EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH

CERN/PS/BD/Note 92-8

PERFORMANCE OF THE PS FFT Q-MEASUREMENT SYSTEM

S. Johnston

ABSTRACT

The swept-filter Q-measurement system of the PS is being replaced with a Digital Signal Processor. This system provides a powerful diagnostic tool which allows not only very rapid Q measurement but also detailed spectral analysis of pick-up data. Fast Fourier Transforms may be performed singly or repeated over large acquisitions to generate 3D pictures of the beam oscillation in time. In this paper an in-depth examination of the performance of the system in terms of computation time and accuracy of results is presented.

Geneva, Switzerland
November 1992

PERFORMANCE OF THE PS FFT Q-MEASUREMENT SYSTEM

S. Johnston

ABSTRACT

The swept-filter Q-measurement system of the PS is being replaced with a Digital Signal Processor. This system provides a powerful diagnostic tool which allows not only very rapid Q measurement but also detailed spectral analysis of pick-up data. Fast Fourier Transforms may be performed singly or repeated over large acquisitions to generate 3D pictures of the beam oscillation in time. In this paper an in-depth examination of the performance of the system in terms of computation time and accuracy of results is presented.

1. INTRODUCTION

The development of a Fast Fourier Transform (FFT) Q-measurement system to replace the swept filter analyser [1] of the PS accelerator began in 1990 [2] with the selection of a VME-based Digital Signal Processing (DSP) card, the VASP-16 Vector & Scalar Processor.

The VASP-16 is a programmable system and consequently it has been developed to perform not only the Q-measurement but also FFT and Sliding-FFT analysis of pick-up data. The theory behind these measurements will be given in chapter 3, following an introduction to the operation of the VASP-16.

Results are then presented of tests designed to find the limits of the system in terms of speed and accuracy and it will be shown that the performance has surpassed the original specification, given below.

- PS Q-measurement : The system should be capable of a measurement repetition-rate faster than every 5 ms with an absolute Q value error less than 10^{-3}.
- FFT analysis : 512 or 1024-point FFT (true magnitude) analysis of pick-up data.
- Sliding-FFT : Acquire as much data as possible for subsequent analysis whereby an FFT window of 256, 512 or 1024 points is moved along the data to generate a 3-Dimensional (3D) description of the beam oscillation.

Examination of the maximum Q-measurement repetition-rate will be shown to be in agreement with the predicted value and the consequence of exceeding this rate will be explained. The accuracy of the Q-measurement is examined under ideal conditions followed by more realistic tests under the influence of damping and noise. These results can be used to obtain a confidence level in a Q calculation dependent upon the input signal conditions. Finally a short look is taken at the FFTs followed by a presentation of results obtained using the Sliding FFT.

2. THE VASP-16 VECTOR AND SCALAR PROCESSOR

The VASP-16 is a multi-processor DSP card aimed at image-processing applications [3]. Running originally under OS9, a new device-driver has been written [4] for the LynxOS operating system.

The card contains five processors : a master Texas TMS320C25 and four slave Zoran ZR34161 Vector Signal Processors (VSPs) (Fig. 1). The TMS320C25 is programmed in C and forms the heart of the system, handling VASP-16/VME communication, acquisitions, data management and function calls to the Zoran VSPs : specialised 16-bit processors with an architecture optimised for rapid DSP operations such as FFTs, convolution, correlation, etc. [5].

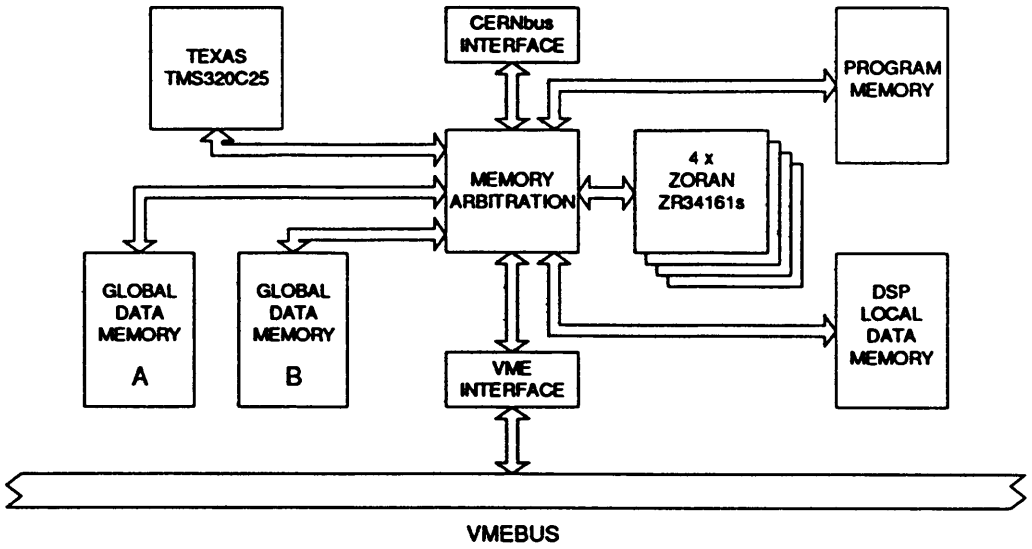

Fig. 1. Block diagram showing the internal architecture of the VASP-16.

Zoran VSPs may perform arithmetic operations with three levels of precision : *integer*, *fixed floating-point* and *block floating-point*. The data which is transferred into the Zorans for processing, e.g. acquisition samples, must be in integer format. FFT analysis of the data uses block floating-point precision which gives better performance than fixed floating-point calculations. The spectrum calculated by the Zorans must however be stored in VASP-16 global memory in integer format, and this results in some precision loss when the floating-point values, i.e. the values representing the amplitudes within the spectrum of the FFT, are truncated.

Data is acquired using a 12-bit ADC located on a separate NIM chassis. Prior to the ADC, the pick-up signal must be passed through a bandpass anti-aliasing filter since the spectrum of the beam extends well beyond the sampling frequency.

Transfer of data into the VASP-16 global data memory is over a 16-bit open-architecture bus, MAXbus, which has been customised and renamed CERNbus. Global memory consists of two, 32 kbyte (16-bit) banks sharing the same address space, access to them being governed by a programmable arbiter. This architecture allows concurrent acquisition in one bank while processing the contents of the other and consequently VASP-16 measurements may be interleaved to obtain a greater measurement frequency.

3. BASIC THEORY

3.1. Measurement of the Q value.

The Q value, defined as the number of betatron oscillations which the particles exhibit as they perform one revolution of the machine [6], is to be analysed using the FFT for hadron and lepton beams. The frequency of these oscillations can be calculated by firstly considering the total energy, E, of the particles in the accelerator (see Table 1).

Table 1. Hadron and lepton energies in the PS.

Particle type	Minimum energy (GeV)	Maximum energy (GeV)
hadrons	1.138	26
leptons	0.6	3.5

Using Table 1, the normalised velocity of the particles, β, can be calculated [7] from

$$\beta = \sqrt{1 - \left(\frac{E_0}{E}\right)^2} \tag{3.1}$$

where

$$E = \text{Total energy} \quad , \quad E_0 = \text{Rest energy} \begin{cases} 0.938 \text{ GeV (hadrons)} \\ 5.110 \times 10^{-4} \text{ GeV (leptons)} \end{cases}$$

This results in the minimum and maximum values for β shown in Table 2.

Table 2. Normalised velocities of particles within the PS accelerator.

Particle type	Minimum β	Maximum β
hadrons	0.56622	0.99935
leptons	1	1

The revolution frequency of the particles around the accelerator, f_{rev}, may be then calculated from

$$f_{rev} = \left[\frac{2\pi R}{\beta c}\right]^{-1} \tag{3.2}$$

where

$R =$ average radius of the PS (100m)

Inserting the values of Table 2 into (3.2) gives the minimum and maximum revolution frequencies for the PS (see Table 3) :

Table 3. PS revolution frequencies for hadrons and leptons.

Particle type	$f_{rev(min)}$ (kHz)	$f_{rev(max)}$ (kHz)
hadrons	270.16	476.82
leptons	477.13	477.13

In the PS, the value of Q is known to lie within the range

$$6.1 \leq Q \leq 6.45 \tag{3.3}$$

The frequency of the betatron oscillation corresponding to Q, f_β, is related to the revolution frequency by

$$f_\beta = (n+Q) f_{rev} \quad \text{where } n = -6 \text{ for the PS} \tag{3.4}$$

From (3.3) we may write $Q = 6 + q$, where $0.1 \leq q \leq 0.45$, and hence (3.4) can be re-written as

$$f_\beta = q f_{rev} \tag{3.5}$$

Using (3.5) and the values in Table 3, the frequency range which the Q-measurement system must be capable of analysing can be calculated,

$$\begin{aligned} f_{\beta(min)} &= 0.1 \times f_{rev(min)} = 27.016 \text{ kHz} \\ f_{\beta(max)} &= 0.45 \times f_{rev(max)} = 214.708 \text{ kHz} \end{aligned} \tag{3.6}$$

The Q-measurement system, because it measures the beam at one position around its orbit, is only capable of calculating the fractional component q and not, as the name implies, the Q value itself [6].

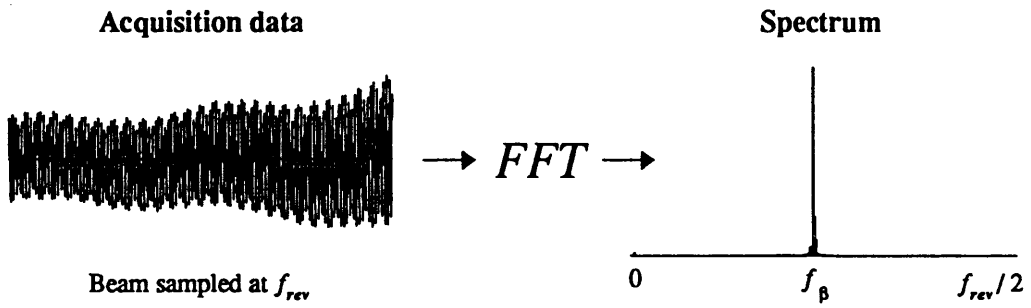

Fig. 2. 512-point FFT analysis of a 1.6 x 10^{13} proton beam, sampled at f_{rev} = 400 kHz, yields a spectrum with f_β around 95 kHz.

Calculation of the fraction, q, is based on the FFT analysis of a 512-point acquisition, using the revolution frequency as the sampling rate (see Fig. 2). The spectrum resulting from the FFT calculation represents the frequency content of the beam from DC to $f_{rev}/2$ and includes f_β, the frequency corresponding to Q. The frequency resolution of the spectrum in Fig. 2 is

$$\frac{f_{sample}}{N} = \frac{f_{rev}}{512} = 781.25 \text{ Hz/bin} \tag{3.7}$$

where N is the number of samples. If the tallest line in the spectrum, located in bin 122, is taken to represent the betatron oscillation, then q could be calculated from

$$q = \frac{f_\beta}{f_{rev}} = \frac{122 \times 781.25}{400 \times 10^3} = 0.2383 \tag{3.8}$$

However, considering that each spectral line represents the frequency range calculated in (3.7), the true value of q may lie between 0.2363 and 0.2383.

The value of q can be more precisely calculated by considering the lines around the main peak. Performing FFT analysis on a sampled sine wave gives one of the two spectra shown in Fig. 3, depending upon the frequency of the signal.

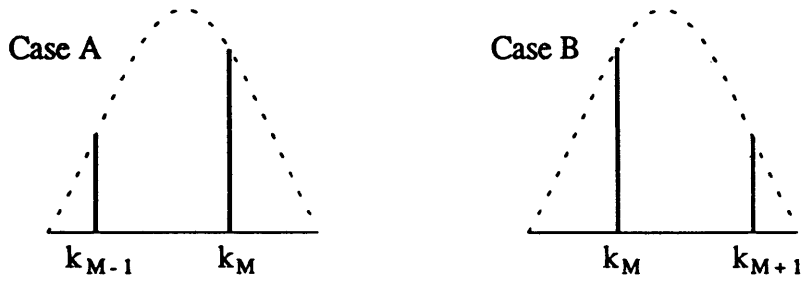

Fig. 3. The FFT analysis of a sine wave shows a peak with sidelobes when the number of periods within the rectangular FFT window is non-integer. The position of the second line within the main lobe indicates whether the true signal frequency is to the left (Case A) or to the right (Case B) of the main line.

The continuous spectrum shown by the dashed line in Fig. 3 has the shape $(\sin x)/x$, and therefore, assuming a sine wave input and for the rectangular FFT window, the location of the peak can be interpolated from the amplitudes of the two highest lines [8]. Applying this calculation to the spectrum in Fig. 2 allows f_β, and consequently q, to be calculated. For the two conditions shown in Fig. 3, the following formulae are used to calculate q :

$$\text{Case A} \quad q = \frac{1}{N}\left\{k_M - \frac{|k_{M-1}|}{|k_M| + |k_{M-1}|}\right\} \tag{3.9}$$

$$\text{Case B} \quad q = \frac{1}{N}\left\{k_M + \frac{|k_{M+1}|}{|k_M| + |k_{M+1}|}\right\} \tag{3.10}$$

N = Number of points in the FFT

In section 2 it was mentioned that the spectral amplitudes calculated by the VASP-16 have integer resolution. For this reason, (3.9) and (3.10) cannot be used exclusively to calculate the q.

The lack of floating-point precision can sometimes mean that both spectral lines around the highest line are truncated to the **same** integer value and the software cannot differentiate between the two lines of the main lobe and the line corresponding to the sidelobe (Fig. 4a).

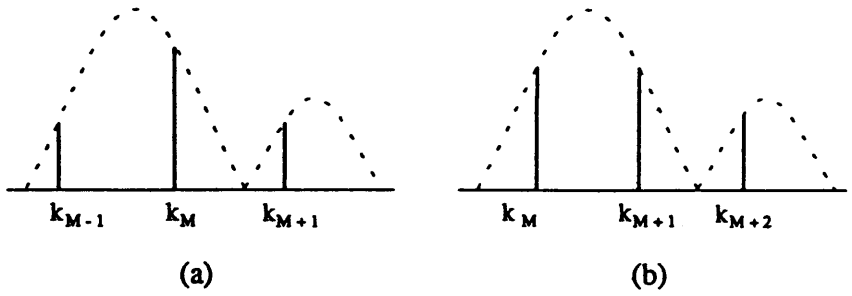

Fig. 4. The effect of integer truncation on floating-point spectra. (a) Both spectral lines around the peak are truncated to the same amplitude. (b) A heavily damped input may give the same amplitude for the two lines within the main lobe.

Under such circumstances, because the truncated lines on either side of the highest are equal, the program decides that the spectral peak does in fact represent the betatron oscillation and the q value is simply calculated from

$$q = \frac{k_M}{N} \tag{3.11}$$

A second exception occurs when the two values within the main lobe are truncated to the same value (Fig. 4b) and hence the software cannot determine which of the two lines is the true peak. If this situation arises, which may occur when the sampled input signal is heavily damped, the software decides that the true frequency is exactly between the two highest lines, and the q is then calculated from

$$q = \frac{k_M + 0.5}{N} \tag{3.12}$$

3.2. VASP-16 FFT analysis.

During the calculation of an FFT, the Zoran VSPs use a complex automatic-scaling mechanism to contain the amplitude of the spectrum within the operating precision of the processors, i.e. < 16-bits. Occasionally an overflow does occur however, and the tallest line in the spectrum is set to zero. In the case of the q-measurement this causes large errors because interpolation is not performed on the true spectral peak, but using the highest remaining line. A method of overriding this automatic scaling, in order to prevent q-measurement errors, has therefore been developed.

The FFT function returns a value, `binexp`, which indicates the amount of automatic-scaling performed during the FFT calculation. Consequently, some function of `binexp` can be used to re-scale the spectrum back to its true amplitude.

Fig. 5a shows the changes which occur to the value `binexp` when the amplitude of the sampled input signal increases. An input with a very small amplitude does not require scaling during the FFT calculation because the spectral amplitude poses no threat of overflow (Fig. 5b). When the input signal amplitude is increased, a threshold is reached at which point the output spectrum seems suddenly to become smaller (Fig. 5c). However, at the same time, the value of `binexp` has been incremented, indicating that the spectral amplitude must be re-scaled, using `binexp`, to obtain the true spectral amplitude.

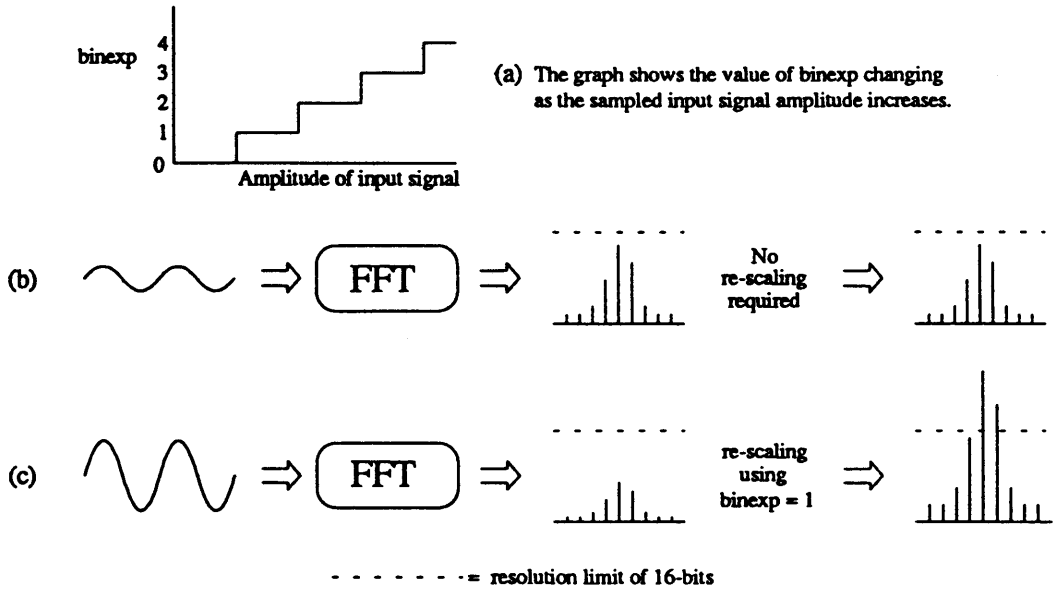

Fig. 5. Explanation of the FFT scaling mechanism. (a) The value binexp, returned by the FFT function, can be used to re-scale the spectrum to its true amplitude. (b) A small amplitude signal may not require scaling, while a larger amplitude signal (c) will require re-scaling using the value of binexp.

To override this automatic scaling mechanism, the Q-measurement software makes use of a second variable, `minscale`. Before calling an FFT function, the software performs some pre-processing of the sampled input data in order to select the value of `minscale`. This value, which may be between 0 and 9, is dependent upon two factors : the maximum amplitude of the sampled input signal, and the number of samples. Correct selection of the value of `minscale` ensures that the FFT function returns a `binexp` value of zero, thus preventing overflows and removing the possibility of q-measurement errors.

To obtain the true amplitude of the spectrum, the value of each spectral line (k_n) still has to be re-scaled, using some function of `minscale`. For the particular FFT size selected, one of the following three formulae should be applied to every point in the spectrum in order to calculate the true amplitude (A_n) of the spectral lines.

256-point FFT $$A_n = \sqrt{k_n} \times 2^{(\text{minscale}+1)} \quad n = 1,2,\ldots,256 \tag{3.13}$$

512-point FFT $$A_n = \sqrt{k_n} \times 2^{(\text{minscale})} \quad n = 1,2,\ldots,512 \tag{3.14}$$

1024-point FFT $$A_n = \sqrt{k_n} \times 2^{(\text{minscale}-1)} \quad n = 1,2,\ldots,1024 \tag{3.15}$$

4. DESCRIPTION OF THE SOFTWARE

The Q-measurement software is designed to run in a continuous loop, corresponding to the PS machine cycle. The loop commences when the VASP-16 receives a list of instructions containing the total number of measurements to be performed, their types, FFT sizes, etc. This measurement list must be received in the order in which they are to be executed because the VASP-16 does not initiate acquisitions but receives triggers from an external timing source.

The program proceeds by checking if the first measurement in the list is the Sliding FFT. This is an exclusive measurement which may not be performed in the same cycle as other measurement types because of the very large volume of data it produces. If selected, the main program loop will repeat after completion of the Sliding FFT calculations and alternatively, if not requested, the program will execute the list of measurements sequentially using concurrent data-acquisition and processing.

The Sliding FFT is used to analyse 52 kbytes of data with 256, 512 or 1024-point magnitude-squared FFTs, the spectra being re-scaled to represent the true amplitude of the sampled input data. A 26 kbyte acquisition is initiated by an external trigger and is stored in data bank A. Upon completion, the program grants access to data bank B and a false trigger, generated by the NIM interface, initiates the second 26 kbyte acquisition, with synchronisation to the same bunch. The number of FFTs specified in the measurement list is then executed and the spectra are returned to the host upon completion of each.

Between FFT calculations, the FFT window is moved along the data by a number of samples as specified by the user. The resultant spectra therefore represent the beam oscillation in frequency and time. A small step can be used to obtain a very detailed picture of the beginning of the acquisition or alternatively a large step can be used to obtain a general view of the evolution throughout the whole acquisition.

Q-measurements and FFTs are interleaved, using background interrupt routines, to give concurrent acquisition and processing. An acquisition is armed for the first measurement in the list and is initiated by an external trigger. Upon completion, the program looks for the subsequent measurement, if any, and arms the next acquisition which may occur concurrently with processing of the first. This continues until the list of measurements is exhausted and the loop is then repeated.

The FFT is the simplest measurement and may contain 512 or 1024-points, magnitude squared. The number of samples acquired, N, equals the number of points in the FFT. This data is pre-processed to prevent arithmetic overflows during the many cross-multiplications of the FFT. Upon completion of the FFT calculation, the spectrum, containing N/2 points, is returned to the VME host along with the scaling information explained in section 3.2. This information is used to re-scale the magnitude-squared spectrum to give the true amplitude spectrum.

The Q-measurement makes use of the 512-point FFT only, and the q-value is calculated by interpolating the peak in the spectrum as described in section 3.1.

5. PERFORMANCE MEASUREMENTS

All acquisition times quoted in this chapter are for sampling rates of 500 kHz, unless specified otherwise. This is approximately the maximum revolution frequency of the PS, as explained in section 3.1.

5.1. Processing Speed

Measurements of the execution times of the FFTs (see Table 4), combined with an understanding of the operation of the VASP-16, allows an estimation of the maximum frequency at which Q-measurements may be made.

Table 4. Execution times.

FFT size (points)	Acquisition duration (ms)	FFT execution time (ms)	Q interpolation (ms)
512	1.024	0.63	1.146
1024	2.048	1.11	-

The duration of the Q-measurement, which uses the 512-point FFT, is 2.8 ms. However, concurrent acquisition and processing is employed to achieve an increased throughput and known acquisition and processing times allow the timing diagram shown in Fig. 6 to be drawn. This diagram suggests that the first result should be ready after 2.8 ms and thereafter every 1.8 ms.

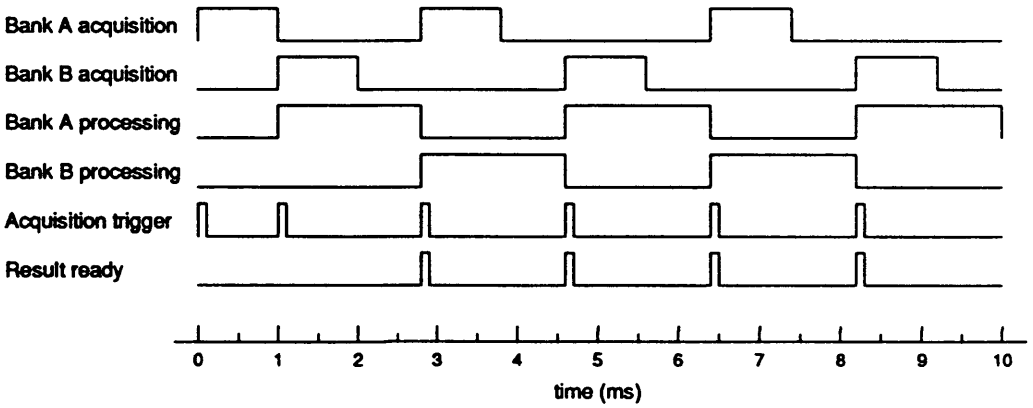

Fig. 6. FFT Q-measurement timing diagram

This estimate is validated by recording the acquisition trigger and the end-of-acquisition signal produced by the VASP-16. For a trigger-rate of 2 ms, the maximum of the PS Q-kicker, the VASP-16 encounters no problems. This is also the case when the trigger-rate is increased to the estimated limit of 1.8 ms (Fig. 7a). Fig. 7b illustrates the consequence however of increasing the trigger-rate to 1.7 ms. The third trigger arrives too soon after the second, at which point processing of the previous acquisition is not yet complete. The third trigger is therefore missed and the third acquisition is actually initiated by the fourth trigger. Since, in practise, the number of triggers will match the number of measurements requested, the VASP-16 will not receive enough triggers to complete the measurements and the system will hang-up, a remote reset of the VASP-16 then being the only option. These results are summarised in Table 5.

Table 5. Maximum repetition rate of measurements.

Measurement type	Maximum trigger rate (Hz)	Period (ms)
Q Measurement	555	1.8
512-point FFT	555	1.8
1024-point FFT	322	3.1

It was foreseen that the Sliding FFT would not be used more than once per cycle and therefore a maximum trigger rate is not applicable. It however requires some additional data manipulation during processing and this extra time has been added to the FFT execution times to allow calculation of the total duration of Sliding FFT measurements (see Table 6). For example, a Sliding FFT of 250, 512-point FFTs will have an acquisition time of 106.5 ms followed by a processing time of 250 x 2.77 ms, i.e. 799 ms in total.

Table 6. Sliding FFT execution times.

Measurement type	Acquisition duration (ms)	Execution time per FFT (ms)
256 point Sliding FFT	106.496	1.62
512 point Sliding FFT	106.496	2.77
1024 point Sliding FFT	106.496	5.19

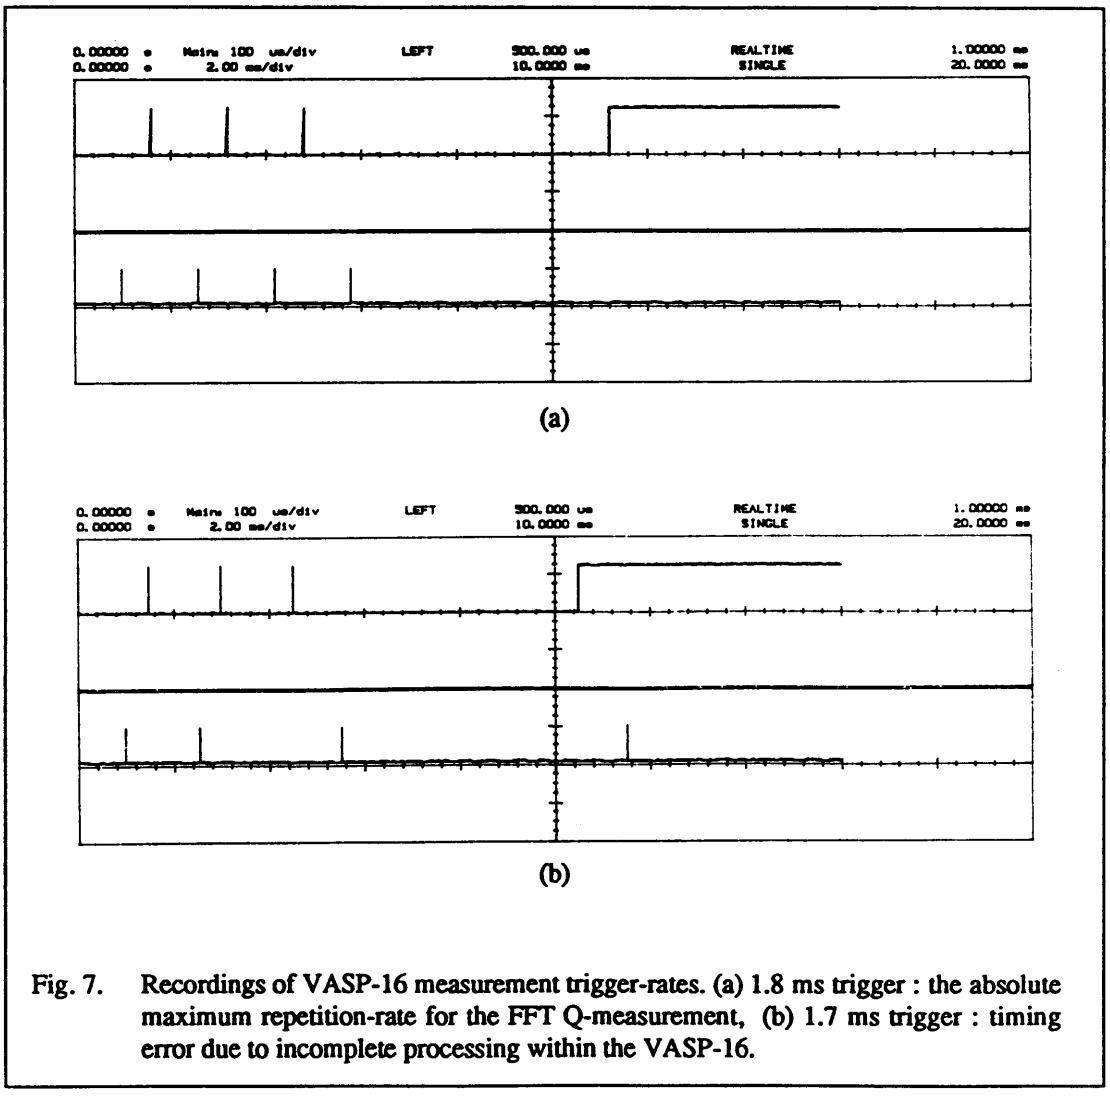

Fig. 7. Recordings of VASP-16 measurement trigger-rates. (a) 1.8 ms trigger : the absolute maximum repetition-rate for the FFT Q-measurement, (b) 1.7 ms trigger : timing error due to incomplete processing within the VASP-16.

5.2. Q-Measurement Accuracy

In this section the accuracy of the FFT Q-measurement is examined by performing measurements on signals simulated by the TMS320C25 processor. Three signals are generated : a pure sine wave, a damped sine wave and a damped sine wave with additional pseudo-random noise.

Measurements have been performed over a frequency range which covers two bins of the 512-point FFT spectrum. This is not the complete range which the FFT Q-measurement system must analyse, but the smaller range still allows all effects due to windowing to be covered, as well as saving much computation time.

5.2.1. Q-measurement analysis of a pure sine wave

To firstly demonstrate the ability of the FFT Q-measurement to interpolate the q value, the pure sine wave described below was simulated :

Input signal	$A \sin(2\pi f T)$
Frequency	$97.656 \text{ kHz} \leq f \leq 99.609 \text{ kHz}$
Amplitude	$A = 1.25$ V (peak - peak)
Sampling rate	$T^{-1} = 500$ kHz

Fig. 8 shows a recording of the absolute error in the value of q, calculated using the original software [2]. The error is seen to rise above the specification limit of 10^{-3} in two regions as the frequency is swept through the simulation range.

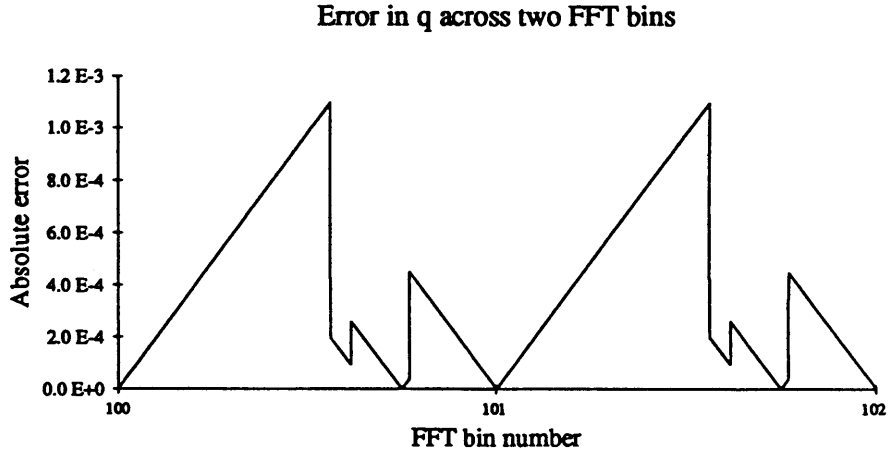

Fig. 8. The graph shows the original software giving q errors > 10^{-3} when analysing a pure sine wave signal. (Frequency range 97.6 kHz to 99.6 kHz)

An investigation into the source of the large errors in Fig. 8 yielded four reasons which were collectively responsible.

As explained in section 3.1, the q value can be interpolated from the magnitudes of the two main spectral lines corresponding to the betatron oscillation. However, the original software used magnitude-squared values, resulting in wrongly interpolated values. Furthermore, the use of equation (3.10) only, and not (3.9), resulted in the linear increase

in error visible in the left half of each bin (Fig. 8). The third source of errors has already been discussed in section 3.2, which was due to the scaling mechanism employed by the Zorans VSPs.

In order to discover the final reason for errors, the same sample of a pure sine wave was repeatedly analysed by the FFT Q-measurement. The q calculation was found to be correct except for a repetitive error which occurred every 3 or 4 FFT operations. After much investigation, this was found to be due to an incorrect value within a hardware register of the Zorans themselves. This register was used as a pointer during the FFT scaling operations and a new routine, written to reset these registers before every FFT operation, eliminated the problem.

Fig. 9 shows the error of the FFT Q-measurement, using the new software, for the same input signal and frequency range as in Fig. 8.

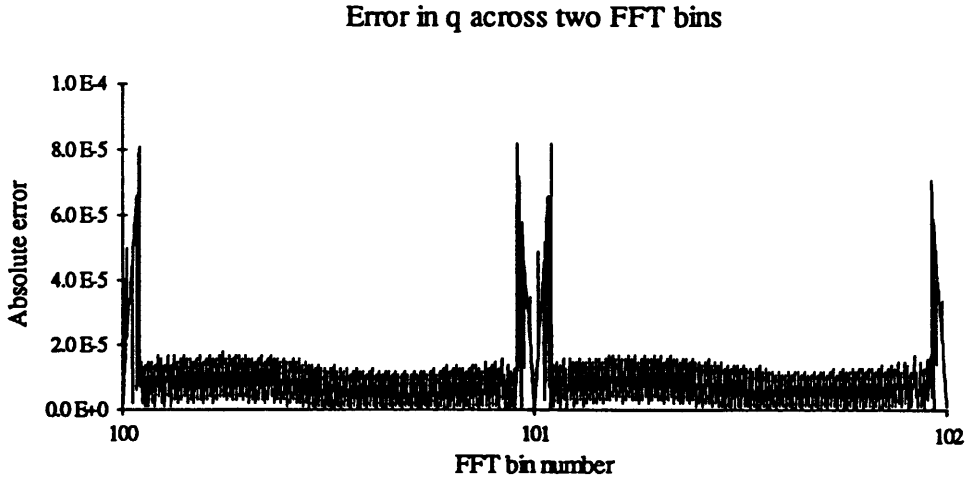

Fig. 9. The absolute q error from the new software, measured during the analysis of a pure sine wave (1.25 V peak-peak) over the frequency range 97.6 kHz to 99.6 kHz.

A comparison of Figs. 8 and 9 shows that the elimination of the four sources of error has significantly reduced the maximum absolute error in q from 1.1×10^{-3} to 8×10^{-5}. The majority of Fig. 9 in fact shows an absolute error well below the maximum, in the region of 10^{-5}.

The shape of the plot in Fig. 9 can be explained by considering the interpolation method described in section 3.1. At 97656 Hz, the sampled sine wave fits exactly into the window of the FFT an integer number of periods. The FFT gives therefore no sidelobes and the software uses (3.11) to calculate the q with zero error. When the sine wave frequency is increased slightly, sidelobes should begin to appear and interpolation, using (3.9) and (3.10), can be used to obtain the correct q. However, because the amplitudes of the sidelobes are initially very small, integer truncation of the values results in significant loss of precision, which in turn gives visible q errors. In the remaining region, where the q error is in the order of 10^{-5}, the sidelobes have a larger amplitude and the effect of integer truncation is reduced to a negligible level.

Having established that the new FFT Q-measurement is more accurate, a second test was performed using the pure sine wave over a range of amplitudes, $A = (2^n - 1) \times 0.61$ mV, where n = 3,4,...,11. At each amplitude, the sine wave was simulated over the 2 bins of the FFT in steps of 1 Hz and the absolute error in the calculated q value was recorded.

From the data, the maximum error found at each amplitude was extracted and this information is shown in Fig. 10. In the range 2^6 (38 mV) to 2^{11}(1.25 V) peak, the maximum q error never exceeds 10^{-4}, a factor of 10 better than required by the specification. Even for an input as small as 2^3 (4.27 mV peak), the absolute error is only 4.2×10^{-4}.

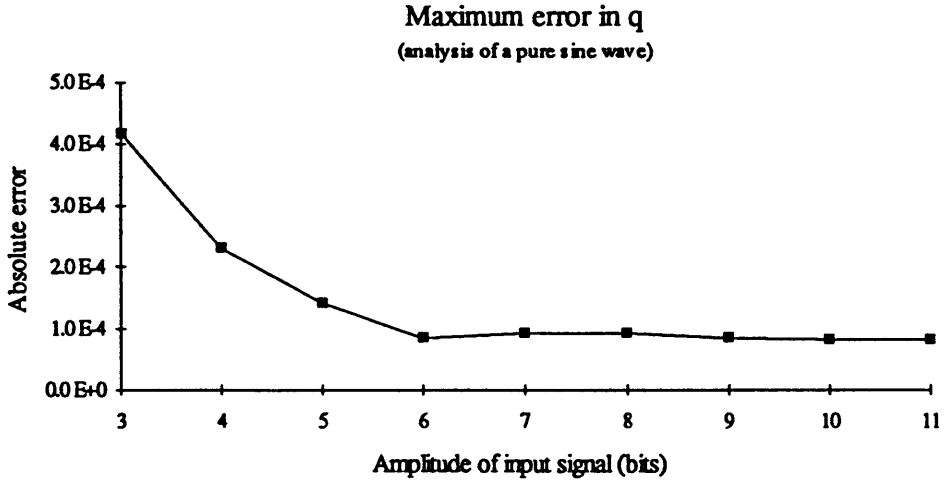

Fig. 10. The maximum absolute error in q is shown for pure sine waves simulated at different amplitudes. At each amplitude, frequencies between 97.6 kHz and 99.6 kHz were simulated in steps of 1 Hz.

5.2.2. Q-measurement analysis of a damped sine wave

The following test attempts to find the effect of different damping rates upon the accuracy of the Q-measurement based on the analysis of a damped sine wave. Simulation has not been restricted to an input signal of one amplitude however. Of two signals, one small in amplitude and the other large, the latter will have a better representation within the FFT window for the same amount of damping and therefore give less q error.

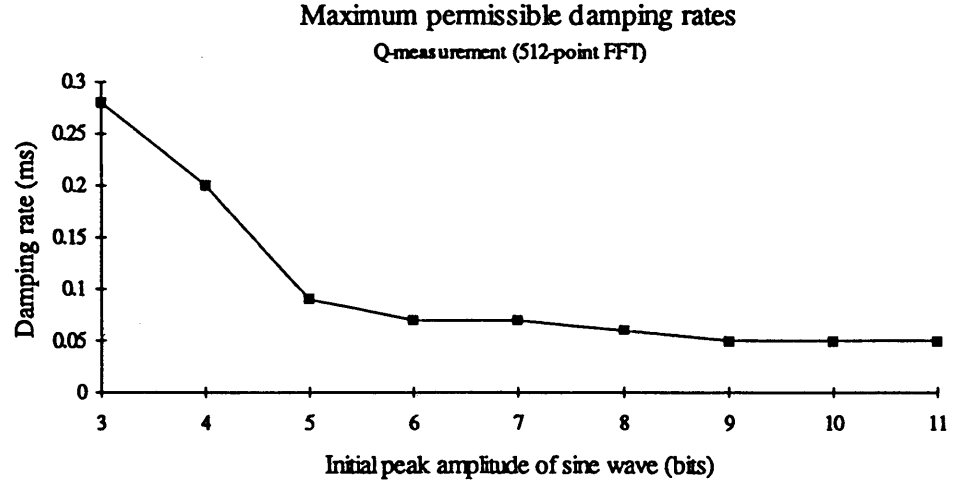

Fig. 11. Fastest damping (τ) for the signal $\sin(2\pi fT)e^{(-t/\tau)}$ which results in a q error below 10^{-3}.

Fig. 11. is a plot of the fastest damping rate which a signal may have, depending upon its initial amplitude, and still be analysed by the FFT Q-measurement with an absolute error less than the specification requirement of 10^{-3}. The graph shows that signals with an initial amplitude greater than approximately 2^5 (19 mV peak) may have damping rates as fast as 0.1 ms and still have a q error within specification.

5.2.3. Damped sine wave with added random noise

A comparison is now made of the accuracy of the FFT Q-measurement for a damped sine wave input signal, with and without additional pseudo-random noise.

The amplitude of the un-damped sine wave was set to 1.25 V peak-peak, corresponding to half of the full-scale representation of the ADC. For a number of different damping rates between 0.1 ms and 10 ms, the damped sine wave was analysed using the FFT Q-measurement twice : before, and then after the addition of noise. The noise was generated by using the ANSI C `rand` function, and was amplitude limited to a maximum of 200 mV peak.

Performing these measurements over the same frequency range as in 5.2.1, a recording was made of the absolute q error at each frequency and the maximum values were extracted to give the graph shown in Fig. 12.

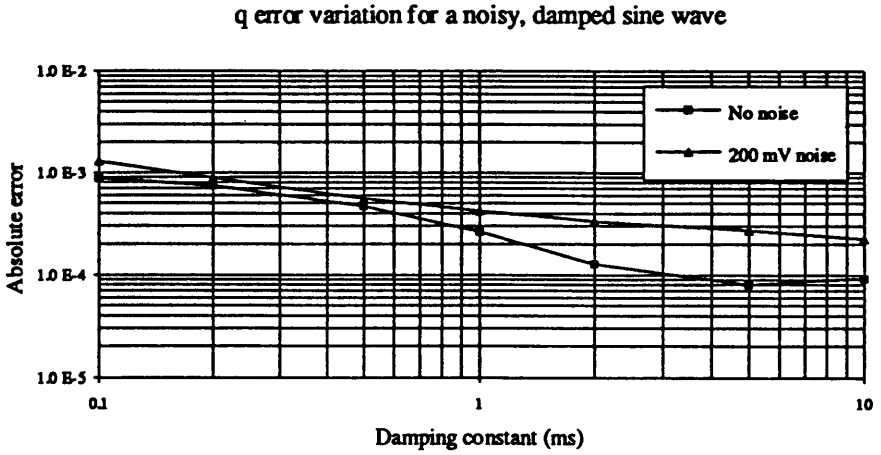

Fig. 12. The absolute q error from the analysis of a damped sine wave, with and without additional noise (200 mV peak).

As expected, there is a slight increase in the q error for the noisy signal, although only for very fast damping rates, at around 0.15 ms, does the absolute error increase beyond the specification limit of 10^{-3}.

Signals with damping rates faster than 0.1 ms may not be reliably measured with the FFT Q-measurement because of the lack of floating-point precision. As explained in section 3.1, damping causes the spectrum to spread out and because of integer truncation of the amplitudes of the FFT lines, it can become impossible to distinguish the true peak from the lines around it (Fig. 4b).

5.3. Block Floating Point FFTs

As mentioned in Chapter 2, the precision of the block floating-point arithmetic used by the Zoran VSPs is lost when the spectra are stored as integers in the VASP-16 global memory. This loss of precision can be seen when measuring the Signal-to-Noise Ratio (SNR) improvement resulting from FFT analysis.

When a pure sine wave has some random noise added to it, a signal is obtained with a certain SNR. FFT analysis results in an improved spectral SNR because the noise content is spread over the entire spectrum while the sine wave becomes more distinct.

Fig. 13 shows the relationship between input and spectral SNRs, based on peak signal amplitudes, for the 512-point FFT. A sine wave with a peak-peak amplitude of 1.25 V was used as the input signal and pseudo-random noise was added to it, the sum of the amplitudes not exceeding the maximum input of the ADC (2.5 V peak-peak).

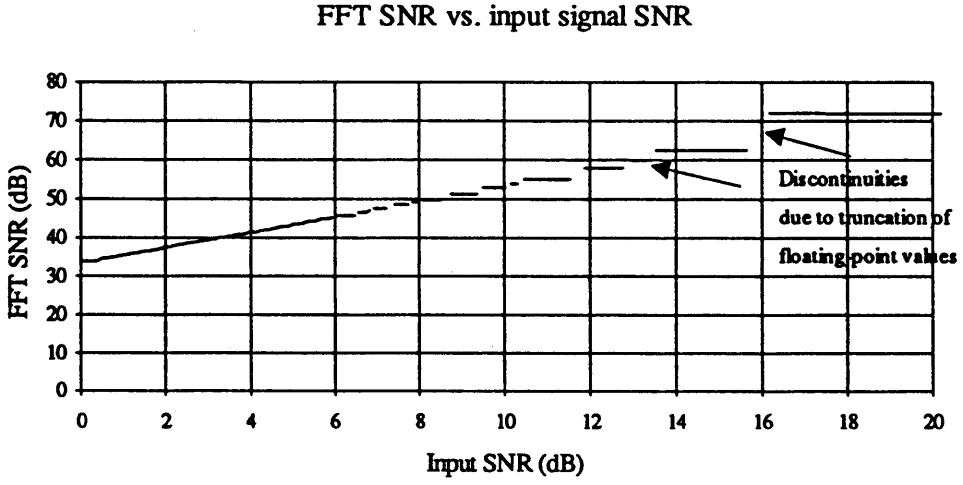

Fig. 13. Discontinuities in the FFT SNR vs. input SNR due to integer truncation of the amplitudes of the spectral lines.

The maximum FFT SNR of 74 dB corresponds exactly to the value quoted by the manufacturer of the Zoran VSPs. At this level, the spectral amplitude of the noise is represented by 1 bit only. An increase in the input SNR therefore causes the spectral noise representation to decrease in amplitude, and integer truncation to zero results in the FFT SNR becoming infinite. Hence the plot cannot be extended beyond 20 dB SNR for a 1.25 V peak-peak input. Floating-point representation of the spectral lines would prevent this effect and the discontinuities, shown by arrows in Fig. 13, would disappear.

The loss of floating-point precision is also evident when errors in the amplitudes of spectral lines are considered. Simulation of a sine wave, of known frequency and amplitude, allows the difference between the FFT amplitude and the actual amplitude to be measured. Fig. 14. shows the percentage error in the amplitude of a spectral line representing a sine wave, fitting exactly into one spectral bin, over the amplitude range 0 to 2.5 V peak-peak. Assuming that we have an oscillation around 1 V peak-peak, the error in the spectral estimation of its amplitude due to integer truncation is around 0.1% only.

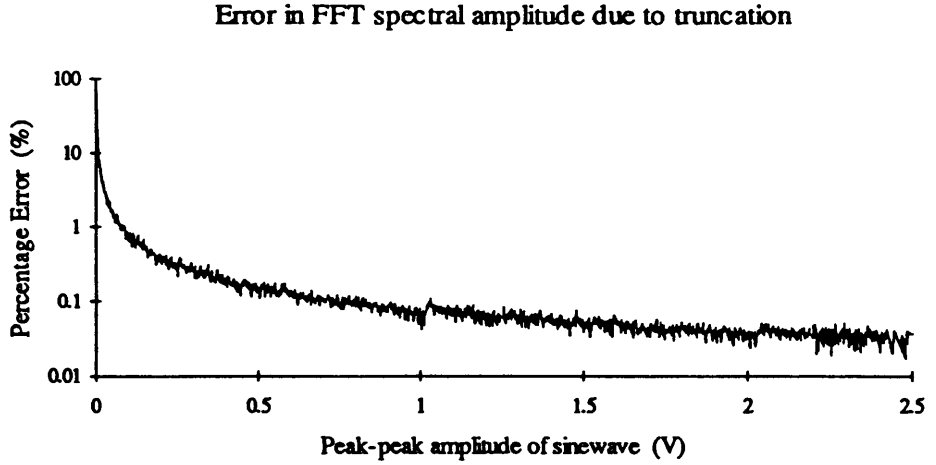

Fig. 14. Error in the amplitude of a single spectral line representing a pure sine wave as the amplitude is varied from 0 to 2.5 V peak-peak.

Before leaving this look at the FFTs, it is worth comparing the performance of the VASP-16 integer FFTs with the floating-point FFT in Matlab (software for the numerical solution of mathematical problems).

The same sample of pseudo-random noise was analysed using both FFTs and the resultant spectra exhibited very good correspondence. A plot of the difference in amplitude between the two does not divulge the reason for the largest errors whereas a plot of the percentage error (Fig. 15), with respect to the amplitude of the lines in the VASP-16 spectrum, shows clearly that the largest errors are associated with the smallest lines, for whom integer truncation is significant.

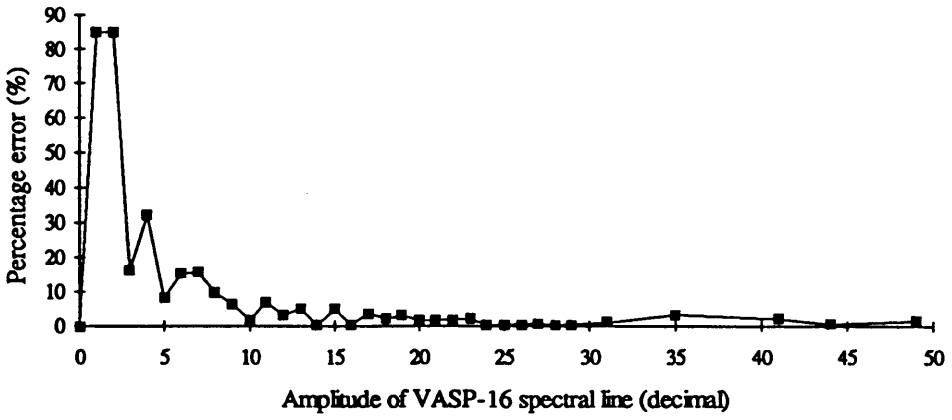

Fig. 15. A comparison of the amplitudes of the VASP-16 integer spectrum and a floating-point FFT output shows that the largest errors are associated with the smallest spectral values.

5.4. Sliding FFT results

The graphics shown in this section are screen-dumps from a display program written specially for the FFT Q-measurement system. The program uses plotting primitives to generate 1D or 3D mountain-range views of Sliding FFT spectra on Tektronix graphics terminals.

Two acquisitions are now analysed using the Sliding FFT to generate a detailed picture of their frequency content with time. The first is actual pick-up data from a 1.6×10^{13} proton beam at injection into the PS. The data was acquired using a 12-bit, 10 MHz ADC sampling at f_{RF} = 8 MHz. Prior to the ADC, the signal was passed through a 2.5 MHz bandpass anti-aliasing filter followed by the digital notch filter of the PS Transversal Feedback system [9].

The second acquisition was acquired with a sampling rate of 500 kHz and is a constant-amplitude sine wave, swept in frequency, produced by a signal generator.

5.4.1. Sliding FFT analysis of a 1.6×10^{13} proton beam

Fig. 16 shows a 1D mountain-range view of the evolution of the beam oscillation during the first 5 ms after injection into the PS, generated using a Sliding FFT of forty, 512-point FFTs. The step size between FFTs is 1318 samples, corresponding to a time interval of 164.75 μs. This short time interval can only be achieved using the Sliding FFT since single 512-point FFT measurements have a maximum repetition rate of 1.8 ms. The graphical plot shown does not include the first 10 bins of the FFT which were removed to allow the graphical scaling to emphasise the oscillations present and not the DC at injection.

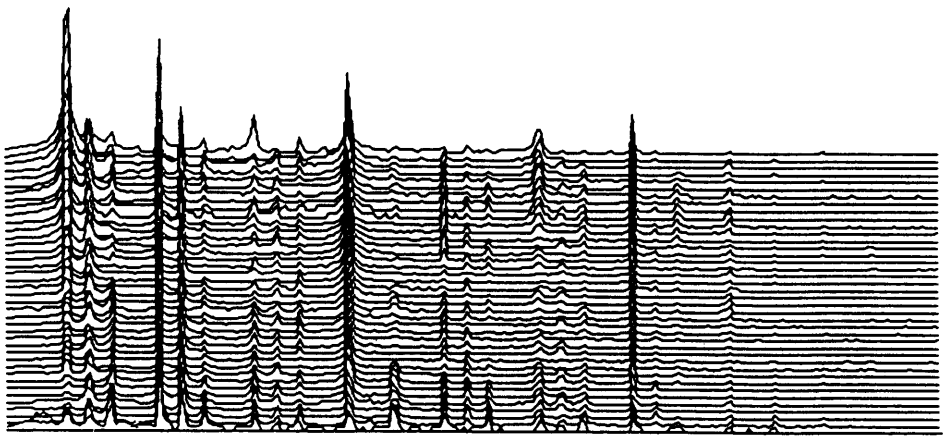

Fig. 16. 1D Sliding FFT analysis shows the frequency content of a proton beam from 156 kHz to 4 MHz (Sampling rate was f_{RF} = 8 MHz).

The 1D mountain-range view in Fig. 16 allows the frequency content to be easily seen but to appreciate amplitude changes the data has to be viewed from a different angle. Fig. 17 was calculated using sixty 1024-point FFTs every 108.75 μs and shows in 3D the beam

oscillation during the whole acquisition. Also visible this time is the DC content and the betatron oscillation corresponding to $q = 0.23$.

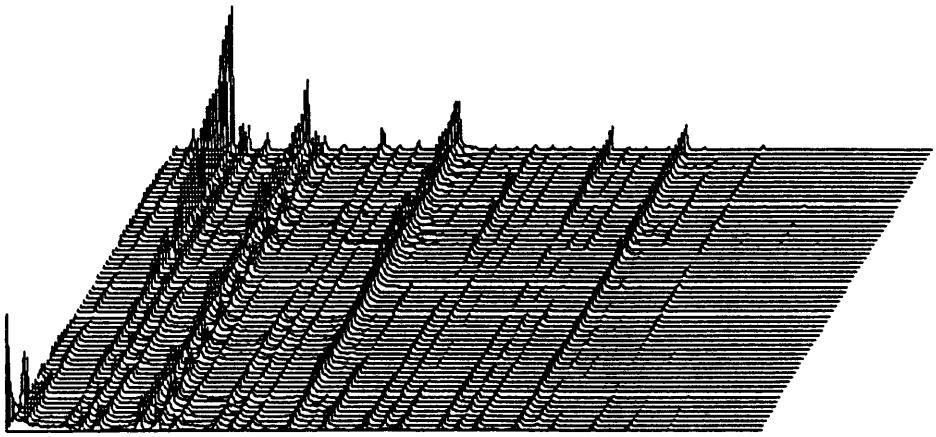

Fig. 17. 3D Sliding FFT analysis shows the frequency content of a proton beam and also allows the amplitude variation to be seen (Sampling rate was $f_{RF} = 8$ MHz).

5.4.2. Sliding FFT analysis of a constant amplitude, swept frequency sine wave

Fig. 18 shows the Sliding FFT spectrum of a constant-amplitude sine wave being slowly swept in frequency. The spectral representation is seen to vary in amplitude as the frequency is swept from low to high. This effect is due to the spectral peak moving across the bins of the FFT and thus the amplitudes of the peak and its sidelobes are changing.

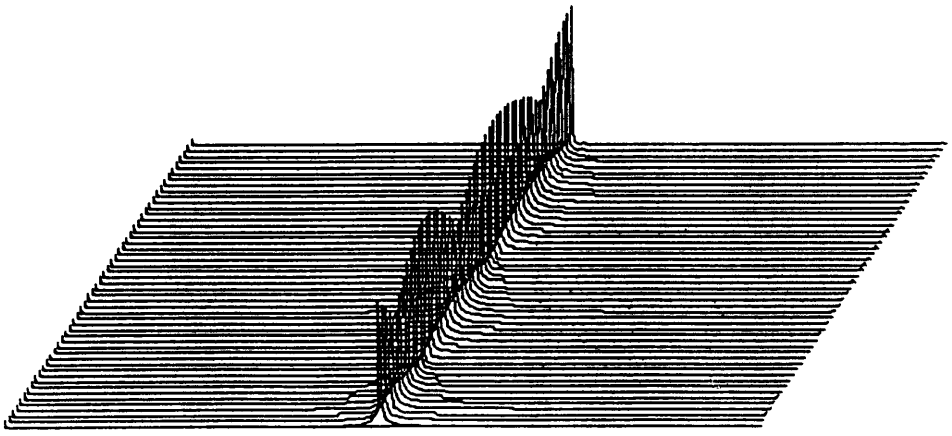

Fig. 18. 3D Sliding FFT analysis, without interpolation, of a constant amplitude, swept-frequency sine wave suggests amplitude variation.

Because the acquisition signal is known to be a sine wave, interpolation may be performed on the two highest spectral lines to determine the true frequency and amplitude of the

signal [10]. The graphics program written for the FFT Q-measurement includes this option and Fig. 19 shows the interpolated mountain-range. Clearly, the amplitude of the input signal is seen to be indeed constant.

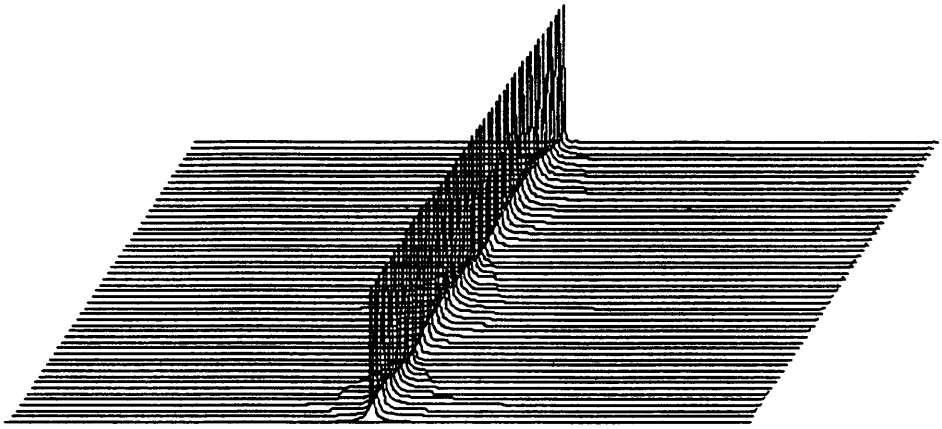

Fig. 19. Amplitude interpolation of the spectra confirms that the sine wave amplitude, swept in frequency, is not changing.

6. CONCLUSIONS AND FUTURE WORK

It has been shown that the worst-case accuracy of the FFT Q-measurement, performing analysis on a pure sine wave input, is around 10^{-4} : a factor of 10 better than required by the specification.

The system can also analyse, with an accuracy greater than 10^{-3}, damped sine wave inputs with damping-rates as fast as 0.1 ms. Even with an additional 200 mV random noise on top of the damped signal, the accuracy is within specification for damping-rates as fast as 0.2 ms.

The use of concurrent acquisition and processing allows Q-measurements to be performed every 1.8 ms : more than twice as fast as required. FFT analysis can also be performed very rapidly, and integer spectral precision instead of floating-point has been shown to be largely insignificant.

Demonstration of the Sliding FFT has shown it to be a very useful aid for appreciating the change in frequency content of the beam during its acquisition. By moving the FFT window over the data in small steps, FFT analysis can, in effect, be performed with a much higher repetition-rate than is possible using real-time methods.

The integration of the system into the PS control protocol is currently proceeding and some further modification to the software is envisaged. A note will be published in the future containing a guide to general programming of the system and a detailed breakdown of the PS Q-measurement software itself.

7. ACKNOWLEDGEMENTS

I would firstly like to thank my immediate colleagues, Messrs. E. Schulte, J. Gonzalez, and J. Belleman, for their help, encouragement and useful discussions throughout this project. I am also indebted to N. de Metz Noblat, who wrote the LynxOS device driver for the VASP-16. Finally, I would like to thank all those others with whom I have discussed this project at some time, and everyone currently concerned with its integration with the control-protocol.

8. REFERENCES

[1] G.C. Schneider, "Q-measurement with swept RC filter", CERN/MPS/SR 73-4, 22.05.73.

[2] C. Kessler, "Betatron Tune Measurement in the PS using the FFT", CERN/PS/PA Note 90-27, 30.08.90.

[3] Future Digital Systems Ltd., "Host User Manual" and "DSP User Manual", 1989.

[4] LynxOS device driver written for the VASP-16 by Nicolas de Metz-Noblat, CERN PS Division, 18.03.92.

[5] Zoran Corporation, "Digital Signal Processors Data Book", 1987.

[6] M. Serio, "Tune measurements", CERN 91-04, pp. 136-160, 08.05.91.

[7] C. Bovet, R. Gouiran, I. Gumowski, K.H. Reich, "A selection of formulae and data useful for the design of A.G. synchrotrons", CERN/MPS-SI/Int. DL/70/4, 23.04.70.

[8] E. Asseo, "Guide méthodique pour la mesure (valeurs corrigées) des oscillations bétatroniques dans LEAR", PS/LEA Note 85-4, April 1985.

[9] J.L. Gonzalez, "Étude et réalisation d'un automatisme qui supprime les instabilités transversales des faisceaux de protons, groupés en paquets de haute densité, dans l'accélérateur PS", Mémoire CNAM, December 1992.

[10] H.J. Schmickler, "Study of the Accuracy and Computation Time requirements of a FFT based Measurement of the Frequency, Amplitude and Phase of Betatron Oscillations in LEP", LEP/BI/Note 87-10, 15th May 1987.