



An ultra-precise fast fourier transform

Manus Henry^{a,b,*}

^a Fluids and Complex Systems Centre, University of Coventry, Priory Street Coventry, CV1 5FB, UK

^b Department of Engineering Science, University of Oxford, Parks Road, Oxford OX1 3PJ, UK

ARTICLE INFO

Keywords:

Fast fourier transform
Prism signal processing
Spectral analysis
Peak detection
Spectral leakage

ABSTRACT

The Fast Fourier Transform (FFT) is a cornerstone of digital signal processing, generating a computationally efficient estimate of the frequency content of a time series. Its limitations include: (1) information is only provided at discrete frequency steps, so further calculation, for example interpolation, is often used to obtain improved estimates of peak frequencies and amplitudes; (2) 'energy' from spectral peaks may 'leak' into adjacent frequencies, potentially causing lower amplitude peaks to be distorted or hidden; (3) the FFT, like many other DSP algorithms, is a discrete time approximation of continuous time mathematics. This paper describes a new FFT calculation which uses two windowing functions, derived from Prism Signal Processing. Separate FFT results are obtained from each windowing function applied to the data set. Calculations based on the two FFT results yields high precision estimates of spectral peak location (frequency), amplitude and phase. This technique addresses FFT limitations as follows: (1) spectral peak parameters are calculated directly, unrestricted by FFT frequency step discretization; (2) the windowing functions have narrowband characteristics which attenuate and localize spectral leakage; (3) the windowing functions incorporate a Romberg Integration mechanism to overcome the discrete/continuous time approximation.

Video to this article can be found online at <https://doi.org/10.1016/j.sctalk.2022.100097>.

Figures and tables




Fig. 1. Prism windowing functions. The two windowing functions used to generate the Prism FFT.

* Corresponding author at: Fluids and Complex Systems Centre, University of Coventry, Priory Street Coventry, CV1 5FB, UK.
E-mail addresses: manus.henry@coventry.ac.uk, manus.henry@eng.ox.ac.uk.




Fig. 2. Example of time series data. This data set is a simulation of a signal with frequency components at 1 Hz, 5 Hz, and 8 Hz.




Fig. 3. Fast Fourier Transform applied to data in Fig. 2. Applying the conventional FFT calculation to the data set in Fig. 2 generates a corresponding spectrum where peaks at 1 Hz, 5 Hz and 8 Hz are visible.




Fig. 4. Second example of time series data. This data set is a simulation of a signal with 13 frequency components, with amplitudes ranging from 1 V down to $1e-9$ V.




Fig. 5. FFT analyses of signal in Fig. 4. The true amplitude and frequency values of the 13 signal components, or tones, are marked by crosses. The blue line shows the results of the basic FFT calculation. While some high tones are identified, lower amplitude tones are obscured by the noise floor of the FFT calculation – this phenomenon is called spectral leakage, and results in hidden tones. The red line shows the results obtained when the well-known Hann windowing function is applied to the data set before performing the FFT calculation. This results in reduced spectral leakage, and hence fewer hidden tones. The results for the Hann windowing function shown in this paper are similar to those that might be obtained with other widely-used windowing functions (Hamming, Rectangular, Blackman etc), to be contrasted with the very high precision results obtained using the Prism FFT technique.




Fig. 6. Plain FFT analysis of three tone simulation. A three tone simulation is used to illustrate the limitations of conventional FFT techniques and to demonstrate the precision of the new Prism method. The three tones are approximately 50 Hz apart, with the outer two tones having amplitudes of approximately 1 Volt, and the middle tone an amplitude of approximately 1 microVolt, as indicated by the cross symbols. The conventional 'Plain' FFT technique locates the two outer tones, but the middle tone is hidden.




Fig. 7. Hann FFT analysis of three tone simulation. This plot extends Fig. 6 by adding the results obtained when applying the Hann FFT technique to the same data set. Spectral leakage is reduced, resulting in a lower noise floor between the true tones. The middle tone remains hidden due to spectral leakage.




Fig. 8. Prism FFT analysis of three tone simulation. This figure includes the corresponding Prism analysis for the same data sets shown in Figs. 6 and 7. The green line shows the low noise floor calculated away from the true peaks, while the blue lines show the calculated Prism peaks alongside ‘dead zones’ around each Prism peak. The low amplitude middle peak is clearly identified using the Prism technique.




Fig. 9. Close up of FFT analyses around first tone. The three FFT analyses of Figs. 6–8 are shown in more detail in the region of the first true tone frequency at approximately 8950 Hz. The sampling rate is 48 kHz and the dataset contains 48,000 samples, so FFT results are generated at 1 Hz steps. On this more detailed scale, it can be seen that both the Plain FFT and Hann FFT techniques only give approximate estimates of the true tone amplitude and frequency. Additional calculation steps, such as polynomial fitting, may be used to improve these estimates. By contrast, on this scale the Prism FFT technique still appears to provide accurate estimates of both the frequency and amplitude of the true tone.




Fig. 10. Close up of Prism FFT analysis around first tone. In Fig. 9, reducing the frequency scale to steps of 1 Hz is sufficient to show the difference between the true tone and the estimates generated by Plain FFT and Hann FFT. Here, the frequency and amplitude scales are reduced to steps of $1e-11$ Hz and $1e-11$ V respectively in order to separate the Prism FFT results from the true tone values, thus demonstrating the high precision of the Prism FFT technique.




Fig. 11. Time series of three-tone example. This plot shows the time series used to generate Figs. 6–10. The start and end values, marked as red circles, have significantly different values, with a difference of approximately 2 V. An FFT calculation assumes that the time series is repeated in an infinitely long sequence. Accordingly, the discontinuity between the start and end points in this data sets introduced a step change in the repeating sequence, which generates a high level of broadband noise in the corresponding frequency spectrum.




Fig. 12. Plain FFT analysis of three tone simulation. The discontinuity illustrated in Fig. 11 is one source of the high spectral leakage (or high ‘noise floor’) observed in the Plain FFT analysis, causing the middle tone to be hidden.




Fig. 13. Hann windowing function. The plot shows the Hann windowing function with length 48,000 samples. This consists of a pure sinusoid, scaled and offset so that it has initial and final values of zero and a peak of 1.0 at the centre of the data window.




Fig. 14. Hann windowing function applied to data set. The Hann windowing function of Fig. 13 is applied to the original data set of Fig. 11 to create the windowed signal shown here. This is achieved by performing a point by point multiplication of each dataset value by the corresponding Hann window value. As the Hann function starts and ends with values at zero, the windowed function is also zero at the end points, so that the discontinuity in the original data set shown in Fig. 11 is removed. This results in reduced spectral leakage.




Fig. 15. Hann FFT analysis of three tone simulation. The conventional FFT calculation applied to the Hann windowed data set of Fig. 14 results in reduced spectral leakage compared to Fig. 12.




Fig. 16. Prism windowing functions. The Prism FFT technique uses two windowing functions, called Gs1 and Gs2.




Fig. 17. Prism windowing function applied to data set. The Prism windowing functions of Fig. 16 are each applied to the original data set of Fig. 11 to create the two windowed signals shown here, where in each case the data set is multiplied point by point by the corresponding windowing function value.




Fig. 18. Raw Prism FFT spectra. Applying the conventional FFT calculation to the Prism windowed data sets of Fig. 15 results in the two spectra shown here. Key features include the low noise floor (compare with Figs. 12 and 15) and the narrow double peaks either side of the true peaks. The middle peak is no longer hidden.




Fig. 19. Raw Prism FFT values around first tone. Over the next few figures, the method of calculating true tone location to high precision is shown, focussing on the first tone at approximately 8950 Hz. The FFT amplitude values are calculated for the Gs1 and Gs2 windowed data sets at 1 Hz intervals. The vertical dashed line shows the location of the true peak. Very low amplitude values (including around the true tone itself), are excluded from the plot.




Fig. 20. Calculation of amplitude ratio. The next step to determine the true tone frequency is to calculate the ratio of the Gs2/Gs1 amplitudes for each FFT-supplied results shown in Fig. 19.




Fig. 21. Calculation of r . A quadratic transformation is applied to the amplitude ratio of Fig. 20 to obtain the value of r , a dimensionless frequency associated with Prism theory. Here, the value of r is proportional to the distance to the true tone frequency, while the slope of r with respect to FFT frequency indicates the direction of the true tone location.




Fig. 22. Calculation of tone frequency. Based on the values of r shown in Fig. 21, estimates of the local true tone frequency are calculated. The horizontal dashed line shows the true value. Good estimates are obtained for a range of the FFT frequency values.




Fig. 23. Tone frequency error. The errors corresponding to the tone frequency estimates in Fig. 22 are shown. Most errors are below $1\text{e-}5$ Hz, and for high FFT amplitude values (Fig. 19) corresponding to FFT frequencies of around 8940 Hz and 8960 Hz, the errors are less than $1\text{e-}11$ Hz.




Fig. 24. Final Prism FFT results. The results of the Prism FFT analysis are shown. The Prism ‘deadband’ around each identified peak replaces the raw FFT double peaks shown in Fig. 18. The level of the deadband is determined from the adjacent noise floor.

Table 1
Prism FFT results based on for 5000 simulations.

	Tone 1		Tone 2		Tone 3	
	mean error	std	mean error	std	mean error	std
Frequency (Hz)	− 8.09e-13	2.14e-12	− 6.10e-07	1.60e-06	− 1.50e-12	2.07e-12
Amplitude (V)	− 7.83e-13	1.55e-12	− 2.83e-13	1.55e-12	− 6.28e-13	1.58e-12
Phase (radians)	2.32e-12	7.29e-12	1.89e-06	5.18e-06	2.35e-12	6.81e-12

This table shows results for 5000 simulations of the three tone example used in Figs. 6–24, where the initial phase of each tone is randomly selected and white noise with standard deviation 1e-10 V is added. For the high amplitude Tones 1 and 3, the mean errors and standard deviations of frequency, amplitude and phase are of the order of 1e-12 in their respective units. For the lower amplitude Tone 2, which cannot be observed using the Plain FFT or Hann FFT techniques, the frequency and phase mean errors and standard deviations are of the order of 1e-6. The Tone 2 amplitude error and standard deviation are still approximately 1e-12, but given the true amplitude is only approximately 1e-6, this gives a relative amplitude error of approximately 1e-6. Overall, the Prism FFT generates results with errors proportional to the local signal/noise ratio. As explained below, this numerical performance is achieved through the use of high attenuation narrowband filtering to minimize frequency leakage, while Romberg Integration is included in the Prism calculation [1] to provide high precision results within the filter passband.

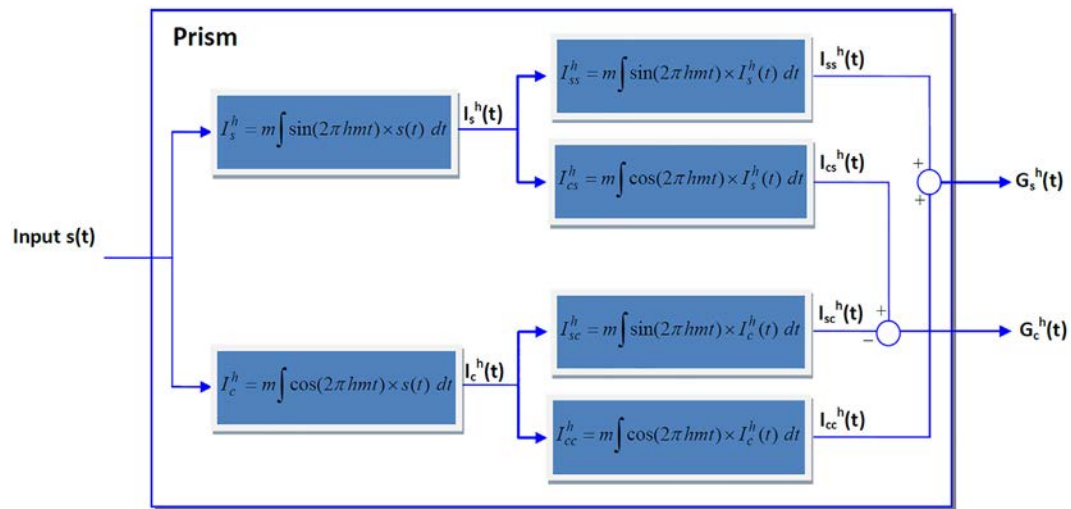


Fig. 25. Prism structure (from [1]). The Prism is a consists of a series of Fourier-style integration blocks. It operates on an input time series $s(t)$ and generates one or two output time series $G_s^h(t)$ and $G_c^h(t)$. The Prism has two configuration parameters, the characteristic frequency m , and the harmonic number h .




Fig. 26. Prism symbol. This symbol can be used to represent an individual Prism in diagrams of Prism networks. The notation is the same as used in Fig. 25.

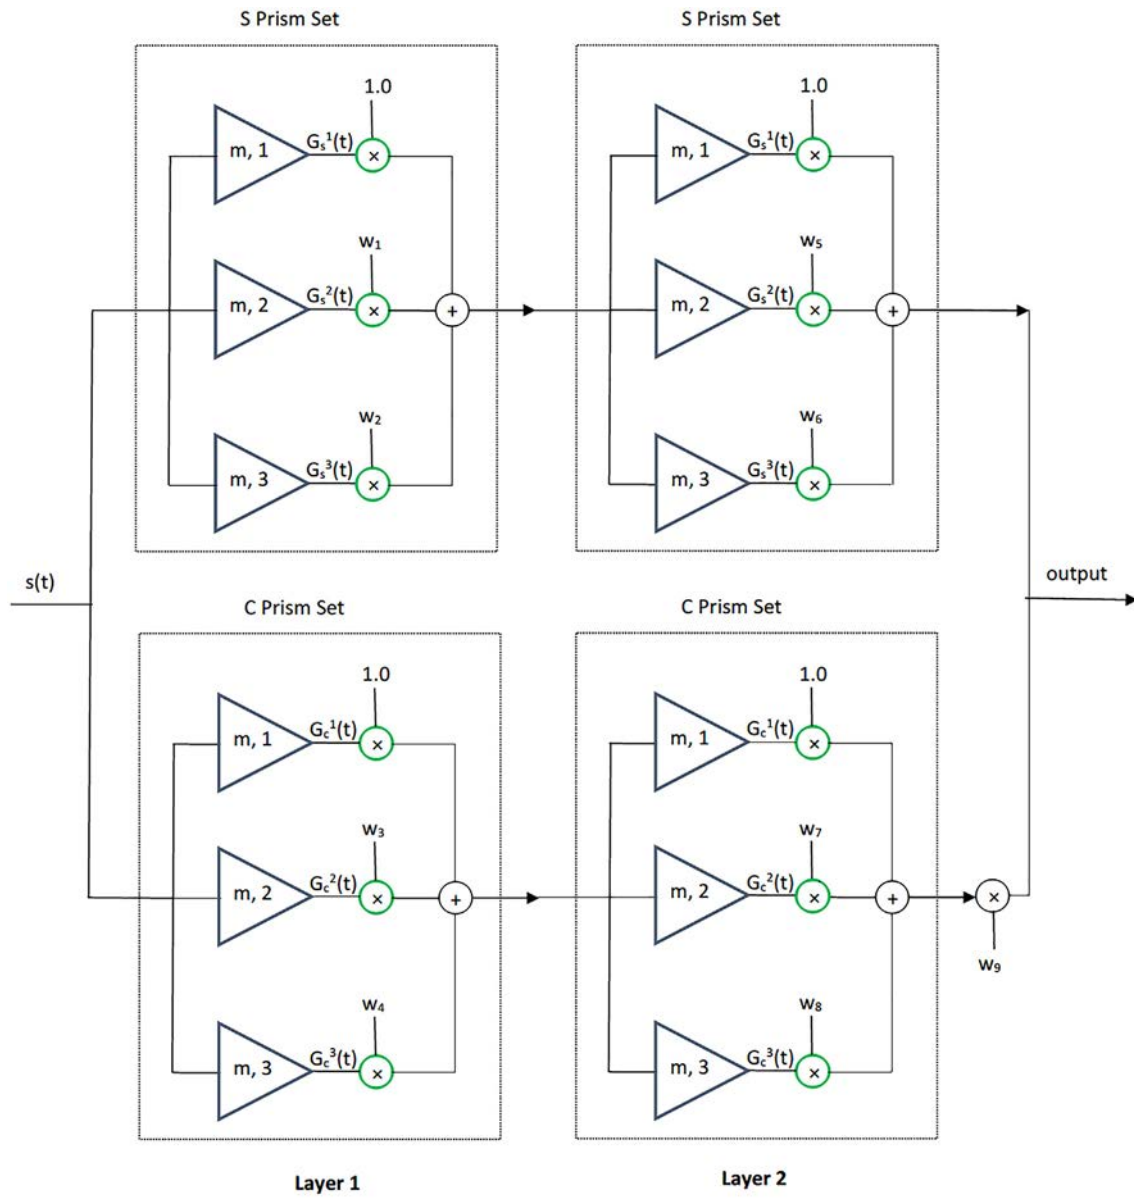


Fig. 27. Prism low pass filter. Prism networks can be used to implement low pass filters. Here twelve Prisms are arranged in two layers to generate a single output from a single time series input. Weights $w_1 - w_9$ are selected to provide the desired low pass frequency response. The same design can be used for any characteristic frequency m .




Fig. 28. Frequency response of Prism low pass filter. This plot shows the frequency response of the 12 layer Prism lowpass filter used to create the Prism windowing functions G_{s1} and G_{s2} in the three tone FFT example of Figs. 6–24. This filter uses the architecture of Fig. 27 extended to create 12 layers, where the weights w_i are optimised across all layers to minimize the gain above the selected value of m , where m is the common characteristic frequency of all Prisms in the network. In this case, the selected value of m is 25 Hz. For all frequencies above m , the attenuation is at least -316 dB.




Fig. 29. Impulse responses of Prism filters. This plot shows the impulse response of two Prism filters, each comprised of the 12 layer Prism lowpass filter of Fig. 28 followed by a single G_s output Prism (Fig. 25) with either $h = 1$ (G_s^1) or $h = 2$ (G_s^2). These impulse responses form the windowing functions used for the Prism FFT calculation of Figs. 6–24.




Fig. 30. Impulse responses of Prism filters on logarithmic scale. This shows the absolute values of the impulse response function of Fig. 29 on a logarithmic scale.

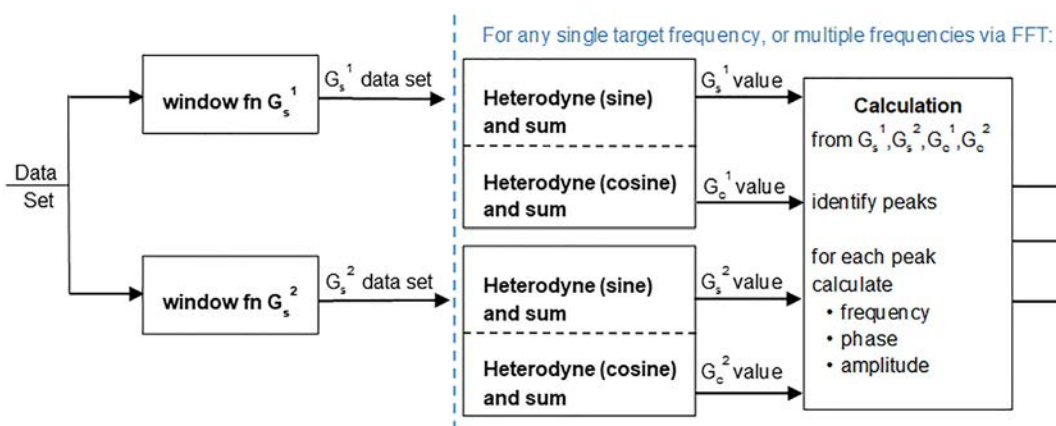


Fig. 31. Summary of Prism FFT calculation. The low pass Prism filter of Figs. 27–30 has a narrow and fixed frequency pass band. This can be adjusted to apply to any target frequency through the technique of heterodyning, which shifts the target frequency into the filter passband. Conventionally in signal processing, heterodyning is applied first, then low pass filtering. The Prism FFT technique inverts this order – the low pass filtering is applied in the form of windowing functions, and then the FFT is used to apply systematic and computationally efficient heterodyning to shift all frequencies into the filter pass band. This inversion of the computational order is possible because the Prism is a linear FIR filter. This calculation results in FFT estimates which have narrowband filtering applied, limiting spectral leakage.

CRediT author statement

Manus Henry: Conceptualization; Investigation; Software; Writing - original draft.

Data availability

The data that has been used is confidential.

Declaration of interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] M.P. Henry, The Prism: recursive FIR signal processing for instrumentation applications, *IEEE Trans. Instrum. Meas.* 69 (2020).



Manus Henry is Professor of Flow Measurement at the Centre for Fluid and Complex Systems, Coventry University and Director of the Advanced Instrumentation Research Group at the Department of Engineering Science, University of Oxford. He is the Editor-in-Chief of the Flow Measurement and Instrumentation. He is a member of the UK Government National Measurement System Programme Expert Groups for Flow Metrology and for Digital Metrology. He has over 130 granted patents, mostly in the field of Coriolis mass flow metering. He developed the Prism in pursuit of improved signal processing for instrumentation applications.