

Zero-deadtime processing in beta spectroscopy for measurement of the non-zero neutrino mass

Benjamin LaRoque¹
(for the Project 8 Collaboration)

¹Pacific Northwest National Lab


Abstract. The Project 8 collaboration seeks to measure, or more tightly bound, the mass of the electron antineutrino by applying a novel spectroscopy technique to precisely measure the tritium beta-decay spectrum. The current system produces a single analog signal, which is digitized and processed in several stages before being saved to local disk storage. Online processing includes two stages, an FPGA connected to the analog to digital converter reduces the data down to the region of interest before shipping over the local network for further processing and storage. A normal CPU-based processing stage applies triggering logic to only save data at times when a signal is present, further reducing the total volume of data which needs to be written to disk or transferred for long-term storage. The next stage of the project will need to process many input channels and will integrate a necessary aggregation and combination step prior to applying the event search and triggering logic. We present the online processing system which has successfully been deployed for the current, single-channel, phase. We also present the status and design for a many-channel platform.

1 Introduction

The Project 8 collaboration seeks to use precision measurement of the tritium beta-decay spectrum to either more tightly constrain or measure the effective mass of the electron-flavor antineutrino. The collaboration has previously demonstrated Cyclotron Radiation Emission Spectroscopy (CRES) as a viable approach to measuring the energy of mildly-relativistic electrons [1]. In CRES, a radioactive source is allowed to decay in a region with an ambient magnetic field, which causes any charged particles produced to produce cyclotron radiation. The frequency of the cyclotron radiation depends upon the magnitude of the magnetic field, and the particle's charge-to-mass ratio, which for a relativistic particle is sensitive to its kinetic energy through the relativistic mass shift. Because of this, a precise measurement of the frequency of cyclotron radiation can be combined with a particle's known rest mass and independently measured magnetic field to determine the particle's kinetic energy.

The initial demonstrations of CRES have been conducted in a configuration where the entire source volume is contained within a waveguide, which acts as a very efficient signal collector, with characteristic dimensions determined by the frequency of the cyclotron radiation to be measured. A photograph of the source containment portion of the waveguide is shown in figure 1. Continuing to pursue CRES requires the ability to observe a substantially larger source volume, which precludes the use of an enclosing waveguide with only one or

Figure 1. The Project 8 Phase II magnet insert's trapping region, the core of which is a waveguide with circular cross section. There are five independent copper coils wound around the waveguide which provide magnetic field perturbations which axially confine electrons while they are being measured. The entire insert is installed within the bore of a superconducting NMR magnet, charged such that it produces an axial background magnetic field of just under 1 T. When a radioactive decay within the waveguide produces a charged particle, the ambient field will cause it to undergo cyclotron motion and produce cyclotron radiation. For electrons with kinetic energies in the low tens of keV (the tritium endpoint is around 18.6 keV, this radiation will couple to the propagating mode of the waveguide, which transports it to a low-noise cryogenic amplifier which is the first component of the analog signal conditioning system. After analog conditioning, the signals are passed to an analog to digital convert and then are processed in the digital signal processing system which is the subject of this manuscript.



few relevant propagating modes. The current plan is to instrument the source volume with a collection of antenna units. Analog signals from these units will be amplified and digitized independently and then combined in order to look for signals of interest.

2 Online signal processing in the current phase

In Phases I and II of Project 8, the entire active source volume is physically contained within a radio frequency waveguide, which is coupled directly to a low noise amplifier. The single-channel analog signal then goes through several analog signal processing stages which amplify the signal such that downstream thermal noise contributions are negligible, mix the signal from an initial frequency around 26 GHz to roughly 1 GHz, and adjust the total system gain such that the total power input to the analog to digital converter (ADC) is appropriate for its dynamic range. The signals of interest have very small amplitudes compared to the amplitude of the thermal fluctuations, but because these signals are coherent over longer time scales, it is still possible to extract them, as illustrated in figure 2. This is done by leveraging two properties of the signal: being narrow-band, which means that all of the signal power is in a narrow part of the power spectrum, while noise is spread out; and being relatively long duration, which means that the signal will remain in consecutive time windows, while noise fluctuations will not. A quite detailed treatment of the signal characteristics has been published previously [2].

After the analog signal processing is complete, the signal is digitized and processed in two online stages. A ROACH-2 platform leverages an FPGA to operate the ADC and do the first processing. The FPGA's bitcode is developed using the open source CASPER tool set




Figure 2. A simplified version of CRES signal extraction for Project 8. On the left all panels are shown in the time domain, the top panel shows a simple sine-like signal with amplitude much smaller than the uniform noise shown in the middle. The bottom panel shows the combination of the two, which is by eye is indistinguishable from the noise alone. On the right, a Fourier transform has been used to move into the frequency domain, where the original signal is clearly visible above the noise.

[3]. The hardware signal processing performance of this system has previously been studied in detail [4]. After digitization at 3.2 GSps (gigasamples per second), the FPGA uses digital down conversion to select three independent frequency regions of interest, each 100 MHz wide. For each region, a Fourier transform is used to convert to the frequency domain; data in both time and frequency domain are shipped over a 10 Gbps (gigabits per second) local area network for CPU-based processing and serialization. Samples are 8 bits long and so this means that after decimation there are 200 MBps (megabytes per second) per frequency region per domain.

A server receives both the time and frequency domain data, caching the time domain while using the frequency domain data to implement triggering logic. Each of the three channels is processed independently by a separate application instance. In the frequency domain, a frequency-dependent power threshold is defined relative to the average measured power level to account for the shape of the system gain. The software is able to be configured with two different trigger levels so that once a signal of interest is identified, power fluctuations do not cause data collection to stop prematurely. It is also possible to configure the system such that multiple threshold violations within a short time interval are required in order to identify a signal, allowing the threshold to be reduced for a given false-trigger rate (or a reduced false-trigger rate at fixed threshold). When a signal of interest has been identified, the corresponding time-domain data are written to local storage, starting at a fixed time interval before the first threshold violation and continuing until there have been no threshold violations for the configured time interval. This allows different or more sophisticated reconstruction to be performed offline.

The computational steps of the trigger logic execute in threads which pass data using buffers. For all relevant configurations, the total data processing speed is limited by the write speed of the SSDs used for initial storage. These support continuous streaming of a single region of interest, in principle up to the full storage volume of the drive, or all three regions when configured with reasonably optimized settings. Data files are written incrementally as trigger events happen, and data may be split into multiple files if a maximum file size is configured. Data management software detects when files are closed and immediately

transfers them to long term storage, allowing local storage to be freed and reducing the total storage required in the on-site system.

3 Plan for Phase III signal digitization and processing

Beginning with Phase III, Project 8 will seek to demonstrate the CRES technique using a source occupying a larger physical volume (of order 100 cm^3 as opposed to of order 1 mm^3 in Phase II). The analog signal collection system will therefore need to make use of many antennas, in order to collect enough of the radiated power to detect the signal above the thermal background noise. Because the source volume will be much larger than the signal's wavelength, simply summing the signals from the antennas will result in interference depending on the path-length between the source and each antenna. In order to be sensitive to the entire source volume, the signals will need to be combined many times using different relative phase delays, which means that the signal from each antenna will need to be digitized independently and the signal combinations completed as part of the digital signal processing. The current baseline design anticipates using a ring of approximately thirty antenna units to surround the source volume. The signal processing required can be thought of in two steps: first the signals from the independent analog channels must be combined to create a single signal which is sensitive to a limited sub-volume, then that combined signal must be processed in a fashion similar to that used in Phases I and II to look for events.

While transporting data from so many digital data streams creates some technical challenges, it also provides some opportunities in terms of data quality. First, because the electrons producing the signals of interest are localized and move through the source region in an understandable way, any source of background which does not behave in this same way is able to be suppressed. Second, a larger source volume means that there is an increased chance that there will be multiple signals present at the same time (in fact, a high-statistics experiment will need this to be the case). By processing signal combinations which are sensitive to only a small physical volume, it is anticipated the signals from simultaneous sources can be disentangled based on their positions.

In order to address the simultaneous challenges of the very large total raw data rate, and the large number of independent channel combinations which need to be calculated, a multi-stage processing system is being planned, with opportunities to scale the amount of computing resources based on performance benchmarks. A schematic of this system is depicted in figure 3 and the rest of this section addresses each stage in turn.

The front-end receiver boards will be FPGA-based systems designed using the CASPER tools. Each analog signal will be digitized by a single ADC, and each FPGA will operate as many ADCs as its capacity allows. For Phase III, there will be only a single region of interest selected for each analog channel and the FPGA will still be responsible for computing Fourier transformations.

The second stage is responsible for collecting and re-packaging the data streams produced by the front-end systems. It will take in the continuous stream of data from each channel and produce blocks of data which include all channels for a limited time duration. This stage is expected to be limited by the bandwidth of available networking hardware, and it may be necessary to have several coordinators acting in parallel, each responsible for a subset of the input channels. Each data block, or collection of a small number of blocks, will contain all of the data required to look for signals of interest during that time interval.

The time-slice processing stage will be responsible for collecting data blocks, computing the phased combinations of the input channels, and performing the signal of interest search on each the output of each combination. If there are N processing nodes, each will be responsible for receiving one out of every N data blocks, reducing the required network bandwidth




Figure 3. A block diagram of the data processing stages anticipated in Phase III of Project 8. Starting from the left, the front-end receivers will host a single FPGA and operate as many ADCs as the FPGA’s capacity allows. The number of front-end systems will be driven by the final channel count and the number of channels per board. Next, a data coordination and dispatch stage will collect streams of data from each independent channel and output blocks of data which include all channels for a limited time interval, with potential for scalability based on limitations of network bandwidth. The processing of data blocks will include both computing a complete set of phased combination of channels, and triggering a signal search on the output of each combination. Finally, after an event has been identified, only those time-intervals and phased combinations with relevant signals will be written to disk.

relative to the coordination nodes and allowing the system to be scaled without increasing the number of copies of the data which must be made and transmitted. The number of nodes will have to be determined by benchmarking the performance after algorithms have been implemented; it is anticipated that GPU-based processing may substantially accelerate this stage. This stage also provides the opportunity to use network-configured multicast to send a single block of data to multiple destinations; providing a means of testing multiple designs in parallel, be that different algorithms, or the same algorithms on different accelerating hardware.

The final processing stage receives identified signal information from the previous stage and determines exactly which data needs to be recorded. This will account for the possibility of an electron moving between the sensitive volumes corresponding to different phased combinations, possibly by computing new phased combinations of the data sensitive to volumes which include parts of adjacent volumes in the initial search. It remains to be determined if this step will be co-located with the prior stage, to further reduce the number of data transmissions required, or with the physical storage where the final data are written, or possibly both.

4 Results and outlook

In Phase II of Project 8, digitized signals are able to be processed in real-time to search for signals of interest and limit the amount of data which must be written to disk. The system is able to continuously process incoming data without deadtime, up to the limit of available local storage for the output. For trigger configurations with sufficiently high data reduction factors, this can mean continuous (deadtime free) data collection on the timescale of experimental operation.

For Phase III this becomes much harder as the amount of required online processing scales not like the increase in number of channels, but like the number of combinations of channels required to observe the entire source volume. In order to continue to achieve deadtime-free operation, a modular processing scheme has been developed. This enables the performance of each stage to be measured independently, and the resources allocated to each stage to be scaled based on that performance. It also provides natural opportunities to deploy hardware acceleration, such as GPUs, if appropriate.

Acknowledgements

This work is supported by the US DOE Office of Nuclear Physics, the US NSF, the PRISMA+ Cluster of Excellence at the University of Mainz, and internal investments at all collaborating institutions.

References

- [1] D. M. Asner et al. (Project 8 Collaboration), “Single-Electron Detection and Spectroscopy via Relativistic Cyclotron Radiation”, *Phys. Rev. Lett.* **114**, 162501 (20 April 2015). doi:10.1103/PhysRevLett.114.162501
- [2] A. Ashtari Esfahani et al. (Project 8 Collaboration), “Electron radiated power in cyclotron radiation emission spectroscopy experiments”, *Phys. Rev. C* **99(5)** (May 2019). doi:10.1103/PhysRevC.99.055501
- [3] J. Hickish, “A Decade of Developing Radio-Astronomy Instrumentation using CASPER Open-Source Technology”, *Journal of Astronomical Instrumentation*. **5** (2016) doi:10.1142/S2251171716410014.
- [4] N. Patel et al., “Characterizing the Performance of a High-Speed ADC for the SMA Digital Backend”, *Journal of Astronomical Instrumentation*. **3** (2014) doi:10.1142/S2251171714500019.