ACHIEVING HIGH-SPEED DATA ACQUISITION FOR REAL-TIME BEAM CONTROL AND MEASUREMENT

R. Soden, Y. Maumary, P-F. Maistre, Agilent Technologies SA, Plan-les-Ouates, Switzerland S. Narciso, G. Hill, Agilent Technologies, Inc. Loveland, CO, U.S.A.

Abstract

Digital data acquisition in real-time applications falls into two categories: digitizing a stream of data without missing a single sample point, and capturing a stream of triggered events without missing a single trigger.

Maintaining these data streams over long periods requires an optimized combination of analog signal conditioning, precise digitization, digital data reduction and high-speed data transfer. This paper describes suggested methods to reduce the amount of measurement data required, reduce the amount of that data that is to be transferred to a measurement where possible, and then transferring this reduced data in the most rapid fashion. Our approach uses a combination of hardware, firmware and software elements that are designed to work together, optimizing performance and managing the data bottlenecks.

New hardware standards and architectures are discussed that improve the capabilities of today's technologies, providing access to higher data and measurement flux. Applications presented in this paper include high trigger rate capture for beam steering and fill pattern monitoring in charged particle accelerators.

THE KEY ASPECTS OF REAL-TIME MEASUREMENT

Measurement in real time requires a constant flow of accurate data. Maintaining the flow and the accuracy of the data transferred requires specific considerations to the techniques used.

Accurate measurement

Measuring once takes half the time of measuring twice! But to have confidence in a measurement means using equipment that is properly specified and meets those specifications.

These specifications need to go beyond the banner specifications of Bandwidth, Sampling rate, Resolution, and into the secondary specifications that talk of distortion and noise, which deteriorate signals and measurement.

Figure 1 shows the how noise and distortion are typically specified for high-speed digitizers. There are a number of surprising interactions between performance aspects that contribute to diminish the accuracy of the final acquisition. The effective number of bits (ENOB) antonymous to signal to noise and distortion (SINAD), provides a good measure of how all other performance aspects add up to the quality of the measurement of the final system. Previous work has discussed noise and distortion in detail[1]. Unfortunately there is no simple answer to the question "Which digitizer should I use in my system?" Instead, each application and device must be reviewed on a case-by-case basis, with a view to optimizing the required measurement.



Figure 1: Noise and distortion specifications.

Data reduction

Smart handling of the acquired data on the card can reduce the amount of data that needs to be transferred to the system.

Segmented acquisition of only required data is the first step (trigger, delay, acquisition of pre-determined length vs, trigger, acquire), thus acquiring only the required data, not the preamble and post-acquisition data.

Data processing within on-board field programmable gate arrays (FPGAs) can be used on the module to reduce data down to a measurement.

Peak detection identifies peaks and troughs in data, outputting only peak heights and peak times, rather than the entire data waveform.

Waveform averaging captures repetitive signals and averages them to remove random noise.

Real time FFT processing reduces a temporal signal into its frequency power spectrum.

Data transfer

Optimized data transfer is achieved through the mix of how the data is transferred, and the bus over which the data is passed. Today, high-speed solutions use a number of different bus systems:

- PCI -up to 532 MB/s low latency 64-bit/66 MHz
- PCIe up to 8 GB/s 16 lane 2.0 (16x)
- VME/VXS 3.125 Gbps with aggregate up to 2.5 GBytes/s
- Optical data ports up to 4.25 Gbps

Burst transfer rate is the key parameter to maximize the measurement throughput.

- Burst transfer rate depends on the size of the transfer.
- Burst transfer rate depends on the chip set payload size.

Figure 2 shows the maximum trigger rate TR_{max} for a segmented acquisition process as a function of the segment size *Seg*.

This example assumes we have a single channel acquisition with 8-bit ADC resolution. Data transfer is accomplished over a PCI Express (PCIe) bus interface with a maximum data rate of 1100 MB/s (PCIe 8x standard). As most high-speed digitizers combine a smaller amount of fast memory a larger amount of slower, but less expensive, acquisition memory, we assume an internal fast memory of 600 kS and an external DDR2 memory of 512 MS. The latency in these memories limits the inter-segment dead time. That is the dead-time between two segments when an acquisition cannot be made. For the purposes of this calculation we assume a dead time of 800 ns for the internal memory (DTe).

Two factors limit the maximum trigger rate in this calculation. The first is this dead-time:

$$if \ Seg < 600kS, \quad TR_{max} = \frac{1}{DTi + \left(\frac{Seg}{SR}\right)},$$
$$if \ 600kS < Seg < 512MS, \quad TR_{max} = \frac{1}{DTe + \left(\frac{Seg}{SR}\right)}$$

where SR is the sampling rate of the acquisition. Maximum trigger rate is the inverse of the period needed to acquire the defined segment length at the defined sampling rate, plus the fixed dead-time between segments, as defined by the requirements of the segment length.

The second factor in the trigger rate is the data rate of the PCIe interface. As we are limited to 1100 MB/s through the 8x PCIe, $Data_{8x}$, our accumulation rate of data cannot exceed this limit without halting the acquisition. This accumulation rate, is simply the product of the segment size Seg and the Trigger rate,

$$Seg \times TR_{max} < Data_{8x}$$
.

With shorter segment lengths, trigger rates are limited by the inter-segment dead-time. As segment lengths Instrumentation increase, the impact of this fixed-length dead-time decreases, most notably for the higher sampling rates.

The backplane bottleneck presented by the PCIe interface limits our acquisition only as we approach these higher sampling rates. From Figure 2 we can see that for all sampling rates over 1 GS/s, the maximum trigger rates are identical for segment lengths over 110 kS.

This is even more apparent in Figure 3, where we use the PCIe x4 standard backplane for our model. Here the data rates are limited to around 500 MB/s, and we see that sample rates over 500 MS/s are limited by this bottle neck for segment lengths in excess of 800 samples.



Figure 2: Maximum trigger rate as a function of segment size in a segmented acquisition process, using a PCIe x8 interface.



Figure 3: Maximum trigger rate as a function of segment size in a segmented acquisition process, using a PCIe x4 interface.

Operational Software

The final aspect to be discussed is the software environment within which the real-time processes are controlled and monitored.

The software code must be optimized for the available system resources. Our low-level kernel takes advantage of hardware configuration and bus to maximise throughput. It uses Scatter-Gather DMA with multibuffer single stream for high speed data transfer. A high level API gives access to the kernel; this is easy to program to, with an entry point at the level of SCPI commands.

REAL-TIME MEASUREMENT EXAMPLES

The use of these measurement principles in high-speed real-time beam instrumentation is proven. Examples include the OASIS system at the Large Hadron Collider at CERN and fill pattern monitoring at the Australian Synchrotron.

Beam control at the LHC, CERN

The performance of the LHC depends on the performance of the injector chain. Along that chain, the Open Analogue Signal Information System (OASIS) acquires and displays analog signals in the acceleration domain.

The signals come from every CERN accelerator and are sampled using various types of digitizers. The digitizers are co-located with front-end computers (FECs) that perform first-level processing. This preserves bandwidth when the acquired data is sent through an Ethernet network for display on a workstation running a dedicated virtual oscilloscope (Vscope) application.

Vscope is a software oscilloscope that takes data from various types of digitizers and displays all waveforms as if they came from the same type of hardware module. Through this scheme, CCC personnel can perform sideby-side comparisons of multiple signals from throughout the injector chain. The OASIS architecture aligns the signals to ensure meaningful comparisons by synchronizing acquisition settings across all of the hardware digitizers.

Figure 4 shows an OASIS Vscope example bunch rotation, similar to that seen in the LHC beam. More than 2,000 individual signals are available to the operators and physicists who use OASIS. They are looking for a variety of indicators from the LHC and its injector chain: instabilities in the proton bunches; quality of the beams; current in the fast-pulsing magnets ("kickers"); phase relationships between the kickers; and the state of the RF signals in the accelerating cavities. All of these characteristics help determine if conditions are optimum for whichever experiments are scheduled for a given shift.

High-performance digitizers are used to perform wideband beam monitoring and monitor forward and reverse RF signals in the accelerator cavities. In a literal sense, the beam monitoring measurements are made possible by Agilent U1064A digitizers, which provide sufficient speed and bandwidth to capture the signals of interest. In the RF measurements, Agilent digitizers measure direct and reflected energy and thereby help determine the final energy provided to the beam during each revolution around the accelerator.



Image courtesy of CERN. © Copyright CERN Geneva.



Real-time measurement of the fill pattern at the Australian Synchrotron

The Australian Synchrotron is a third generation light source, operating at 3.0 GeV with an expected electron bunch width of 20 ps. The storage ring's harmonic number is 360 and its period of revolution is 720.5 ns.

To ensure a high-quality beam, the synchrotron uses fill-pattern monitoring (FPM) to measure real-time intensity distribution of electron bunches in the storage ring[2]. Knowledge of the electron fill-pattern profile is important for experiments that require precise spatial and temporal control. For experiments that are spatially and temporally sensitive, a source with a known temporalintensity profile makes it possible to analyze the effects of the radiation source on the results.

The combination of an ultra-fast photo diode, suitable electronics and the Agilent Acqiris digitizer enables realtime measurements of fill patterns with a resolution of one RF bucket. This use of the FPM is now an integral part of the control system software at the Australian Synchrotron.

NEXT GENERATION HARDWARE ARCHITECTURES FOR REAL-TIME CONTROL AND MEASUREMENT

New advances in hardware interfaces allow high-speed serial buses with data transfer rates beyond that of older parallel bus architectures. These new bus architectures are implemented in a number of different form factors that lend themselves to different applications.

PCI Express, or PCIe, is a high-speed serial replacement for the PCI and PCI-X bus.

Data rates of PCIe currently reach up to 8 GB/s for 16 lane PCIe v2.0, this will increase to 16 GB/s with v3.0. However the latency of PCIe is actually higher than some of the older architectures.

The transfer of short bursts of data can be dominated by bus "handshaking." It is often wiser, when possible, to collect bursts of data together in bank and read the entire bank out in one go. We have found that for data packets of less than around 30kB, such as pulse data, standard parallel PCI (64-bit/66 MHz) can outperform the serial 4 lane PCIe bus. To truly take advantage of these improved data rates, we need to combine them with on-board FPGA data processing that can make the measurement on the card.

AXIe is an open standard based on AdvancedTCA, designed to create a robust ecosystem of components, products and systems for general purpose instrumentation and semiconductor test.

The hardware standard provides a larger footprint than both cPCI and VME, and as such can provide greater performance density per inch of 19" rack space than those standards. Board footprints provide space for high performance signal conditioning and sensing, with high power/cooling capacity for high-performance analog and digital subsystems. Multi-module synchronization is possible using the AXIe trigger buses, star trigger, sync, and clock resources. AXIe also has a local bus between adjacent modules, using LVDS pairs, for high speed module-to-module communications.

CONCLUSIONS

This article describes a number of considerations and techniques needed to optimize high speed data acquisition in real-time applications. Application of these to beam instrumentation enables continuous streams of accurate data in raw data flow or high trigger rate applications. Finding the right combination that meets the requirements of the system in question will always require specific knowledge—not only of the application, but also of the digitizing capabilities of the hardware.

REFERENCES

- R. Soden, "Interpreting Quoted Specifications when Selecting Digitizers", Agilent Measurement Journal 2004, No. 4,
- [2] D. J. Peake; M. J. Boland, R. Dowd, G. S. LeBlanc, M. J. Spencer, Y.-R. E. Tan, R. P. Rassool, "Measurement of the Real Time Fill-Pattern at The Australian Synchrotron", Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, Volume 589, Issue 2, 1 May 2008, pages 143-149.