EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH
ORGANISATION EUROPEENNE POUR LA RECHERCHE NUCLEAIRE

CERN PS DIVISION

PS/BD/Note 2002-216(Tech.)

PERFORMANCE TESTS OF A NEW FAST DIGITISER FOR BEAM DIAGNOSTIC APPLICATIONS.

Maria Elena Angoletta

Abstract

A new type of PCI-based fast digitisers has been deployed to implement new beam diagnostic systems and as a prototype for a new family of applications. The modules selected for the first tests and applications are the Acqiris DC265 fast digitiser boards, characterised by a high sampling speed, a large amount of memory per channel (2 MSamples per channel as the chosen option) and a PCI bus interface.
This note details the tests carried out, and the results obtained, to ascertain the DC265 board and crate suitability to general beam diagnostics applications.

Geneva, Switzerland
18 December 2002

PERFORMANCE TESTS OF A NEW FAST DIGITISER FOR BEAM DIAGNOSTIC APPLICATIONS.

Maria Elena Angoletta

INDEX

1. INTRODUCTION

In 2001 the PS/BD group decided to purchase a new type of PCI-based fast digitisers to implement new systems and as a prototype for a new family of applications [1]. The Acqiris DC265 fast digitiser boards, selected for the first tests and applications, are characterised by a high sampling speed, a large amount of memory per channel (2 MSamples per channel as the chosen option) and a PCI bus interface. The first development based on this new hardware, i.e. the PS Transverse Phase Space Measurement System [2], was carried out successfully and within a very short time. This note describes a series of tests (including the results) carried out on the DC265 board, crate and associated software driver, to ascertain the suitability to general beam diagnostics applications.

2. FAST DIGITISER OVERVIEW

The fast digitiser system purchased by the BD group and used for the tests is composed of a 6U CompactPCI crate with five slots. One slot is taken by a PCI extension interface board, two remain unused and available for future usage, while the last two are taken by two Acqiris DC265 fast digitiser PCI boards.

The PCI extension interface board is connected via a short cable to a PCI host card plugged into the local PC. This allows a software application running on the PC to control and exchange data with the digitiser boards, by means of a software driver. The PC software application developed for this purpose, named CaP (from Control and Processing) [3], carries out a customised data processing and reduction. In this way only a reduced data set (as opposed to the whole set of raw data) has to be transferred over the network for storage, visualisation or post-processing purposes. CaP also sets up the digitiser boards and arms them for an acquisition. The data acquisition phase is then carried out autonomously by each board, according to its setup and depending on externally connected trigger or clock signals. When the data acquisition is completed, CaP retrieves the digitised data via the PCI interface. The time required by this operation depends on the total amount of memory to be retrieved and on the form in which data are retrieved. In fact, data can be specified as "char" (ADC values) or as "floating point" (voltage values), the former being much faster than the latter.

Each DC265 board has four channels that can be simultaneously sampled at a frequency of up to 500 MHz, internally or externally generated. An external input channel is available for use as a trigger or as an external clock. Several hardware settings can also be selected for each input channel, such as the impedance (50 Ω or 1MΩ), coupling (DC or AC) and full-scale value (from 50 mV to 5 V in a 1, 2, 5 sequence). Level and slope can also be selected for the trigger input.

The internal sampling rate can be programmed and is selectable in a 1, 2, 2.5, 4, 5 sequence. This means that the higher internal sampling rates, expressed in decreasing order, are 500 MHz, 400 MHz, 250 MHz, 200 MHz, 100 MHz and so on. In the case of external clock, the external clock signal is internally converted to a square wave, and that will be the signal effectively used as a clock. The frequency of the so-generated clock is therefore user-selectable. However, the manual specifies that the external clock signal frequency should be of at least 10 MHz, therefore this is the lower limit for the clock frequency. There are also additional caveats on the external clock signal characteristics, as detailed in paragraph 6.1.

The ADC resolution is 8 bits, its aperture uncertainty is of ± 0.5 ps and the differential linearity is of ± 0.7 LSB. The 3 dB attenuation bandwidth of the input signal is 150 MHz. The default memory size connected to each channel is 128 kSamples, but deeper memories are also available. The fast digitiser system used for the tests described in this note has a 2-MSamples memory bank connected to each channel. Once the digitiser has been armed for an acquisition, the memory can be filled in different, user-selectable ways: continuously after a trigger (Single Acquisition Mode), in segments (Sequence Acquisition Mode) or in burst mode (Start/Stop External Clock Mode).

In the Single Acquisition Mode the user specifies the number of samples to be acquired. The digitiser starts sampling the input signals after receiving a trigger. The sampling rate is also user-selectable. The acquisition is then continued until the specified memory is filled, and at that moment it is stopped.

In the Sequence Acquisition Mode the acquisition memory is divided into a user-selectable number of segments of a certain length. Each segment will be acquired following a trigger. The maximum number of segments depends on the memory size per channel and it is 4000 for the 2 Msamples memory option installed on the board,. Note that the Single Acquisition Mode can be seen as a special case of the Sequence Acquisition Mode where the number of required segments is 1.

In the Start/Stop External Clock Mode the clock is generated outside the board and it is not constant. Once the board is armed, the trigger and the clock coincide and a new sample is acquired for each pulse. The user specifies the acquisition memory size during the setup phase and the board stops acquiring data when the acquisition memory is filled, disregarding any additional clock pulse it receives.

The two DC265 boards are linked by a proprietary, 1 GHz AutoSynchronous Bus System (ASBus) bridge that allows all channels to be driven by the same clock and trigger. This connection, shown in Figure 1, is located at the front of the PCI crate, it is optional and easily removable. When using the ASBus, the digitiser boards in a crate are divided between master (one board) and slaves (all the other boards). Such roles can be automatically chosen or can be forced by the user, depending on the software function used for initialising the board. In case of automatic roles allocation for the two digitiser boards case, the master board is the PCI bus board nearest to the PCI extension interface board.
The sampling clock used by all ASBus-linked boards is the clock generated by the master digitiser board, which distributes it to the other cards through the ASBus with a time delay as low as 100 ps. This is valid for both an internally generated clock and a clock derived from an external signal (external clock or burst mode). The external clock signal can be connected to the master digitiser board only.
The trigger signal can be connected to any board, independently on whether that board is the master or a slave. In case it is connected to a slave board, the trigger is sent via ASBus to the master board thus maintaining synchronicity between different boards.

A function included in the board software driver combines all ASBus-linked boards to make them appear as a single device with more channels. As a consequence, Channel 1 to 4 as seen from the driver correspond to Channel 1 to 4 in the lower board, while channels 5 to 8 correspond to Channel 1 to 4 of the upper board.

Figure 1 shows the 5-slots PCI crate containing two DC265 boards linked by the ASBus.

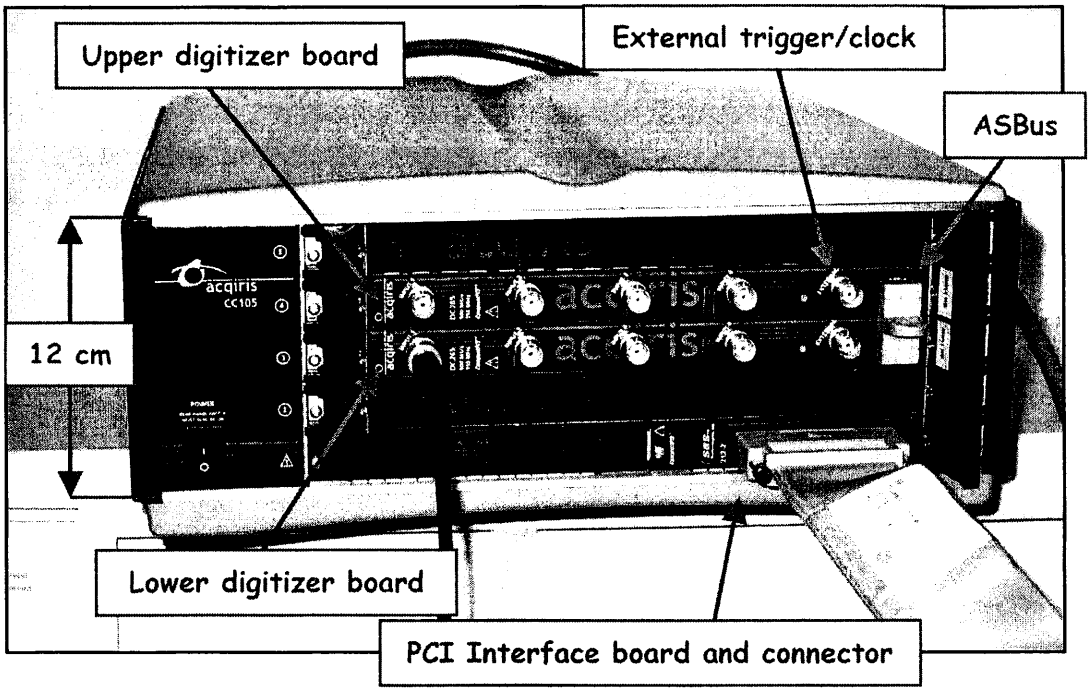

Figure 1: *Five-slot PCI crate with two DC265 fast digitiser boards and the PCI interface card installed.*

3. TESTED PERFORMANCE AND FEATURES.

The tests described in this note focus on three different aspects of the digitiser, two of them are performance-related and one is a feature. These aspects are relevant for current and future applications in LHC-operation relevant machines, cover a wide scope of beam diagnostics and include both hardware and software characteristics.

First, the ASBus should allow quasi-synchronous data acquisition (with an application-dependent tolerance on DAQ triggering and clock) on channels located on different digitiser boards, hosted in the same crate. This is essential in case of inter-channel measurements, for instance if some signal cables coming from the same pick-up are connected to more than one digitiser board.

Second, when working in the Sequence Acquisition Mode, the dead time should be as short as possible and typically up to 1.5 μs for PS applications. One must recall that the dead time is the minimum interval between the end of a segment acquisition and the beginning of the next segment acquisition. The dead time corresponds to the trigger rearm time of the digitiser and depends on the sampling frequency. A short dead time is essential for instance when one wants to fill a segment for each beam revolution period. This will be possible only if the dead time is sufficiently small compared to the revolution period.

Third, the board should support "burst mode" operations, i.e. acquisitions clocked by a non-continuous train of pulses, generated outside the PCI crate. This requirement is dictated by the need to study the evolution of a certain phenomenon over as long a time as possible. The observed object can be for instance the shape of a bunch circulating in a machine, described by signals coming from a pick-up. The bunch will pass in front of the pick-up only once per

revolution period. As a consequence, digitising the pick-up signal with a continuous clock (Single Acquisition Mode) would mean filling partially the memory with irrelevant data, corresponding to the bunch being located in the machine far from the pick-up. The Segment Acquisition Mode allows triggering an acquisition immediately before the bunch arrives at the pick-up. However the memory usage might still be inefficient, especially in case of a very short bunch. In that case, in fact, the number of observable turns would be limited by the maximum allowed number of segments, even if the memory is not completely filled. The Burst Mode allows, on the contrary, to use the entire memory allocated to each channel to store meaningful data. The drawback is that external hardware has to be provided and connected to the digitiser boards.

Paragraphs 4 to 6 describe the tests carried out, the methods used and the results obtained. In paragraph 4 and 6 the *CaP* program is used to set-up the fast digitiser, arm it for an acquisition and retrieve data from the its memory. This program runs under Windows 2000, is written in VisualC and uses the manufacturer-provided version of the driver. In addition, it implements also the processing capabilities needed for beam diagnostics instruments, such as for the system detailed in [2]. On the other hand, the AcqirisLive program [4] was used for the tests described in paragraph 5. This is a program provided free of charge with any Acqiris digitiser and allows testing and setting up the fast digitiser, capturing data, displaying and archiving them.

4. QUASI-SYNCHRONOUS DAQ ON DIFFERENT CHANNELS

4.1 Tests description

The same 10-ns wide pulse at 1 MHz frequency was fed via a splitter and same-length cables to 2 channels located on different digitiser boards. In the first test, described in paragraph 4.2, a 500 MHz internal clock was used. In the second test, described in paragraph 4.3, a 400 MHz external sinusoid was used as external clock.
The output data were acquired and displayed by using the CaP program, as shown in figures 2 to 5. The horizontal axis represent the sample index, counted from the start of the acquisition. The vertical axis represent the ADC count corresponding to the voltage applied to the input channel. Subsequent samples are linked by lines. Acquired data were zoomed-in to determine to which extent the data acquisition had been performed synchronously on both channels. For the purpose of the PS Phase Space Measurement System a delay as low as even 1 ns was acceptable.

4.2 Test with Internal Clock

The digitiser was setup to generate internally a 500 MHz clock after receiving the trigger. An external trigger was applied to the External Input on the lower board in the digitiser crate, and data from channels located on both DC265 boards were acquired.
Figure 2 shows the input pulse train as it was acquired by channel 2, sampled at 500 MHz. As a consequence, two consecutive samples are spaced by 2 ns.

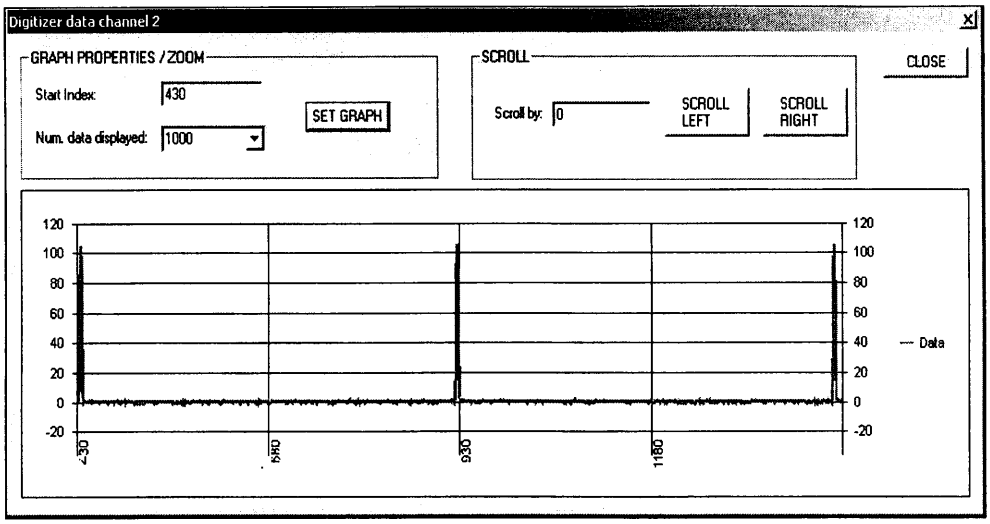

Figure 2: *Input pulse signal digitised at a 500 MHz frequency internally generated, as acquired by Channel 2. The horizontal axis represents the sample index, and the vertical axis represents the ADC count. A thousand samples are shown, corresponding to a 2 μs window.*

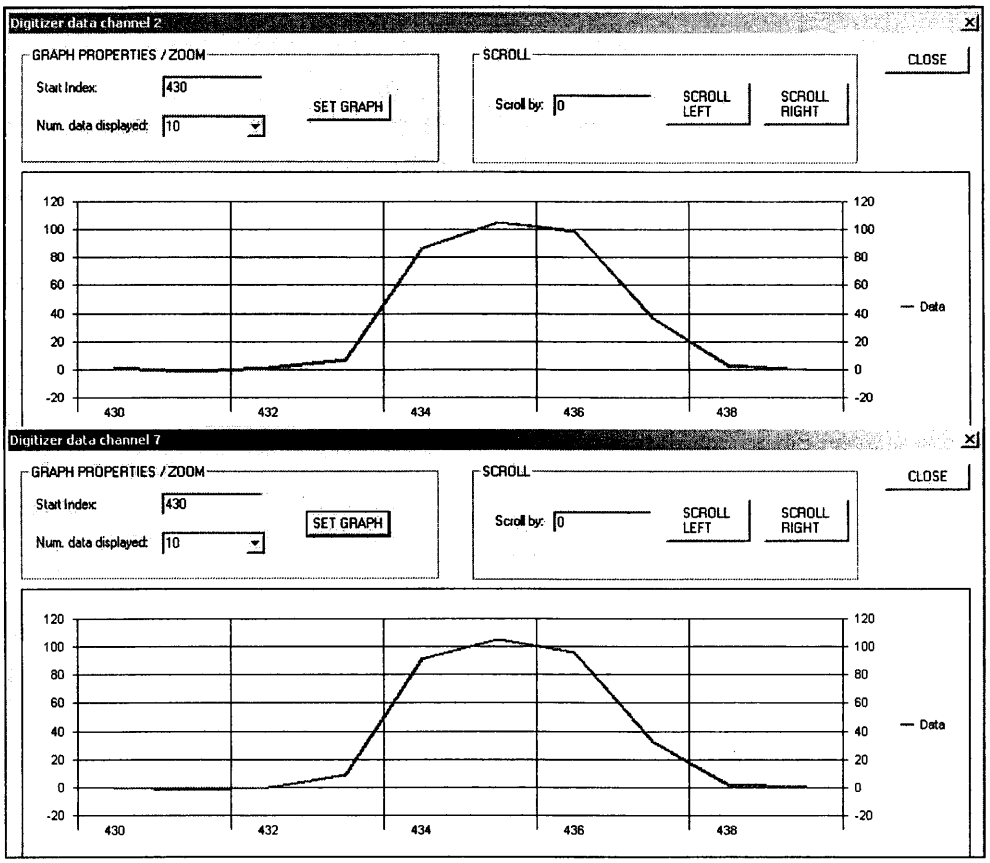

Figure 3: *Zoom on digitised data acquired on channels 2 (upper picture) and 7 (lower picture). Subsequent samples are connected by lines.*

Figure 3 shows a zoom on the same pulse as digitised by Channel 2 (located on the lower board) and by Channel 7 (located on the upper board). The two curves represent the digitised version of the same signal and show a time lag lower than about half a horizontal unit, i.e. 1 ns.

4.3 Test with External Clock

The digitiser was setup to accept an external signal as clock. The signal used was a 400 MHz, 3 V peak-to-peak sine wave applied to the External Input connector of the crate's lower board. Figure 4 shows the input pulse train as it was acquired by channel 4, sampled at 400 MHz. As a consequence, two consecutive samples are spaced by 2.5 ns.

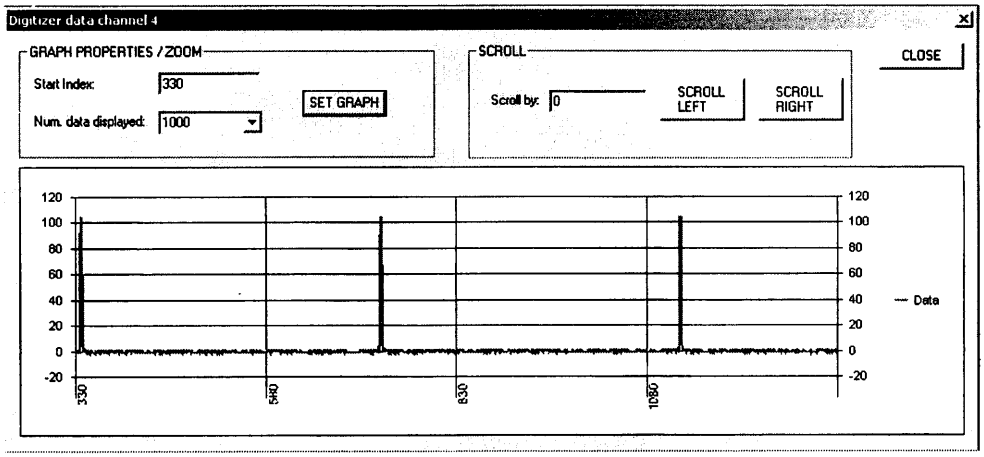

***Figure 4**: Input pulse signal digitised at a frequency of 400 MHz externally generated. A thousand samples are shown, corresponding to a 2.5 μs window.*

Figure 5 shows a zoom on the same pulse as digitised by Channel 4 (located on the lower board) and by Channel 8 (located on the upper board). Similarly to the previous case of Figure 3, the two curves show a time lag lower than about half a horizontal unit, i.e. 1 ns.

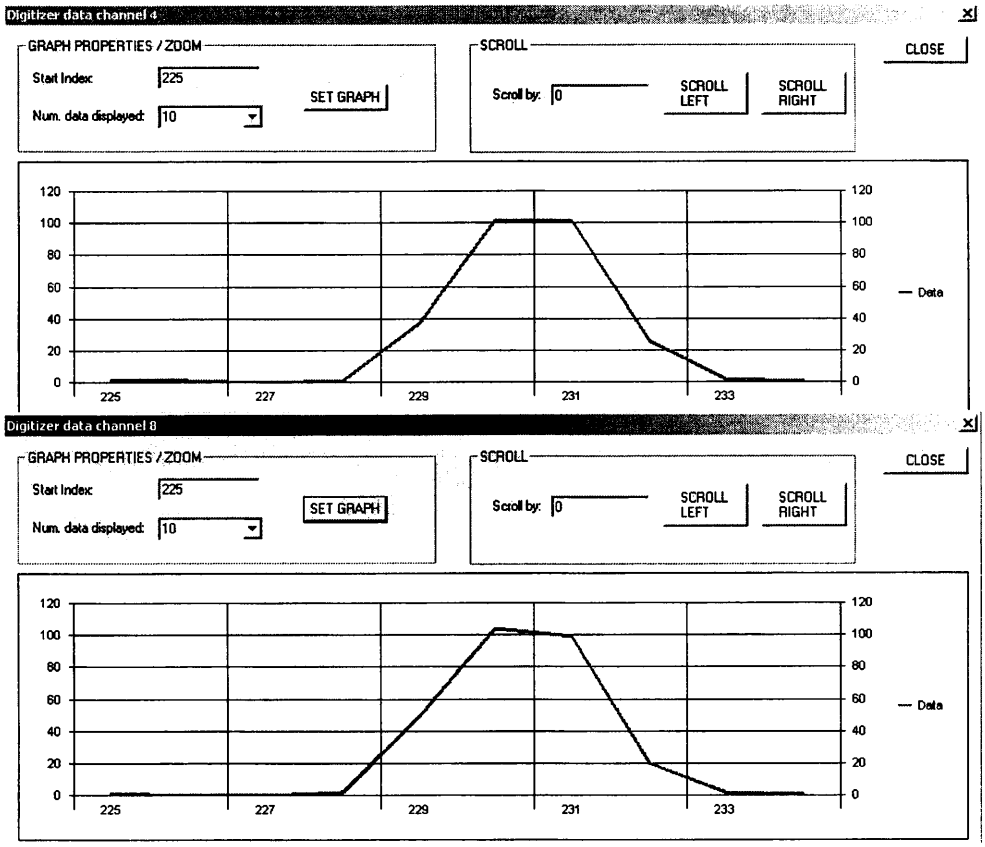

***Figure 5**: Zoom on digitised data on channels 4 (upper picture) and 8 (lower picture). Subsequent samples are connected by lines.*

5. DEAD-TIME IN SEGMENT MODE

The Acqiris Digitisers User Manual [4] specifies in paragraph 3.3.3 the dead time as being less than 500 ns. However, what the manual does not mention is that this value refers to a sampling frequency of 1 GHz, and that the dead time is greater at lower sampling rates. Since the DC265 fast digitiser can sample at frequencies of up to 500 MHz, it is necessary to measure the dead time upper bound at various sampling rates. As mentioned in paragraph 3, this test makes use of the AcqirisLive program.

5.1 Test description

The AcqirisLive program was setup to use the board as a Transient Recorder. A complete acquisition was specified as composed of 20 segments, each one made of 160 samples. The sampling rate during a segment was changed in different tests and ranged from a maximum of 500 MHz to a minimum of 100 MHz.
A sinusoidal signal was fed to Channel 1; the acquisition trigger was set on the same channel, with an input voltage offset equal to 0 V and with positive slope.
The period of the sinusoid was selected so as to be initially much higher than the total acquisition time of a sequence.
A complete sequence acquisition was taken, and the timestamps for each segment were inspected via AcqirisLive. The averaged time interval between the start of subsequent

segments should equal the sinusoid period, if the sinusoid period is sufficiently large compared to the segment DAQ time.
This situation is depicted in Figure 6 for a sinusoid with frequency 730 kHz sampled at 500 MHz. The sinusoid period of 1.37 μs is sufficiently long as to allow acquiring consecutive segments at consecutive sinusoid periods. This can be seen in the "Segment Timestamps" window, where the average time interval between consecutive segments is 1.37 μs. The Waveform display shows the data acquired in 20 segments. Each segment contains 160 samples of the input sinusoid, starting when the sinusoid passes through the 0 V value with positive slope. The samples of the same segment are inter-spaced by 2 ns.

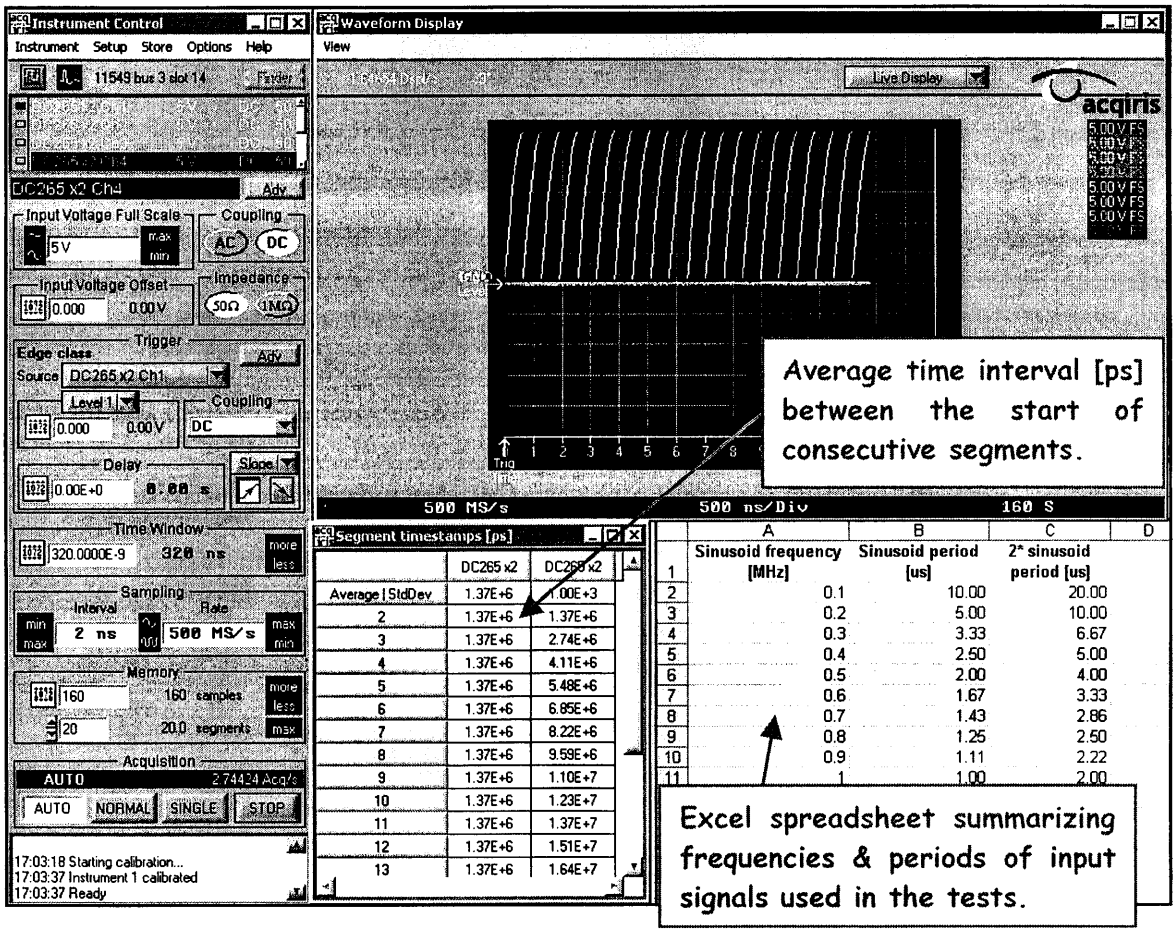

***Figure 6**: Acquisition in Sequence Mode by the AcqirisLive program of a 730 kHz frequency sinusoid, sampled at 500 MHz. In the lower right corner an Excel spreadsheet summarises the different sinusoid frequencies used during the tests and their corresponding periods.*

From this situation, and keeping constant the sampling rate, the frequency of the input signal was increased in constant-length steps until the average interval between two consecutive segments jumped to the double of the input sinusoid period. This occurrence meant that an acquisition trigger was missing because of the dead time.
There was a frequency region where a certain number of segments were acquired at one sinusoid period time interval, others at two periods distance. This is shown in Figure 7, where a zoom into the Segment Timestamps window is shown for a 750 kHz input signal (corresponding to a period of 1.33 μs). The sampling frequency is still 500 MHz.

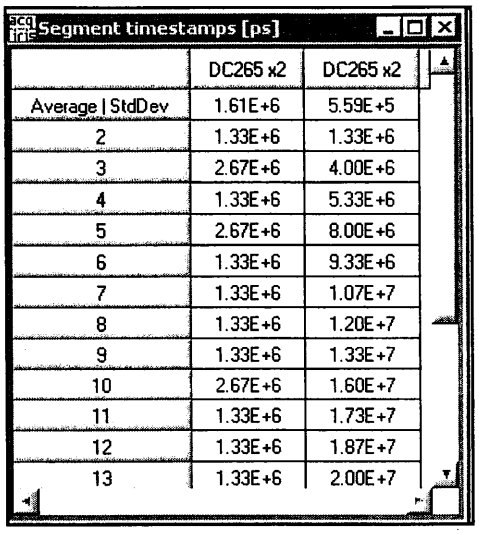

	DC265 x2	DC265 x2
Average | StdDev	1.61E+6	5.59E+5
2	1.33E+6	1.33E+6
3	2.67E+6	4.00E+6
4	1.33E+6	5.33E+6
5	2.67E+6	8.00E+6
6	1.33E+6	9.33E+6
7	1.33E+6	1.07E+7
8	1.33E+6	1.20E+7
9	1.33E+6	1.33E+7
10	2.67E+6	1.60E+7
11	1.33E+6	1.73E+7
12	1.33E+6	1.87E+7
13	1.33E+6	2.00E+7

***Figure 7**: Zoomed-in AcqirisLive "Segment Timestamps" window for a 750 kHz sinusoidal signal. The time interval between segments corresponds sometimes to one period of the sinusoid input (1.33 μs), other times to two periods (2.67 μs).*

The dead time can be measured by using equation 1, with *Points per Segment* always equal to 160 in the tests carried out.

*Dead Time [μs] = Sinusoid Period [μs] - (Sampling Interval [μs] * Points per Segment)* (1)

By substituting in (1) the longer sine wave period that allows consistently acquiring a segment every period, one obtains a lower bound on the dead time. Similarly, one obtains an upper bound by substituting in (1) the shorter sine period, which allows to consistently acquire a segment every other periods.

5.2 Test results

Table 1 details the ranges of dead time values taken in correspondence to different sampling frequencies. A numerical example will detail how the numbers were obtained.

For a sampling frequency f_S = 500 MHz, corresponding to a sample period T_s = 2 ns, 160 samples per each segment would give an acquisition window of 0.32 μs/per segment.

By incrementing the sine wave frequency by 10 kHz consecutive steps, and by looking at the averaged time interval between segments, two frequencies $f_{SIN,H}$ and $f_{SIN,L}$ were found.

$f_{SIN,H}$ is the higher sinusoidal frequency that allows to acquire a segment for each sinusoid period, without missing any period. The value found was 730 kHz (i.e. $T_{SIN,H}$ = 1.37 μs).

$f_{SIN,L}$ is the lower sinusoidal frequency that consistently allows to acquire a segment only every other sinusoid period. The value found was 790 kHz, (i.e. $T_{SIN,L}$ = 1.26 μs).

From (1) one can derive the upper (2) and lower (3) values defining the range were the dead time lies for the 500 MHz sampling frequency.

$$\text{Dead Time}_{\text{UPPER VALUE}} = (1.26 - 0.32)\ \mu s = 1.05\ \mu s \quad (2)$$

$$\text{Dead Time}_{\text{LOWER VALUE}} = (1.36 - 0.32)\ \mu s = 0.95\ \mu s \quad (3)$$

Table 1 was filled by proceeding in a similar fashion for different sampling frequencies. It is worth pointing out that the measured dead time ranges are significantly lower than the nominal dead times for several commercially available digital oscilloscopes.

Sampling frequency [MHz]	Dead Time Range [μs]	
500	1.05	0.95
400	1.16	1.03
250	1.49	1.36
200	1.70	1.58
100	2.95	2.57

***Table 1**: Measured dead time ranges as a function of sampling frequency.*

As expected, Table 1 shows that the dead time ranges increase as the sampling frequency decreases. In the case of the Transverse Phase Space System the measured dead time ranges allow segment acquisition with a sampling frequency as low as 200 MHz. In fact, the PS revolution time is about 2.1 μs and the bunch length can be of up to 200 ns.

As mentioned in paragraph 2, the 2 MSamples option used for these tests allows a maximum of 4000 segments. This maximum number of segments is fixed by the board driver, depends on the amount of installed memory per channel and does not allow for an optimal memory usage, because of the memory overhead needed by the driver itself. As a consequence, lowering the sampling frequency or reducing the number of samples per segment might not result in an increased number of segments acquired.

6. BURST-MODE OPERATION TESTS

6.1 INTRODUCTION

For operation in the Burst-Mode, an in-house built burst generator module [5] is used to generate the burst clock. Figure 8 gives an overview of the system set-up used to carry out the burst mode tests. The burst generator module produces a train of pulses after receiving a trigger. The number of pulses generated for each burst, their frequency and the time delay between the trigger arrival and the generation of the first pulse are selectable. The output of this module was connected to the External Input of the lower digitiser board. The burst generator was set-up by the CaP program via the PC parallel port, and received the triggers from a trigger generator.

The digitiser drifts into saturation if the time interval between the start of two bursts lasts for more than 100 ns. The digitiser has to recover from this saturation hence the first sample of

each burst will be invalid i.e. discarded. This has to be accounted for when specifying the number of pulses needed for each burst.

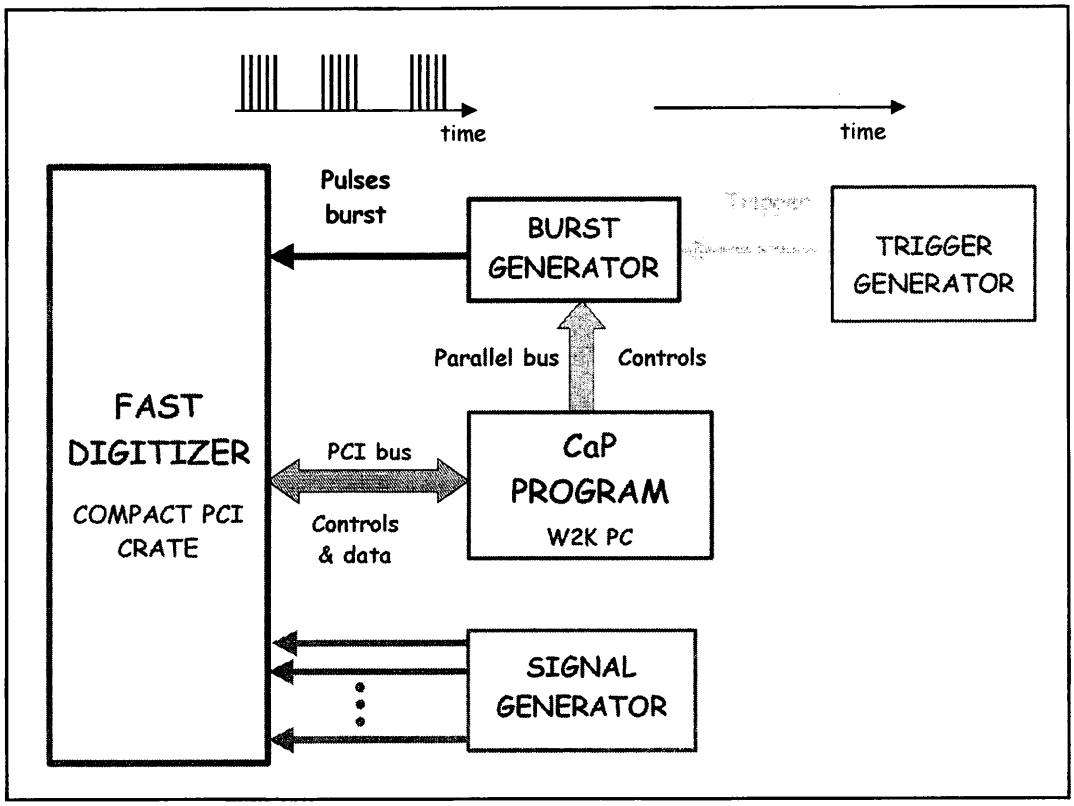

***Figure 8**: Overview of the hardware set-up used when carrying out tests of burst mode operation.*

The number of clock pulses to be generated in order to complete the acquisition must be higher than the number of samples to acquire. This is due to the internal software driver management routines. Once the data acquisition is completed the board will not react to any additional trigger, until it is re-armed again.

It is not possible to specify the whole memory available as acquisition memory for the burst mode, since a buffer at the end of the memory has to be left for additional software activities carried out by the driver. However, when retrieving digitised data from the digitiser memory, the additional points acquired will not be visible and only as many data as were specified can be read from the digitiser board.

The manual specifies that a voltage signal should have amplitude of at least 500 mV peak-to-peak if used as clock, both for burst and external clock modes. It turns out that the amplitude needs to be much larger than 500 mV peak-to-peak to reliably trigger acquisition at high frequencies. In fact, it was not possible to perform reliable data acquisition with an 800 mV peak-to-peak square burst signal at frequencies over 100 MHz. This is because the clock is generated by the digitiser from the external signal by converting it into a square wave. It is very important for this to happen correctly that the threshold level is set to the mean value of the incoming signal. Because the offset of the input clock range is not calibrated, it may exhibit some 'natural' offset that can be as high as 300 mV, with the consequence that the threshold level might fall on the edges of the input signal. To avoid this problem, the

manufacturer verbally recommended using a clock signal with at least 1.2 V peak-to-peak amplitude in order for the clock to be reliably generated at frequencies approaching 500 MHz.

6.2 MEMORY WRAPPING-UP PROBLEMS

A first set of tests carried out during the spring 2002 made clear that the driver version could not appropriately support burst mode operations. Since then, the board manufacturer upgraded appropriately the driver to support burst mode, as detailed in paragraph 6.3. It is however useful to describe how the problem was detected and documented in a simple but effective way.

The burst generator was set up to produce 128 pulses per burst. The acquisition memory size was set to 512, corresponding exactly to 4 bursts. The input signal voltage range in the digitiser was set to 5 V peak-to-peak. The signal to digitise was a DC voltage whose value was changed from burst to burst, thus allowing to track back when acquired data had been generated. Table 2 details the order in which bursts were produced (from 1 through 7), the DC value of the input that was present during each burst and the corresponding ADC count. Every burst was manually triggered.

Burst number	Peak-to-peak input signal amplitude [mV]	ADC count
1	38	2
2	512	26
3	890	46
4	1170	60
5	1300	67
6	1400	72
7	1800	92

***Table 2**: Bursts generated during the test, input signal DC voltage and corresponding ADC count.*

Since the memory was set to 512 samples, the acquisition of 4 bursts only was foreseen. As mentioned in paragraph 6.1, more bursts (in our case bursts number 5, 6 and 7) had to be provided for the acquisition to be completed.

Figure 9 shows the results obtained. As expected, 512 samples were retrieved from the digitiser memory. Data corresponding to bursts number 1 through 4 were digitised; however, it appears that also burst number 5 was partially acquired and actually overwrote several points belonging to the first burst. This memory-wrapping problem was caused by a malfunction of the software driver rather than by the filling up of the whole memory attached to each channel. In fact, the specified acquisition memory per channel (512 samples) was much less than the whole memory available per channel (2 MSamples).

An improved version of the driver was afterwards released. This version could manage correctly the burst mode up to 256 pulses per burst, but presented the same problem for longer bursts.

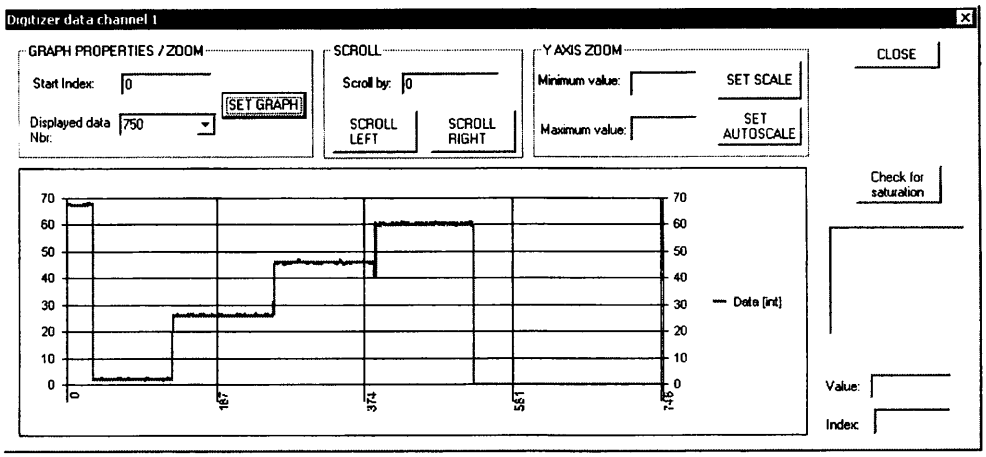

***Figure 9**: Digitised data showing the wrapping-up effect. The horizontal axis reports the sample index, and the vertical axis reports the ADC count corresponding to different DC input voltages. Sampled points are connected by lines.*

6.3 SUCCESSFUL BURST-MODE TESTS

During the summer 2002 a new version of the digitiser driver was delivered, that could successfully manage the burst mode even for a number of pulses per burst higher than 256. Tests such as the one presented in paragraph 6.2 were successfully carried out. It was also possible to determine that the number of additional pulses needed to complete an acquisition was between 500 and 750 for a 500 MHz sampling frequency.

As a further step, the burst mode operation was tested in the case of a short time interval between bursts and with a two million samples acquisition memory. These in fact are the conditions that had to be satisfied to acquire data for a short interval every revolution period in a short circular machine (revolution periods of a few microseconds) and for as many turns as possible. Of course, in such situations it was not possible anymore to manually trigger each burst.

During the first test, a 30 kHz triangle signal was used as input signal. The burst generator was set-up to generate 512 pulses for each burst; two subsequent pulses in a burst were spaced by 2 ns. The trigger generator was programmed to generate 3908 triggers spaced 2 μs apart. The digitiser acquisition memory was set to $2 \cdot 10^6$ samples. Under such circumstances, a little over 3906 bursts were needed to fill the acquisition memory, and a sufficient number of additional pulses was available for terminating the acquisition.

Figure 10 shows the first 10000 samples acquired, connected by lines. The foreseen drift in the digitiser owing to the time interval between bursts is clearly visible as the first sample for each burst is clearly wrong. One should also be aware that in burst mode there are always three additional samples placed at the beginning of the digitised array. They are caused by the internal driver mechanism and should be disregarded.

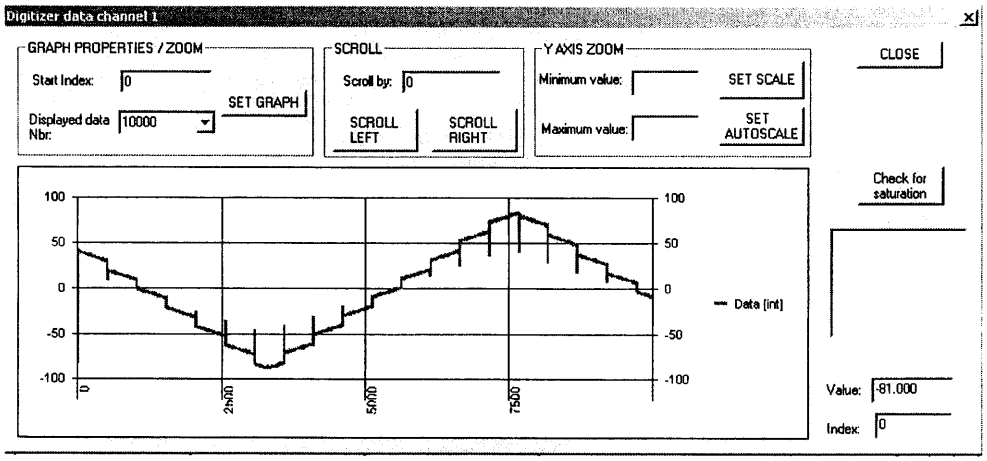

***Figure 10**: Triangle signal of 30 kHz frequency, acquired in bursts of 512, 2 ns-spaced pulses. Samples are connected by lines and bursts are started every 2 μs. The starting point of each burst can be recognised by its value, which deviated from the digitised triangular signal because of the ADC drift.*

A second test was made, using a higher number of pulses per burst. The same 30 kHz triangular signal was selected as input and the acquisition memory size was still set to $2 \cdot 10^6$ samples. Each burst contained 1024 pulses and all bursts were spaced by 5 μs. Exactly 1954 triggers were needed to complete the acquisition, of which a little over 1953 were used for actually digitising the input signal and the final part of the last burst was used to complete internal software tasks of the driver.

Figure 11 shows the last 1700 samples acquired; the ADC drift value acquired as the first point of each burst is now hidden in the vertical line. The value corresponding to the first point of the last burst is highlighted in Figure 11 and its index and value are reported in the right lower corner. It is easy to check that no samples were lost. In fact, the initial point of the last burst had index 1999875. If one takes into consideration also the 3 additional points inserted at the beginning of the digitised sequence, the real index is 1999872, i.e. corresponding to burst number 1953 (1999872 / 1024 = 1953).

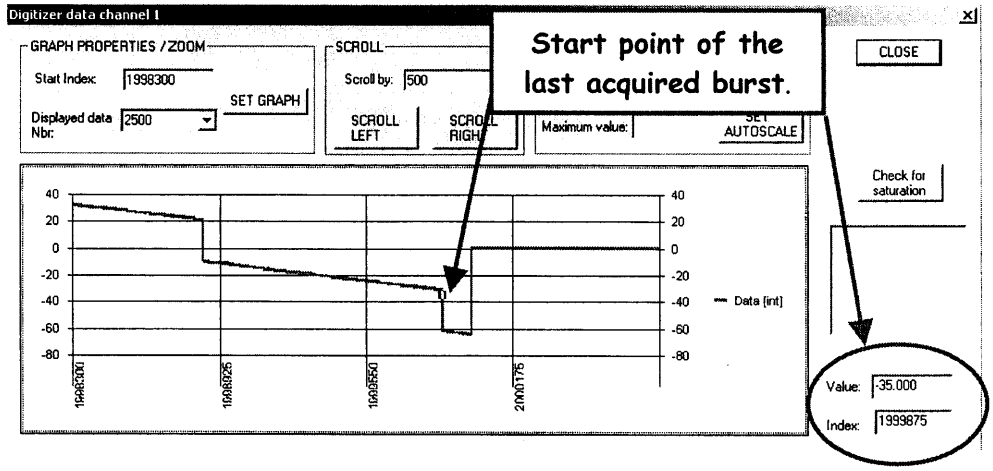

***Figure 11**: Final samples of a 30 kHz triangle signal, acquired in bursts of 1024, 2 ns-spaced pulses. Sampled points are connected by lines and bursts are started every 5 μs. The value corresponding to last burst's first point is highlighted, and its index and value are reported.*

7. CONCLUSIONS

The PCI crate with DC265 fast digitiser boards and latest version of the software driver satisfies all needs detailed in paragraph 3. Under Sequence Acquisition Mode the dead time is short compared to nominal values for several oscilloscopes on the market. An additional useful feature is the possibility of carrying out a customised data reduction on the local PC hosting the CaP program, and then of transferring a reduced data set thus created through the network to the data storage, visualisation and post-processing hosts. Because of these performance and features, this digitiser looks attractive for use on several beam diagnostics applications.

8. ACKNOWLEDGEMENTS

Several tests for the Burst Mode Operations were carried out in collaboration with M. Ostapowicz (technical student), J.L. Gonzalez and J. Belleman from PS/BD.

9. REFERENCES

[1] V. Chohan, *Purchase of a fast digitiser system for beam studies*, PS/BD/Note 2001–005 (Tech).

[2] M. E. Angoletta et al., *Analysis of a multi-turn beam position measurements in the CERN PS*, EPAC 2002 Proceedings, June 2002, Paris.

[3] M. E. Angoletta, *CaP User's Manual - Control And Processing program for the Acqiris DC265 Digitiser*, in preparation.

[4] Acqiris, *Acqiris Digitisers User Manual*, ZM020010B Rev. A, March 2002.

[5] M. Ostapowicz, *Burst Generator*, PS/BD/Note 2002-162.