Denoising for Supernova Detection with DUNE

Brenton Outangoun

Advisor: Jovan Mitrevski

August 2023

ABSTRACT

The Deep Underground Neutrino Experiment (DUNE) plays a crucial role in detecting neutrinos, which are key to understanding supernovae within the larger field of multi-messenger astronomy. Processing the vast amounts of data, filled with important neutrino signals and distracting background noise, is a significant challenge. Current methods, which mainly rely on Central Processing Units (CPUs), face issues of slow processing and high energy use, especially as data volumes grow. To improve this, we propose a new approach using In-Storage Computing, Field Programmable Gate Arrays (FPGAs), and 2D Convolutional Neural Networks (CNNs). FPGAs, due to their flexible design, can handle specific tasks efficiently. By moving much of the processing directly to where data is stored using FPGAs, we can cut down on data movement, resulting in faster data access, quicker processing, and lower energy consumption. With the AMD/Xilinx Vitis platform, we've optimized FPGAs to prepare data for the 2D CNNs, which act as filters to separate valuable signals from noise. Our early tests, including simulations and visual analyses, show the potential of our system to prepare the right signals from complex data sets. In conclusion, our combination of FPGA-based processing with 2D CNNs offers a

promising direction for supernova detection, with increased speed and accuracy. Future work will focus on real-world testing of the system, further refining our 2D CNN model, and potentially exploring the integration of a 1D CNN for a deeper analysis of space events.

ACKNOWLEDGEMENTS

This work was supported in part by the U.S. Department of Energy, Office of Science, Office of Workforce Development for Teachers and Scientists (WDTS) under the Community College Internships Program (CCI).

I. INTRODUCTION

When a star reaches the end of its life, it erupts into a supernova, releasing vast numbers of neutrinos. These small particles, known for their ability to pass through most matter, act as essential informers of cosmic events. DUNE is an advanced facility that is able to capture and analyze these neutrinos, placing it at the forefront of supernova discovery. DUNE's design allows it to efficiently capture these particles, providing insights into their origins, particularly from supernovae. When DUNE detects potential supernova signals, it acts as an alert system for a wide array of observatories worldwide. Together, these observatories employ a method called multi-messenger astronomy to collaboratively study the event.

However, the sheer volume of data, comprising both neutrino signals and background noise, can be daunting to process. This rich data stream contains crucial scientific insights, but its size and complexity pose significant processing challenges. Efficiently handling this data is critical. Delays or inefficiencies might cause scientists to miss the early moments of a supernova, which could hold essential clues about such cosmic events. Therefore, enhancing data processing efficiency remains a pressing need to ensure timely responses to supernova detections.

II. CURRENT DATA PROCESSING METHOD

Modern data processing and Data Acquisition (DAQ) tasks have historically been anchored around the use of the Central Processing Unit (CPU). This traditional methodology, which forms the basis of the majority of computational undertakings, unfortunately is not without its setbacks. One glaring concern arises in the realm of processing efficiency. As tasks become more complex and data-driven, the CPU-centric model struggles to maintain optimal performance levels.

Another closely tied challenge pertains to energy consumption. Every transaction involving the transfer of data between the storage system and the CPU not only introduces a discernible amount of latency but also consumes energy — often more than is sustainable in computationally heavy scenarios. This energy expenditure is not just a matter of operational cost but also environmental responsibility. With DUNE being built far underground, dissipating energy will also be a problem. The lack of ventilation is just one of the many issues that would come with this method.

Additionally, with the exponential growth in data volumes, driven by our increasingly digital world, these limitations and inefficiencies risk amplifying. The cavern that holds DUNE limits the amount of bandwidth available to transfer to the surface, making sending data to the surface for processing quite slow. The continual shuttling of vast datasets back and forth can lead to system bottlenecks, which in turn can hamper both the speed and quality of data processing, especially in critical, time-sensitive applications.

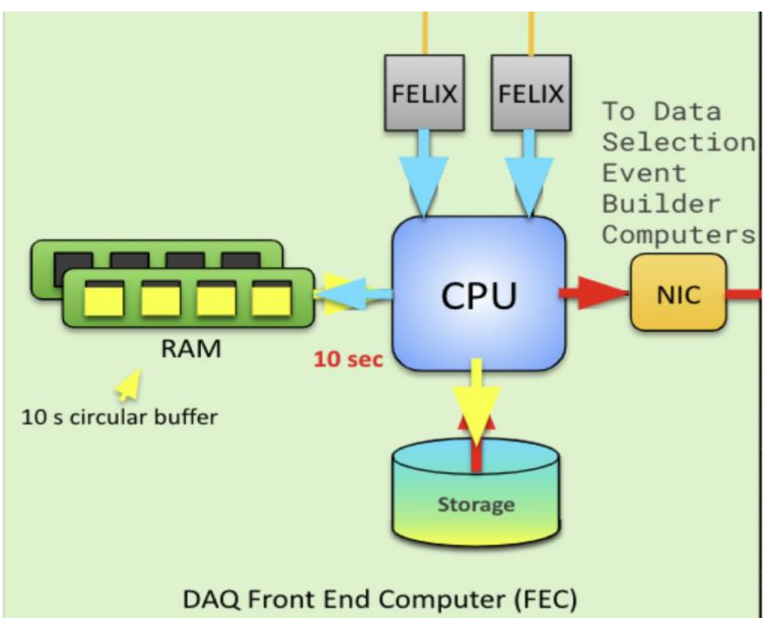

Figure 1: CPU data processing method [1]

III. PROPOSED METHOD: ADVANCED IN-STORAGE COMPUTING WITH FPGAS

In-Storage Computing, specifically harnessing the power of FPGAs, is a solution to this problem. Field-Programmable Gate Arrays, or FPGAs, are integrated circuits that stand out due to their unique adaptability. Their customizable design allows them to be tailored to specific tasks, offering an unmatched level of flexibility in hardware operations. Furthermore, FPGAs can perform multiple operations in parallel, making them highly suited for intricate computational demands. By moving a significant portion of the processing load from the CPU to an FPGA, which is integrated directly within the storage system, we open the door to a more streamlined computational process. This reconfiguration means that processing occurs closer to where the data is actually stored.

Such proximity offers a plethora of advantages. Firstly, it reduces the constant shuttling of data between storage and the CPU. This direct approach leads to a noticeable decrease in latency, ensuring that data is processed and accessed more quickly. Secondly, there's a tangible increase in processing speed. Without the added lag from data transfer times, systems can perform computations more immediately, directly accessing and processing data from its source. Lastly, this computational approach has the added benefit of being energy efficient. Less movement of data means less power consumed in the transfer process. This not only decreases operational costs but also moves us toward a more sustainable and environmentally friendly computational model. By strategically redistributing processing tasks and leveraging the capabilities of FPGAs within storage systems, we find ourselves on the verge of a promising data processing method, leading to significant improvements in latency, speed, and energy usage.

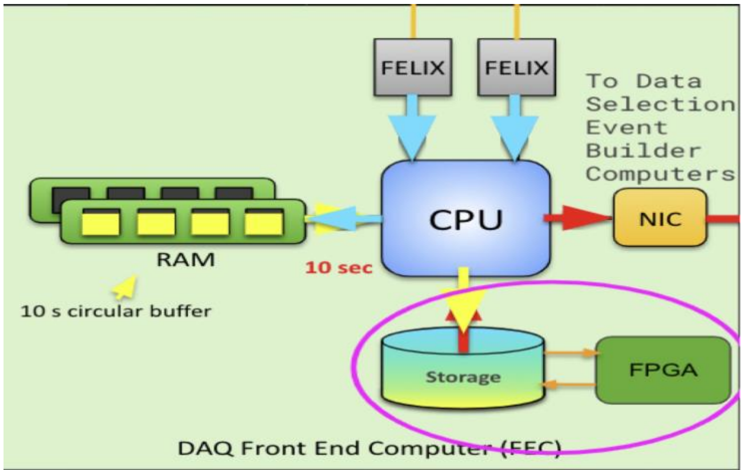

Figure 2: In storage computing with FPGAs [1]

IV. PROGRAMMING FPGAs WITH 2D CNNS

To fully leverage the capabilities of these FPGAs, we utilize AMD/Xilinx Vitis, a specialized development environment designed for FPGA programming. This environment enables the FPGAs to operate as high-speed accelerators. The kernel code we develop for the FPGA essentially prepares data and runs that data through a 2D CNN. A 2D CNN is a powerful, complex machine learning model capable of interpreting and analyzing multidimensional data. This makes it a strong choice for the task of assisting in the detection and interpretation of the signals we aim to identify.

For our purposes, the 2D CNN is used as a background rejector to help process the raw data collected by the sensors of the DUNE and give us signals. To make the implementation of the 2D CNN on our FPGAs easier, we utilize High-Level Synthesis for Machine Learning (hls4ml), a tool specifically designed to convert machine learning models into hardware-ready code. Using hls4ml, the 2D CNN is fine-tuned and optimized for deployment on the FPGA, ensuring efficient resource use and optimal performance.

Once operational, the 2D CNN, guided by the kernel code, sifts through the processed data, identifying and extracting the desired signals necessary for supernova identification. This combination of sophisticated hardware and machine learning techniques presents a promising future for supernova detection.

V. RESULTS

In the pursuit of enhancing data analysis capabilities for supernova detection, our kernel code yielded noteworthy outcomes. Optimized to closely interface with the FPGA, the code demonstrated its effectiveness in streamlining the operations of the 2D CNN. Our first significant step was the successful simulation and synthesis of the data preparation kernel code into FPGA programmable logic using Vitis High-Level Synthesis (HLS). During the data processing phase, we identified three distinct collection planes: the U, V, and Z planes. For the purpose of our research and ensuring comprehensive data interpretation, our focus centered on processing both sides of the Z plane. This specific processing approach was key to validating our computational model's strength.

With the foundational results established, our subsequent focus was directed towards refining the kernel code. We resolved multiple violations, leading to a significant decrease in the computational resources the code needed. In addition to this, to provide a tangible representation of our kernel code's capabilities, we leveraged data visualization tools. This methodical approach enriched our insights into the kernel code's performance metrics.

Further fortifying our confidence in the kernel code, we subjected it to a representative data file—mirroring the typical inputs it would be tasked with in genuine scenarios. Impressively, the code consistently generated the anticipated signals, emphasizing its readiness to merge seamlessly with the 2D CNN, even in the face of data complexities.

Transitioning our attention to the 2D CNN, we were able to use a prototype model in our code. The integration of this with our main kernel code was an intricate process, culminating in a comprehensive kernel program adeptly able to process DUNE data for the 2D CNN. Validating its readiness, the system advanced through the simulation, synthesis, and co-simulation stages in Vitis HLS.

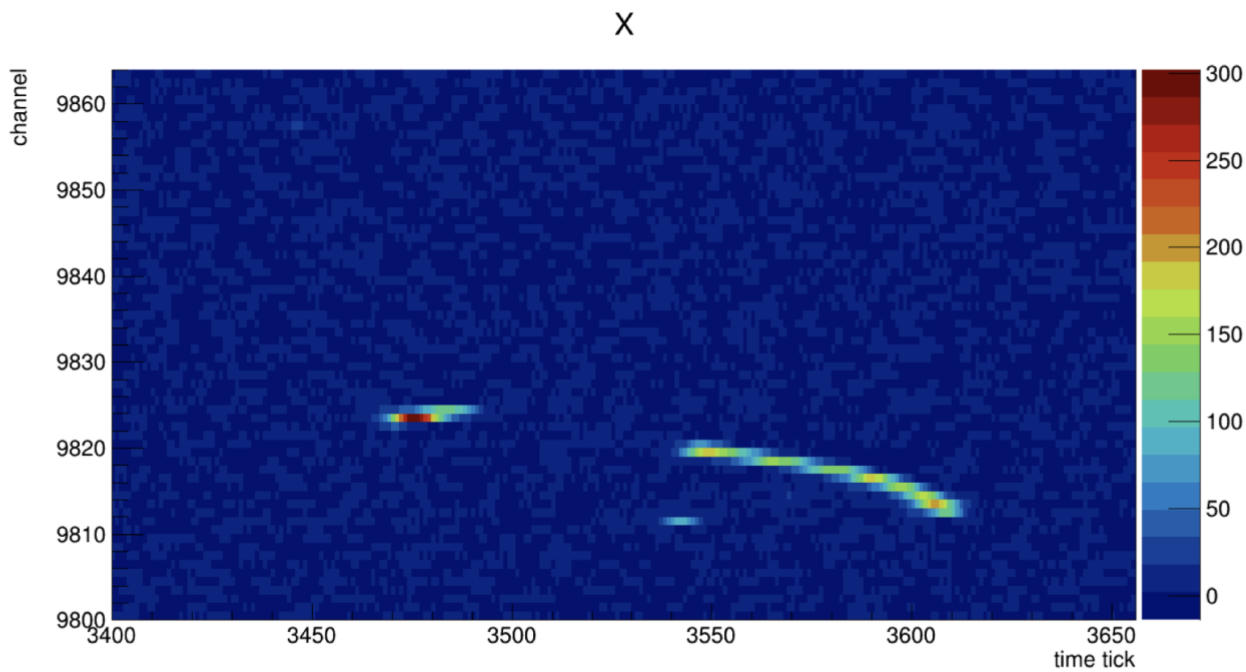

Figure 3: Event signal found by kernel code.

VI. CONCLUSION AND FUTURE DIRECTIONS

In summarizing our research, we've found that using FPGA-based In-Storage Computing, together with advanced programming and 2D Convolutional Neural Networks (CNNs), provides a new direction for supernova detection. This method differs from our usual techniques, aiming to improve both the efficiency and accuracy of our detections. Our current results suggest that this combination holds promise. They show potential advantages and point out areas needing more work.

Looking ahead, we plan to increase our simulations and make further adjustments to the kernel code. While simulations give us a good starting point, real-world testing is crucial. Thorough testing on actual FPGA platforms is needed to better gauge our system's real-world applicability. There are also many optimizations that can be added to make the code much more efficient. In addition to this, a Vitis project is needed to merge host and kernel codes for better coordination. This host code will be written in OpenCL and will be able to call our kernel code.

Our initial 2D CNN model is just a prototype, and we are currently developing a more refined version that can effectively sift through background noise. This refined version is nearly ready to be used with our kernel code. In addition to this, the data needs to go through another step after being passed through our 2D CNN. A potential 1D CNN is being built and will dive deeper into the results, offering a more thorough analysis.

On a broader scale, this research doesn't just focus on supernova detection. These methods may be used to improve large-scale data analysis in other areas too. Better detection techniques can lead to new research opportunities and a deeper understanding of space. In closing, our research is about bridging advanced technology and space study in a practical way. The 2D CNN is just the beginning of our data processing steps, but it sets the stage for a more detailed understanding of space events in the future.

REFERENCES

[1] Benjamin Hawks et al., Exploring FPGA in-storage computing for Supernova burst detection in LArTPCs.