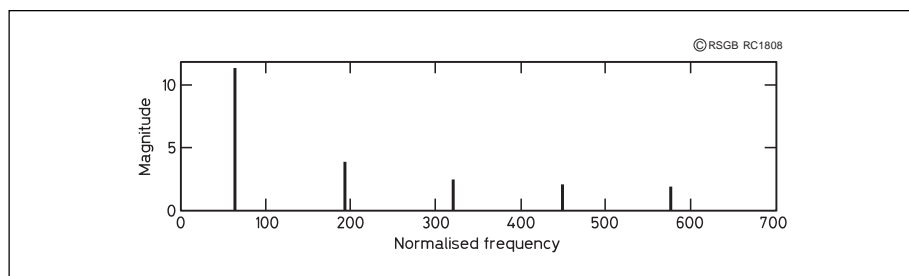
APPENDIX

THIS SIMPLE THEORY of Amplitude Modulation (AM) serves to illustrate the principle of Fourier Analysis, but by a different route.

Suppose we have a carrier at a frequency f_c and a single modulating frequency f_m. These are represented by the formulæ:

$$V_c = \sin(2\pi f_c t) \text{ and } V_m = \sin(2\pi f_m t)$$

The modulation process effectively multiplies these two together, and introduces the depth of modulation, m [m = 0 for zero modulation, and m = 1 for 100% modulation]. The resultant modulated signal thus becomes:

$$V = [1 + m \sin(2\pi f_m t)] \sin(2\pi f_c t)$$

For simplicity, let us assume that m = 1. The expression reduces to:

$$V = \sin(2\pi f_c t) + \sin(2\pi f_m t) \text{ x } \sin(2\pi f_c t)$$

Recourse to an A-level book of trigonometrical formulæ will produce an equivalent form for the product of two sines:

$$\sin(A) \text{ x } \sin(B) = \tfrac{1}{2} \cos (A-B) - \tfrac{1}{2} \cos (A+B)$$

Applying this to the previous equation, we have:

$$V = \sin[2\pi(f_c)t] + \tfrac{1}{2} \cos [2\pi (f_c - f_m) t] - \tfrac{1}{2} \cos [2\pi (f_c + f_m)t]$$

This gives the same result that Fourier theory would derive - that the radiated signal has *three* discrete frequency components, at frequencies $(f_c - f_m)$, f_c and $(f_c + f_m)$. These frequencies have been put in round brackets in this expression to facilitate their identification.

REFERENCES

[2] See part one of the feature in the June *Radcom*.

[3] *The Satellite Experimenter's Handbook*, by Martin Davidoff, K2UBC, ARRL, 1st ed, p10-2, 2nd ed, p13-2.

[4] The programme FOURIER.EXE, which runs under DOS, was written some years ago by the author for use with sixth form students. It shows how rectangular, triangular and sawtooth waves are made up from their individual components, using any number of terms. A copy of it may be found in the C:\FOURIER directory of GB7BAS. #32.GBR.EU, and can be obtained via a YAPP download for those who can access the BBS directly, or in smaller segments by using the REQDIR and REQFIL commands on this same directory.

[5] *Radio Engineers' Handbook*, by Frederick E Terman, McGraw-Hill (New York) 1943, p20.

[6] The Windows programme *Mathcad* has been available for several years now, and is currently supplied by Adept Scientific (see current computing press for details). ♦

Fig 7: Fourier transform of the waveform of Fig 6. See the text for an explanation of the frequency axis.