



# High performance storage technologies for the DAQ of the DUNE experiment

Adam Abed Abud (Univ. of Liverpool / CERN)

October 11, 2021
tCSC 2021

# The DUNE experiment

## Detector

- Long baseline neutrino detector located 1300 km from the source and 1.5 km underground
- Detector composed of a time projection chamber with a photon detector system
- 4 module detectors with 17-kton LAr
- Installation and data taking starting from 2025/2026

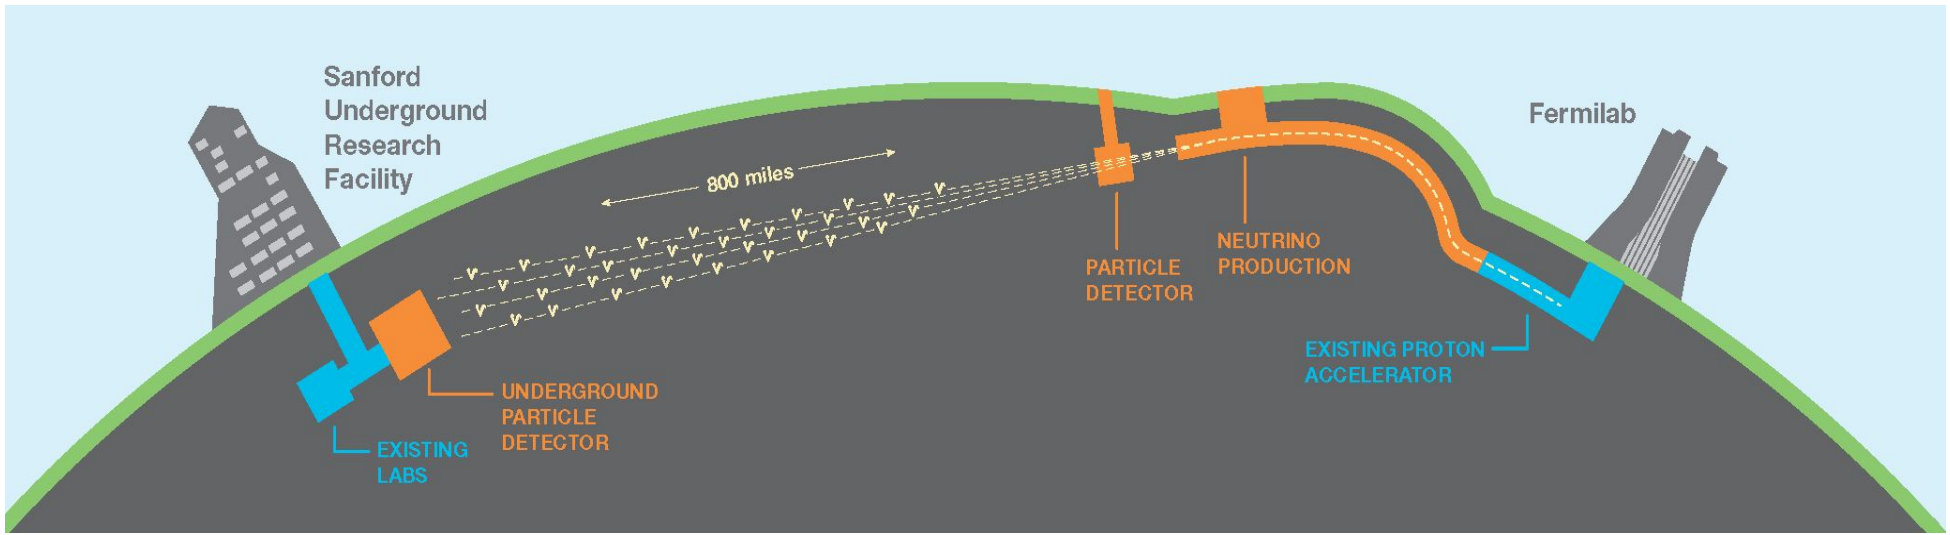

# Objective and requirements

- **Goal**:
  - One of the physics goals of DUNE is to store Supernova Neutrino Burst (SNB) events
  - Detection of **rare**, **low energy** and **distributed** signatures
- **Requirements:**
  - Sustain **high data rates**: 10 GB/s for each DAQ readout unit
  - **High data volumes:** minimum capacity of 1 TB per readout unit (~100 s of recording time)
  - 150 total number of readout units
- **Solution**:
  - Use modern storage devices to meet the performance demands of the DUNE experiment

# Supernova Neutrino buffer

## Persistent memory devices

- The supernova storage buffer needs to store **critical data** and sustain a **high bandwidth**
- Ideal technology candidate to store the data: **persistent memory devices**
    - Memory devices capable of permanently storing data
    - High performance devices: O(10) GiB/s
    - Large storage capacity and endurance

- Implementation of the supernova storage buffer with persistent memory devices could be a better solution than DRAM devices or solid-state drives

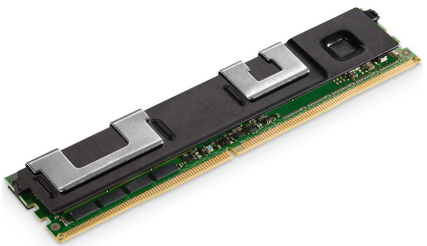

Device with DRAM-like properties!

# Results with PMEMs

- Use of persistent memory devices:
  - Measured the write throughput of persistent memory devices with a workload similar to the DUNE use case
  - Sustained 80 % of target throughput !
- Optimize software stack:
  - Use of efficient modern libraries (PMDK)
    - Not enough performance
  - Low level application: create memory mapped files and persist them using the MOVNTI instruction (~20% gain)
  - Understand the performance bottlenecks of the SW (use of VTune Analyzer)

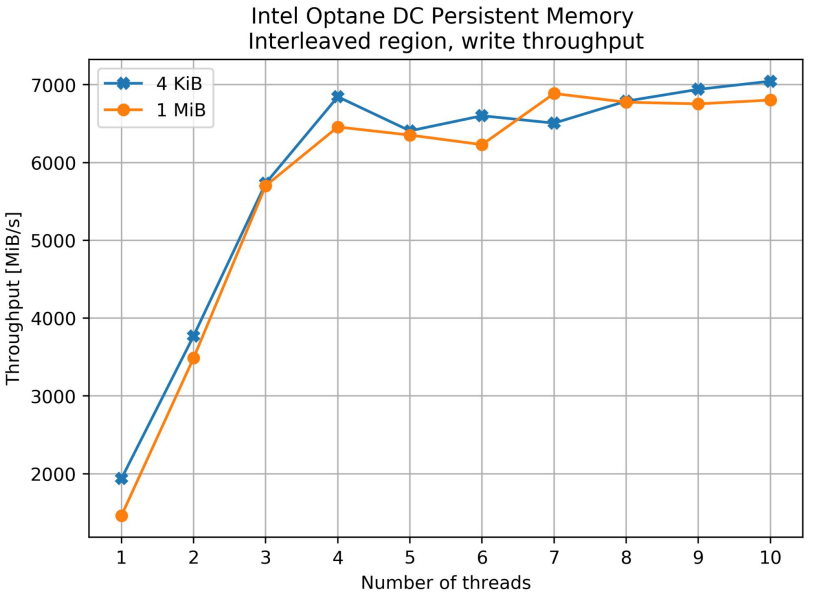

# Micron X100 as a high-performance drive

- The Micron X100 is an example of PCIe  NVMe SSD based on the 3D XPoint technology
- **Synthetic benchmarks**
  - Evaluate the sequential write throughput with  a single thread
  - Perform asynchronous file operations with using native kernel AIO
  - Use of o-direct flag to bypass operating system cache
  - Max. write throughput achieved is **8.5 GiB/s**
    - Compatible with hardware specification
- Micron X100 device is capable of sustaining 85% the throughput of a single readout unit
- Promising technology:
  - **Example**: combine 3 Micron X100 SSDs for 4 readout units

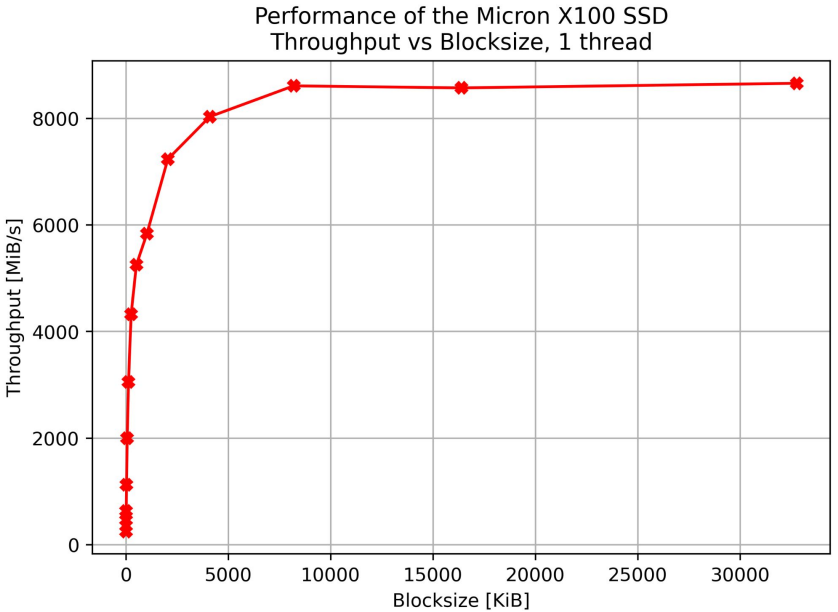

# Other solutions

- Take advantage of PCIe 4 storage adapters
    - Connect multiple SSD drives together: up to 4 x PCIe 4.0/3.0 M.2 Slots
- Newly available fast NVMe SSDs provide a seq writing throughput of ~4.5 GB/s
- Connecting 4 drives with poor man's raid provides a theoretical throughput greater than 10 GB/s
    - Interesting solution for the DUNE DAQ local storage buffer
    - Preliminary tests show good performance results