Signal Phase Cannot Be Determined from a Fourier Transform of Sampled Data

The Discrete Fourier Transform is a well-established tool for obtaining a frequency spectrum from a set of sampled data. The coefficients obtained from the DFT are complex-valued and so both phase and amplitude can be extracted. However, ***David Gibson*** *asserts that if the sampled data originated in the real world (say, as a recording from a broadband radio antenna) the phase of the spectral components cannot be determined. This, he claims, is an 'obvious' observation, but one that is not emphasised in textbooks. The details of the problem are left for the reader to work out, as an interesting 'examination question' in DSP.*

The Fourier Transform

The Fourier transform is a mathematical equation that converts a function of time (i.e., a 'signal') into the frequencies that make it up, similarly to how a musical chord can be expressed as the amplitude of its constituent notes. So says the Wikipedia, which goes on to say that the Fourier transform of a function of time is a complex-valued function of frequency, whose absolute value represents the amount of that frequency present in the original function, and whose complex argument is the phase offset of the basic sinusoid in that frequency.

We can consider this further… A Fourier Transform is one of a class of so-called integral transforms that work on functions evaluated from $-\infty$ to $+\infty$, and that produce a function of a different variable, expressed over the same interval. In our applications, we work with finite time intervals and there is an underlying assumption, in the maths, that our finite signal will repeat itself from $t=-\infty$ to $t=+\infty$. With this proviso, the maths 'works'; and we obtain a set of discrete frequencies rather than a continuous spectrum. For example, we can envisage a square wave that exists from –0.5 ms to +0.5 ms, changing state at $t = 0$ and repeating itself endlessly outside this interval. Its **Discrete Fourier Transform** will be an infinite set of discrete frequencies, at 1kHz intervals

A second modification to the basic FT is that we are usually dealing with sampled data waveforms. If the square wave in the above example were sampled at 1µs intervals, we would have 1000 samples. We know that, within this sampled-data environment, the highest frequency we can represent is 500kHz so, when we calculate the FT of the waveform, we do not need to consider frequencies higher than this.

A brute-force algorithm for calculating the coefficients of a 'Discrete Fourier Transform of Sampled Data' can take a large amount of computer time. Various short-cuts exist, the most common of which is known as the Cooley–Tukey algorithm, which gives rise to the term ***Fast Fourier Transform***. For highest efficiency, an FFT requires the number of samples to be a power of 2; and so, for example, a 1024-point time sample is transformed into 1024 complex values of frequency. The top 512 of those are 'essentially' repeats of the lower set. They do not correspond to anything physical, but are an artefact of the algorithm. (The same is not true if the inputs are complex-valued. Obviously, time itself is 'real' but the FT is an abstract tool).

To summarise; if we take a set of time samples, we can apply a Discrete Fourier Transform and produce a complex-valued dataset, from which we can derive the amplitude and phase of the sine waves that comprise the signal.

However, this is not as straightforward as it sounds and, in certain situations, the phase information cannot be relied on. This is one of those observations that is 'obvious' when one considers it in detail but, because it does not seem to be mentioned in popular textbooks, it makes an interesting conundrum for electronics undergraduates to consider!

Cave Radio Using a Quadrature Phase Space

This problem with Fourier Transforms came to light while I was studying the results of the tests with my channel sounder, which I had built as part of my PhD work on subsurface radio in the early 2000s. It had occurred to me that the parts of the spectrum that are occupied by broadcast stations must, were it not for the station itself, be the quietest areas of the spectrum. Therefore – paradoxically – this might be the best part of the spectrum to use for cave radio.

For an AM broadcast station, the two sidebands are in phase with the carrier, and there is no signal in phase quadrature. An AM envelope demodulator is insensitive to phase, so the region occupied by a broadcast station needs to be free of interference at all phases. It is therefore clear that, if we could transmit a signal that was in phase quadrature to the broadcast carrier, we could detect it, using a phase-sensitive receiver, in such a way that we did not pick up the sidebands of the broadcast AM signal.

I mentioned this briefly in my PhD thesis [Gibson, 2003] – now reprinted with some corrections as [Gibson, 2010] – and I also discussed it in the CREG Journal [Gibson, 2002a]. My initial idea was to transmit a narrow-band FM or a double-sideband AM signal with the carrier in quadrature to the broadcast carrier. Obviously, the true situation is subtler than that because it is at the cave radio *receiver* that the quadrature phase relationship is required; and this might require some 'training' of the communications system. However, as I explained in the CREG article, if the 'phase-space' is quiet in the quadrature phase, then it is as equally quiet for an SSB signal as it is for a DSBSC or FM signal; and SSB transmissions do not require the receiver carrier to be synchronised. We can therefore transmit using *any* phase of the hijacked broadcast carrier, and we receive using the *local* quadrature carrier.

Moreover, the SSB receiver design is simplified because we can use a DSB demodulator. This will cause the two sidebands to be superimposed and, normally, we cannot do this with an SSB signal because there may be interference in the unwanted sideband. However, as we have already stated, this area of the spectrum is asserted to be free from interference.

I went on to describe the operation of a 'phase-skewed' DSB transmission where we advance the phase of the modulation in one sideband and retard it by the same amount in the other. This scheme allows us to utilise the necessarily quiet region in quadrature phase space under an existing AM broadcast in a rather elegant fashion.

Measuring Phase With a Fourier Transform

My channel-sounding work had resulted in a number of datasets, which I presented and discussed in a CREG Journal article [Gibson, 2002b]. The radio receiver signal was sampled at 500kHz, capturing 2^{17} samples per dataset and resulting in spectral data at 3.8Hz intervals. The 16-bit ADC was an AD7721 from Analog Devices, featuring an over-sampling anti-aliassing digital filter with a 3dB bandwidth of 244kHz (16MHz clock) and an attenuation of 72dB at 276kHz.

To evaluate whether my quadrature phase space technique was likely to work, I wanted to look at the amount of interference that was present in the spectrum, in the area occupied by the BBC Radio 4 broadcast transmitter at 198kHz.

On the face of it, you would think this was easy – after all, the data has been sampled and, although sampling does remove information, it is still "all there", from the point of view of applying sampling theory to reconstruct the waveform. However, as I looked at the problem more closely, I realised that there were subtle difficulties. Consider two possible methods of analysis…

Firstly, we could simply compare the amplitude of the corresponding points on either side of the carrier. Since the two sidebands should be identical, this "ought to" give us a measure of the interference. However there are a number of obvious difficulties – we need to look at the phase as well; and it is clear that there is no unique solution – we could work out the difference in the signals but we cannot know which sideband (or sidebands) to attribute it to.

A more workable approach would be to synthesis the operation of the proposed radio receiver itself. That is, we extract the phase of the broadcast carrier (it is simply the coefficient of the FT corresponding to 198kHz) and use it to demodulate the sampled data in such a way as to cancel the AM signal and reveal only the interference – that is, the component of the signal that is in phase quadrature to the broadcast carrier. It sounds straightforward, but when I investigated this I obtained some strange results, which caused me to take a closer look at the problem.

Using MATLAB, I ran a number of simulations and discovered that 'real-world' data sets behaved differently to data sets that I created 'on paper'. Eventually (and rather too slowly!), I realised what was causing the effect, and it is embarrassingly 'obvious' that, contrary to the simplistic notions we might have about how the Fourier Transform works, it does *not* allow us to extract the phase of 'real world' sampled data. The fact that the phase is arbitrary is not the problem, because we can always find some time reference to relate it to. Rather, the salient point is that the phase information produced by the Discrete Fourier Transform of a sampled data signal is, in effect, 'random'.

This is, on the face of it, an astonishing claim and, as I said earlier, one that I have not seen mentioned in any textbooks – although perhaps I have just been reading the wrong books? I am going to take the unusual and provocative step of not explaining this result, since it is such a good puzzle for engineers and students to think about.

Practical Evaluation of a Quadrature Phase Radio

Unfortunately, the problem with the Fourier Transform means that we 'probably' cannot use a sampled dataset to investigate my concept of a quadrature phase radio. However, if we perform the quadrature demodulation in hardware *before* we sample the signal, the resulting sampled data will be adequate for us to do a spectral analysis of the interference.

One reason why I have not done this in the intervening years since I proposed the concept is that I think my original assertion that the concept has merit is probably wrong. Although the interference may well be at too low a level to affect a strong broadcast signal, it could still be strong enough to affect a cave radio. My reasoning is related to the behaviour of the Loran signal, which is a strong source of interference to cave radios. Loran's carrier is at 100kHz, and the spectral content at 80kHz is very low. The problem is that cave radio signals are also very low – and so, in relative terms, the interference is significant. Additionally, if the radio receiver pre-amp does not have a flat phase or amplitude response, this is going to lead to leakage of the broadcast signal into the quadrature phase space. Thus, I have my suspicions that piggybacking on a broadcast carrier might not be a workable idea after all. Although, clearly, the field is still wide open for someone to do some experimental work.

Concluding Remarks

I have described the concept of a 'quadrature phase space' cave radio, and asserted that it cannot be analysed by using a Discrete Fourier Transform of sampled data. I am claiming that the reason for this is 'obvious' but that, because it is under-reported in textbooks, it makes a good puzzle for the reader to consider!

The essential point is that you cannot extract phase information from a sampled 'real-world' data signal. But *why* does the Discrete Fourier Transform of such a signal not contain any useful phase information? The salient point is not the sampling 'as such'; it is the fact that we are using a Discrete Fourier Transform.

But – you may ask yourself – why is 'real-world' data the problem, rather than any theoretical sample set? What is the difference? And why should phase be affected, and not the amplitude? After all, the FT is merely an abstract device and it does not 'know' what amplitude and phase are – they are just similar aspects of a single entity, the complex number obtained by transforming the time-series data. (In fact, amplitude *is* affected but we usually process the sampled data in such a way that the effect on amplitude is minimised).

My earlier notes did not specifically refer to this problem, because, firstly, phase is not usually of interest to us in studies of the frequency spectrum; and secondly, I did not spot it till later. However, the clues are on those notes, just as they are in the article you are reading now.

The chances are that, if you use the DFT in your work, you will immediately recognise the salient point that I have been avoiding mentioning in this article, and you will agree with me that it is 'obvious'. But I hope you will also agree that the specific point of the phase being 'random' is under-reported in textbooks. Hopefully, you will also agree that, even though I have expressed the problem in somewhat provocative terms – i.e. that the DFT *does not* allow you to extract phase information – it is a real and actual problem and a fascinating conundrum.

I probably need to write a follow-up note to this article, explaining what is going on; and readers may like to discuss it on the CREG forum at **bcra.org.uk / cregf**.

References

Gibson, D. (2002a). *Cave Radio Notebook: 51 - Cave Radio Using a Broadcast Carrier*. CREGJ **49**, 23-24

--- (2002b). *A Channel Sounder for Sub-Surface Communications: Part III - Spectrum Survey: Preliminary Results*. CREGJ **49**, 14-22

--- (2003). Channel Characterisation and System Design for Sub-Surface Communications (PhD Thesis). School of Electronic and Electrical Engineering. Leeds: University of Leeds.

--- (2010). Channel Characterisation and System Design for Sub-Surface Communications. ISBN 978-1-4457-6953-0. Available at **lulu.com / content/5870557** [retrieved 22 Feb 2016]

CREG