Excerpt From:

R.B. Randall, Frequency Analysis, Bruel & Kjaer, 1987.

4.4. ZOOM FFT

The FFT algorithms discussed so far result in a so-called "Baseband analysis", where the frequency range extends from zero up to the Nyquist frequency f_N, and the frequency resolution is determined by the number of frequency lines up to f_N (normally half the number of original data samples). In certain situations it is desirable to obtain a considerably finer resolution over a limited portion of the spectrum, and the so-called "Zoom-FFT" procedure permits this. It can be considered as "zooming in" on a limited portion of the spectrum with a resolution power corresponding to the number of lines normally used for the whole spectrum (Fig.4.14).

In fact there are two main procedures used for digital zoom, each having certain advantages and disadvantages, so it is interesting to compare the two methods.

To understand the fundamental difference, it is necessary to examine the factors determining the resolution of an FFT analysis:

The sampling interval $\Delta t = 1/f_s$
and record length $T = N\Delta t = N/f_s$
Thus, analysis resolution $\Delta f = 1/T = f_s/N$ (4.1)

Consequently, the two ways of reducing the resolution Δf, are either:

(1) Reduce the sampling frequency f_s. This corresponds to so-called "Real-time zoom".

(2) Increase the record length N. This corresponds to so-called "Non-destructive zoom".

A detailed discussion follows.

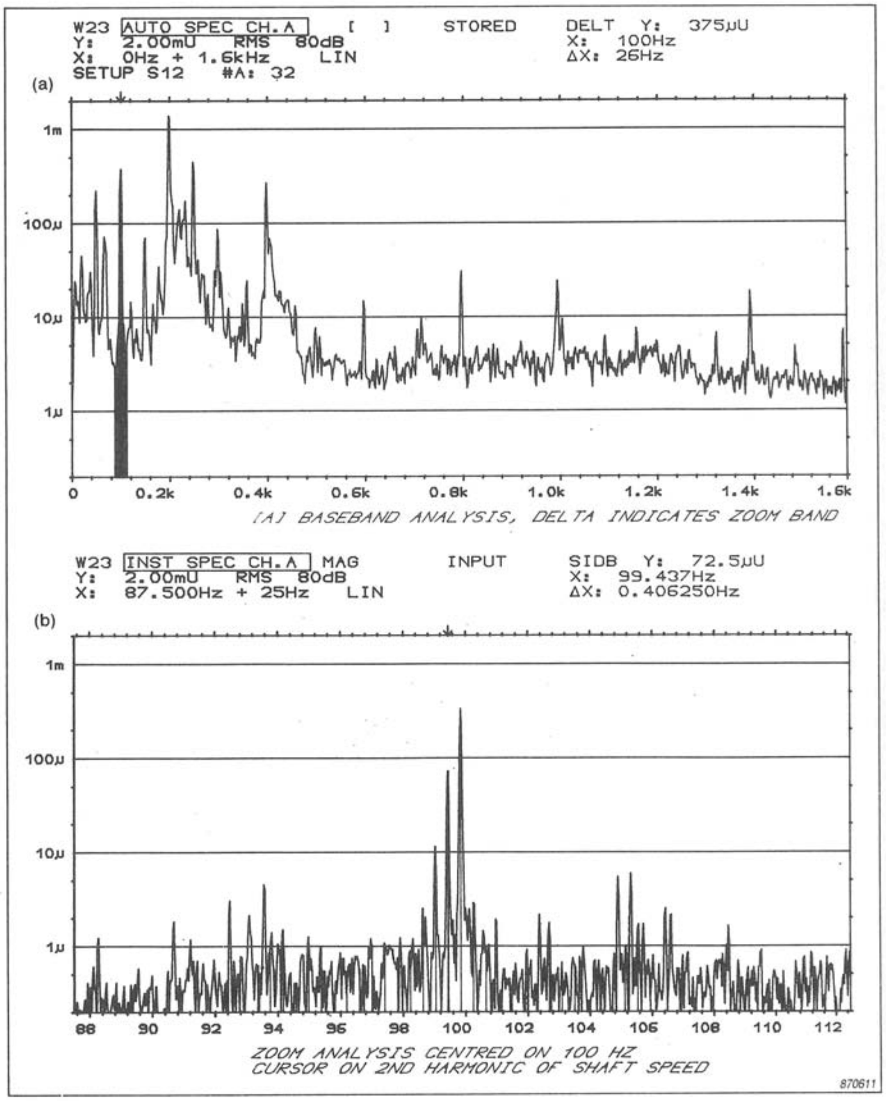

Fig. 4.14. (a) Original baseband spectrum
(b) Shaded section of (a) "zoomed" by factor 64:1

4.4.1. Real-time Zoom

This technique derives its name from the fact that the signal must be processed in real-time by a zoom processor, in order to shift the frequency origin to the centre of the zoom band, low-pass filter the signal, and resample as for the digital filtering described in Section 3.5. The final FFT operation does not have to be in real-time.

The basic principles can be understood by analogy with the discussion of Section 2.2.1, on the way in which the Fourier integral (Eqn.(2.10)) functions. As described there, multiplication by a rotating unit vector $e^{-j2\pi f_k t}$ effectively shifts the frequency origin to frequency f_k. The component at frequency f_k is stopped in the position it occupied at time zero, and virtually becomes a new DC component (although in general it is complex). The positive and negative sampling frequencies $\pm f_s$ are likewise moved by an amount f_k, as illustrated in Fig.4.15. (This may introduce aliasing in the negative frequency region, as the new negative Nyquist frequency $(-f_N + f_k)$ may lie higher in frequency than the lowest frequency component.) Note that even if the original time signal $g(t)$ were real, the modified signal would be a sequence of complex values.

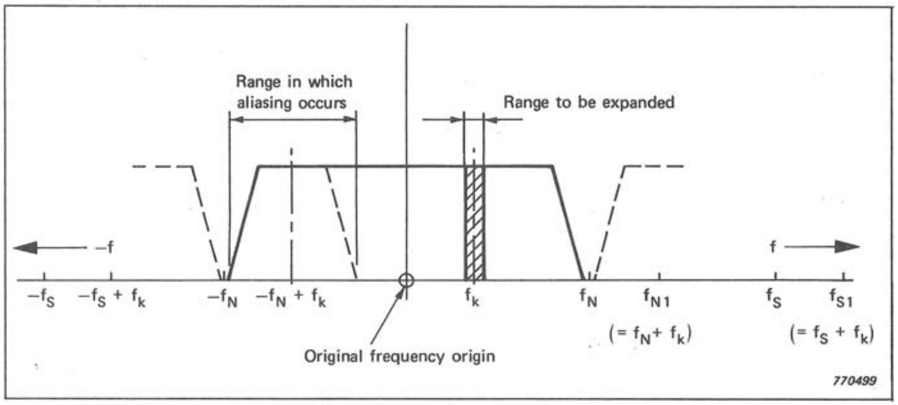

Fig. 4.15. Frequency shift caused by multiplying signal by unit vector rotating at $-f_k$

It is thus possible by multiplying any time signal $g(t)$ by a unit vector rotating at $-f_k$ to change its frequency origin to frequency f_k. The complex signal, thus modified, can then be low-pass filtered (using a digital filter) to remove all frequency components except for a narrow band around f_k, as illustrated in Fig.4.15. Note that at the same time this lowpass filtration would generally remove the portion of the spectrum where aliasing may have occurred. The narrow frequency band remaining after lowpass filtration (the shaded area in Fig.4.15) is shown to a larger scale in Fig.4.16 where it is also made apparent that it is now possible to reduce the sampling frequency while still complying with the sampling theorem. For example, if the total bandwidth after filtering is less than 1/10 of the sampling frequency, it is possible to reduce the sampling frequency to 1/10 without overlapping in the vicinity of the new Nyquist frequency. In a similar manner to the digital filters discussed in Section 3.5, the reduction in sampling frequency is achieved simply by retaining a reduced number of samples, in this case every tenth, the rest being discarded. The resampled sequence of lowpass-filtered complex samples can be transformed by a (complex) FFT transformation to give the required "zoomed" spectrum.

The above discussion applies strictly to the general complex FFT transform. If the baseband analysis system is designed to produce $N/2$ spectral values from N real data values (Section 4.2), the data memory will only hold $N/2$ complex data values. On the other hand, the complex forward transform (Eqn.(2.18)) gives $N/2$ complex results which are now all valid, because in general there is no symmetry about the new zero frequency (the original frequency f_k). Thus the number of lines resolution in the zoomed spectrum is unchanged. Note that as mentioned in the footnote on p.28 the second half of the frequency spectrum obtained represents the negative frequencies (i.e. the original frequencies below f_k) which should be moved to their correct position before the first half prior to display.

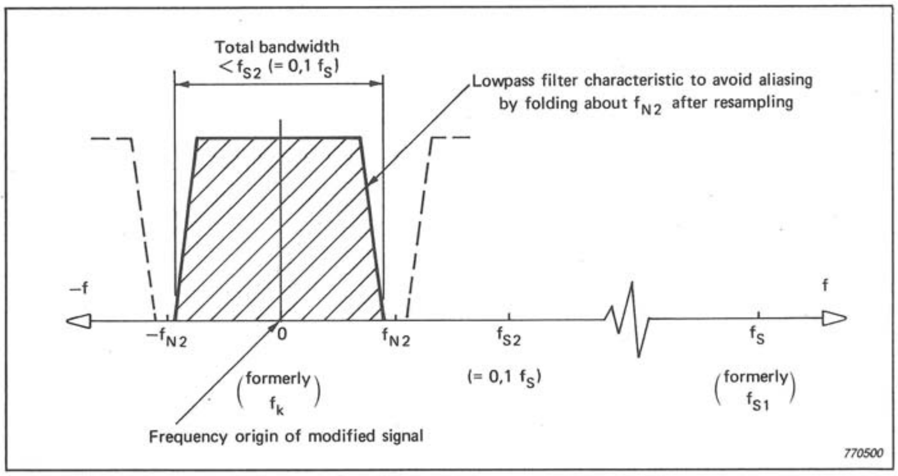

Fig. 4.16. Detail of range to be expanded after resampling

In this case, in order to achieve a zoom factor of 10, the sampling frequency would have to be changed by 20:1, the first 2:1 zoom virtually achieving the same effect as the procedure described in Section 4.2. Note that in the Type 2032 Analyzer (Fig.4.1) both procedures can be used to obtain the same spectral data (i.e. "Baseband" analysis, or "Zoom" in the centre of the baseband range). Even though the spectral result is the same, the data treatment is different, and in particular the analytic time function will be different in the two cases.

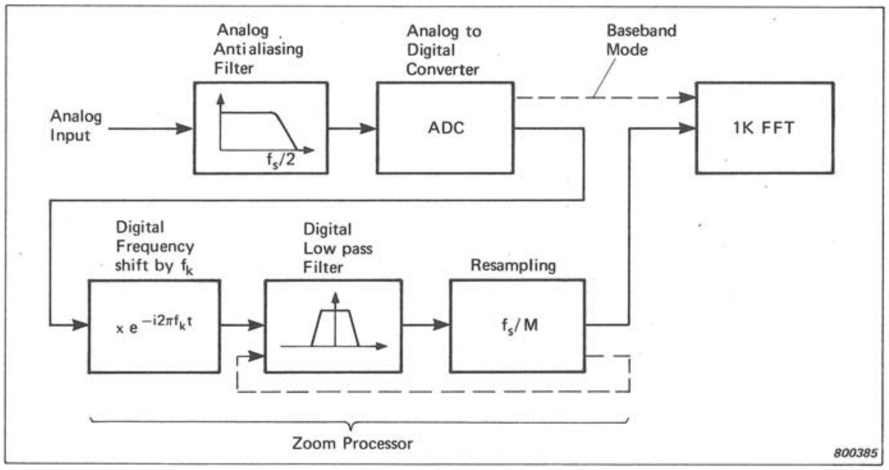

Fig. 4.17. Flow diagram for real-time zoom

In practice, as with the digital filters of Section 3.5.3, it is usually most efficient to low-pass filter and resample in cascaded octave (2:1) steps, meaning that the zoom factors would normally be powers of 2. Fig.4.17 shows a block diagram of the procedure for real-time zoom. Two features should be noted:

(1) In order to zoom around a new centre frequency f_k, it is necessary to re-process the signal through the zoom processor. This may not be possible unless the signal has been stored.

(2) The memory buffer required for the time signal is no longer than for baseband analysis. The extra length of time signal required to obtain the finer resolution is achieved by reducing the sampling rate.

4.4.2. Non-destructive Zoom

Non-destructive zoom is simply a way of achieving a large transform size (e.g. 10K or 10240 samples) by repeated application of a smaller transform (e.g. 1K or 1024 samples). For a zoom factor *M*, it requires a data buffer *M* times longer than the transform size *N*.

Fig.4.18 illustrates the basic principles, for a 10K transform. The samples in the data record are numbered 0 through 10239. Ten 1K transforms are performed on records obtained by taking every 10th sample from the original 10K record, first Nos. 0, 10, 20, ..., 10230, then Nos. 1, 11, 21, ..., 10231 etc. until all data values have been transformed. Because of the linearity of the DFT, the sum of the transforms of the undersampled records must be equal to the transform of the sum of the records. This sum is only equal to the original 10 K record when compensation is made for the small time displacements of each of the undersampled records (except for the first). In the Fourier spectra this represents a simple linear change of phase proportional to both the time displacement and the frequency of the component in question (Ref.4.7).

Fig.4.19 is a block diagram of the process, for comparison with Fig.4.17. Note that the zoom accumulator is of limited size, and thus only a selected 400 line section of the resulting high resolution spectrum is generated at any time. If a 4000 line buffer were available it would be possible to generate all 4000 lines from one set of transforms, because even though only 512 complex spectrum results are produced by each partial transform, the entire 5120 lines can be deduced by periodic repetition (the undersampling of the part records leads to this periodicity by aliasing as illustrated in Fig.4.18). It should be noted that any errors introduced by this aliasing cancel out in the final summing operation (by the linearity properties of the DFT).

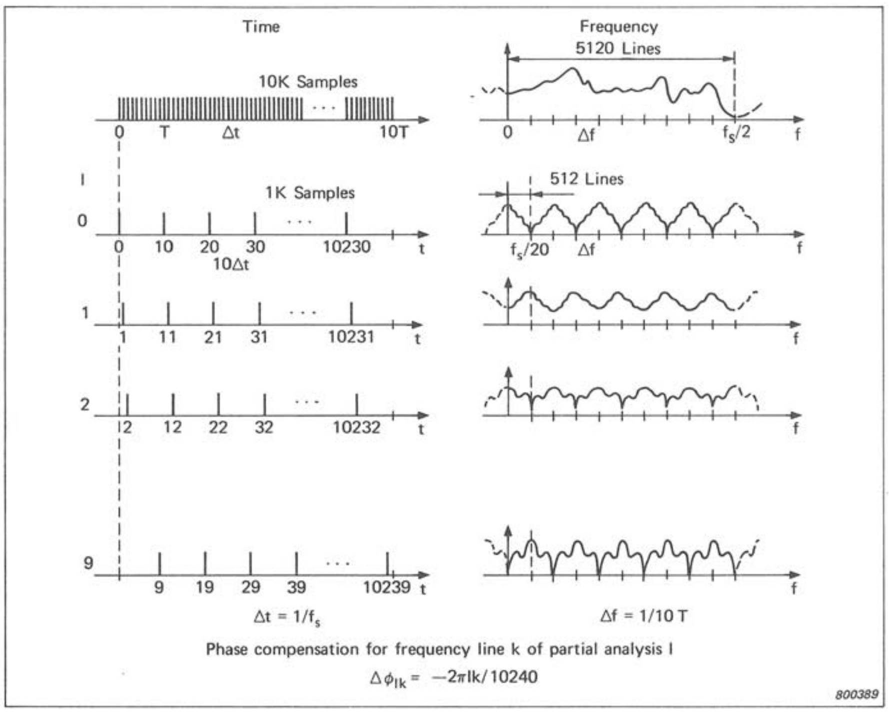

Fig. 4.18. Implementation of a 10 K transform using a 1 K transform ten times

The block diagram of Fig.4.19 also makes the following two features evident:

(1) Zoom in different frequency regions is based on exactly the same data record.
(2) A long data buffer is required and the zoom factor is limited by the length of memory.

4.4.3. Comparison of Zoom Techniques

The main advantages of non-destructive zoom accrue from the fact that it is based on exactly the same data record. Thus, it is most valuable for (single channel) signal analysis because even very stable signals vary slightly from one record to another. A typical example would be in gearbox analysis, where one might first want to zoom around the first harmonic (i.e. fundamental) of the toothmeshing frequency and then around the second and third harmonics, knowing that because the same data record is used, there is an exact integer 1,

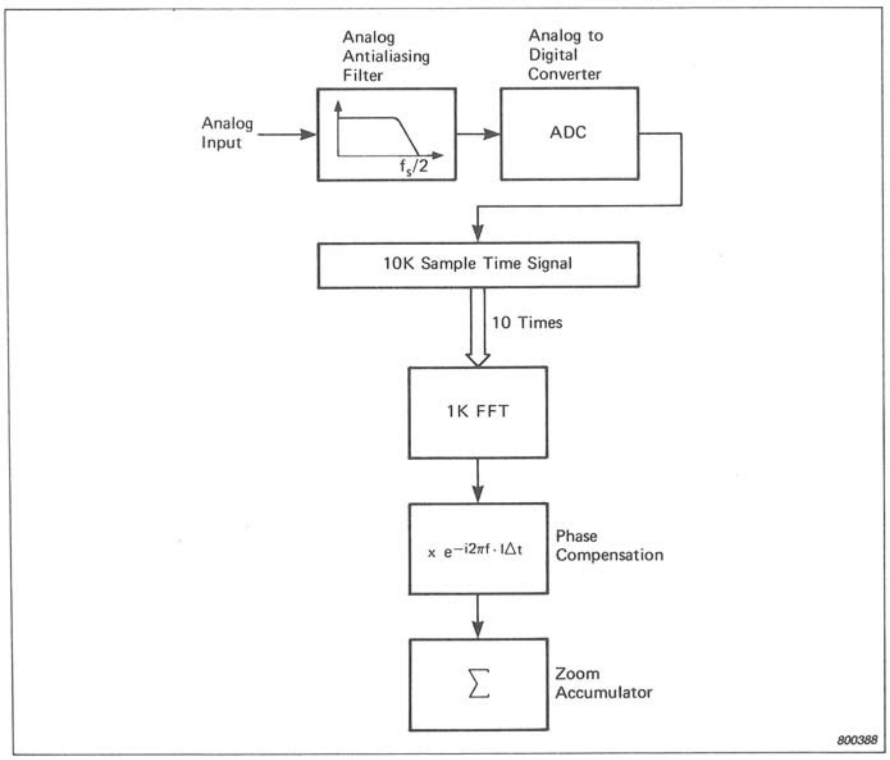

Fig. 4.19. Flow diagram for non-destructive zoom

2, 3 relationship between the frequencies of these components. The same cannot be said for real-time zoom, because for zooming around each harmonic, a different data record is processed.

It is only with non-destructive zoom that frequency component definitions to an accuracy of 1:20000 (as described at the end of Section 4.3.3) can be achieved.

The disadvantage, that zoom factors are limited by the record length, is rarely a restriction in signal analysis, because it is rarely that individual components are so stable in frequency as to justify a resolution better than 1:4000 or so, for example when tied to machine speed.

A possible exception is equipment with servo-controlled speeds such as Hi-Fi audio equipment.

Larger zoom factors are more relevant for (dual channel) system analysis, where for example frequency response functions can be expected to be stable even if the individual excitation and response signals vary somewhat. Thus, even the results of real-time zoom in contiguous (or overlapping) bands can be expected to match up at the intersections. Zoom is often required in frequency response measurements to achieve adequate resolution of lightly damped resonance peaks. (See Chapter 7.)