Available online at www.sciencedirect.com
ScienceDirect
Nuclear Instruments and Methods in Physics Research A 568 (2006) 350–358
NUCLEAR INSTRUMENTS & METHODS IN PHYSICS RESEARCH Section A
www.elsevier.com/locate/nima

Current trends in developing digital signal processing electronics for semiconductor detectors

W.K. Warburton*, P.M. Grudberg

XIA LLC, 31057 Genstar Road, Hayward, CA 94544, USA

Available online 4 August 2006

Abstract

A major recent trend in semiconductor (SC) detectors has been to combine very good energy resolution with high counting rates or position sensitivity, which results in arrays of many detectors, detectors with many electrodes, or both. This produces a requirement for large numbers of high performance but relatively inexpensive processing channels. Digital processing electronics are particularly effective in these circumstances because of their high performance, flexibility, and ease of interfacing to a computer control system. Driven by consumer demand for low power, high-performance communication devices, the electronics industry has developed an increasingly powerful set of digital components, including ADCs, FPGAs, and DSPs that can be applied to meet this requirement. This paper reviews these trends and also briefly examines how increased digital processing power can be applied both to extract additional information from the detectors and also to carry out in real time what were previously post-processing operations.

PACS: 29.30.Kv; 29.40.Wk

Keywords: Spectrometer electronics; Digital signal processing

1. Introduction

In 1996, when XIA first commercially introduced the DXP-4C CAMAC module that offered digital signal processing electronics for multichannel X-ray detector arrays, the kinds of semiconductor (SC) detectors that were available had not changed very much in 20 years and primarily consisted of small planar HPGe or Si(Li) detectors for X-ray or low-energy γ work and larger co-axial HPGe detectors for γ-rays. The major recent development was the introduction of the planar detector array to meet the high-speed counting requirements of synchrotron X-ray fluorescence (e.g., EXAFS) experiments. XIA's DXP-4C was developed specifically to support these kinds of detectors and experiments.

In the last 10 years, however, there has been enormous progress in detector design and fabrication, due both to an improved understanding of induced charge effects and increased access to lithographic production methods. Resultant designs range from large segmented co-axial detectors that can locate γ-ray absorption events with a few millimeter precision to Peltier-cooled silicon drift detectors (SDDs) that match the energy resolution of LN_2-cooled detectors at 10 times the count rate. Further, SDD arrays are now appearing that may have hundreds of elements, either operating independently or as parallel strips in two-dimensional (2D) imaging detectors with energy resolution.

Detectors require processing electronics to be useful and these may be either integrated circuits (ICs) or discrete components (DCs). The ICs offer higher densities, low production cost and low power per channel, but have restrictions on their performance and flexibility and typically have large development costs. DCs, on the other hand, provide the highest performance and flexibility, but are physically much larger and also have much higher production costs. While there has been some convergence between the two approaches in recent years, as the power of both custom-designed and commercially supplied ICs

*Corresponding author. Tel.: +1 510 401 5760 × 12; fax: +1 510 401 5761.
E-mail address: bill@xia.com (W.K. Warburton).

0168-9002/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.nima.2006.07.021

has increased, there are still several detector classes where the DC solution remains superior. These classes are primarily those that require either uncompromising energy resolution or complex signal processing (typically of induced charge signals) or both. They include arrays of small, independent detectors, crossed-strip Si(Li) or HPGe detectors, and large, co-axial HPGe detectors with many electrodes. Their principle requirement has been for increasing numbers of high-speed processing channels at a lower price per channel. An important secondary requirement has been for more front end signal processing "intelligence" to maximize real-time data reduction and minimize data storage. With modern bus and backplane speeds approaching 100 MB/s, it has become easy to record 10 TB of data a day if one is not careful.

In this paper we will first examine some of these detectors' needs (Section 2), review developments in digital-processing ICs (Section 3), and then describe the instruments and a few example applications (Section 4).

2. Trends in energy resolved SC detectors

2.1. Spectroscopy detector arrays

Higher counting rates are the first force pushing for arrays of multiple spectroscopy detectors. Modern synchrotron beamlines can easily deliver 10^{13} or more X-rays/s to a sample, which can then easily emit 10^8–10^{10} fluorescent X-rays/s into 2π. Because a single detector can handle fewer than 10^6 X-rays/s with good energy resolution, researchers are now pressing for arrays, both to use the X-rays more efficiently and also to reduce experiment times. While the early standard Canberra 13 element HPGe array is still a workhorse in many laboratories, 30 element arrays have become common and a few 100 element arrays exist. There are occasional rumors of people considering 1000 element arrays.

A desire for larger solid acceptance angle is the second force pushing toward arrays, particularly when count rates are limited. This commonly occurs on astronomical satellites missions, where SDDs are particularly advantageous since they do not require LN_2 cooling and can be arrayed lithographically. The improved count rate capability of an SDD array would also be very useful in synchrotron applications below 15 keV, where absorption efficiency is good.

2.2. Crossed-strip detectors

This class of detectors is being developed in two modes. In the simpler mode, the crossed strips simply define pixels, so all the electronics need to do is look for top–bottom strip coincidences, measure their energies and record positions. In the more complex mode, the induced charge signals on that appear on "spectator" strips neighboring the charge-collecting strips are analyzed to obtain sub-pixel resolution. In ongoing work at XIA, for example, the difference between the spectators' maxima will be processed in real time to create 100 Hz frame-rate images for small animal cardiac studies. Our resolution goal is 200 μm from 2 mm strips at 10^6 count/s, using a 1 cm thick HPGe detector. Others are developing similar designs for Compton camera applications.

2.3. Arrays of large co-axial detectors with multiple contacts

Several large nuclear physics experiments are developing detector systems that will array large numbers of segmented HPGe co-axial detectors in order to cover a large fraction of 4π solid angle. The large detectors are required to obtain useful stopping power in the MeV energy range, but better spatial location is required to refine the photons' momentum values. These arrays may use hundreds of individual detectors and have thousands of electrodes as a result. Typically, in addition to energies, it is also important to measure photon arrival times, using constant fraction discriminators, and to capture spectator signal traces, where the traces from different electrodes must be precisely aligned in time at the nanosecond level. [1] In a first step toward minimizing data storage requirements, the electronics are required to identify charge-collecting electrodes and only save spectator signals from their local spectator electrodes. The detector developers' goal is to eliminate trace capture entirely by moving online signal risetime and spectator shape analysis directly into the data collection electronics as soon as appropriately accurate algorithms can be developed.

2.4. The bottom line

The upshot of these detector development trends is a need both for growing numbers of processing channels and an increase in the sophistication of the analyses they can perform. Naturally, since large numbers of processing channels are desired, the price per channel should be reduced as well.

3. Trends in digital processing components

3.1. Processing issues

Before reviewing advances in digital components, we first take a brief look at the reasons that we still use discrete processing components at all, rather than developing custom ASICs, for example. The major issues are cost and performance. ASIC solutions are typically expensive to develop and generally sacrifice some measure of performance in order to meet specific engineering requirements, such as physical size or allowed power. As a result, they win out when there are extremely large numbers of detectors, so that the ASIC development cost per detector is small, or when the engineering constraints are more important than the final degree of performance. As the costs of developing discrete, digitally based solutions are

much smaller, they can be readily designed or modified for specific applications. It is also important to remember that typical digital components have been exquisitely engineered by companies with large development budgets to deliver the highest possible performance at the lowest cost and smallest volume using the latest technology. Except in rare cases (e.g., rad-hard components) it is not possible to match them with home-grown designs. Because each component is individually packaged, there is a space penalty to be paid in using these components, but in most applications this is not a significant issue. Moreover, package sizes are also falling rapidly.

3.2. The important digital components

The four digital-processing components that are critically available for digital-processing designs are memory, analog-to-digital converters (ADCs), field programmable gate arrays (FPGA), and digital signal processors (DSPs). Anyone who pays much attention to computers is probably aware that the size of memory chips has increased as predicted by Moore's law over the last decade, while the price per part has actually fallen. Similar advances have occurred with ADCs, FPGAs, and DSPs, but the details are much less well known. Here we summarize some of the key points in Tables 1–3.

3.3. ADCs

In 1992, when we started designing our first digital X-ray processor, our only serious choice was a 10 bit, 20 MHz part that cost about $40. 12 bit, 20 MHz parts were only available as hybrids and cost $200–400, which made them too expensive for a high-speed X-ray processor and forced us to invent a ramp subtraction method to handle the 5 V preamp output ranges of the period. This was the part used in our first commercial product, the 4-channel DXP-4C CAMAC module. Shortly thereafter, in 1995, a 12 bit part with the same price and speed became available and was the part we used in our first standalone single-channel product, the DXP-X10P.

Table 1 summarizes ADC development progress. There are several important trends. First, about 2 bits in resolution get added every 5 years. Second, output speed at a given resolution doubles about every 2–3 years. Price and power draw at constant performance halve every 2–3 years. Some of these trends may not continue much longer, particularly the push to higher speeds and number of bits, since the available parts are now starting to be completely satisfactory, from a specification point of view for the majority of needs. The trends that will continue are the drop in package size and price per part since these lead directly to ever smaller and cheaper commercial products. In particular, Table 1, with its emphasis on high-performance parts does not show the explosion of small, low-power ADCs (e.g., octal 12 bit, 2 MHz parts that use 10–20 mW/ADC and cost less than $10). The table does indicate the trend to multiple ADCs on a single die, even at high performance levels. This is being hastened by the transition from parallel to serial output, even at high digitization rates, which leads to much smaller package sizes. The quad 12 bit 65 MHz part, for example, comes in a 7 mm × 7 mm package with differential inputs and outputs. The octal version is a 14 mm × 14 mm package.

3.4. FPGAs and DSPs

FPGAs and DSPs have followed the same trends as ADCs, as may be seen by reference to Tables 2 and 3. For example, XIAs first commercial digital X-ray processor,

Table 1
ADC property changes over 25 years

Year	N_c	B (bit)	R (MHz)	$N \times B \times R$	$
1979	D-1	10	20	200	3500
1986	D-1	12	20	240	~2500
1988	H-1	12	10	120	~800
1990	M-1	10	75	750	~450
1992	M-1	10	20	200	40*
	M-2	10	10	200	80
	H-1	12	20	240	~200
1995	M-1	12	40	480	40*
1997	M-1	8	200	1600	~100
2000	M-1	14	105	1470	~200
2002	M-1	12	40	480	40*
	M-1	12	210	2520	~150
2005	M-1	12	210	2520	100
	M-1	14	75	1050	60*
	M-1	14	125	1470	100
	M-1	16	100	1600	100
	M-4	12	65	3120	40
	M-8	12	50	4800	60*

D = discrete, H = hybrid, M = monolithic; * identifies XIA used part.

Table 2
FPGA property changes over the last decade

Year	Xilinx series	Clock (MHz)	Gates (K)	RAM (kbit)	XIA product
1992	XC4000	20	8	7	DXP-4C
1995	Spartan	20	20	15	DXP-X10P A
1998	Spartan XL	40	40	30	Saturn
2002	Spartan 2	100	200	125	Pixie-4
2003	Spartan 2e	100+	400	300	xMAP
2004	Virtex 2	100+	1000	900	Pixie-16 A
2005	Spartan 3	100+	1500	800	Pixie-16 B

Table 3
DSPs used in XIA products

Year	Part	Bits	Clock (MHz)	RAM (kbit)	XIA product
1992	NEC	16	20	4	DXP-4C
1995	ADI	16	40	80	DXP-2X Saturn, etc.
2003	ADI	16	160	160	XMAP
2004	ADI	32/64	100	4000	Pixie-16A

the DXP-4C CAMAC module used a single Xilinx XC4000 FPGA per channel, a device with 8 K gates. While processing included a fast timing channel, slow energy filter, pileup inspection, and baseline correction, it all had to be carefully floorplanned and laid out schematically using custom accumulator designs in order to all fit in. The digital filters could only have sum lengths of 32, so that five designs were required to handle peaking times from 0.25 to 64 μs. By 2002, when Spartan 2's became available (XIA usually works about 1–2 years behind the newest part introductions where the performance per dollar value peaks) we were able to implement the digital filters as 128-sample, double-precision designs in order to achieve 0.1% energy resolution at 1.3 MeV and also fit in an 8000-sample FIFO, eliminating an extra part for trace capture. Currently, in 2005, we are fitting four such channels into a single Spartan 3 in the most recent Pixie-16 design.

XIA digital X-ray processor designs have never made serious demands on DSPs. The bulk of our processing is carried out in the FPGAs, with the DSPs doing final value corrections at the output data rate and carrying out a variety of housekeeping chores such as computing and maintaining baseline values and histories. Table 3 shows only the parts that we have actually used. For example, although they were available much earlier, we did not put an Analog Devices (ADI) SHARC processor into a design until 2004 when its ability to read 32 bit words became critical to maintaining throughput. There are currently available, and always have been, far faster parts. However, they are also much more expensive, a critical concern when total instrument cost has to be minimized.

Figs. 1 and 2 encapsulate this progress. Fig. 1 shows the 10 bit, 20 MHz ADC and 8 K gate FPGA from XIA's first card, the DXP-4C. Fig. 2 shows the four 12 bit, 100 MHz ADCs and the single Virtex 2 FPGA that processes them

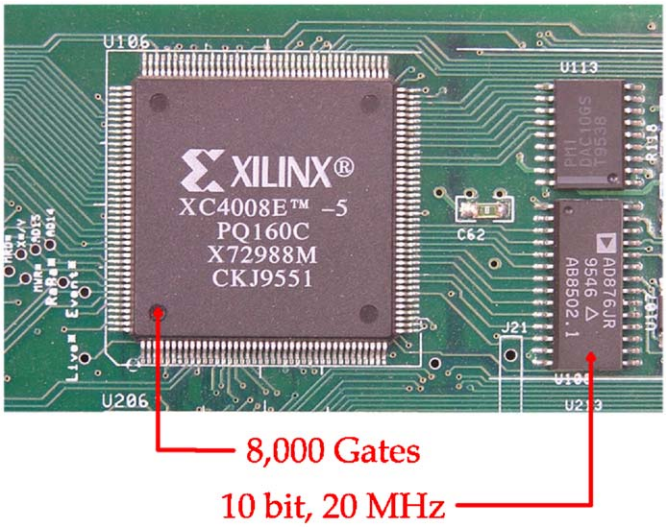

Fig. 1. ADC and FPGA from XIA's first commercial card, the DXP-4C, which was 10 bit at 20 MHz.

that are installed on our latest card, the Pixie-16. The two images are at the same physical scale, showing that we are getting over 20 times the performance in the same physical space.

4. Specific instrument capabilities

4.1. Standard processing capabilities

All of XIA's standard processors share a common basic topology—at the input, a DC-coupled analog section matches the input preamplifier signal range to the ADC's input range and applies Nyquist filtering to limit the bandwidth to less than half the ADC's sampling frequency. After digitization, all other steps are carried out digitally to

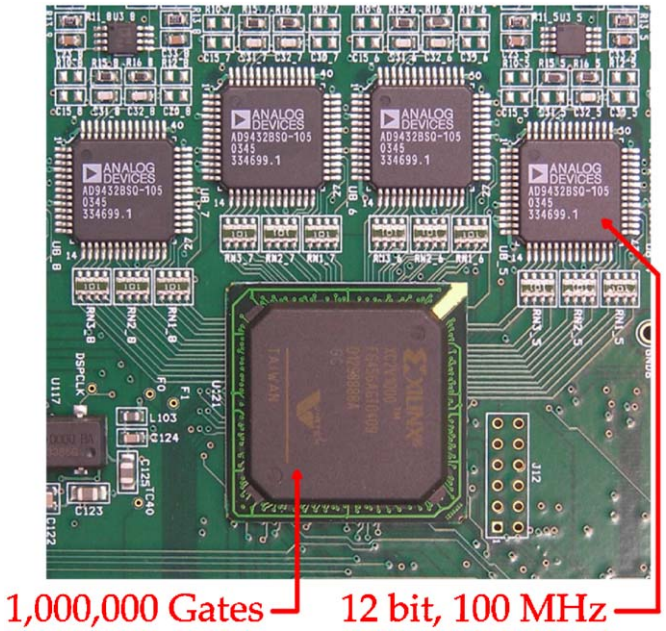

Fig. 2. Shared four ADCs and one FPGA from XIA's most recent card, the Pixie-16, which is 12 bit at 100 MHz.

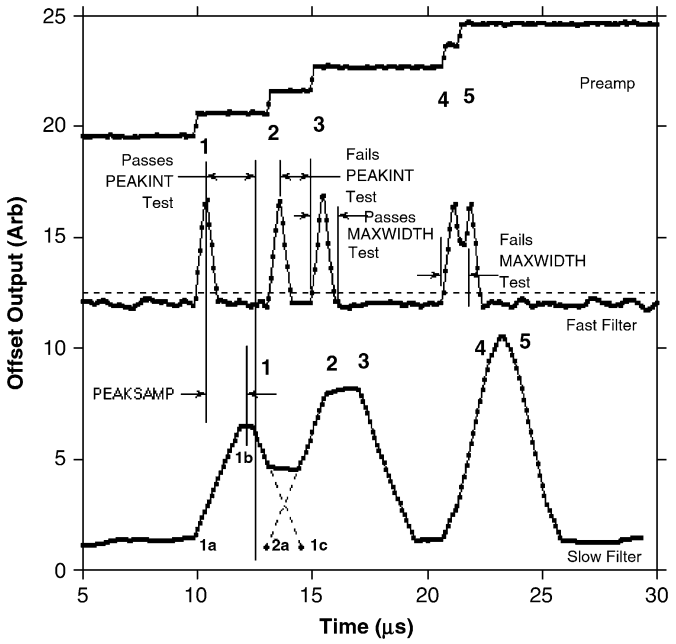

Fig. 3. Sketch indicating methods of XIA pileup inspection. Pulses are detected when the fast filter crosses threshold and their energies sampled PEAKSAMP clock cycles later. Pulses that are too close together "pile up" and produce distorted outputs from the slow energy filter. PEAKINT and MAXWIDTH test for pileup conditions in the slow and fast channels, respectively.

produce at least the following functions: fast-channel pulse detection, slow-channel energy filter, pileup inspection, baseline capture, average and subtraction, and deadtime correction. The fast channel is typically a short time constant triangular filter with a 50–200 ns peaking time. Pileup inspection is carried out as indicated by Fig. 3. Times between sequentially detected pulses are measured and compared to the value PEAKINT to inspect for slow energy filter pileup. The widths of pulses detected in the fast channel are also measured and compared to MAXWIDTH to inspect for fast-channel pileup, with the result that deadtime corrections are typically accurate to 0.5% or better, even at very high counting rates.

The energy channel measures pulse amplitudes using quasi-trapezoidal filters of the form

$$E = A_0 \sum_{i=-2L-G}^{-L-G-1} y_i + A_g \sum_{i=-L-G}^{-L-1} y_i + A_1 \sum_{i=-L}^{0} y_i \quad (1)$$

where y_i is the ADC sample at time i, and the lengths L and G determine the peaking time $\tau_p = L\tau_s$ and flat-top or "gap" time $\tau_g = G\tau_s$, where τ_s is the ADC's sampling time. Peaking times can range from 100 ns to 100 μs, though the range implemented in a particular firmware design typically only spans a multiplicative range of 5–10. For reset preamplifiers $A_0 = A_1$ and $A_g = 0$, to give a purely trapezoidal filter, while for resistive feedback preamplifiers A_0 is adjusted to correct for the RC decay constant and A_g is adjusted to compensate for the effects of ballistic deficit. This class of filter is particularly effective at higher counting rates where one operates in the regime where energy resolution is proportional to $1/\tau_p^{0.5}$.

Baseline corrections have been found to be critical for avoiding peak position drift and loss of energy resolution with increasing count rates. XIA uses a novel quarter length baseline filter that continues to be able to capture baseline samples even when the energy filter is heavily piled up (see Fig. 4). By making the baseline filters values L_b and

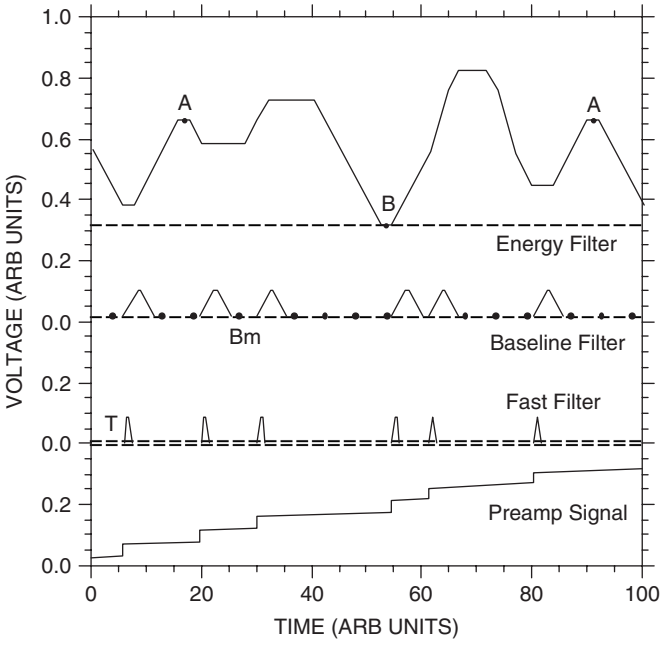

Fig. 4. Outputs of the fast timing, slow energy, and middle baseline filters for pulses with different time separations under high data rate conditions.

G_b exactly 1/4th those of the energy filter's L and G, the baseline output values of the baseline filter are exactly 1/16th those of the energy filter and trivially scaled to correct the energy filter's outputs. Our spectrometers in which this is implemented are thus able to show essentially no peak position drift or significant energy resolution degradation even when running at 95% deadtimes, compared to their low rate values. This quarter length filter also makes an effective pulse detection circuit when working with very soft X-rays, since it has much lower noise than the fast filter, yet still is enough faster so that it can inspect for energy-channel pileup. As a result, our designs perform superbly in the soft X-ray regime, as exemplified by the data shown in Fig. 5. Since this result has not been duplicated using analog spectrometers, we speculate that its success is due to the fact that XIA's RC firmware provides ballistic deficit correction that compensates for the different signal risetimes associated with X-rays absorbed in different locations in the SDD.

4.2. Pixie-4: clover detector and PSPMT

The Pixie-4 is a 4-channel PXI card intended for very high energy resolution detectors in coincident γ-ray detection applications. It has 14 bit, 75 MHz ADCs and runs 22 bit energy filters of which the 15 MSBs are read out, allowing 32 K channel MCA spectra. Spectrometer noise, linearity, and temperature drift per 10° are all less than one channel in 32 K and pulse pair resolution can be as short as 50 ns. Timing jitter is sufficiently low as to allow time resolutions below 1 ns to be achieved in fast scintillator/PMT measurements using coincident γ-rays from ^{22}Na. The Pixie-4 is intended for small laboratory detector arrays and has the capability to share triggers, clocks and multiplicity data across both channels and modules. Fig. 6 shows a 32 K spectrum from a mixed set of nuclear sources taken using the Pixie-4. The FWHM values of some well-known lines are indicated to show the very high resolution that was attained. In particular, better than 0.1% is achieved at the ^{208}Tl 2616 eV line (see the right-hand inset). The left-hand inset shows how the same detector can also resolve the Pb K shell X-ray fluorescence lines from local shielding and, in addition, that they may be clearly seen in the same spectrum when it has 32 K channels.

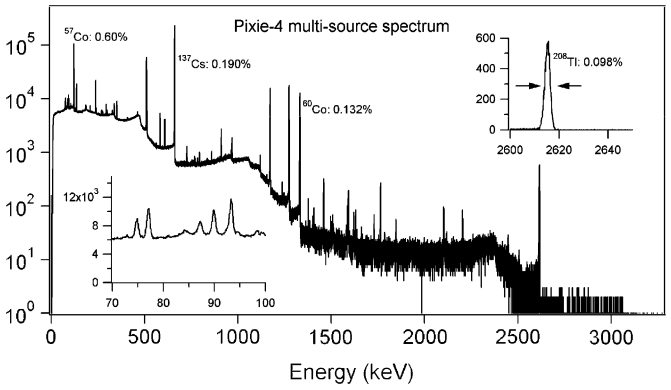

Fig. 6. Spectrum of a mixed set of sources taken using the Pixie-4.

Figs. 7 and 8 show two simple examples of how the Pixie-4's ability to combine spectroscopy with coincidence timing may be employed experimentally. Fig. 7 is a sketch of a clover detector, comprised of four co-axial HPGe detectors, and the four preamplifier output signals. Clovers are used like large single co-axials except that they have better energy resolution in the quadrants, and both better timing capability and lower cost because of the smaller detector size. However, if the electronics are smart and can recognize coincident events between the quadrants, then events such as the one shown that scatter within the compound detector can be identified and their energies recovered in an "add-back" spectrum. In the shown case, the Pixie-4 would note the coincidence between detectors

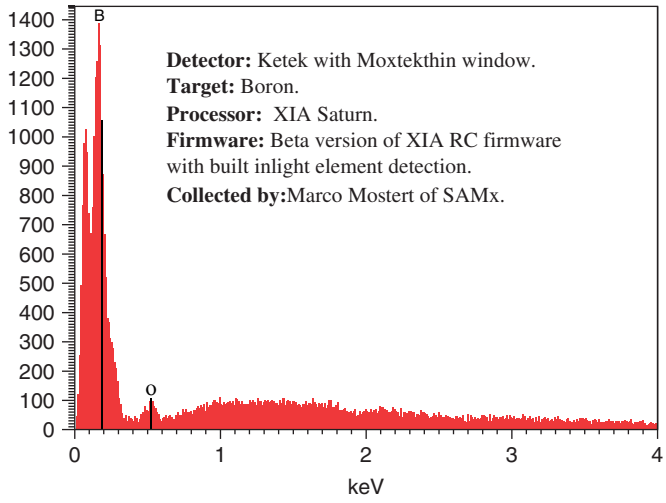

Fig. 5. Soft X-ray spectrum of B target taken using XIA Saturn processor and B target in SEM.

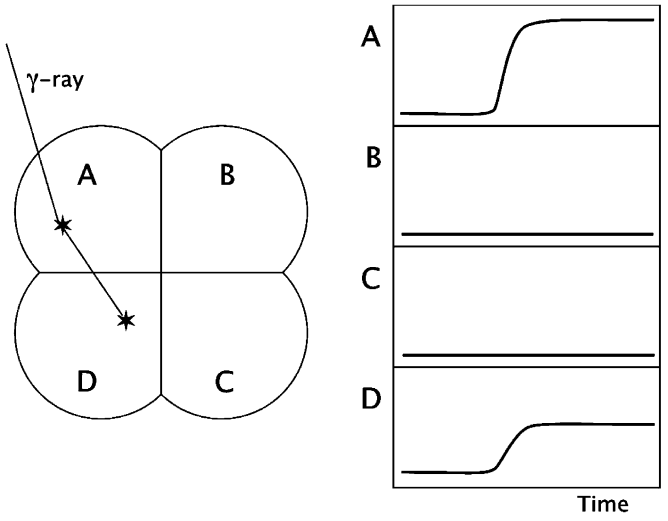

Fig. 7. Clover detector showing coincident detector signals.

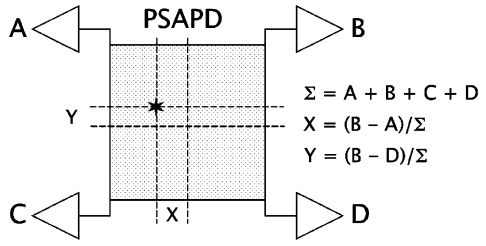

Fig. 8. PSAPD showing signal combinations required to recover incident light location.

A and D and add the separately found energies E_A and E_D to recover the energy of the initial γ-ray. Experimentally, we find that this method has better energy resolution than performing the addition off-line, better peak to Compton values and much higher full energy efficiency. Only a single Pixie-4, is required for implementation, replacing the traditionally required collection of amplifiers, gates, timing modules and multiple MCAs.

A position-sensitive avalanche photodiode (PSAPD) or position-sensitive PMT (PSPMT) can be handled similarly using only a single Pixie-4. As sketched in Fig. 8, the four outputs share the amplified charge produced by an absorption event. If these are integrated by charge-sensitive preamplifiers, then the event location can be estimated by the shown equations. In this case the Pixie-4 notes that an event has occurred, captures the four electrode energies and then performs the indicated computations. Because of the geometry, the resultant image typically has very bad barrel distortion. Using a lookup table, this can also be corrected in real time on an event-by-event basis. Or, if preferred, the Pixie-4 can build an image in its on-board memory and then perform the distortion correction only once, on a per pixel basis, at the end of the data collection period.

4.3. xMAP: high-speed mapping

Fig. 9 shows a picture of the xMAP four-channel X-ray processor developed particularly for high-speed scanning X-ray mapping applications at synchrotron research facilities. In these experiments, the focused X-ray beam is held fixed and the sample is scanned in front of it in a 2D raster pattern with a fixed dwell time per pixel. Fluorescent X-rays emitted from the sample are collected by an N element array detector (per section 2.1), with each element being processed by a dedicated xMAP channel that produces a real-time energy spectrum on a pixel-by-pixel basis. The xMAP has 0.1–100 μs peaking times with full dead time and ballistic deficit correction up to 1 Mcps per channel. Because the xMAP's PXI interface can transfer data at 100 MB/s, the data readout time for a 1000-channel spectrum (32 bit deep) is only 40 μs. Thus, when $N \times 40\,\mu s$ is less than the dwell time per pixel (e.g., 520 μs for a 13-element detector), the spectra can simply be read out on a pixel-by-pixel basis. For shorter dwell times, up to 32 regions of interest (ROIs) can be assigned in the energy spectra and stored in the xMAP's on-board 4 MB memory, which is then read out once per row of pixels in a "ping-pong" mode that allows one row to be read out while the next is being collected.

Fig. 9. xMAP high-speed digital X-ray spectrometer.

4.4. Pixie-16: 2D crossed-strip detector and large segmented HPGe detector

The most recent and most powerful XIA processing card, shown in Fig. 10, is the Pixie-16, 16-channel coincidence spectrometer that was initially designed to

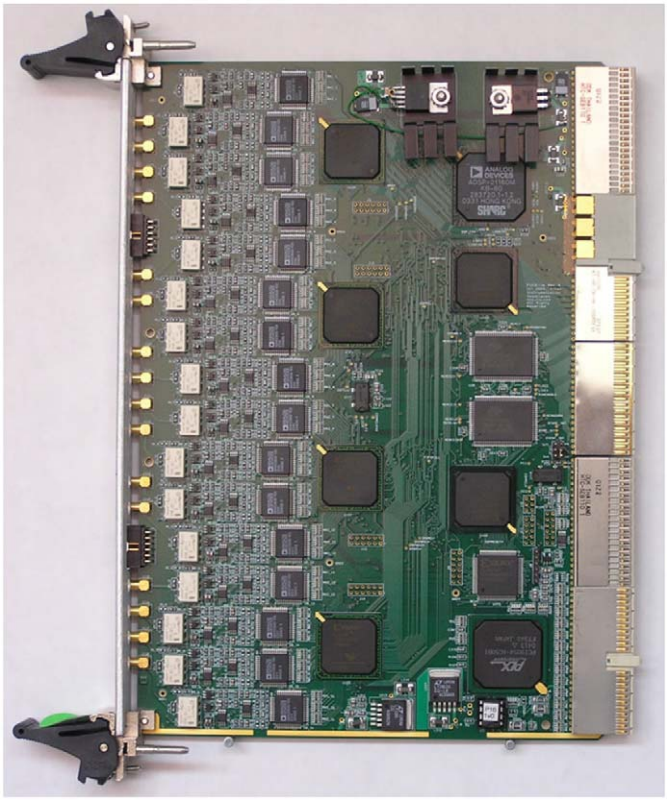

Fig. 10. Pixie-16 module, in 6U format.

implement a 128×128 crossed-strip Si detector for a radioactive ion implantation experiment that required antipileup inspection to detect the small fraction of implantations that produced proton decays. This board is now in Revision B, which added inter-card communication features that will allow it to carry out event descriptor building. By taking advantage of shared back-panel trigger and communications lines, the module can be used to instrument quite large array detectors. Here we describe two recent applications.

4.4.1. Crossed-strip interpolating image detector

This detector, shown schematically in Fig. 11, has 10×10 strip electrodes at 2.5 mm × 25 mm and is being developed at XIA to use spectator signal processing (per section 2.2) to achieve 10^6 cps imaging at 35 keV with 200 μm spatial resolution. The strip signals will be processed two ways—first, directly by a Pixie-16 channel to detect pulses and measure energy and time of arrival. Coincidences detected between top and bottom strips will roughly identify points of interaction. Second, analog differences between pairs of spectator signals will be applied to a resettable peak capture and hold. This circuit will output a step pulse of about 500 ns duration each time it sees transients of the indicated shape. When energy filtered by the Pixie-16, the recovered "energy" will then be the difference $a-b$ of the transients' pulse heights. The strip signals will be connected in order to the Pixie-16's even channels ($2N$) and their associated spectator difference signals to neighboring odd channels ($2N+1$). Thus, when an even channel detects an energy event it can trigger its neighbor to capture the associated value of $a-b$ as well and then form the ratio $(a-b)/E$ using a lookup table of $1/E$ for the division. The resultant pair of values (E, δx) will be given an event number and then passed to a "Director" Pixie-16 module that matches it with its coincident (E, δy) pair and then places a count in the associated (E, δx, δy) pixel in the developing image. With the exception of the analog sample and hold conditioning circuits, all the data processing required to go from pulses to image can be carried out by the Pixie-16's.

The second application is to large arrays of segmented HPGe co-axial detectors (Fig. 12). This application will be slightly more sophisticated than the crossed-strip detector

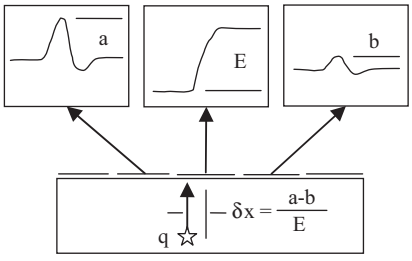

Fig. 11. Sketch of crossed-strip detector showing charge collecting strip and its two spectator strip signals.

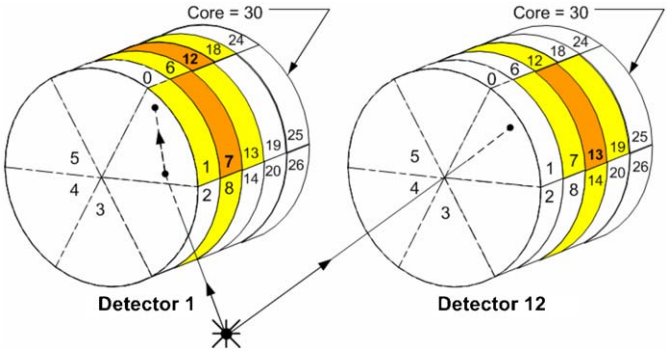

Fig. 12. The pair in an array of HPGe segmented detectors that absorb two γ-rays in coincidence from a nuclear decay.

system above. Here there will be as many as 18 detectors, where each detector has 32 surface electrodes and 1 core electrode. The core electrodes will be attached to a neighboring pair of Pixie-16's and used to generate system level triggers whenever two or more γ-rays are absorbed (multiplicity greater than 1). For valid events, the desired data set should contain: (a) all the core energies and pulse shapes; (b) all the energy values and pulse shapes for charge-collecting surface electrodes; and (c) pulse shapes for the nearest-neighbors (spectator electrodes) to all electrodes in (b). Notice that, if there is Compton scattering within the detector, it may have more than one surface electrode that collects charge, which is a critical feature of the experiment.

In previous designs, whenever a core electrode triggered a discriminator, traces were captured from all its surface electrodes and the energy measurements and pulse shape analyses required to identify charge-collecting electrodes were carried out off line. However, in a large array, this requires excessive storage as typically only about 6 of the 32 electrodes actually carry useful information. Using the Pixie-16's, the charge-collecting electrodes will be identified on line and then a lookup table used to identify their spectators and tag them for data capture, where the spectator electrodes may connect to a different Pixie-16 than the one making the identification. While this is a sophisticated operation, it is well within the module's capabilities and reduces data storage and transfer times by a factor of about 5.

In addition, rather than just time stamping all data and leaving it to the off-line computer to sift all the data from 38 Pixie-16 modules to build the event, the Pixie-16's can actually build an event descriptor to bypass this step. Each module detecting a charge-collecting electrode can record its label and those of its spectators and then pass these values together with an assigned event number to a central "Director" Pixie-16. The Director receives these label sets from both the modules processing the core electrode signals and from the modules processing the surface electrode signals and can then generate a data packet completely describing the event. The data-processing computer can then read the descriptors for a set of events and assign

space in a data table for them so that when the data are finally read from the modules they can be placed into preassigned storage locations. This means both that the data can now be read from the modules in any order and also that the data readout can be made completely independent of the data capture, which will increase throughput by a factor of 2 or more. Our initial estimate is that, with an average multiplicity of 4 (with 8 signal captures 1.2 μs long at 10 ns/sample per γ-ray) that the data collection system should be able to operate at between 5000 and 10,000 events/s, a factor of about 10 faster than before.

Finally, even though the Pixie-16 is a 12 bit design, we find that its energy resolution is nearly identical to that of the optimized Pixie-4.

5. Conclusions

Digital signal processing still has advantages in situations where complex information has to be processed with a high degree of accuracy. In this paper we have shown how progress is digital components, driven by the communications industry, has made it possible to achieve ever more complex processing operations with higher accuracy and a decreasing cost per channel.

Acknowledgments

We acknowledge with pleasure the many fruitful discussions over the years with other XIA staff members that have lead to the development of these ideas and their translation into working devices.

References

[1] The Gamma-ray Tracking Coordinating Committee, A National Plan for Development of Gamma-ray Tracking Detectors in Nuclear Science, 2002 ⟨http://www.pas.rochester.edu/~cline/grtcc/GRTCC-report.pdf⟩.