

# MICROBOONE-NOTE-1030-PUB:

## The Continuous Readout Stream of the MicroBooNE Liquid Argon Time Projection Chamber for Detection of Supernova Neutrinos

---

The MicroBooNE Collaboration

*E-mail:* [microboone\\_info@fnal.gov](mailto:microboone_info@fnal.gov)

ABSTRACT: The MicroBooNE Continuous Readout Stream is a parallel readout of the MicroBooNE liquid argon time projection chamber (LArTPC) which enables detection of non-beam events such as those from a supernova neutrino burst. The low energies of these neutrinos and the intense cosmic-ray flux due to the near-surface detector location makes triggering on these events very challenging. Instead, MicroBooNE has pioneered a novel approach for detecting supernova neutrinos relying on a delayed trigger generated by SNEWS (the Supernova Early Warning System). The continuous readout of the LArTPC is achieved through the use of real-time compression algorithms (zero suppression and Huffman compression) implemented in an FPGA (field-programmable gate array) in the readout electronics. We present the results of the optimization of the data reduction algorithms, and their operational performance. In order to demonstrate the capability of the continuous stream to detect low-energy electrons, a sample of Michel electrons from stopping cosmic-ray muons is reconstructed and compared to a similar sample from the lossless triggered readout stream.

---

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>The MicroBooNE Continuous Readout Stream</b>	<b>2</b>
<b>3</b>	<b>Data reduction algorithms</b>	<b>4</b>
3.1	Zero suppression	4
3.2	Huffman encoding	5
<b>4</b>	<b>Configuration of zero suppression parameters</b>	<b>6</b>
<b>5</b>	<b>Data compression results</b>	<b>7</b>
<b>6</b>	<b>Analysis of Continuous Readout Stream data</b>	<b>9</b>
6.1	Event building	10
6.2	Signal processing	10
6.3	Michel electron reconstruction	12
<b>7</b>	<b>The Continuous Readout Stream as a development platform</b>	<b>18</b>
<b>8</b>	<b>Conclusion</b>	<b>20</b>
<b>A</b>	<b>Dynamic-baseline algorithm</b>	<b>31</b>
<b>B</b>	<b>Physics-driven plane-wide zero-suppression thresholds</b>	<b>31</b>

---

## 1 Introduction

Liquid argon time projection chamber (LArTPC) detectors produce images of particle interactions with exquisite spatial and calorimetric resolution at the expense of large data volumes. For a typical ADC sampling rate of 2 MHz and an ADC resolution of 12 bits, each continuously read-out channel generates 3 MB/s without compression. In order to achieve good spatial resolution, most detectors use a wire pitch of a few millimeters spanning several meters, leading to several thousands of readout channels. Consequently, front-end electronics data rates can reach several GB/s. For the acquisition of events from a neutrino beam, these data rates are manageable since the readout of the detector is driven by the accelerator beam spills, which occur with a known maximum frequency of  $\mathcal{O}(10)$  Hz, determined by the accelerator repetition rate, and only requires an acquisition window of milliseconds, determined by the electron drift time across the TPC. For non-beam events such as supernova neutrino interactions or nucleon decays, since they cannot be anticipated, processing such amount of information in a continuous fashion, either for

generating triggers based on the TPC patterns or for data acquisition is a formidable task. Furthermore, with the upcoming multi-kiloton-scale detectors like the DUNE Far Detector modules [1] read out by hundreds of thousands of channels, the challenge becomes more acute as data rates at the front-end will be of the order of TB/s.

Within this scenario, MicroBooNE, as the unique LArTPC in operation with continuous readout, has a great opportunity to spearhead the development of the required technology and inform future experiments. The implementation of the continuous readout of the MicroBooNE LArTPC is described in section 2. The data reduction algorithms required to achieve the continuous readout within the MicroBooNE DAQ constraints are presented in section 3. The configuration of the main algorithm for zero suppression is discussed in section 4. The data compression results achieved are presented in section 5. The performance of the MicroBooNE Continuous Readout Stream in detecting and reconstructing Michel electrons from stopping cosmic-ray muons, which are similar to the ones expected from electron-neutrino interactions emitted by a core-collapse supernova, is discussed in section 6. The potential of the MicroBooNE Continuous Readout Stream as a development platform for other experiments such as DUNE is briefly discussed in section 7.

## 2 The MicroBooNE Continuous Readout Stream

The primary goal of the Continuous Readout Stream is to enable the acquisition of the burst of supernova neutrinos by MicroBooNE, should a nearby core-collapse supernova happen during the lifetime of the experiment. For this reason, the Continuous Readout Stream is also known as the Supernova Stream (SN Stream). Detection of supernova neutrinos in a LArTPC is especially interesting due to the higher sensitivity to the electron-neutrino flux through the  $\nu_e + {}^{40}\text{Ar} \rightarrow e^- + {}^{40}\text{K}^*$  channel [2]. The expected number of interactions in the MicroBooNE active volume is  $\mathcal{O}(10)$  for a supernova burst at 10 kpc (based on the prediction for DUNE [1]), spread over  $\sim 10$  s and with energies in the range  $\sim 5 - 50$  MeV. Due to the large cosmic-ray rate ( $\approx 5.5$  kHz [3]) resulting from the close-to-surface location of the detector, MicroBooNE cannot rely on self-triggering on these events. Instead, the TPC data is continuously saved to disk on the DAQ servers and the alert issued by the Supernova Early Network System (SNEWS) [4] can be used as delayed trigger. The latest data is kept for more than 48 hours, allowing the collaboration to react to the alert. We defer the discussion about the reconstruction and selection of the supernova neutrino interactions in the MicroBooNE detector to future work.

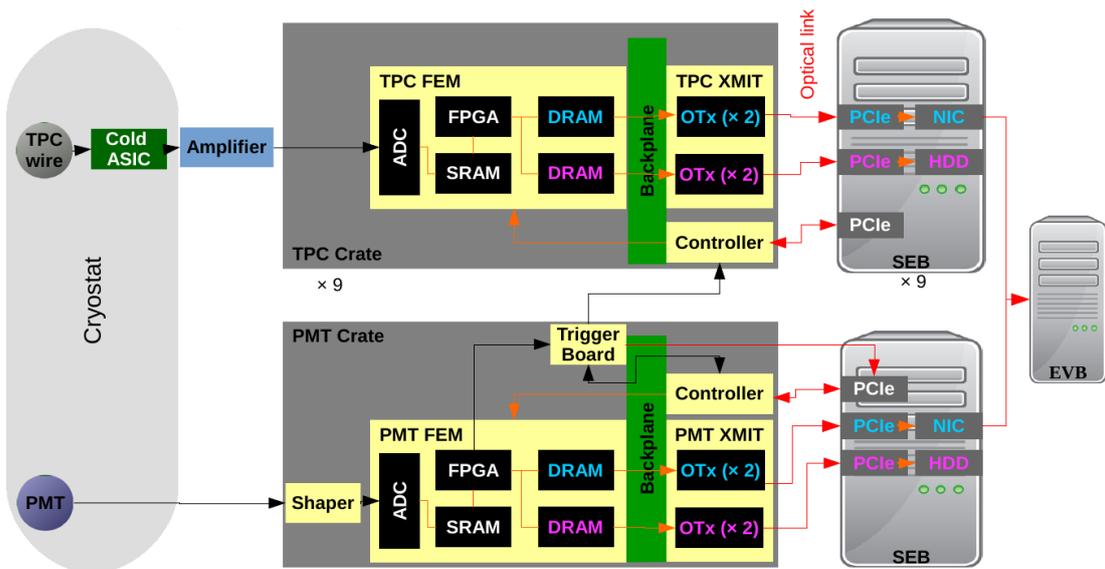
The MicroBooNE TPC is read out by three consecutive wire planes. The first two planes crossed by the drifting electrons are configured as induction planes; each having 2400 wires oriented at  $\pm 60^\circ$  from the vertical. The last plane is configured as a collection plane, with 3456 vertical wires. Figure 1 illustrates the dataflow in the MicroBooNE readout, from a TPC wire to the DAQ server. First, the signal in each TPC wire is preamplified and shaped by ASICs immersed in the liquid argon and then extracted from the cryostat through feedthroughs. The signal is further amplified by warm electronics immediately upon extraction to condition it for transmission using shielded twisted-pair cables to the readout digital electronics on a platform above the detector [5]. The readout

of the 8256 TPC channels is distributed among 9 crates, each one connected to a dedicated DAQ server, known as a Sub-Event Buffer (SEB). Seven of these crates (labeled as crates 02 – 08) are loaded with 15 Front-End Modules (FEMs), each reading out 64 channels, consisting of 16 wires from the first induction plane (plane U), 16 from the second induction plane (plane V) and 32 from the collection plane (plane Y). Crate 01 is loaded with 11 FEMs, reading exclusively induction channels from the first plane. Crate 09 is loaded with 14 FEMs and reads out mostly induction channels (720) from the second induction plane, with some channels (96) from the collection plane and some channels (48) from the first induction plane. In each FEM, the waveforms are digitized by an array of commercial 12-bit ADCs [6] at 16 MS/s. An FPGA (Altera Stratix III [7]) downsamples the waveforms to 2 MS/s and writes them in time order to a 1 M × 36 bit 128 MHz static RAM (SRAM) configured as a ring buffer. The TPC data is read back from the SRAM by the same FPGA, but now ordered by channel and split into two parallel streams. The Trigger Stream is only read out upon a trigger, generated by the Trigger Board and distributed through Controllers to the FEMs. For a detailed description see [5]. The Supernova Stream is continuously read out. The data of each FEM is sent to a transmitter board (XMIT) through a crate backplane with bandwidth up to 512 MB/s. The backplane dataway is shared between both streams, with the Trigger Stream given priority over the Continuous Stream using a token-passing scheme. The FEM has a dynamic RAM (DRAM) for each stream to buffer the data waiting for its turn to be transferred. The XMIT has 4 optical transceivers (OTx's), each rated to 3.125 Gb/s. Two are used for the Trigger Stream, and two for the Continuous Stream. Finally, the data from each stream is read out by custom PCIe ×4 cards in each SEB. The triggered data from each SEB is sent to the Event Builder (EVB) DAQ server using a network interface card (NIC), while the continuous stream is written to a 15 TB local hard disk drive (HDD) in each SEB, awaiting a SNEWS alert to be further transferred to offline storage. If no SNEWS alert is issued and the disk occupancy reaches 80%, the oldest data is permanently deleted until the occupancy falls below 70%.

The large data rate read out by the front-end electronics in each crate,  $\sim 4$  GB/s, prevents the continuous acquisition of the TPC waveforms by the back-end DAQ without compression. In particular, the bottleneck of the Continuous Stream is found in the disk-writing speed of the local hard drive system, which is in the range of 50 – 200 MB/s. In order to achieve such data rates, the FPGA applies data reduction algorithms described in section 3 to achieve a  $\sim 20 - 80$  compression factor.

The Continuous Stream data is arranged into frames corresponding to 1.6 ms of detector readout (the MicroBooNE drift time at the design electric field). Each FEM creates its own frame record, writing the TPC data in a payload which is preceded by a header consisting of twelve 16-bit words indicating the FEM address, a word count of the data in the payload, a sequential identifier, the frame number, and a simple checksum of the payload data. The TPC data consists of 16-bit words. The data from each channel is preceded by a channel header and a timestamp.

A continuous readout of the PMT system, based on the more conventional out-of-beam-spill discrimination signals described in [5], also exists but is not used for any of the results of this work.



**Figure 1.** Simplified diagram of the MicroBooNE readout dataflow. The Trigger Stream components are highlighted in blue text and the Continuous Readout Stream components are highlighted in magenta.

### 3 Data reduction algorithms

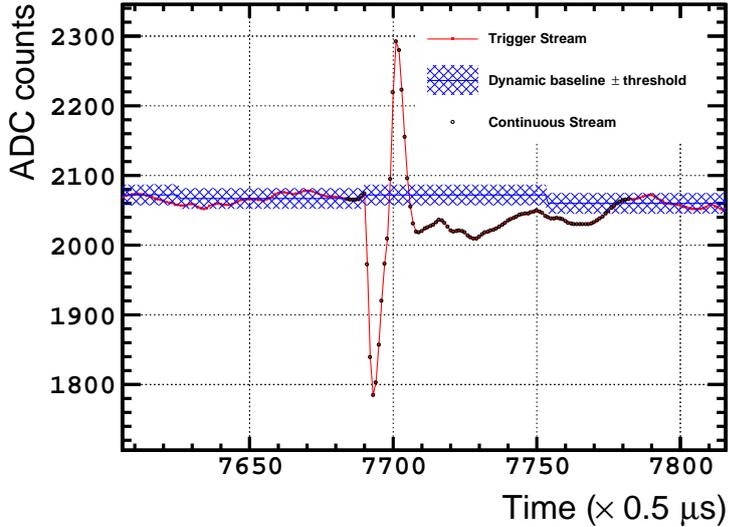
The Continuous Readout Stream FPGA firmware applies two data reduction algorithms sequentially. The first one is a zero-suppression scheme in which ADC samples within a channel not passing a configurable threshold with respect to an estimated baseline are discarded. The second data reduction algorithm is a fixed-table Huffman compression, in which consecutive ADC samples which differ by less than 4 ADC counts are encoded using a reduced number of bits.

#### 3.1 Zero suppression

The zero-suppression algorithm aims at removing the samples which do not carry any signal. For this, it checks whether an ADC sample passes an amplitude threshold after subtracting the channel baseline. The sign of the threshold can be chosen to be positive (passing samples are greater than the sum of the baseline and threshold), negative (passing samples are smaller than the baseline minus the threshold) or either. The threshold value and sign are configurable for each channel. The different methods used to determine the thresholds are described in section 4 and appendix B. In addition, a number of samples preceding the first one to pass the threshold (presamples) and following the last sample that passed the threshold (postsamples) are retained in order to better capture the waveform. The set of samples that pass the threshold, plus the presamples and postsamples, is considered a Region-Of-Interest (ROI). The numbers of presamples and postsamples are configurable for each FEM (64-channel block). They have been set to 7 presamples and 8 postsamples, the maximum values allowed by the current FPGA firmware.

Two versions of the zero-suppression firmware have been produced. One version uses a sliding-window algorithm in the FEM FPGA to estimate the baseline dynamically (see appendix A for details). Another version uses a static baseline per channel which is configured at the beginning of the run. In this version, the baseline value for each channel is extracted from a previously chosen reference run from the DAQ Trigger Stream, taking the mode of the raw ADC distribution.

An illustration of the effect of the zero suppression (with dynamic baseline) on a waveform is shown in figure 2. The results from each configuration are further discussed in section 5.



**Figure 2.** Example of data from a test stand at Nevis Laboratories showing the samples of a zero-suppressed waveform in the Continuous Readout Stream (empty black circles), superimposed on the same waveform from the Trigger Stream (red). An emulation of the dynamic baseline used by the FPGA is shown as a blue line, with the threshold shown as a blue cross-hatched band. Only the samples with ADC values out of this band are saved, plus a number of samples preceding them (presamples) and following them (postsamples). The threshold and baseline estimation tolerance values used in this figure do not correspond to the ones used in MicroBooNE.

### 3.2 Huffman encoding

After the digitized waveform has been zero suppressed, it is run through a Huffman encoding [8] stage in which successive ADC samples differing by no more than  $\pm 3$  ADC counts relative to the predecessor ADC sample value are encoded according to table 1. This stage reduces the memory footprint of the saved waveform by attempting to store more ADC samples in the same memory space it would take to store a single uncompressed ADC sample.

The readout electronics data format uses 16-bit words. Non-Huffman-encoded ADC words use the lowest 12 bits to store the 12-bit ADC value, and use the rest as header to

**Table 1.** Huffman encoding table relating the value of the difference between the current ADC sample and the preceding one,  $\Delta\text{ADC} = \text{ADC}_i - \text{ADC}_{i-1}$ , and the Huffman binary code.

$\Delta\text{ADC}$	Code
0	1
-1	01
+1	001
-2	0001
+2	00001
-3	000001
+3	0000001

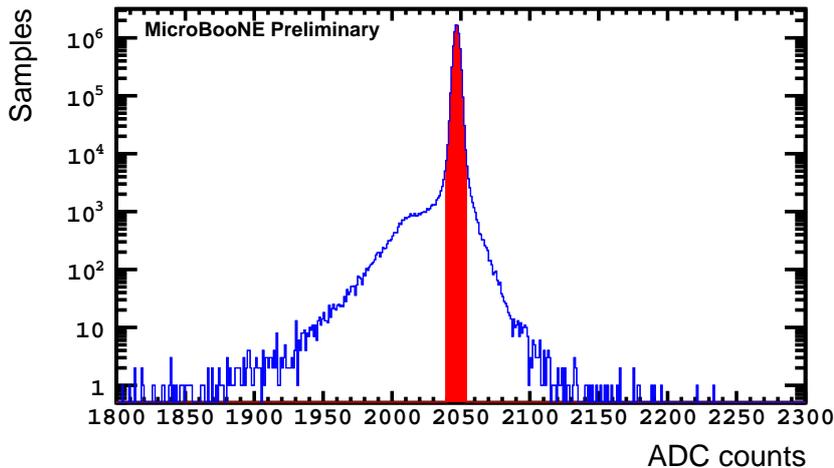
identify the word as such. Huffman-encoded words have the sixteenth bit (most significant bit) set to 1 to identify the word as Huffman-encoded. The other 15 bits are available to contain ADC information using the codes shown in table 1. If there are no more samples to be encoded in the Huffman word (because the next ADC difference is larger than  $\pm 3$  ADC counts) or the required code does not fit in the available bits, the unused least significant bits are filled with zeros. In the latter case, a new Huffman-encoded word will be created to continue storing the ADC differences.

#### 4 Configuration of zero suppression parameters

This section describes the current method we are using to determine the channel-wise thresholds for zero suppression of TPC waveforms. Our initial approach, now deprecated, was to use a single physics-motivated threshold for each TPC plane (a plane-wide threshold) to separate signals from noise and is described in appendix B. As will be discussed in section 5, the initial thresholds were too conservative and resulted in data rates well below the original goal of 50 MB/s. This allowed the setting of lower thresholds not driven by the separation of signal from noise, but by exploiting the bandwidth of the readout electronics, as described next. The motivation is that any data that is zero-suppressed online will be lost forever, while we can add additional higher thresholds offline if needed.

In order to set the threshold values per channel, we analyze the Trigger Stream ADC distribution for each channel. This distribution consists overwhelmingly of noise. We find the ADC values defining the shortest symmetrical interval around the mode of the ADC distribution, containing more than 98.5% of the ADC distribution, corresponding to a zero-suppression compression factor greater than 67. An example of the method is shown in figure 3. The mode of the distribution is also taken as the baseline ADC value for the zero-suppression firmware that uses static baselines as described in section 3.1. The ADC values of the integration limits found, after subtracting the ADC baseline value, are taken as the zero-suppression thresholds

The channel-wise thresholds and baselines used for the static-baseline firmware are shown in figure 4. Almost every channel has a threshold lower than their plane-wise threshold counterpart. The average threshold is 3.6 times smaller for U plane channels,



**Figure 3.** Examples of raw ADC distributions from run 18468 from the Trigger Stream used to determine the zero-suppression thresholds and static baseline for channel 2000 from the first induction plane. The red area shows the data that is zero-suppressed, found by integrating 98.5% of the distribution symmetrically around the maximum (taken as the baseline value). The limits of the red area denote the position of the thresholds.

2.2 times smaller for V plane channels, and 5.2 times smaller for Y plane channels, allowing the recording of more data and ensuring a higher charge-collection efficiency for low-energy signals. Furthermore, noisy channels are effectively masked by setting higher thresholds for them.

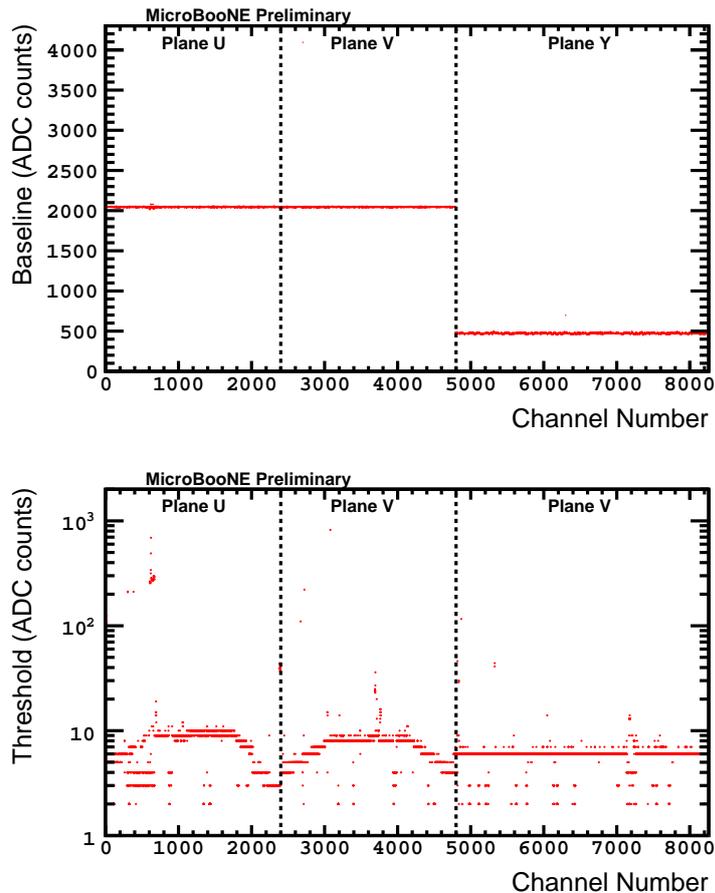
## 5 Data compression results

Figure 5 shows the compression factors achieved in each SEB for the three zero-suppression configurations tested of the Continuous Readout Stream, summarized in table 2. The Continuous Stream data dominant physics contribution is cosmic-ray muons crossing the MicroBooNE detector. The compression factor is computed as the ratio of the expected data rate without compression and the measured data rate.

**Table 2.** Summary of zero-suppression configurations tested in the MicroBooNE detector.

Thresholds \ Baseline	Dynamic	Static
Physics-driven plane-wide	SN Run Period 1	Not used
Bandwidth-driven channel-wise	SN Run Period 2	SN Run Period 3

From November 2017 to July 2018 (SN Run Period 1), the zero suppression configuration described in appendix B was used. This resulted in data rates well below the 50 MB/s target, except for SEB06. The cause of the high data rates and variation observed in SEB06 was traced back to the large number of noisy channels (due to ASIC miscon-

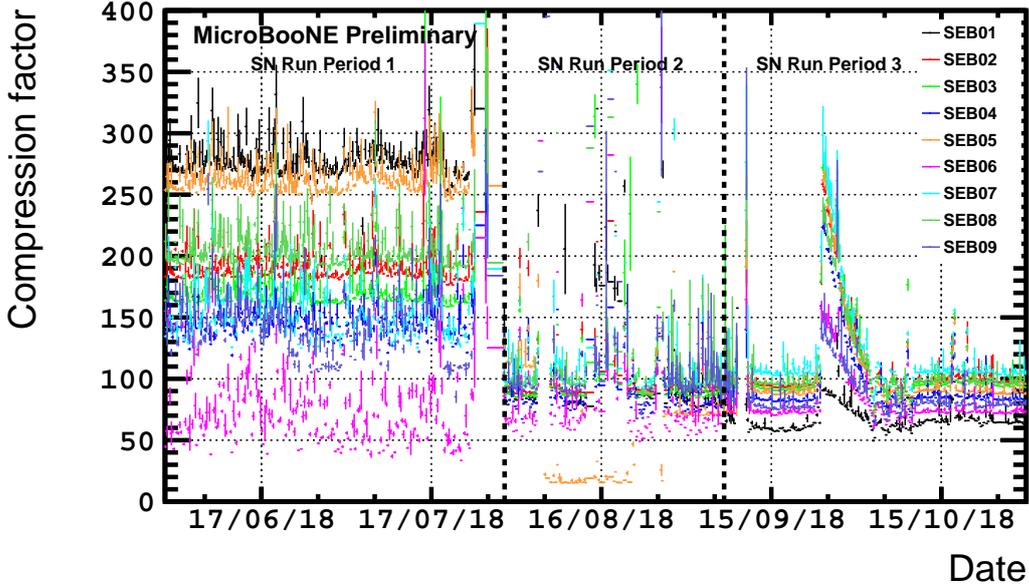


**Figure 4.** Channel-wise baselines (top) and thresholds (applied bipolarly, bottom) used for zero suppression with the FPGA firmware with static baselines.

figuration) read out by that crate, which prevented the dynamic baseline algorithm from establishing an accurate and stable baseline upon which to execute the zero-suppression. The dispersion in data rates is also seen in other SEBs (e.g. SEB07 and SEB09, but with a reduced magnitude). Moreover, the low data rates in the rest of the SEBs suggested it was actually feasible to lower the thresholds to gain efficiency for low-energy signals, even at the expense of recording some noise which could be eliminated during the offline reconstruction. This was the motivation to deprecate the plane-wide thresholds in favor of (mostly lowered) individualized thresholds adjusted to produce data rates closer to the target goal.

Beginning in August 2018 (SN Run Period 2), the zero-suppression per-channel thresholds described in section 4 were deployed, keeping the dynamic baseline estimation. Because most of the thresholds were below the plane-wide values, this resulted in an increase of data rates for most of the SEBs, while the raising of thresholds for a few especially noisy channels decreased the fluctuations in SEB06, but not completely. The lack of an accurate and precise baseline during large portions of the run was a major concern. This motivated

the replacement of the baseline estimation algorithm with the static configuration version, beginning in September 2018 (SN Run Period 3), in order to have a baseline value for zero suppression from the beginning of the run, regardless of the noise conditions. As seen in the figure, this latest zero-suppression configuration using static baselines and channel-wise thresholds achieved the compression-factor target and resulted in better stability for all the SEBs, and has been adopted as the default running mode.



**Figure 5.** Compression factors achieved in the Continuous Readout Stream for the 9 TPC DAQ servers (SEBs) with the three zero suppression configurations used so far. Date format is day/month/year. Until end of July 2018 (marked with a heavy dashed line) the configuration used the plane-wise thresholds and dynamic baselines. During August 2018, the lower channel-wise thresholds were tested keeping the dynamic baseline. Beginning in September 2018 (marked with a heavy dashed line), the static baselines were introduced. Each point shows the mean of the data disk-write rate over 6 h. Error bars show the standard error on the mean. The low compression factor for SEB05 during mid-August was caused by a misconfigured front-end ASIC after emerging from one power outage, and returned to the proper configuration in the following power outage. The ramp starting on September 24th, 2018, corresponds to the filling of the cryostat with a batch of lower-purity argon, followed by a period of high-voltage instabilities.

## 6 Analysis of Continuous Readout Stream data

This section describes the offline analysis of the Continuous Readout Stream data using LArSoft [9] and uboonecode [10], with the final goal of assessing the sensitivity to electrons with energies in the supernova neutrino range (few to tens of MeV).

## 6.1 Event building

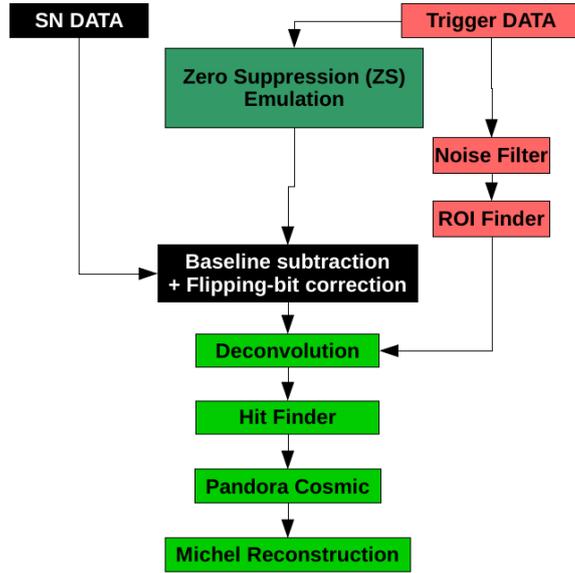
The raw binary data from the Continuous Stream is written to files on each of the nine SEB DAQ servers, each of which has a 15 TB RAID array which temporarily stores the data for hours (typically more than 48 h) before it is permanently deleted. Events are only built on demand. Parallel processes independent of the primary DAQ run on each SEB, retrieving the requested frames from a given run number and sending them via TCP (transmission control protocol) connections to a central server. Data from all TPC SEBs and the PMT SEB are assembled for each frame and written to a MicroBooNE-format file. There is no automated trigger to respond to the SNEWS alert. Instead, the runs surrounding the alert timestamp would be marked as such by collaborators to prevent deletion in the hours following the alert, and subsequently assembled to search for supernova neutrinos.

The encapsulation of the Continuous Stream is slightly more complex than the regular Trigger Stream since the data consists of ROIs occurring randomly in time. In the case of a core-collapse supernova, the neutrino burst is spread out over tens of seconds and there is no clear start time, so the event concept must be defined. For the Michel electron analysis described in section 6.3, the events were defined as 6400-samples long, the size used for the Trigger Stream. Each event is formed by taking a 1.6 ms-long frame (3200 samples), and appending the preceding last 0.8ms (1600 samples) from the previous frame and the following first 0.8ms (1600 samples) from the next frame. This ensures that objects crossing frame boundaries can be well reconstructed by the pattern recognition algorithms. For all the results from the continuous stream shown in section 6.3, only the Michel electrons with the decay vertex in the central frame are considered to avoid double-counting.

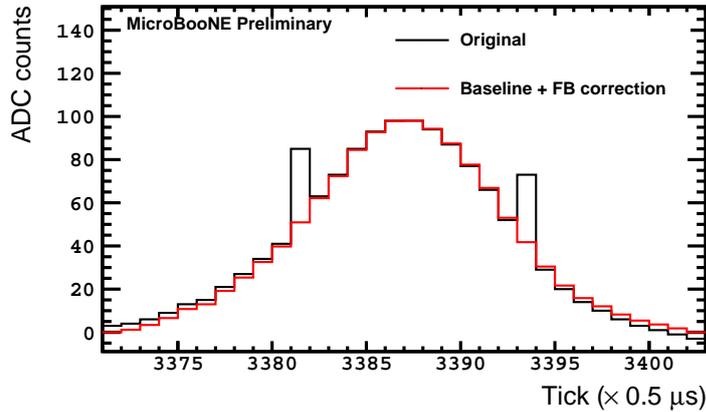
## 6.2 Signal processing

The ROIs produced by the zero suppression in the Continuous Readout Stream follow a reconstruction chain similar to the one in the Trigger Stream (see figure 6). No noise-filter stage is used as the zero-suppression removes the baseline regions where this filter is effective. No software ROI finder stage occurs either, since the zero-suppression in the FPGA is already producing ROIs. An additional challenge not present in the Trigger Stream is that 4% of the ADC samples of the Continuous Readout Stream exhibit flipping bits randomly (see figure 7). This results in one or more of the bits encoding the ADC value to switch from 0 to 1, or vice versa, shifting the original ADC value by a combination of powers of 2. Since this has not been observed in the Trigger Stream, it must occur during digital processing after the SRAM, when the two streams separate (cf. figure 1). While the investigation of the origin of the flipping bits continues, we have mitigated their effect offline as described next.

A tailored offline baseline subtraction is done for the Continuous Stream using a linear interpolation between one of the presamples and one of the postsamples which are acquired during the zero-suppression. Due to the occurrence of the flipping bits, the ADC value of the pre/postsamples needs to be checked before using it for interpolation. For this, we compute the median ADC value for the presamples and the postsamples, separately. This provides a first estimation of the baseline on each side of the pulse, since most of the



**Figure 6.** Illustration of the reconstruction chains used for the different data sets: the stages unique to the Continuous Readout Stream data (SN data) are shown as black boxes, the stages unique to the Trigger Stream data are shown as red boxes. A zero-suppression simulation stage emulates the real-time digital processing of the Continuous Stream data in the FPGA and allows the Trigger Stream to be converted into Continuous-Stream-like data for direct comparisons. The light green boxes show common reconstruction stages for which the processing is identical.



**Figure 7.** Example of the flipping-bit filter on a waveform from a collection plane channel. The baseline-subtracted waveform after the flipping-bit filter is shown in red. The original waveform is shown in black (ADC counts have been shifted by the average baseline value to fit in the same axis range). ADC samples affected by flipping bits can be seen at ticks 3381 and 3393 of the 2 MHz clock.

samples are not affected by flipping bits. Then we compare the presamples (postsamples) to the median and choose the earliest presample (latest postsample) within 15 ADC counts (absolute difference) of the median as the reference points for interpolation. The 15 ADC counts cut was chosen to reject samples affected by a flipping bit in the fourth bit or higher, as that is the minimum deviation which could be identified as flipped bits as opposed to noise fluctuations or ionization charge signal-related deviation.

Once the baseline has been subtracted, the waveform goes through a filter for flipping bits. This is done by comparing each ADC value to a linear interpolation built using the preceding and the following samples. If the difference between them is bigger than 32 ADC counts, the ADC value is replaced with the interpolation. Flipping bits with a shift smaller than 32 ADC counts are difficult to tell from actual signals and we do not try to correct them. The ADC cuts were chosen by hand-scanning waveforms and identifying the spikes which could be attributed to flipping bits. Unlike in the case of the baseline estimation, identifying flipping bits in the rising or falling edge of a pulse is more challenging, which motivated the choice of more conservative values than for the baseline selection. An example of this algorithm is shown in figure 7.

Finally, the waveform is deconvolved using the same 1-D deconvolution tool [11] that is used to deconvolve the Trigger Stream waveforms.

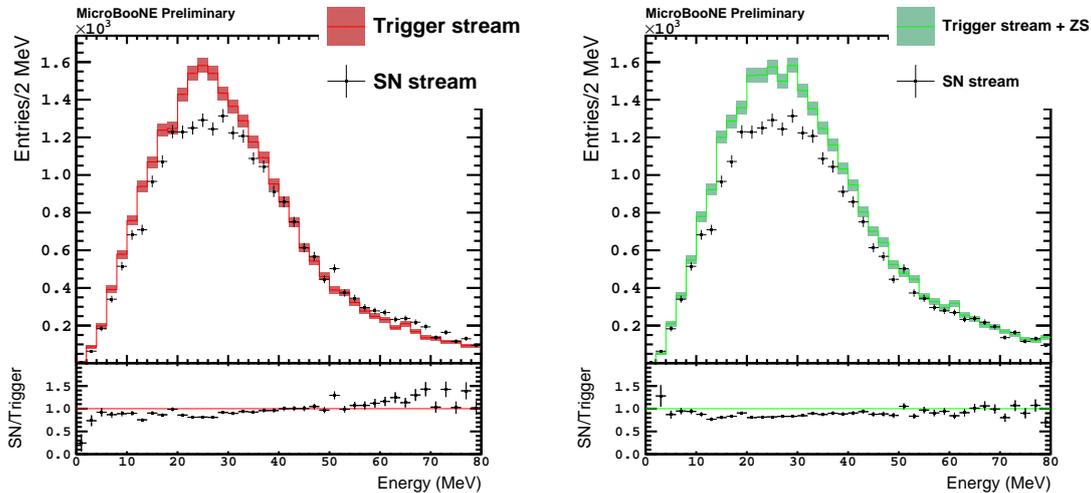
### 6.3 Michel electron reconstruction

In order to show the performance of the Continuous Readout Stream, we use Michel electrons from stopping cosmic-ray muons as a case study. The Michel electron spectrum spans an energy range very similar to the electrons resulting from charged-current interactions of electron neutrinos from a core-collapse supernova. The reconstruction and selection is done following the MicroBooNE publication [12], but extending its application to the induction planes as well. For each plane, the Michel reconstruction is run independently, for the purposes of studying the effect of zero suppression on induction (bipolar) signals vs. collection (unipolar) signals.

The Continuous Stream data set used corresponds to 1999022 frames (53.31 minutes) taken on September 21<sup>st</sup>, 2018. To provide a reference for comparison, we use a data set from off-beam zero-bias triggers from the Trigger Stream corresponding to 1102845 events (58.82 minutes) taken between December 1<sup>st</sup>, 2017 and July 7<sup>th</sup>, 2018, after applying data quality criteria for the detector operating conditions. This data set is processed following the standard reconstruction for the Trigger Stream. In addition, we process 97.9% of the Trigger Stream data set through a zero-suppression (ZS) emulation that reproduces the FPGA algorithm and reconstruct it using the same tools as the Continuous Stream.

The Michel electron spectra from the three TPC planes are shown in figures 8 and 9. They show the total energy of the Michel electron candidates, summing over all the hits of the ionization and radiative clusters as in reference [12] (see an example in figure 10). In order to avoid ambiguities when reconstructing Michel electrons with their radiative components overlapping, we reject the events with more than one Michel electron candidate. The ADC count-to-MeV calibration constants for the U, V and Y planes are  $9.00 \times 10^{-3}$  MeV/ADC,  $8.71 \times 10^{-3}$  MeV/ADC and  $9.24 \times 10^{-3}$  MeV/ADC, respectively,

following from [13]. In the three spectra, a rate and shape discrepancy between the Continuous Stream and the Trigger Stream is observed. The shape discrepancy is caused by the zero suppression, as evidenced by the good agreement (i.e. flat ratio) with the Trigger Stream when we simulate the zero-suppression (see also the shape comparison in figure 11).

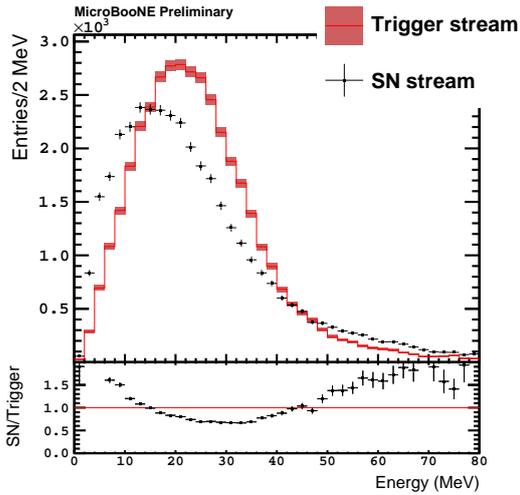


(a) Plane Y (standard Trigger Stream).

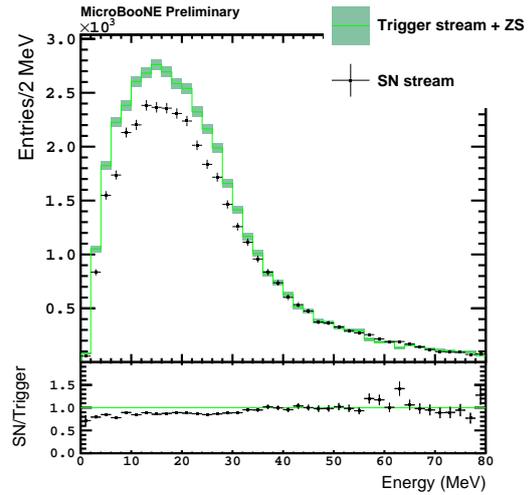
(b) Plane Y (Trigger Stream with ZS emulation).

**Figure 8.** Michel electron candidate energy spectra reconstructed using only the collection plane. The black points in the upper panels show the same spectra from the Continuous Readout Stream (SN stream). (a) shows the Trigger Stream data processed through the standard reconstruction (red histogram). (b) shows the Trigger Stream data processed through the zero-suppression (ZS) emulation and the continuous stream reconstruction (green histogram). Both Trigger Stream spectra are normalized to the exposure of the Continuous Stream. The bottom panels show the ratio between the Continuous Stream and the Trigger Stream data points. The error bars and bands show statistical uncertainties.

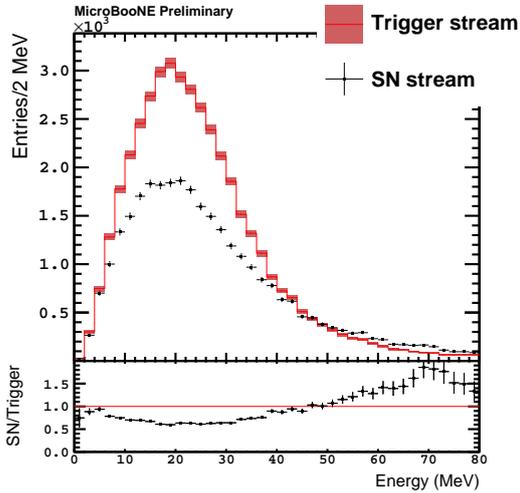
Table 3 shows the Michel electron rates, considering only the candidates with energy below 80 MeV as in [12]. The rate and shape discrepancy in the middle induction plane when compared to the non zero-suppressed Trigger Stream is especially remarkable. The middle induction plane has signals which are more vulnerable to not passing the zero suppression due to their smaller amplitudes and bipolar symmetric shapes, which favors cancellations. Excluding this plane, table 3 shows that the Michel electron rates between the Continuous Stream and the Trigger Stream agree within 10%. It is important to take into account that the Continuous Stream data set corresponds to 53.31 minutes of actual run time while the Trigger Stream data set is spread over 218 days of run time, and hence are subject to different fluctuations. In particular, the average surface temperature during the Trigger Stream data set was 6 degrees Celsius, while the average temperature for the Continuous Stream was 26 degrees Celsius. It is known that the seasonal temperature variations induce a modulation on the cosmic muon ray flux, as they change the density of the atmosphere in which the muons are produced. Determining the exact effect on the stopping muons inside the MicroBooNE detector is out of the scope of this work, but there



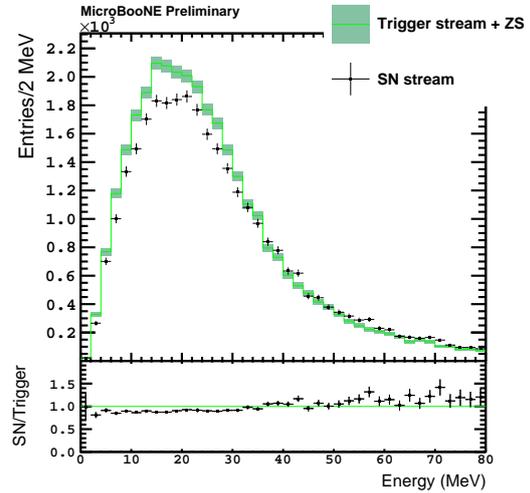
(a) Plane U (standard Trigger Stream).



(b) Plane U (Trigger Stream with ZS emulation).



(c) Plane V (standard Trigger Stream).

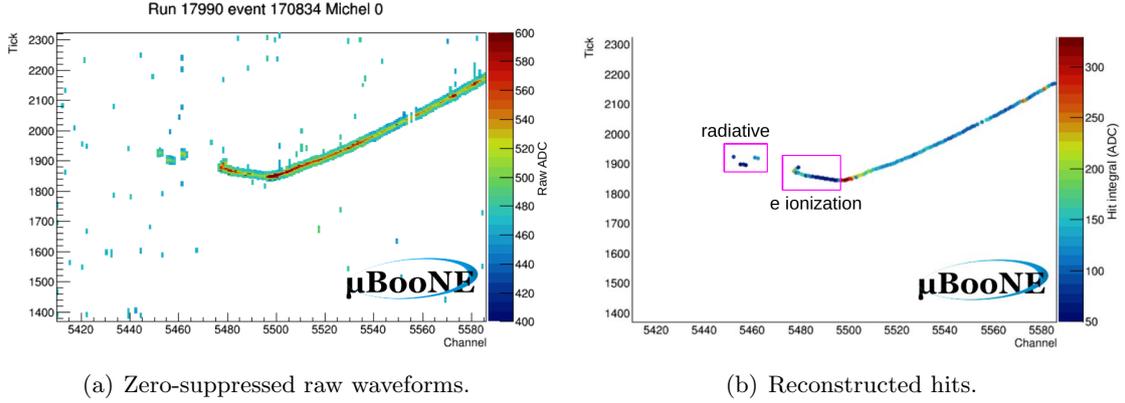


(d) Plane V (Trigger Stream with ZS emulation).

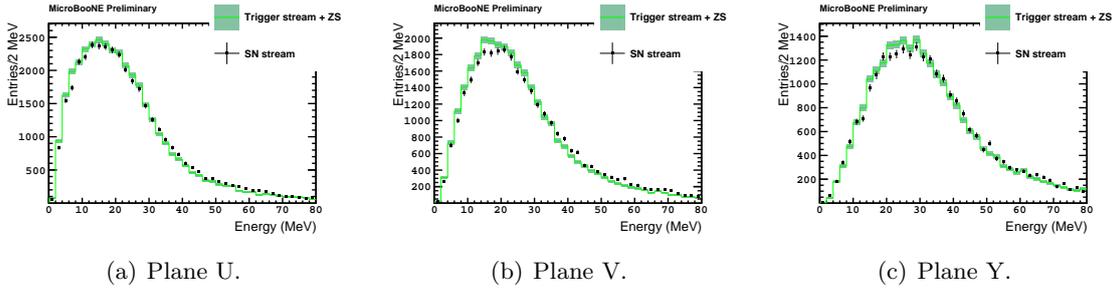
**Figure 9.** Michel electron candidate energy spectra reconstructed using only one of the induction planes. The top row shows the first induction plane (plane U), the bottom row shows the second induction plane (plane V). For each row, the markers and colors follow the same convention as figure 8.

are on-going efforts to measure the cosmic muon rates in MicroBooNE, and the estimated impact on the rate is 3 – 10%. The slight increase in the Trigger Stream rates when simulating zero suppression over the Trigger Stream, and the associated decrease in the ratios of the rates (except for the V plane discussed above), can be understood as events from the overflow bin (above 80 MeV, not included in the rate measurement) migrating into lower energies when adding the effect of the zero suppression.

The effect of the zero suppression is further studied by separating the ionization and



**Figure 10.** Michel electron candidate event display on the collection plane. The white background areas on (a) show the channel readouts which have been zero suppressed. The pink boxes on (b) illustrate the electron ionization and radiative components.

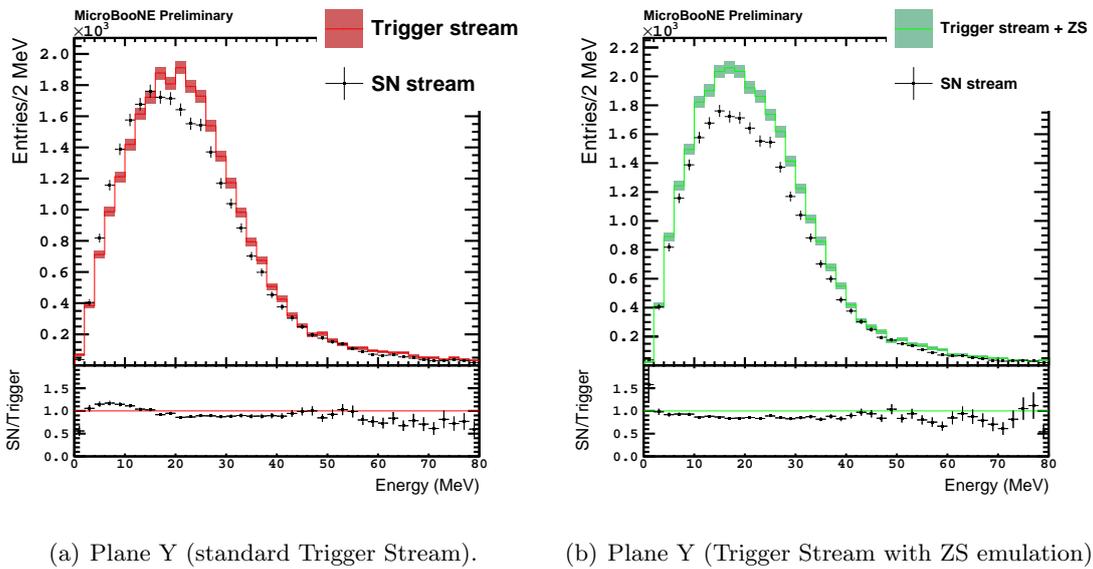


**Figure 11.** Michel electron candidate energy spectra from the Continuous Readout Stream (black points) overlaid on reference spectra from the Trigger Stream with simulated zero suppression (green histogram) normalized to the same area. Error bars and bands display statistical uncertainty.

**Table 3.** Michel electron candidate rates measured in the continuous stream (SN), the Trigger Stream (Trigger) and the Trigger Stream with simulated zero suppression (Trigger + ZS) on the three TPC planes. The last two rows show the ratio between the continuous stream and the Trigger Stream (without and with simulated zero suppression). Uncertainties are statistical only.

Michel e rate	Plane U	Plane V	Plane Y
SN ( $s^{-1}$ )	$11.58 \pm 0.06$	$9.49 \pm 0.05$	$7.63 \pm 0.05$
Trigger ( $s^{-1}$ )	$12.22 \pm 0.06$	$13.10 \pm 0.06$	$8.20 \pm 0.05$
Trigger + ZS ( $s^{-1}$ )	$13.01 \pm 0.06$	$10.11 \pm 0.05$	$8.79 \pm 0.05$
SN/Trigger	$0.948 \pm 0.007$	$0.724 \pm 0.005$	$0.930 \pm 0.008$
SN/(Trigger + ZS)	$0.890 \pm 0.006$	$0.939 \pm 0.007$	$0.868 \pm 0.008$

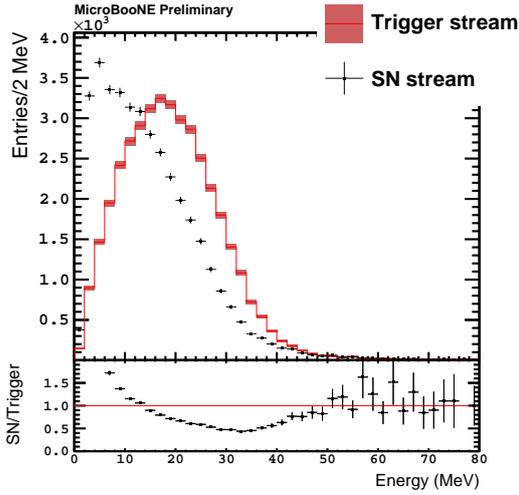
radiative components of the Michel electrons. The zero suppression is found to cause a shift to lower energies of the ionization component (see figures 12 and 13), and an excess of the radiative contribution at high energies (see figures 14 and 15). We interpret this effect as Michel electron tracks being split into segments by the zero-suppression. This leads to shorter reconstructed ionization components, confirmed explicitly by measuring the length of the ionization component in figures 16 and 17, and analyzing the hit multiplicity of the ionization component in figures 18 and 19. A consequence is the detached ionization segments being reconstructed as radiative components, increasing the radiative hit multiplicity as shown in figures 20 and 21. These effects compensate each other when computing the total energy of the Michel electron by summing over all the ionization and radiative components.



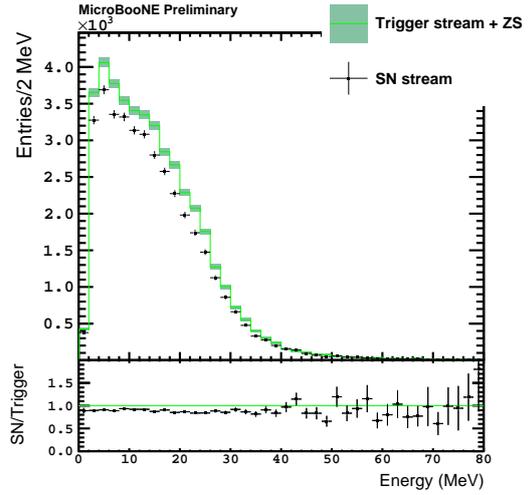
**Figure 12.** Energy spectra of the Michel electron candidate ionization component reconstructed using only the collection plane. The markers and colors follow the same convention as figure 8.

The impact of the zero suppression is stronger for the induction planes, as they feature smaller signals and higher thresholds (cf. figure 4), resulting in a relative shift of the distribution peaks to lower values when compared to the same distribution for the collection plane. Nevertheless, a similar shift is also seen in the Trigger Stream spectra (e.g. see figures 9(a) and 9(c) with respect to figure 8(a)). This points to a higher inefficiency in collecting the charge on the induction planes which will also contribute to this shift.

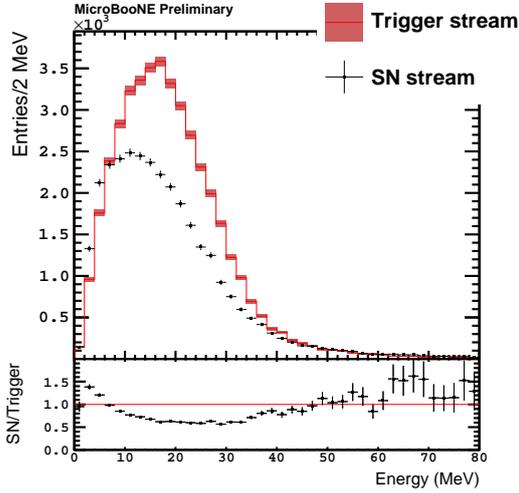
Among the induction planes, the U plane shows more extreme shifts to lower values due to zero suppression. This is expected since the U plane thresholds are higher than the ones for the V plane, but also because the charge from the slow-rising induction component of the U plane signals is not fully captured by the limited number of presamples allocated in the zero suppression. This effect can be further investigated by analyzing the energies of the individual hits that form the ionization (figures 22, 23) and radiative (figures 24, 25) clusters. In particular, the U-plane hit-energy spectrum (figure 23(a)) shows an increase



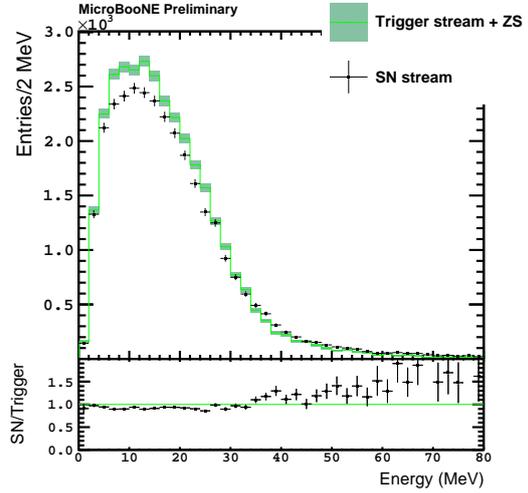
(a) Plane U (standard Trigger Stream).



(b) Plane U (Trigger Stream with ZS emulation).



(c) Plane V (standard Trigger Stream).

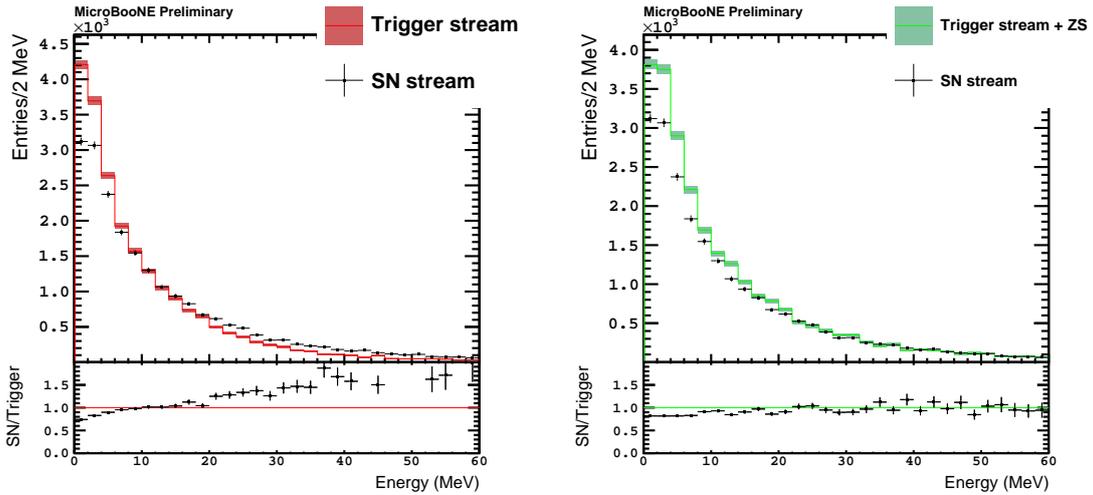


(d) Plane V (Trigger Stream with ZS emulation).

**Figure 13.** Energy spectra of the Michel electron candidate ionization component reconstructed using only one of the induction planes. The top row shows the first induction plane (plane U), the bottom row shows the second induction plane (plane V). For each row, the markers and colors follow the same convention as figure 8.

of the “Compton-like” tail with respect to the V-plane hit-energy spectrum (figure 23(c)).

The good agreement shown by the Trigger Stream with simulated zero suppression and the Continuous Stream hit-energy spectra allow us to anticipate that the impact of flipping bits on calorimetry after signal processing is very small. Figures 22, 23, 24 and 25 show the flipping bits as a small peak at 0.1 – 0.2 MeV. This is understood to be caused by the flipped bits which escape correction by the filter and distort the waveform, forcing the hit finder to allocate extra hits to fit the waveform shape. Because only small shifts



(a) Plane Y (standard Trigger Stream).

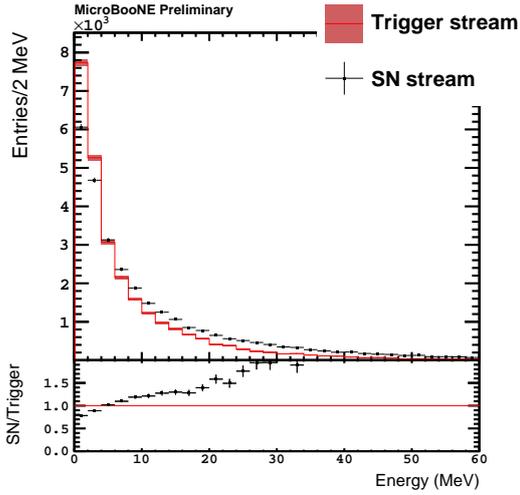
(b) Plane Y (Trigger Stream with ZS emulation).

**Figure 14.** Energy spectra of the Michel electron candidate radiative component reconstructed using only the collection plane. The markers and colors follow the same convention as figure 8.

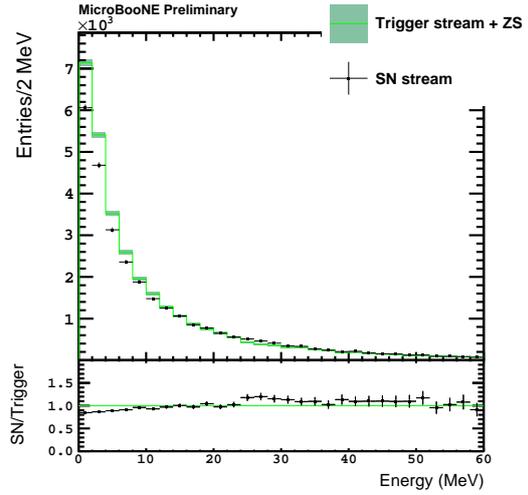
in ADC counts escape correction, these additional hits have small amplitudes. Using a data-driven simulation of the flipping bits we have evaluated the impact on the hit energy resolution to be  $\sim 10\%$ . Note these additional hits are effectively rejoined when summing over all the hits within the cluster to estimate its energy. Due to the currently dominant  $\sim 20\%$  resolution caused by the failure to reconstruct very low energy photons radiated by the electrons [12], we deem this additional contribution acceptable, even though the investigation of the origin of the flipping bits continues. In addition, not all the hits contributing to the  $0.1 - 0.2$  MeV peak come from flipping bits, as the zero suppression is found to also create additional radiative-like hits in that region, shown by the peak found in the Trigger Stream distributions with simulated zero suppression in figures 24(b), 25(b) and 25(d), for which flipping bits were not simulated.

## 7 The Continuous Readout Stream as a development platform

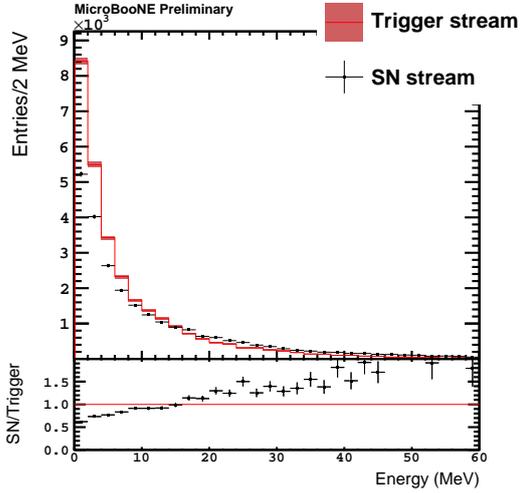
The MicroBooNE Continuous Readout Stream is the first realized stage toward a self-triggering LArTPC based on TPC information. Its successful operation enables both physics measurements as well as its use as a platform to develop and test TPC-based data selection algorithms for current and future detectors. For example, the zero-suppressed ROIs can be processed offline to extract features to be used as trigger primitives, which then would be clustered and processed through pattern recognition. See reference [14] for an example using the DUNE trigger primitive prescription [15]. A detailed discussion is beyond the scope of this work, but our analysis of Michel electrons already suggests that the energy bias observed in the reconstruction of the ionization component is an effect to be taken into account when defining the thresholds for such a trigger. The possibility of



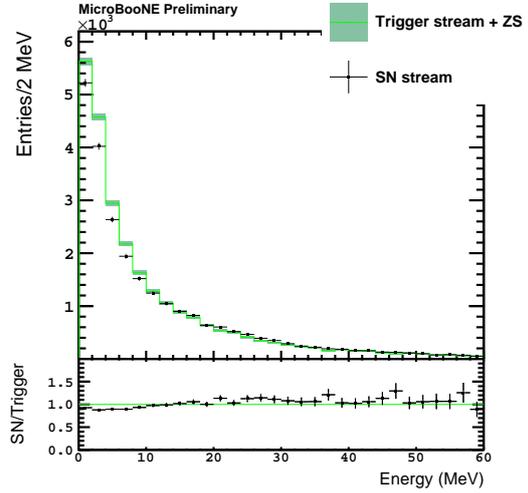
(a) Plane U (standard Trigger Stream).



(b) Plane U (Trigger Stream with ZS emulation).



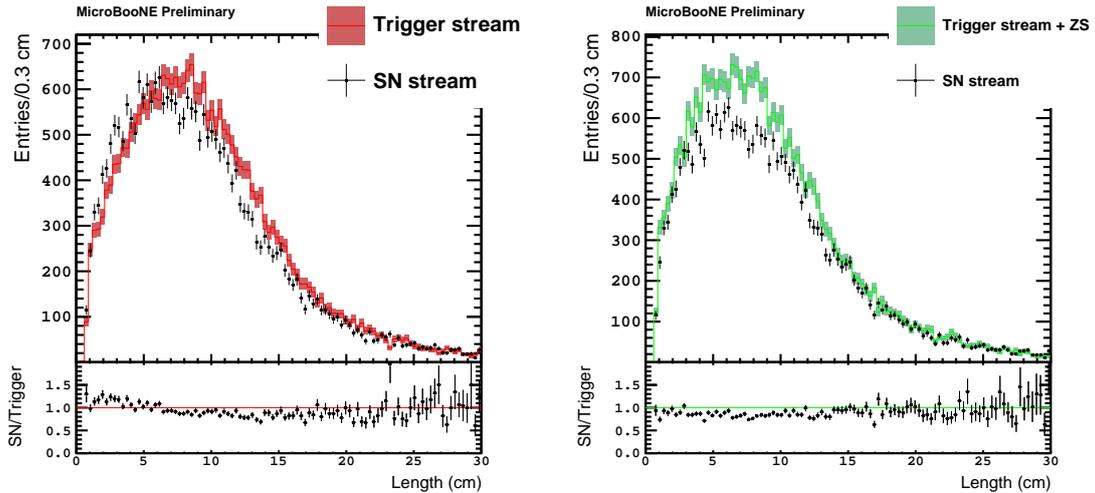
(c) Plane V (standard Trigger Stream).



(d) Plane V (Trigger Stream with ZS emulation).

**Figure 15.** Energy spectra of the Michel electron candidate radiative component reconstructed using only one of the induction planes. The top row shows the first induction plane (plane U), the bottom row shows the second induction plane (plane V). For each row, the markers and colors follow the same convention as figure 8.

triggering using any of the TPC planes, especially for enabling an online 3-D hit reconstruction based on matching time and wire coordinates, which can reduce the impact of noise and ambiguities when clustering, is very attractive. In this regard, we observe a large loss of reconstructed Michel electrons on the middle induction plane caused by the zero suppression, which would decrease the trigger efficiency for this plane. The middle induction plane is shielded by the first induction plane, resulting in signals with smaller amplitudes, which are more difficult to separate from the electronics noise, and symmet-



(a) Plane Y (standard Trigger Stream).

(b) Plane Y (Trigger Stream with ZS emulation).

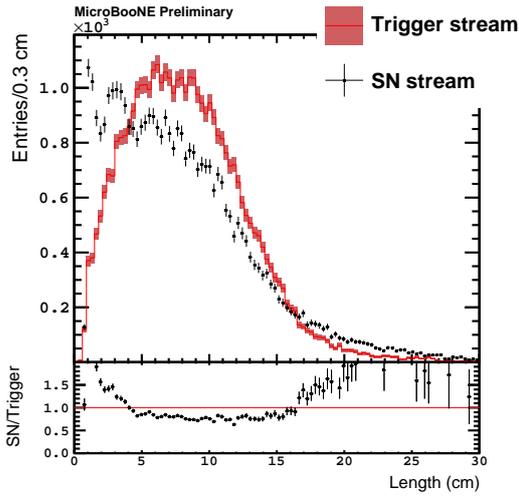
**Figure 16.** Length of the Michel electron candidate ionization component on the collection plane. The markers and colors follow the same convention as figure 8.

rical, which makes them prone to cancellation due to destructive interference. While the long induction rising edge of the signals on the first induction plane may be challenging to capture, the asymmetrical nature of the pulses on this plane (mostly unipolar negative as shown in figure 28) makes them more suitable for triggering. The inefficiency of a shielded plane can be partially mitigated if a peak-signal-to-noise ratio large enough is achieved (for reference, MicroBooNE U and V planes have 18.1 and 13.1, respectively [16]), which would allow the setting of lower zero-suppression thresholds.

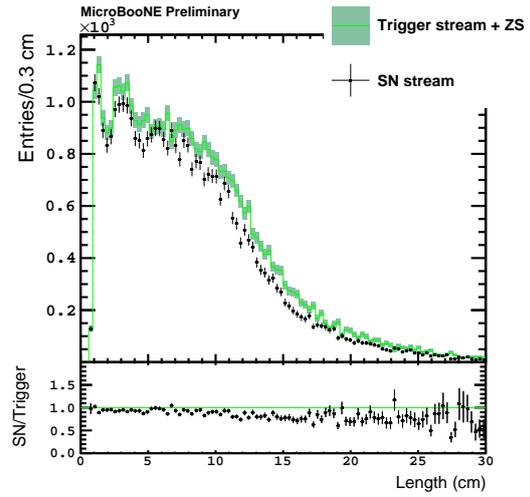
## 8 Conclusion

MicroBooNE has been the first liquid argon time projection chamber experiment to successfully commission and operate a continuous readout, opening a new way to look at data from the MicroBooNE detector. This novel data stream grants MicroBooNE the possibility to detect the burst of core-collapse supernova neutrinos using the SNEWS alert as a delayed trigger, expanding the physics program of the experiment. We defer the discussion about how to reconstruct and select those neutrinos to future work. After one year and a half of successful operation, during which we have tested multiple FPGA-based compression algorithms, the goal of reaching a stable compression factor  $\sim 80$  with sensitivity to supernova neutrino energies has been accomplished on all the TPC planes, including the induction planes where the pulse shapes are more challenging. The best performance was found for a zero-suppression algorithm that employs static baselines and (bandwidth-driven) individual thresholds for each channel.

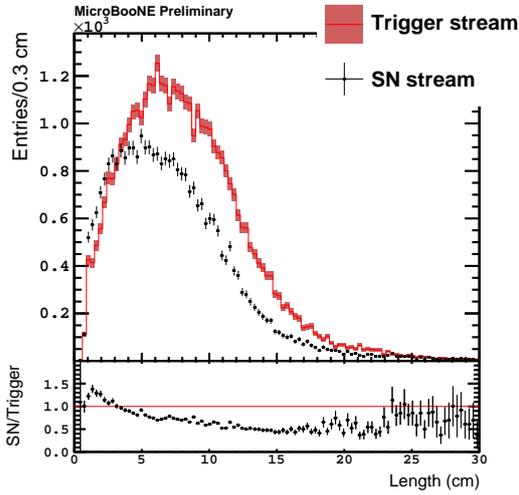
Based on the rate of Michel electrons from stopping cosmic-ray muons reconstructed in the Continuous Readout Stream, relative to the rate observed in the Trigger Stream, we



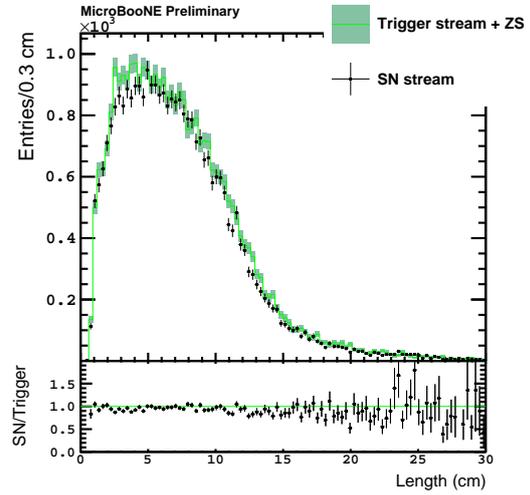
(a) Plane U (standard Trigger Stream).



(b) Plane U (Trigger Stream with ZS emulation).



(c) Plane V (standard Trigger Stream).

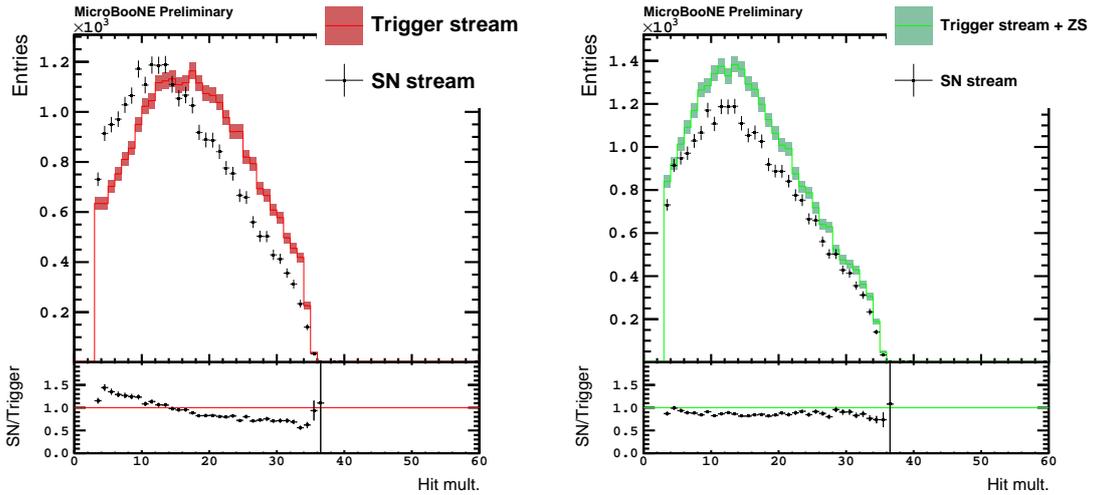


(d) Plane V (Trigger Stream with ZS emulation).

**Figure 17.** Length of the Michel electron candidate ionization component on the induction planes. The top row shows the first induction plane (plane U), the bottom row shows the second induction plane (plane V). For each row, the markers and colors follow the same convention as figure 8.

estimate a relative detection efficiency of  $(93.0 \pm 0.8)\%$  on the collection plane,  $(72.4 \pm 0.5)\%$  on the middle induction plane, and  $(94.8 \pm 0.7)\%$  on the first induction plane, where the uncertainties are only statistical, and do not include systematic effects such as the expected difference in muon rates between the two data samples due to seasonal variations.

An unexpected challenge on the Continuous Readout Stream is the appearance of flipping bits, affecting  $\sim 4\%$  of the ADC samples. Their origin is still being investigated, but their effect has resulted in an acceptable loss of resolution ( $\sim 10\%$ ) due to several mitigation steps in the offline reconstruction.



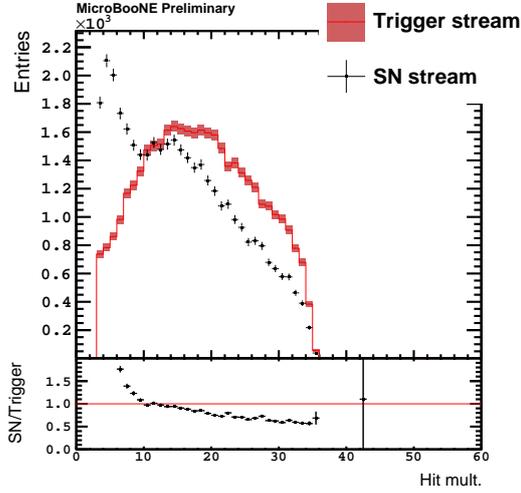
(a) Plane Y (standard Trigger Stream).

(b) Plane Y (Trigger Stream with ZS emulation).

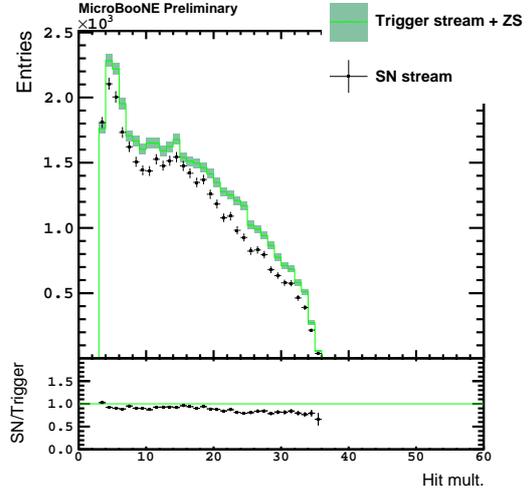
**Figure 18.** Hit multiplicity of the Michel electron candidate ionization component reconstructed using only the collection plane. The markers and colors follow the same convention as figure 8.

The continuous readout of a LArTPC is the first stage towards eventually developing a trigger based on the ionization patterns observed in the TPC. Analysis of the Michel electrons reveals that the zero suppression impacts the reconstruction of the ionization component of the electrons, resulting in reconstructed lower energies. This effect is well reproduced by our readout simulation. While this energy bias would have little impact on the measurement of the energy of supernova neutrinos, where the total energy is computed by including radiative-like components, it is an effect that would have to be accounted for when setting thresholds for an hypothetical TPC-based trigger that uses the zero-suppressed ionization clusters. We also observe a large inefficiency in the reconstruction of Michel electrons on the middle induction plane caused by the combination of the smaller signals due to screening by the first induction plane and the zero suppression.

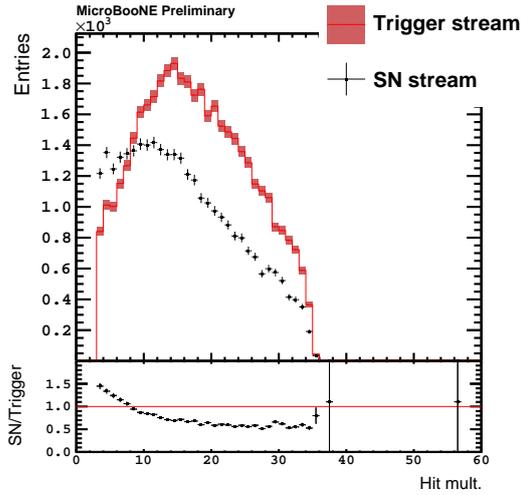
While this work focuses on using the Continuous Stream to detect low-energy electrons such as those produced by supernova neutrinos, this stream enables the study of other off-beam physics such as nucleon decay (proton decay, neutron-antineutron oscillation, etc). MicroBooNE cannot make competitive searches due to its size, by it can be used as a platform to develop analyses for future detectors or study backgrounds [17].



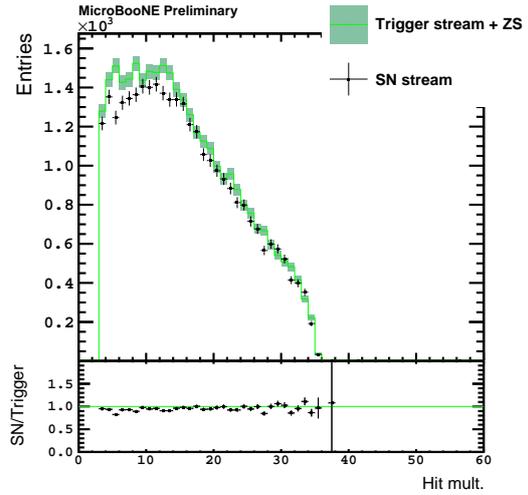
(a) Plane U (standard Trigger Stream).



(b) Plane U (Trigger Stream with ZS emulation).

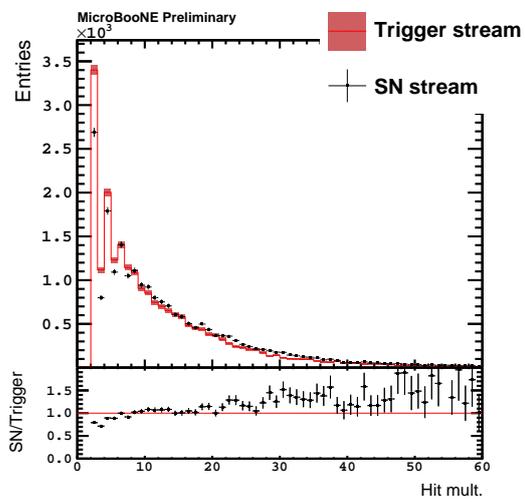


(c) Plane V (standard Trigger Stream).

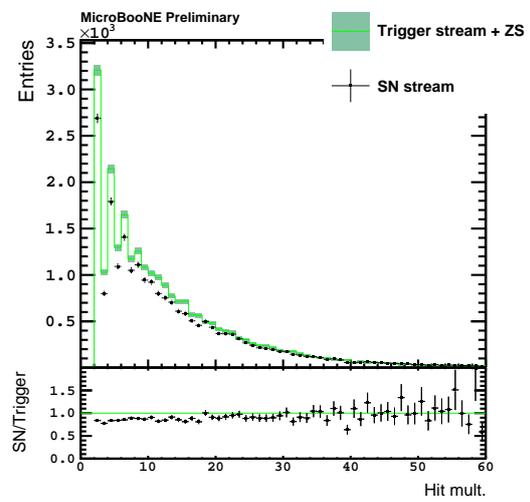


(d) Plane V (Trigger Stream with ZS emulation).

**Figure 19.** Hit multiplicity of the Michel electron candidate ionization component reconstructed using only one of the induction planes. The top row shows the first induction plane (plane U), the bottom row shows the second induction plane (plane V). For each row, the markers and colors follow the same convention as figure 8.

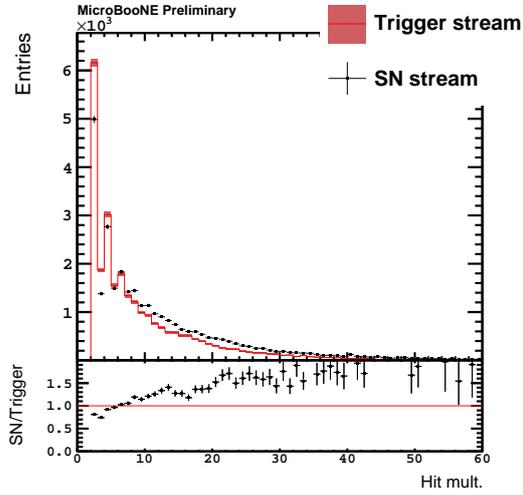


(a) Plane Y (standard Trigger Stream).

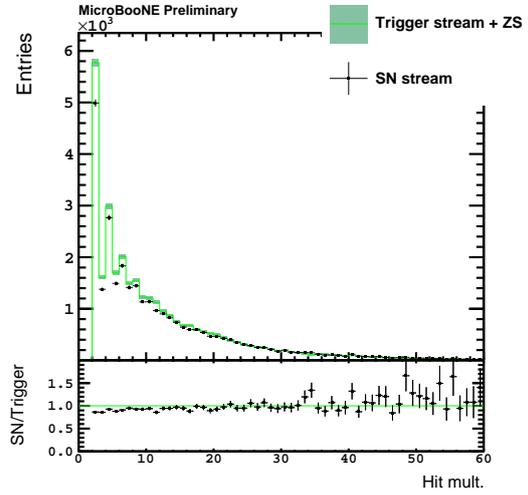


(b) Plane Y (Trigger Stream with ZS emulation).

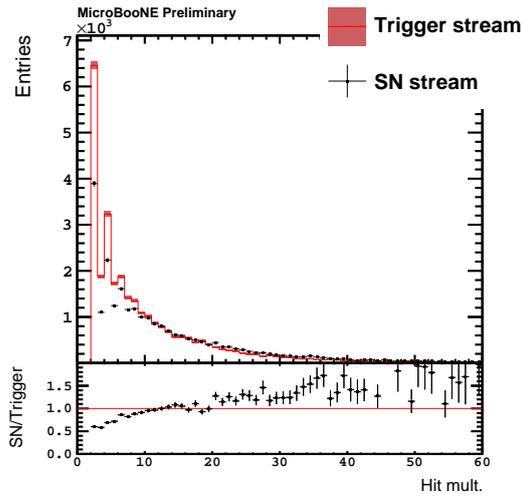
**Figure 20.** Hit multiplicity of the Michel electron candidate radiative component reconstructed using only the collection plane. The markers and colors follow the same convention as figure 8.



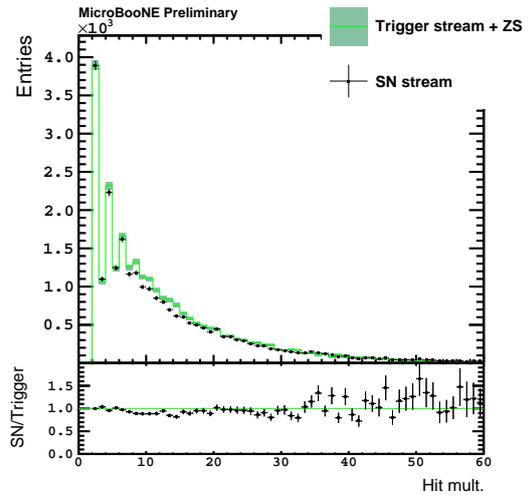
(a) Plane U (standard Trigger Stream).



(b) Plane U (Trigger Stream with ZS emulation).

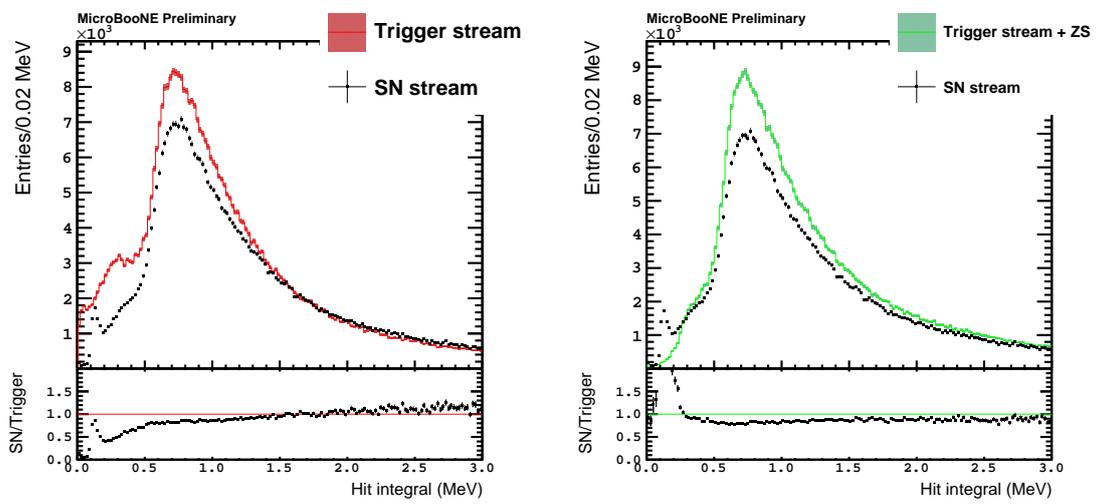


(c) Plane V (standard Trigger Stream).



(d) Plane V (Trigger Stream with ZS emulation).

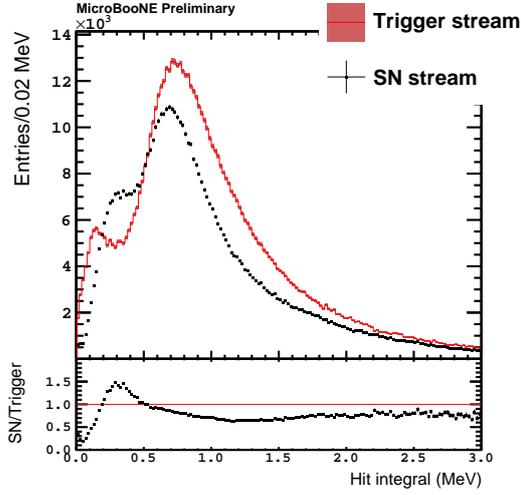
**Figure 21.** Hit multiplicity of the Michel electron candidate radiative component reconstructed using only one of the induction planes. The top row shows the first induction plane (plane U), the bottom row shows the second induction plane (plane V). For each row, the markers and colors follow the same convention as figure 8.



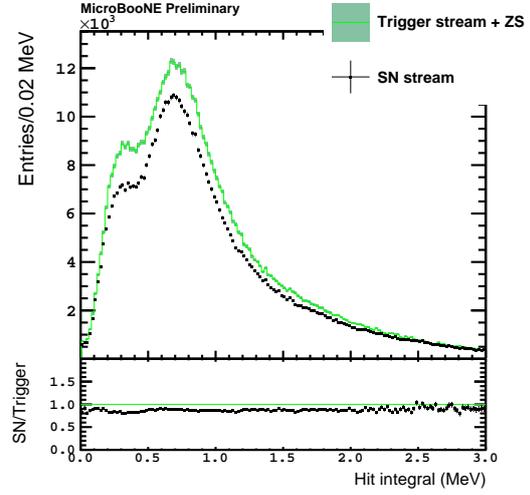
(a) Plane Y (standard Trigger Stream).

(b) Plane Y (Trigger Stream with ZS emulation).

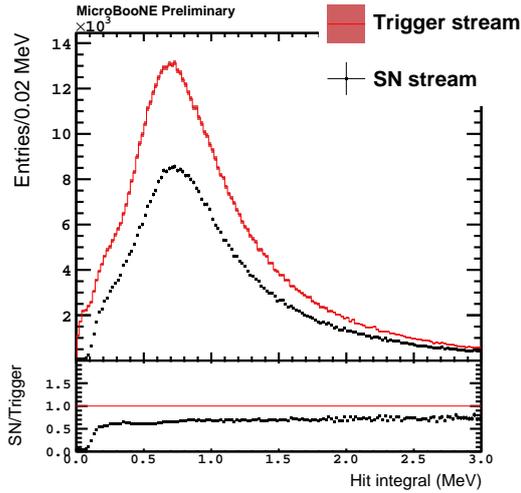
**Figure 22.** Hit energy spectra of the Michel electron candidate ionization component reconstructed using only the collection plane. The markers and colors follow the same convention as figure 8. The peak at 0.1 – 0.2 MeV found in the Continuous Readout Stream (SN stream) data is dominated by additional hits caused by flipping bits affecting the ADC words.



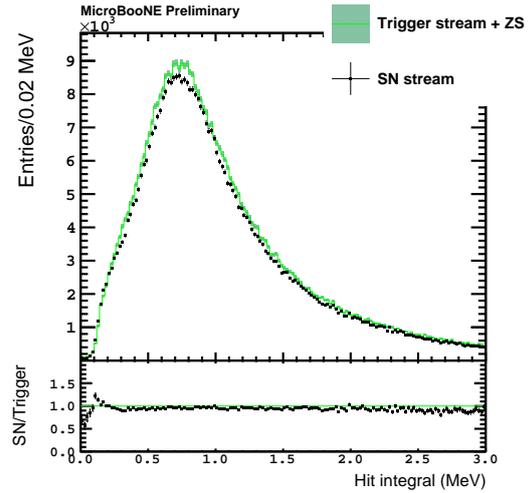
(a) Plane U (standard Trigger Stream).



(b) Plane U (Trigger Stream with ZS emulation).

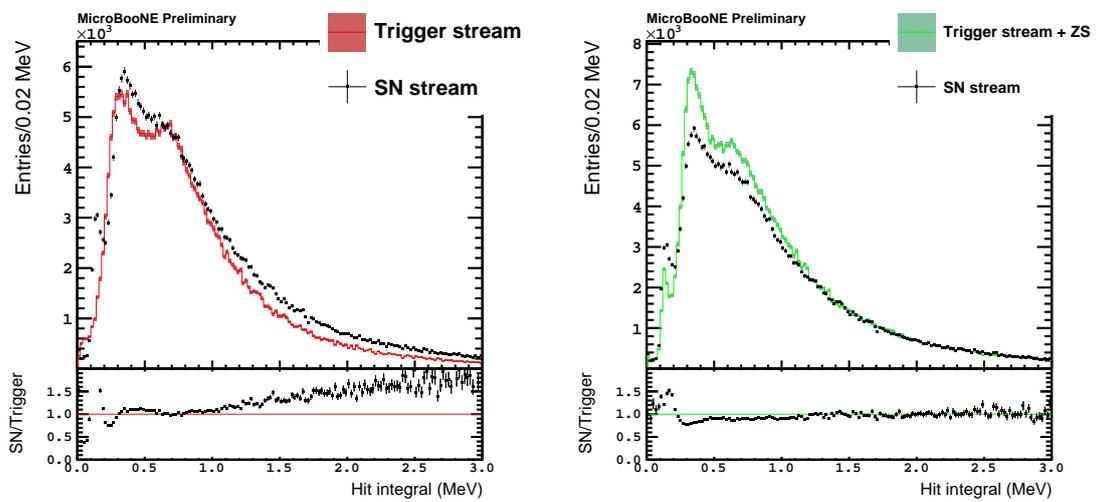


(c) Plane V (standard Trigger Stream).



(d) Plane V (Trigger Stream with ZS emulation).

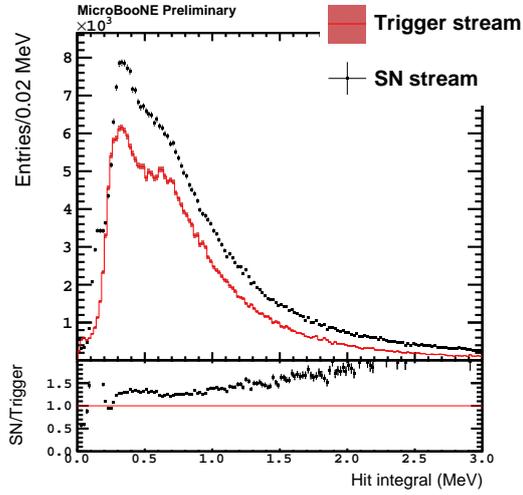
**Figure 23.** Hit energy spectra of the Michel electron candidate ionization component reconstructed using only one of the induction planes. The top row shows the first induction plane (plane U), the bottom row shows the second induction plane (plane V). For each row, the markers and colors follow the same convention as figure 8. The Continuous Readout Stream (SN stream) data in (a) shows a low-energy tail above the Trigger Stream reference caused by the incomplete acquisition of the slow-rising pulses of the first induction plane resulting from the limited number of samples below threshold available in the FPGA implementation of the zero suppression.



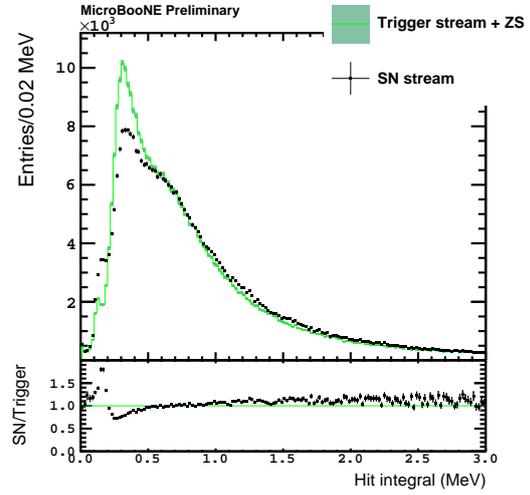
(a) Plane Y (standard Trigger Stream).

(b) Plane Y (Trigger Stream with ZS emulation).

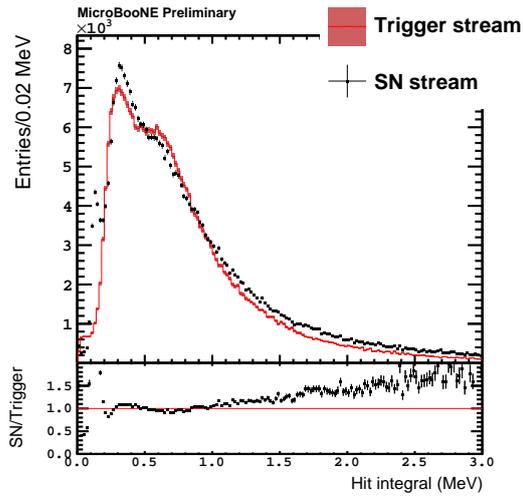
**Figure 24.** Hit energy spectra of the Michel electron candidate radiative component reconstructed using only the collection plane. The markers and colors follow the same convention as figure 8. The peak at 0.1 – 0.2 MeV in the Continuous Readout Stream (SN stream) data is found to be caused by a combination of the zero suppression and flipping bits affecting the ADC words.



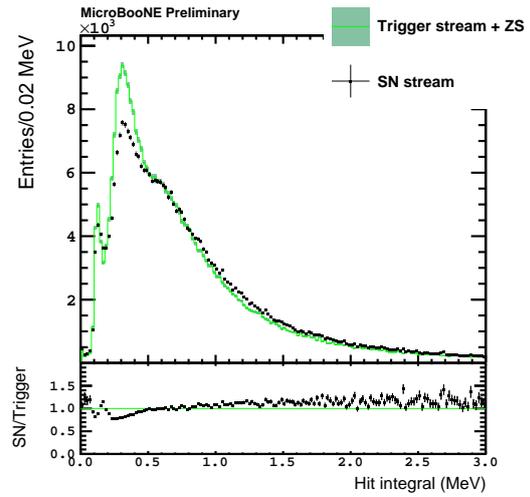
(a) Plane U (standard Trigger Stream).



(b) Plane U (Trigger Stream with ZS emulation).



(c) Plane V (standard Trigger Stream).



(d) Plane V (Trigger Stream with ZS emulation).

**Figure 25.** Hit energy spectra of the Michel electron candidate radiative component reconstructed using only one of the induction planes. The top row shows the first induction plane (plane U), the bottom row shows the second induction plane (plane V). For each row, the markers and colors follow the same convention as figure 8.

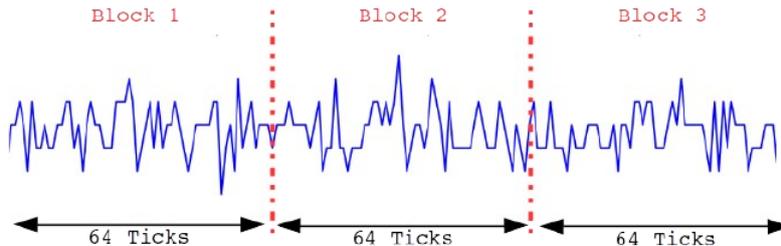
## References

- [1] DUNE collaboration, B. Abi et al., *The DUNE Far Detector Interim Design Report Volume 1: Physics, Technology and Strategies*, [1807.10334](#).
- [2] A. Bueno, I. Gil Botella and A. Rubbia, *Supernova neutrino detection in a liquid argon TPC*, [hep-ph/0307222](#).
- [3] MICROBOONE collaboration, R. Acciarri et al., *Measurement of cosmic-ray reconstruction efficiencies in the MicroBooNE LArTPC using a small external cosmic-ray counter*, *JINST* **12** (2017) P12030, [[1707.09903](#)].
- [4] P. Antonioli et al., *SNEWS: The Supernova Early Warning System*, *New J. Phys.* **6** (2004) 114, [[astro-ph/0406214](#)].
- [5] MICROBOONE collaboration, R. Acciarri et al., *Design and Construction of the MicroBooNE Detector*, *JINST* **12** (2017) P02017, [[1612.05824](#)].
- [6] “AD9222 Octal, 12-Bit, 40/50/65 MSPS Serial LVDS 1.8 V A/D Converter.” <http://analog.com/media/en/technical-documentation/data-sheets/AD9222.pdf>, 2011.
- [7] “Stratix III Device Handbook.” [https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/stx3/stratix3\\_handbook.pdf](https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/stx3/stratix3_handbook.pdf), 2011.
- [8] D. A. Huffman, *A method for the construction of minimum-redundancy codes*, *Proceedings of the IRE* **40** (Sep., 1952) 1098–1101.
- [9] “larsoft 6.80.00. LArSoft Collaboration — Software for LArTPCs.” <https://larsoft.org/>, 2019.
- [10] “uBooNE code.” <https://cdcvns.fnal.gov/redmine/projects/uboonecode>, 2019.
- [11] MICROBOONE collaboration, C. Adams et al., *Ionization electron signal processing in single phase LArTPCs. Part I. Algorithm Description and quantitative evaluation with MicroBooNE simulation*, *JINST* **13** (2018) P07006, [[1802.08709](#)].
- [12] MICROBOONE collaboration, R. Acciarri et al., *Michel Electron Reconstruction Using Cosmic-Ray Data from the MicroBooNE LArTPC*, *JINST* **12** (2017) P09014, [[1704.02927](#)].
- [13] MICROBOONE collaboration, C. Adams et al., *Calibration of the charge and energy response of the MicroBooNE liquid argon time projection chamber using muons and protons*, [1907.11736](#).
- [14] C. E. Hinrichs, “Development of a Michel Electron TPC-Based Trigger for LArTPC Detectors.” [https://www.nevis.columbia.edu/reu/2019/reports/Hinrichs\\_report.pdf](https://www.nevis.columbia.edu/reu/2019/reports/Hinrichs_report.pdf), 2019.
- [15] DUNE collaboration, B. Abi et al., *The DUNE Technical Design Report, In preparation* .
- [16] MICROBOONE collaboration, R. Acciarri et al., *Noise Characterization and Filtering in the MicroBooNE Liquid Argon TPC*, *JINST* **12** (2017) P08003, [[1705.07341](#)].
- [17] J. E. T. Hewes, *Searches for Bound Neutron-Antineutron Oscillation in Liquid Argon Time Projection Chambers*. PhD thesis, Manchester U., 2017. 10.2172/1426674.

## Appendix

### A Dynamic-baseline algorithm

This appendix describes the algorithm implemented in the FEM FPGA to estimate the baseline in real time. The estimation of the channel baseline is performed using 3 contiguous blocks of 64 samples each (each block corresponding to  $32 \mu\text{s}$  of the waveform, see figure 26). A rounded mean ADC value for each block is computed by summing the ADC values of all 64 samples and then dropping the 6 least significant bits (equivalent to an integer division by 64). A truncated ADC variance for each block is computed by summing the squared differences between the ADC value of each sample and the rounded mean computed above, and then dropping the 6 least significant bits. If the absolute value of the difference between an ADC sample and the rounded mean is greater than or equal to 63 ADC counts, the contribution of that sample to the variance is fixed to 4095 ADC counts to prevent arithmetic overflows. The rounded mean and the truncated variance of each block are compared to the ones from the other two blocks to avoid choosing a baseline corresponding to an actual signal or a non-representative fluctuation. If the three rounded mean differences and the three truncated variance differences between blocks are within the configurable tolerance values, the rounded mean of the central block is taken as the new baseline and applied for zero suppression for ADC samples beginning after the third block. The local baseline estimation is applied continuously as a sliding window from the beginning of the run, dropping the oldest 64-sample block and adding a newer block. The baseline tolerance parameters (rounded mean and truncated variance differences) are configured at FEM level (in groups of 64 channels). If the baseline conditions are not satisfied (e.g. the difference between blocks never meets the tolerance) in a channel, the zero suppression algorithm does not produce output data for that channel.



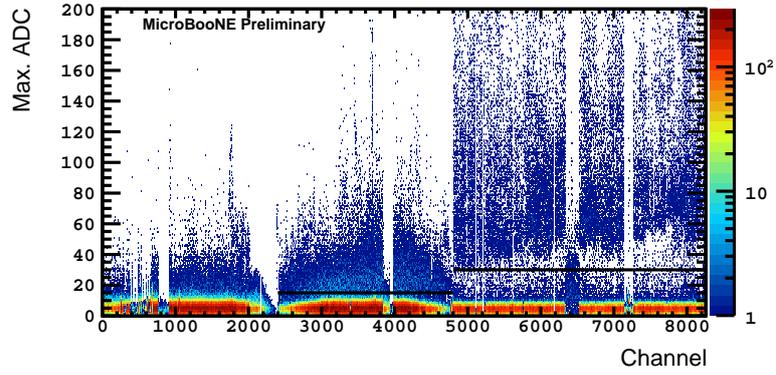
**Figure 26.** Cartoon showing how the waveform is split into blocks for the dynamic baseline estimation.

### B Physics-driven plane-wide zero-suppression thresholds

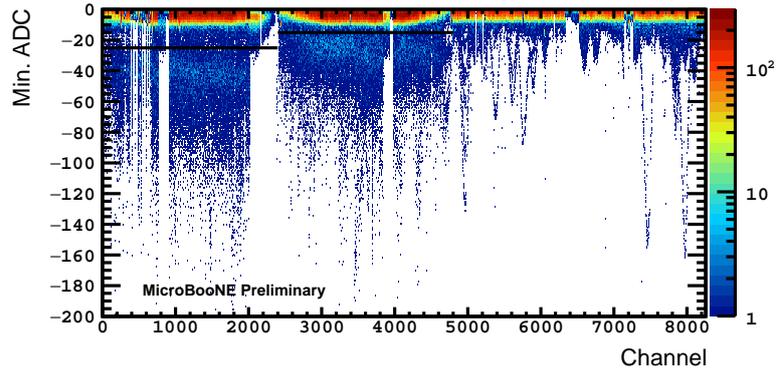
This appendix describes the first method used to determine the settings for the zero suppression of TPC waveforms. This method aims at establishing a single threshold for each TPC plane to separate signals (dominated by cosmic-ray muons) from electronics noise.

Using events from off-beam zero-bias triggers from the Trigger Stream during previous runs with good electron lifetime, a first pass with a software emulation of the zero-suppression algorithm with a threshold of  $\pm 5$  ADC counts (chosen arbitrarily to reduce the raw data while keeping the signals and some noise fluctuations) and dynamic baseline was done. ADC spectra of the maximum and minimum ADC values found in each zero-suppressed waveform, after offline baseline subtraction using the first sample, were produced as shown in figure 27. For these spectra, the peak closest to the origin is interpreted as noise, while the peak furthest from the origin is interpreted as the signal, dominated by near-MIP cosmic-ray muons. The ADC spectra also reflect the polarity of the typical signals within each plane (see figure 28). The FPGA firmware only admits one threshold value per channel and its polarity. For plane U, the negative valley was used to set the threshold at  $-25$  ADC counts (unipolar negative threshold). For plane V, the two valleys are found at approximately the same absolute ADC value. Hence, a threshold of  $\pm 15$  ADC counts (bipolar threshold) was established. For plane Y, the valley in the positive ADC distribution was used to set the threshold at  $+30$  ADC counts (unipolar positive threshold).

These amplitude thresholds were tested with the dynamic-baseline zero suppression, with the baseline tolerance parameters set to 2 ADC counts for the rounded mean and  $3 \text{ ADC}^2$  counts for the truncated variance based on the MicroBooNE noise levels [16]. In the end, these amplitude thresholds were deprecated in favor of the bandwidth-driven thresholds described in section 4, which allow the recording of more data.

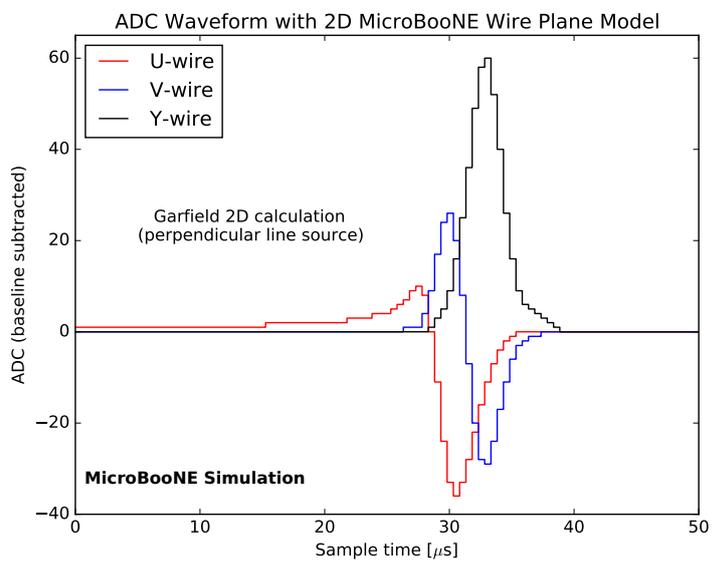


(a) Maximum ADC values.



(b) Minimum ADC values.

**Figure 27.** Maximum (a) and minimum (b) ADC values in the regions of interest produced by the emulation of the zero-suppression algorithm with a  $\pm 5$  ADC count threshold to decimate the Continuous Stream raw data, as a function of channel number. The horizontal black lines mark the location of the chosen thresholds for the three TPC planes:  $-25$  ADC counts for plane U (channels  $0 - 2399$ ),  $\pm 15$  ADC counts for plane V (channels  $2400 - 4799$ ) and  $+30$  ADC counts for plane Y (channels  $4800 - 8255$ ).



**Figure 28.** The digitized signals from a central wire from each plane that are induced by an ideal MIP track in a 2-D model of the MicroBooNE TPC. From [16].