From Analog to Digital DAQ Transition in Physics Applications

4th School for Particle Detectors and Applications at KNU

(SPDAK2024)

CAEN
Tools for Discovery

Electronic Instrumentation

Slide 2

The goal of this class is to explain the different types of readout chains for the detectors used in nuclear physics applications. The presentation begins with a brief description of the fundamental concepts related to detector readout and then focuses on the comparison between the traditional readout chain, mainly based on a combination of different analog modules, and the digital chain, based on a single multi-purpose module called “Waveform Digitizer”.

Slide 4

Let’s start from a very basic definition and specifically we do define Spectroscopy the study of the interaction between matter and radiation with the aim to get information about the energy distribution of the source. When we talk about radiation we do refer to charged (α, β, light nuclei) or neutral particles (photons – X and γ in our case – and neutrons).

Slide 5

There is a multitude of detector types for nuclear physics, each with specific characteristics that make them suitable for different spectroscopy applications. They can be broadly divided into three main categories:

1. High-resolution energy spectroscopy detectors: typically, these are semiconductor detectors such as silicon or high purity germanium detectors (HPGe) that produce weak charge signals when traversed by a particle. With this type of detector, a low-noise preamplification stage is always necessary, usually a Charge Sensitive Preamplifier (CSP), often integrated directly inside the detector to maximize the signal-to-noise ratio.
2. Medium-resolution detectors: this category includes the most common scintillators. In this case, the passage of a particle through the crystal produces a light pulse that is then converted into an electrical charge using a photo-converter, such as a Photomultiplier Tube (PMT) or a Silicon Photomultiplier (SiPM). In certain cases, the converter already has sufficient gain to produce an electrical signal of hundreds of millivolts, so a preamplifier may not be necessary.
3. Low-resolution energy detectors: once again, these are scintillators coupled with a PMT or SiPM. Due to their limited energy resolution, they are often used for counting, active shields, anti-Compton shields, etc.

The electronics used to read them often consist of simple discriminators, counters, and, in some cases, time-to-digital converters.

Slide 6

When a particle passes through and interacts with the detector, it produces a charge pulse, whose amplitude, duration, and shape depends on both the detector characteristics and the type of incident particle. Generally, the charge pulse is very weak and requires a low-noise, highly sensitive preamplifier, which is often placed very close to the detector or even inside it to minimize noise effects. There are various types of preamplifiers, but they can be broadly grouped into two categories:

1. Charge-sensitive preamplifiers: primarily used for high-resolution energy spectroscopy applications. Their output is typically slow, on the order of microseconds, and does not faithfully reflect the input signal's shape generated by the detector.
2. Fast (or current) preamplifiers: in this case, the preamplifier has a wide bandwidth and generally does not alter the signal shape. It produces a very fast output signal, on the order of nanoseconds, and is typically used in applications where timing information is crucial.

In certain cases, the detector itself has sufficient gain to generate signals in the tens or hundreds of millivolts range, eliminating the need for a separate preamplifier.

This is, for example, the case with some PMTs. Since these detectors are fast, their signals can be considered similar to those produced by fast preamplifiers.

The amplified signal is then read by the readout chain, which aims to identify the pulses and acquire the necessary information for the application. Typically, the signal amplitude (in the Pulse Height Analysis, PHA) or the pulse area (in the charge-to-digital conversion, QDC) is acquired to obtain energy information. In most cases, timing information is also acquired (in the form of time of arrival, rise time, time over threshold, etc) using a TDC . In other cases, information about the signal shape is needed to identify the type of particle or interaction.

All this information is temporarily stored in dedicated memory buffers within the acquisition chain, waiting to be transferred to the computers running the Data Acquisition and Analysis software (DAQ), which ultimately produces the desired results for the application (spectra, lists, statistics and output files).

Slide 7

After the data has been acquired, it undergoes further processing to extract significant physical information, including pulse shape, event correlation, and interaction position. This analytical approach is widely known as multiparametric analysis, which yields physical quantities that characterize the interaction between matter and radiation.

Slide 9

The waveform digitizer samples the signal and produces sequences of waveforms, similar to any digital oscilloscope, and saves them into a memory buffer.

Slide 10

The difference from commercial oscilloscopes, is that the waveform digitizer designed for nuclear physics applications implements a multitude of online data processing algorithms, that process the waveforms acquired by the analog-to-digital converters (ADCs), and extract the required synthetic information, such as energy spectra, time distribution, signal waveform variation and more.

Slide 11

In the traditional acquisition chain, the detector signal, after being amplified by a preamplifier and brought to an amplitude level typically ranging from tens of millivolts to several volts, enters a series of mainly analog modules that condition the signal. These modules typically include shaping amplifiers, peak stretchers, charge integrators, discriminators, time-to-amplitude converters, etc.

The output is usually either a quasi-static analog signal that can be converted by a slow Analog-to-Digital Converter (ADC) or a logical signal that can be counted by a scaler or measured by a TDC. In this scheme, a significant portion of the acquisition chain relies on analog signal processing, with digital conversion taking place only at the end of the chain.

As a result, there are multiple interconnected modules, each serving a specific function. Consequently, the system lacks flexibility, as one or more modules need to be changed when the required measurement type changes. Additionally, a setup constructed in this manner is often complex, requiring intricate wiring and manual settings that are difficult to manage and, most importantly, replicate.

Slide 12

The "fully digital" acquisition chain takes an almost specular approach compared to the analog chain. While a preamplifier is still necessary to bring the weak detector signal to an appropriate amplitude level, the conversion from analog to digital is performed right after the preamplifier, without adding any analog signal processing, thus preserving the complete signal information. Sampling is carried out using a fast flash ADC, typically with a sampling frequency starting from 100 Mega samples per second and exceeding 1 Giga sample per second. The choice of sampling frequency depends on the signal speed, i.e. the detector response, as well as the type of preamplifier and the required measurement. The number of bits in the ADC also varies depending on the application, typically ranging from a minimum of 10 to a maximum of 16. A good compromise covering a wide range of nuclear physics applications is a 500 Mega sample per second digitizer with 14-bit resolution.

The analog signal is continuously sampled and converted into a digital format without interruptions. The output of the ADC is a stream of data that is typically read by an FPGA. This continuous data stream is digitally processed using dedicated algorithms implemented in the FPGA firmware, aiming to extract the required information and measurements for the application. Therefore, we can say that the digital acquisition chain is capable of providing the same information as the analog chain, such as Pulse Height, Charge, Time Stamp, Pulse Shape, counting, etc., but in a much more flexible, efficient, and compact manner. Unlike the analog chain, which consists of multiple separate modules, the digital chain has a single acquisition module, namely the waveform digitizer, within which various processing blocks are implemented, each providing specific information. The number of wiring connections is greatly reduced, as well as the physical footprint. Additionally, the functional blocks can be modified at any time through firmware upgrades. Parameter settings are done through registers, making them fully controlled by software and reproducible at any time.

Slide 13

As mentioned, the flash ADC continuously samples the signal and outputs a high-rate data stream. For example, a 16-bit ADC with a sampling rate of 500 Mega samples per second has an output data throughput of 1 Giga byte per second.

There are very few applications where it is necessary to acquire and save such a massive amount of data to disk. In most cases, especially when dealing with multiple acquisition channels, a triggering system needs to be implemented to define regions of interest and acquire data only within a certain time window.

Slide 14

There are two different triggering scenarios:

1. Common Trigger: in this scenario, there is a trigger that is common to all channels (often from an external source) and defines an acquisition window of a certain length. All channels receive the common trigger and save a certain number of samples around this trigger in a local memory buffer. Typically, in this mode, called Oscilloscope Mode, the acquisition is limited to saving portions of waveforms, so sequences of raw, unprocessed samples.
   The readout software then reads the waveforms and saves them to disk or computer memory buffers, where they are processed off-line to extract the necessary information for the specific application. The data throughput to the computer is now much lower than the output from the ADCs, but it can still be high, especially if the trigger rate is high or if the acquisition window is wide.
2. Self-Trigger: Each channel independently acquires data, creating its own self-trigger when it detects the presence of an incoming pulse. An algorithm for pulse identification such as baseline restorer, digital discriminator with a threshold, etc, is necessary. Once the pulse is identified, a local trigger (valid only for that channel) is activated, opening a certain acquisition window that can also be

retroactive, thus saving a certain number of samples that arrived before the trigger. In many cases, the required information for physics analysis is not the complete waveform recording but only specific characteristic parameters such as pulse height, charge, time stamp, rise time, etc. Therefore, it is possible to implement algorithms in the FPGA that process the pulse waveform and extract these parameters. This acquisition mode, called Digital Pulse Processing (DPP), allows for list mode acquisition, where only the list of parameters extracted from the pulse processing is saved in the board's memory buffers, without saving the raw waveforms or saving only a part of them. By acquiring in list mode, the data throughput from the board is significantly lower compared to reading the waveforms, making it possible to acquire pulses at very high rates.

Slide 15

Of course, there are also intermediate situations between the oscilloscope mode and the DPP mode.

For example, it is possible to utilize DPP algorithms only for pulse identification and generating a channel trigger, which is then combined with triggers from other channels using appropriate trigger logic (coincidences, multiplicities, etc.) to generate a global trigger that opens the common acquisition window to save waveforms on all channels simultaneously.

Similarly, it is possible to acquire in list mode with DPP, where channels acquire on independent self-triggers, but with a common validation signal for all channels or groups of channels, allowing to save the list data only if it belongs to a certain time interval. This approach enables the implementation of coincidence logic, veto logic, etc.

Slide 16

Let's now examine in more detail the types of DPP algorithms implemented in the FPGA of digitizers, noting the parallels with the traditional analog chain.

Slide 18

The first algorithm we will look at is called DPP-PHA and it is designed to acquire the pulse amplitude from a charge-sensitive preamplifier.

Typically, the PHA algorithm is used with semiconductor detectors such as silicon, CZT or HPGe detectors, but it also works well with signals from scintillators coupled with PMTs, as long as they are not too fast.

In general, we can say that the DPP-PHA algorithm works with exponential signals, typically with a rise time of less than 100 nanoseconds and a decay time usually in the range of microseconds but can go down to a few tens of nanoseconds. Therefore, there is no need to sample the signal at very high frequencies; typically, a sampling rate of 100 Mega samples per second is more than sufficient. It is, however, important to have an ADC with good resolution, at least 14 bits.

Slide 19

The waveform acquired by the flash ADC is processed online continuously and without interruption by the FPGA. There are two branches of processing:

1. The first one is the Fast Timing Filter and Discrimination algorithm. The purpose of this block is to identify pulses and generate a self-trigger. It consists of a differentiator that operates on the voltage steps, producing a bipolar signal, and a small integrator, called smoothing, that reduces the noise level. The output signal from the timing filter is compared to a threshold to arm the trigger, which is then generated at the zero-crossing point.

2. The second branch is the trapezoidal algorithm for energy measurement. This block serves the same purpose as the shaping amplifier and peak sensing in the traditional analog chain. Its aim is to transform the exponential signal into a trapezoidal shape, from which the pulse amplitude is extracted at the flat top.

Slide 20

The duration of the trapezoid can be programmed, similar to the shaping time of the shaping amplifier. A longer duration, that is a greater integration, leads to better energy resolution. However, the trapezoid duration also affects the probability of pile-up between two trapezoids and therefore impacts the dead time of the acquisition.

It's important to note that the dead time is not related to the ADC conversion, which occurs continuously without interruption, but it is related to the processing algorithm, that is the mathematics of the trapezoid formation.

The DPP-PHA algorithm is capable of acquiring the time stamp of the pulses, thanks to the self-trigger generated by the timing filter algorithm, their amplitude through the trapezoidal algorithm, the input and output count rates (ICR and OCR respectively) of pulses before and after pile-up rejection, and if necessary, even the raw waveforms of the sampled signals.

However, the transmission of raw waveforms incurs a significant data throughput, and it is typically used only for specific applications or during system tuning to monitor the behavior of signals, including those generated internally by the digital filters.

Slide 22

The figure shows some examples of trapezoid formation.

The blue trace represents the output of the preamplifier, which is the analog signal sampled by the ADC and processed by the DPP-PHA algorithms. A sequence of stacked pulses, P1 to P7, can be observed.

The red trace shows the bipolar output of the fast timing filter and the self-trigger generated at the zero crossing. As can be seen, the first five pulses have been correctly identified and will be counted in the ICR.

However, pulses P6 and P7 are so close to each other that the timing filter, although fast, is unable to discriminate them. The filter will perceive them as a single pulse (as there is only one zero crossing), resulting in a deficit in the ICR count. This deficit can be statistically recovered in post-processing, considering the known dead time of the fast filter.

The green trace shows the trapezoid formation of the pulses and the corresponding point on the flat top where the signal height is measured, providing the energy information.

As can be observed, the energy can be determined for the first three pulses, and even for P2 and P3, despite the overlapping trapezoids. However, for P4 and P5, the distance between them is too short, causing their trapezoids to overlap in a way that makes it impossible to determine their heights. This condition triggers the pile-up rejector. Consequently, for P4 and P5, we will have time and ICR count information, but no amplitude information. Therefore, they will not be used for energy spectrum and OCR count and will increase the dead time. In the case of P6 and P7, since the fast filter does not discriminate them and perceives them as a single pulse, the energy filter considers them as a single trapezoid with an amplitude equal to the sum of the two overlapping trapezoids. In the energy spectrum, these events contribute to the sum peaks.

The only way to attenuate (but not completely eliminate) this effect is to analyze the shape of the output pulse from the fast filter, as it can be easily recognized as a double pulse, despite having a single zero crossing.

However, this requires a rather complex waveform analysis that typically cannot be done online by the FPGA. That's why there are cases where saving waveforms is important.

Slide 24

The second group of algorithms we are going to analyze is composed by the DPP-QDC and DPP-PSD. They are often embedded in the same FPGA firmware.

The DPP-PSD/QDC firmware is designed to capture charge pulses, typically much faster than those used for the DPP-PHA. In fact, in most cases, these are anode signals from PMTs and do not pass through a charge-sensitive preamplifier. The pulse generally has a duration of a few nanoseconds, sometimes a few tens, but rarely exceeds a hundred nanoseconds.

Slide 25

Energy information is contained in the charge of the pulse, that is, in its area. For this reason, the digital algorithm implements the "gated integrator" function, which sums the samples within a gate that is automatically generated by the algorithm through a leading-edge discriminator or a constant fraction discriminator (CFD).

Slide 26

It is also possible to disable the self-gating function and use an external gate, or program coincidence or anti-coincidence between external and internal gates. Thanks to the FPGA's memory buffer, it's possible to introduce a delay on the signal, allowing integration of the pulse component that arrives before the discriminator triggers. All parameters for integration, such as discriminator threshold, gate duration, and position, are programmable.

Slide 28

In addition to energy measurement, which represents the total charge of the pulse, the algorithm performs a shape analysis, calculating the ratio between the fast and slow components of the signal. For this purpose, the integration algorithm computes a double integral: one over a "long gate" to capture the total charge of the pulse, and another over a "short gate" to capture only the fast component of the signal. Consequently, the slow component (the tail of the signal) is derived through subtraction between the total charge and the fast charge.

Finally, the "PSD" parameter is obtained as the ratio between the charge in the tail and the total charge. The "PSD" parameter provides information about the signal's shape and can be used for discriminating the type of particle that interacted with the detector. For instance, it allows discrimination between gamma rays and neutrons in a liquid organic scintillator.

Slide 29

The digital discriminator (both leading edge and CFD) is used not only to generate the integration gate but also to calculate the timestamp (so the arrival time) of the pulses, thus completing the multiparametric analysis composed of three quantities: energy (so the charge), shape, and time.

To achieve good timing resolution and minimize walk, which refers to the threshold crossing point shifting with signal amplitude, the technique of Constant Fraction Discriminator (CFD) is employed. Its operation in the digital version is much like the analog version: an attenuated copy of the signal is created (according to a certain fraction), along with a delayed copy. The second copy is then subtracted from the first one, resulting in a bipolar-shaped signal with a zero crossing that occurs at a constant time, independent of the signal's

amplitude. In the digital version, the delay is achieved using memory cells, and attenuation is a simple multiplication by a programmable parameter.

To enhance resolution, interpolation between ADC samples is used to pinpoint the zero crossing position with much higher precision than that given by the sampling period. This gives rise to a "coarse time," determined by the sampling clock, and a "fine time" that indicates the exact position within the clock period.

Slide 30

Naturally, the ADC's sampling frequency should be appropriate for the speed of the signal being read. As a general rule, for effective interpolation and resolution, you need at least 3 points on the leading edge, preferably 5. This implies that if you want to read a pulse with a leading edge of 5 nanoseconds, you require a digitizer with a sampling rate of 1 Giga sample per second. However, through suitable calibration techniques, it's possible to optimize the interpolation algorithm even in cases where signals have edges equal to the sampling period or even faster.

For instance, signals produced by a barium fluoride detector (which have edges below 2 nanoseconds), when read by a 14-bit, 500 Megasample per second ADC, with proper calibration done on the signal shape, a better resolution of 100 picoseconds can be achieved.

Slide 32

In this section we will discuss about two advanced zero suppression algorithms, the DPP-ZLE and the DPP-DAW.

In many applications, it is necessary to acquire the "raw waveform" of signals from detectors. In fact, there are cases where synthetic parameters (such as height, charge, time stamp) are not sufficient to retrieve the required information. The raw waveform preserves the complete signal information, making it possible for offline analysis to extract the desired parameters. However, this approach generates a very high volume of data, typically not sustainable by the system's readout bandwidth, thus causing dead-time and data loss.

Therefore, waveform processing algorithms have been developed, no longer focused on parameter extraction as in PHA or PSD algorithms, but rather on identifying regions of interest, allowing for the suppression of unnecessary data.

Slide 33

In the case of the ZLE firmware, the acquisition is "triggered," meaning there is a common trigger shared among all channels of the system, which acquire waveforms simultaneously within the same acquisition window. Typically, not all channels are "fired," so the first level of data suppression involves eliminating channels that were not triggered from the outgoing data stream. A "suppression threshold" is defined, which is entirely independent of any trigger threshold used to generate the global trigger.

Channels that, even though triggered by the global trigger, never exceed the suppression threshold within the window are discarded. An additional step involves searching for regions of interest within the window of triggered channels. Only portions that exceed the suppression threshold, plus a programmable extra portion before and after the threshold crossing, are saved. Parts containing only the baseline are suppressed. Each saved chunk has a timestamp relative to the start of the acquisition window.

It should be noted that, both with the Scope and ZLE firmware, it is possible for the common trigger to "cut off" a pulse when it spans the edges of the window. This issue is not present in the DAW firmware, which we will discuss in the next slide.

Slide 34

The DAW firmware is designed to acquire waveforms independently on a per-channel basis, meaning it does not have a common acquisition window like the ZLE firmware. It is, therefore, a triggerless type of acquisition, in the sense that there is no system-wide trigger; each channel acquires based on its own self-trigger, meaning it only records data when a signal is detected at its input.

Additionally, the length of the waveform is not fixed but adapts automatically to the duration of the signal, that is, as long as it remains above a set threshold, plus a programmable extension in duration. Each acquired waveform has its own timestamp, allowing for the reconstruction of the relative positions of each chunk through software.

Slide 35

With the DAW, unless there are issues related to excessive data throughput leading to data loss, it can be considered free of dead-time, in the sense that there are no losses of regions of interest in the acquisition, as seen in the case of cut-off events with the ZLE. Compared to the ZLE, the DAW is less suitable when searching for sparsely correlated events across different channels. With the DAW, very small pulses that do not exceed the trigger threshold may be lost, whereas with the ZLE, thanks to the global trigger, it is possible to set a much lower suppression threshold than the trigger threshold. This is based on the assumption that in a coincidence event, there is at least one channel with a pulse sufficiently high to exceed the trigger threshold.

Slide 37

Flexibility: A waveform digitizer is a "general-purpose" tool that is adapted to the application through the reprogramming of algorithms. During the acquisition phase, the only thing to evaluate is the ADC's sampling frequency, which should match the signal's speed. The acquisition mode, signal processing type, and output data can all be modified at any time through firmware and/or software updates. The latest digitizer models, thanks to "open FPGA" functionality and software like SciCompiler, allow users to develop algorithms freely, making the system even more flexible and dynamic. On the other hand, the analog chain is a "hard-wired" system whose functionality is determined by the individual hardware modules and their connections. There is therefore limited room for adaptation and reuse.

Multi-parametric: The digital solution produces various types of information as output parameters that characterize the pulse (height, charge, arrival time, shape, etc.). The number and type of generated parameters are determined by the algorithm loaded into the FPGA and can be modified at any time. In the analog system, adding a new type of information requires adding a new acquisition module, as each module produces a specific output.

Dead-time: Waveform digitizers are based on flash ADCs that continuously sample the input signal and have no conversion dead time, meaning there is no period where the system is blind. However, there can be dead time in the signal processing algorithm. For example, the trapezoidal filter requires time to reach the flat top of the trapezoid before accepting a new pulse. Generally, though, the digital system has less dead time than the analog one and can sustain higher trigger rates as long as "raw waveforms" don't need to be read. In this case, the data volume becomes so high that the acquisition bottleneck becomes the communication link to the computer, resulting in dead time due to internal memory occupation, drastically reducing the sustainable trigger rate.

Trigger logic: In systems with many channels, it's often necessary to create trigger logic by combining self-triggers from individual channels to implement coincidences, anti-coincidences, multiplicities, etc. Here again, the multiple I/Os present on digitizers and the ability to modify FPGA firmware allow customization of this logic without resorting to dedicated logic boards and complex wiring. Furthermore, as "time-stamped list

files" can be produced as output, coincidences and event building can also be performed in post-processing by the software.

Complexity: A waveform digitizer with DPP algorithms has many parameters to control. Consequently, the managing software and its user interface can be quite complex. The learning curve is steeper than that of an analog system. However, the digital system offers many facilities to aid users. Firstly, there's the possibility of "signal inspection" via software and the "waveform" mode, eliminating the need to connect tools to the setup. Signal inspection extends not only to input and output signals but also to individual stages of digital filters, such as the trapezoid filter in PHA. Additionally, unlike analog systems, the digital system is easily replicable and maintainable.

Cost: The cost per channel of analog and digital systems depends on many variables, making it impossible to say beforehand which one is more cost-effective. Generally, if the signals to be acquired are fast and require a sampling frequency of at least a gigasample per second, then the digital solution tends to be more expensive. For spectroscopy applications where signals come from charge-sensitive preamplifiers and are relatively slow, digitizers of around 100 Megasaples per seconds can be used, which have a very low cost per channel, also thanks to the possibility of creating dense boards (for example 64 channels per board).

Slide 39

A separate case is presented by digitizers based on switched capacitor arrays. This technology allows for high sampling frequency (higher than 1 Gigasample per second), high density, low power consumption, and low cost. However, they have two drawbacks: high dead time and a fixed simultaneous acquisition window across all channels. These characteristics make them inadequate for certain applications.

Slide 40

In this table, we have attempted to summarize the entire scenario regarding which functionalities are available for the different sampling rates offered by CAEN digitizers

Slide 42

Many applications in nuclear physics, especially high-energy physics, require thousands or even tens of thousands of channels. The high channel density goes hand in hand with other requirements, such as low power consumption, compact size, short cables, and low cost per channel. To meet these demands, dedicated ASIC-based readout electronics are commonly used.

Often, the ASIC implements the entire analog chain, as described earlier, integrating many channels (typically from 8 to 64) into a single chip. The A/D conversion takes place at the end of the chain, either inside the ASIC or with an external ADC. In both cases, the readout cannot be considered "digital" as in the case of waveform digitizers with DPP algorithms.

In other cases, the ASIC implements a digitizer of the "switched capacitor array" (SCA) type, which acquires "raw" waveforms at very high sampling frequencies and high channel densities. However, this type of digitizer typically has limitations due to its small memory depth (i.e., a short acquisition window) and the potential for dead-time that can reach tens of microseconds per trigger. Systems based on SCA-type ASICs are to be considered as digital chains, very similar to digitizers based on flash ADCs.

Generally, ASICs are highly specialized for a particular detector and application, so there are many different ASICs developed by companies or research institutions to cover most applications.

Slide 43

There is, therefore, a strong motivation to develop a common infrastructure that makes it easy and fast to integrate various ASICs while maintaining almost the same user interface and ease of use. This was the driving force behind CAEN's development of the FERS, which is a front-end readout system that integrates different types of ASICs.

It is a small board that houses one or more ASICs and all the necessary electronics to make it work. The board provides control and readout interfaces, making it possible for the user to use a particular ASIC without the need to develop hardware, firmware, and software.

CAEN S.p.A.
Via Vetraia 11
55049 - Viareggio
Italy
Phone +39 0584 388 398
Fax +39 0584 388 959
info@caen.it
www.caen.it

CAEN GmbH
Eckehardweg 10
42653 - Solingen
Germany
Phone +49 212 254 40 77
Fax +49 212 254 40 79
info@caen-de.com
www.caen-de.com

CAEN Technologies, Inc.
1 Edgewater Street - Suite 101
Staten Island, NY 10305
USA
Phone: +1 (718) 981-0401
Fax: +1 (718) 556-9185
info@caentechnologies.com
www.caentechnologies.com

CAENspa INDIA Private Limited
B205, BLDG42, B Wing,
Azad Nagar Sangam CHS,
Mhada Layout, Azad Nagar, Andheri (W)
Mumbai, Mumbai City,
Maharashtra, India, 400053
info@caen-india.in
www.caen-india.in

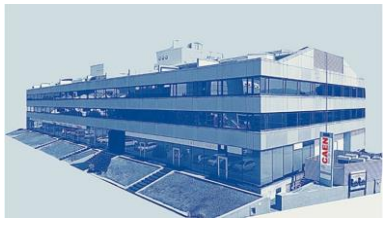

CAEN
Tools for Discovery

Electronic Instrumentation