Computing for LUX-ZEPLIN (LZ)

Maria Elena Monzani

U.S. DEPARTMENT OF ENERGY | Stanford University

SLAC NATIONAL ACCELERATOR LABORATORY

Direct DM searches in 2021

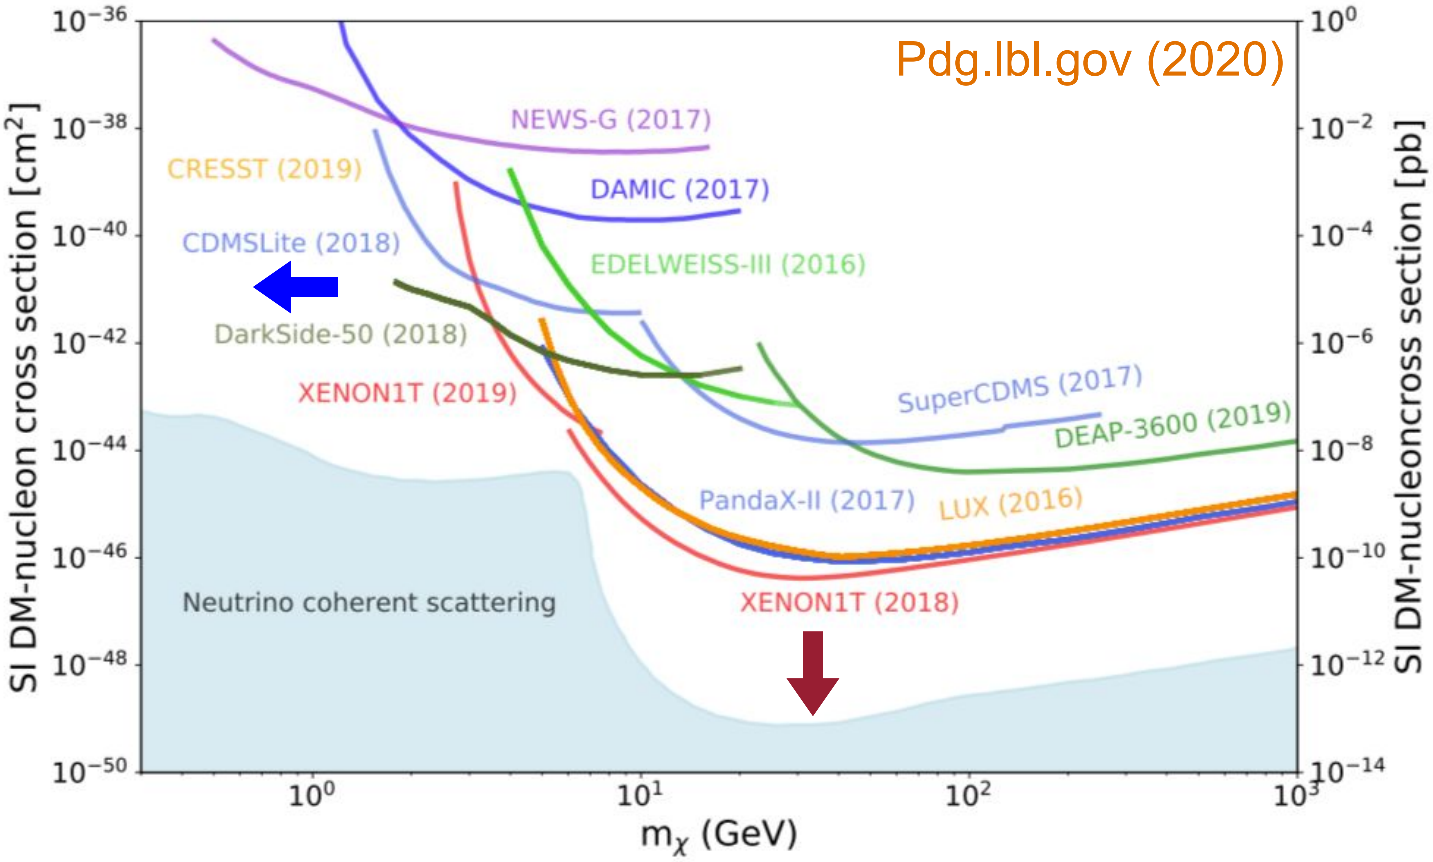

Main Thrusts:

- Enhance Sensitivity
- Expand the mass range
- New detector technologies
- Computing challenges!

Direct DM searches: “G2” experiments (previous Snowmass)

SuperCDMS SNOLAB sensitivity

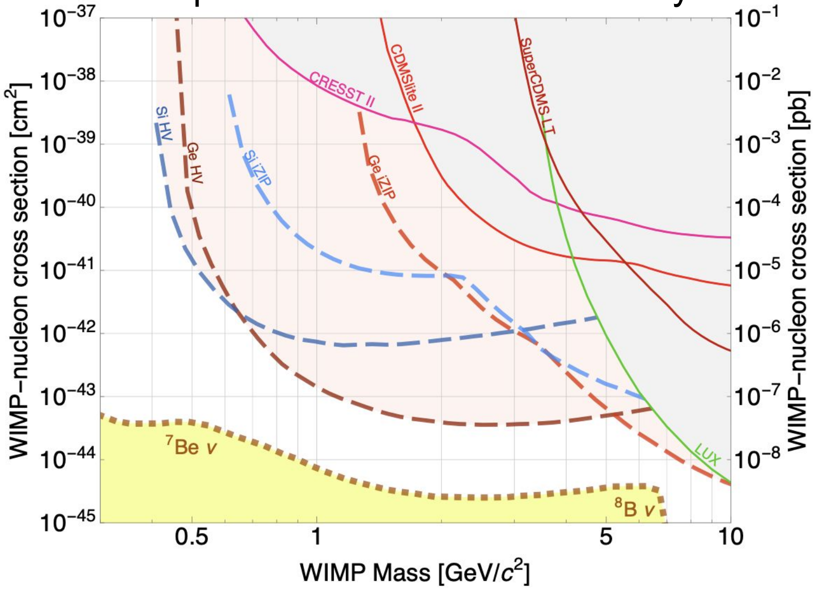

SuperCDMS status: under assembly. Data is stored and processed at SDF at SLAC. Other solid state DM experiments include: CRESST, Edelweiss, DAMA, ANAIS, etc.

LUX-ZEPLIN sensitivity

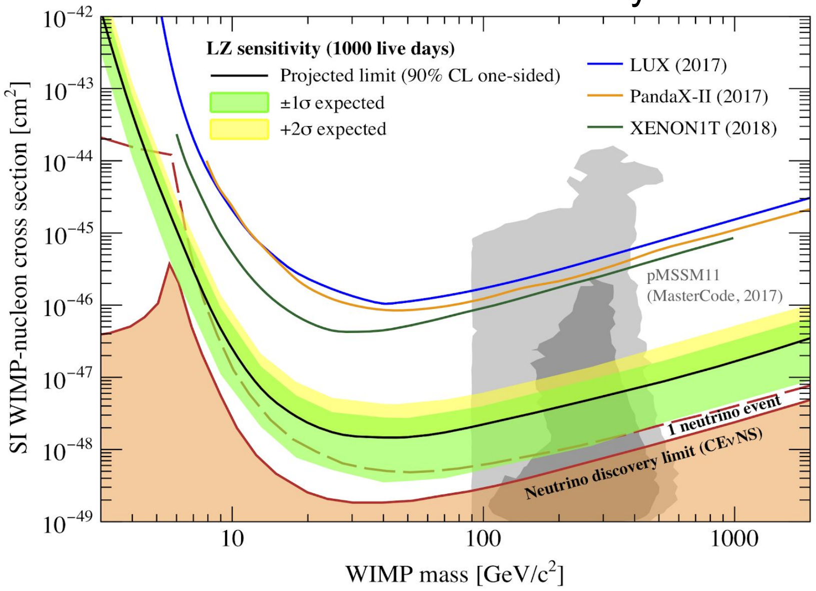

LZ status: commissioning. Data is stored and processed at NERSC and on GridPP (UK). Noble liquid experiments include: Xenon-nT, PandaX, DarkSide, DEAP, (COUPP), etc.

Example Future Dark Matter Detectors from CF1 LoIs

Cryogenic Carbon (Diamond + Si C)

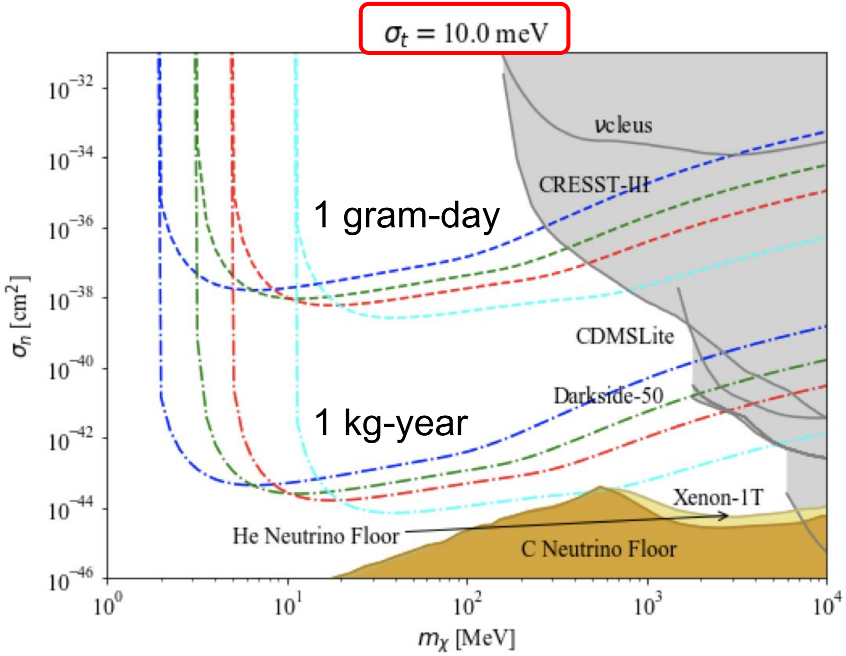

liquid helium (blue), diamond (green), silicon (red), xenon (cyan)

HydroX (Hydrogen in Xe)

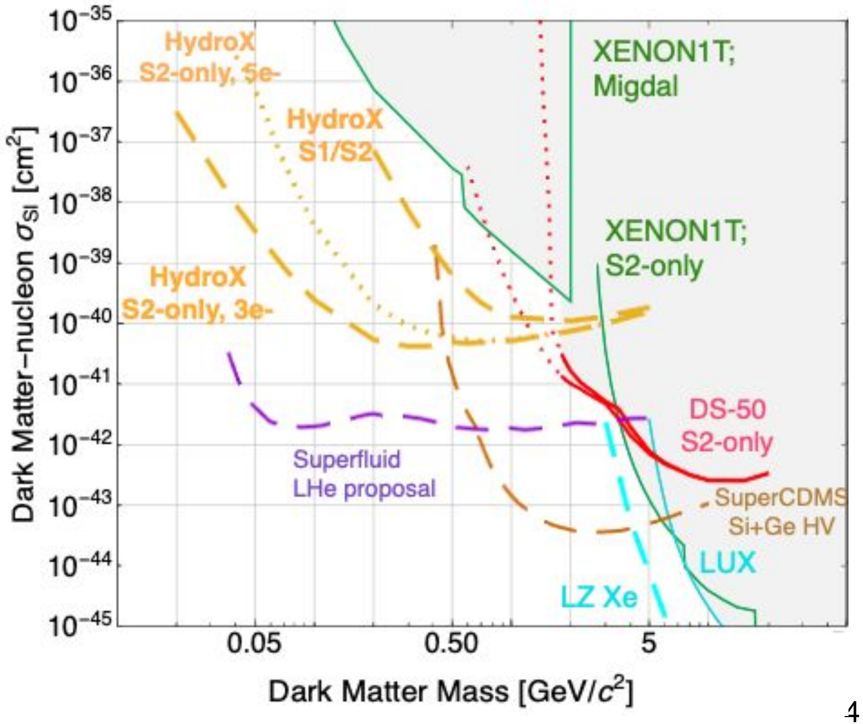

Noble Liquid Time Projection Chambers: Detection Principle

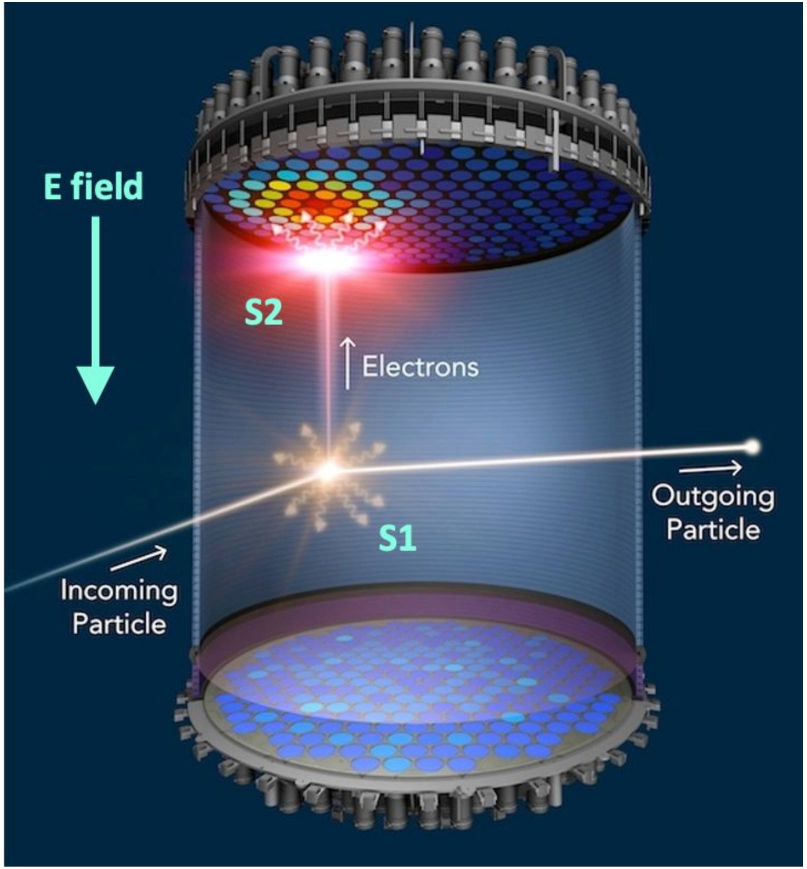

- WIMP-induced nuclear recoils: few keV energy
  - S1, S2 → event energy
  - S2 image → xy coordinate
  - S1-S2 timing → z coordindate
  - S2/S1 (Xe) → recoil type
  - S1 PSD (Ar) → recoil type
- No long-lived isotopes (Xe) and self-shielding
- LZ: 10 tons LXe (6 tons fiducial WIMP target)
- Fully active outer detector system (n-tagging)
- Next decade: 50-100 ton-scale detectors?

What are the computing challenges?

- Extreme "needle in a haystack" problem (LZ example):
  - Expected background events in nominal 1k-days run: 5 billion
  - Expected signal events in the same timeframe: a handful
- Challenges for data storage, reconstruction and analysis:
  - LZ will write 1PB/year (which is x3 Fermi-LAT and 1/3 ATLAS)
  - How do we make sure that a handful of DM events is not lost?
- Challenges for detector and accurate physics modeling:
  - Accurate/large-scale simulations of detector physics (custom!)
  - Ability to handle rare, unmodeled, anomalous backgrounds

Large-scale computing: challenges & approaches

U.S. DEPARTMENT OF ENERGY | Stanford University

SLAC NATIONAL ACCELERATOR LABORATORY

How do we deal with this massive amount of data?

SLAC

The direct detection community does not have a culture of large-scale software development. DOE/HEP is not keen on funding a large-scale software effort.

1. Use standard HEP tools as much as possible (Geant4, ROOT, Gaudi, etc.)
   - ↦ Challenge: it is not obvious that "standard HEP tools" are going to be supported for the long run (example: DOE does not fund Geant4 anymore)
2. Pool resources with DM and neutrino experiments with similar needs
   - ↦ Challenge: competition. Few experiments believe in public data/software
3. Use infrastructure at DOE's supercomputing centers (NERSC, ANL, etc.)
   - ↦ Challenge: no control over architecture. Little to no resources to adapt existing software framework(s) to a constantly evolving paradigm

NERSC: Computing Model Evolution

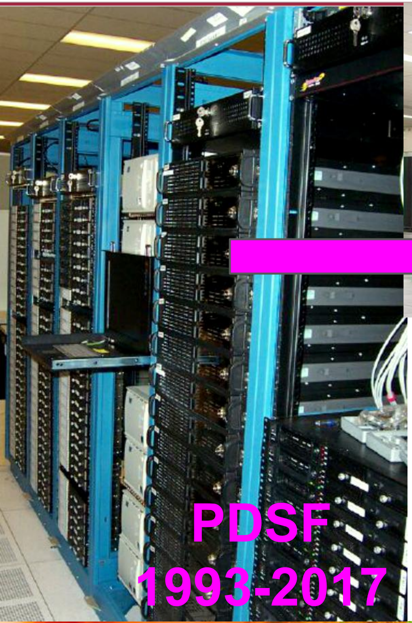

- Mixed architecture, optimized to support complex ML workflows:
  - 12 cabinets: CPU-only nodes
  - 12 cabinets: GPU-accelerated
- Prepare for a new architecture!
  - NERSC Exascale Science Applications Program (NESAP)
  - NERSC fellows, early access, hackathon with NVIDIA, etc.
  - Recruiting/retention challenges

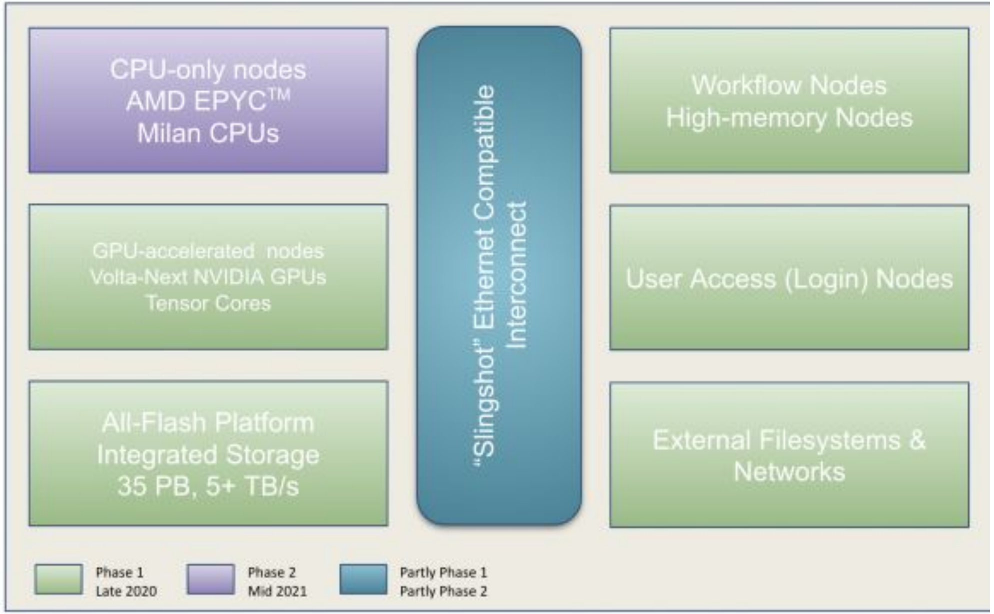

Anderson, Tyler	Brodsky, Jason	Buuck, Micah
Cartaro, Tina	Davies, Gavin S.	Diamond, Miriam
Fan, Alden	Ippolito, Valerio	Jillings, Chris
Kravitz, Scott	Krezko, Luke	Monzani, Maria Elena
Ostrovskiy, Igor	Psihas, Fernanda	Renshaw, Andrew
Riffard, Quentin	Roberts, Amy	Sander, Joel
Sangiorgio, Samuele	Stifter, Kelly	Tunnell, Christopher
von Krosigk, Belina	Wright, Dennis	Reto Trappitsch

- DANCE Community (Dark-matter And Neutrino Computation Explored)
  - Meetings in 2019, 2020 (ML)
  - Whitepaper in preparation

- Common needs across DM and neutrino experiments:
  - DAQ and triggers; data management, data processing; data preservation; simulations; ML; analysis tools; sw engineering, developer tools, training
  - Goal: share resources and cross-train; advocate for funding of common tools

Simulation, reconstruction, analysis: “custom” vs. community tools

Detector Modeling with Geant4

- Low-background experiments are designed, built and analyzed using Geant4
  - This requires a high-level of customization: SOURCE ≡ DETECTOR
  - Each low-background experiment develops a customization of Geant4
- DOE terminated its support of Geant4 in 2019
  - Geant4 development supported by CERN only
  - Without G4 or a well-validated replacement, we will not be able to design future experiments!!!
- Several ideas for new simulation frameworks!
  - How many person-millennia have gone into the validation of Geant4 over the last 25+ years???

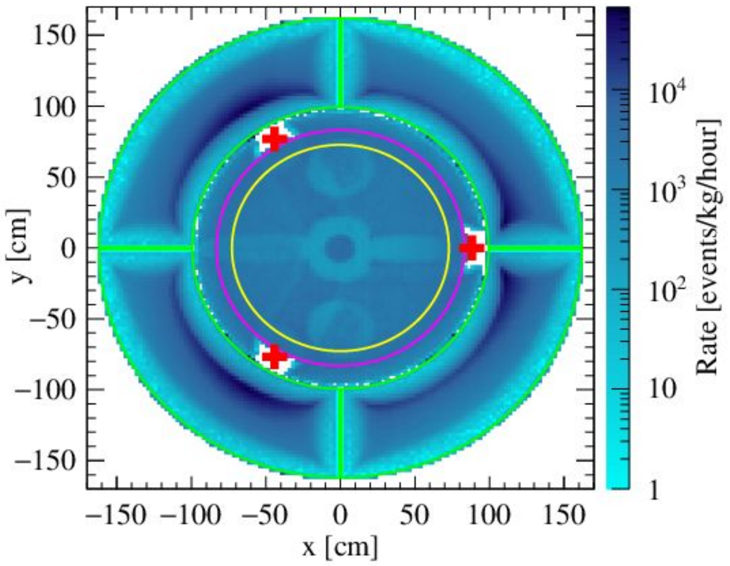

The Noble Element Simulation Technique (NEST)

- NEST simulates the scintillation, ionization and electroluminescence processes in noble elements
- It is validated for liquid, gas, and solid Xenon
- Material response is provided as a function of energy and electric field
- Used in DM (LUX, LZ, XENON, DARWIN, LAr DM Collaboration) and neutrino (EXO/nEXO, DUNE, MicroBoonNE, ICARUS, SBN, RED, COHERENT), plus PET scans and other radiation detectors
- Example of a community tool used by several “competing” collaborations (see LOI here)

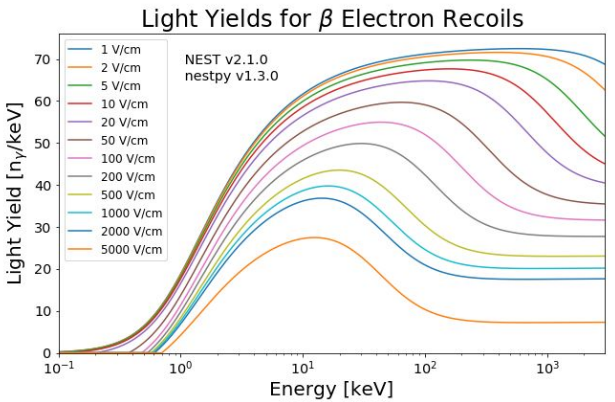

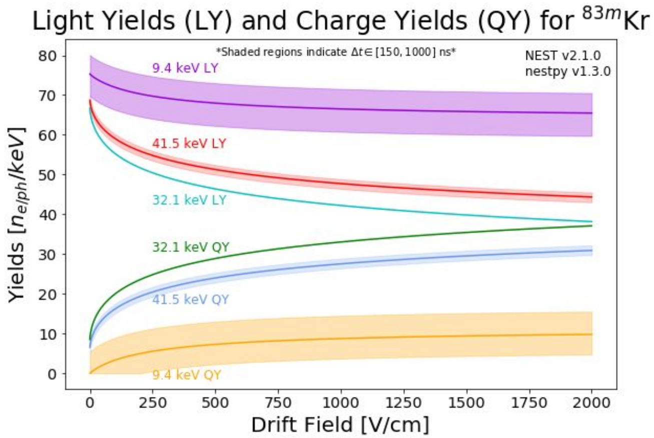

Leveraging Perlmutter GPUs for simulations

- Propagation of optical photons is the most computing-intensive aspect of simulations in noble liquid experiments (>99.0% of CPU time)
  - Idea: offload the raytracing of photons to GPU
  - Use NVIDIA's OptiX library to track photons
  - And open-source Opticks interface to Geant4
  - See LOIs here and here
- Pilot program for an overall simulation approach?
  - "disassemble" Geant4 into a set of independent libraries, that can run on multiple architectures

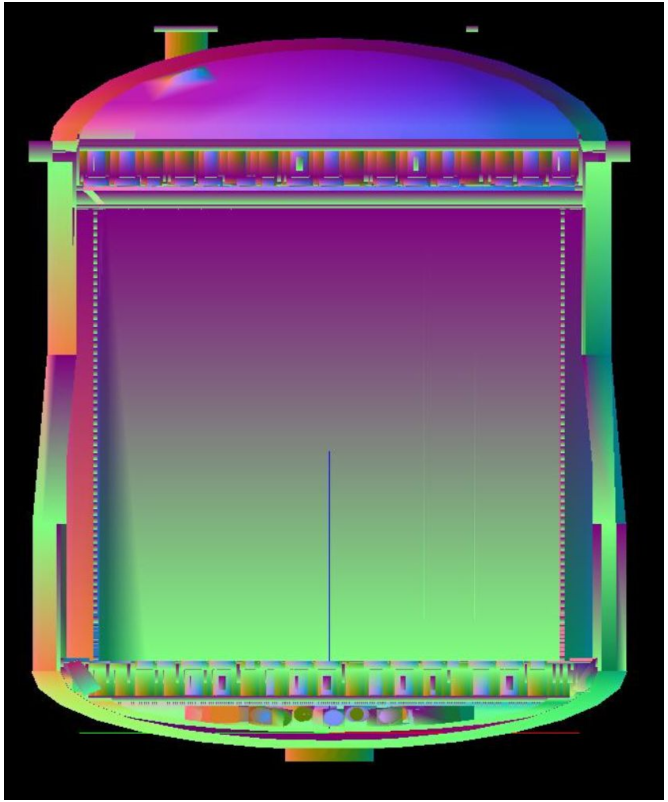

Event processing and reconstruction. Data analysis

- Event processing and reconstruction:
  - **LZap (LZ analysis package):** doesn't actually do analysis. Custom framework based on **Gaudi**
  - Spatial/energy reconstruction: custom code, based on **ANTS2** (see webpage here)
- Data Analysis:
  - **ALPACA (Analysis Lz PACkAge):** does the actual analysis. Custom code, based on **ROOT**
  - Key feature: allows to **correlate between events** when doing analysis

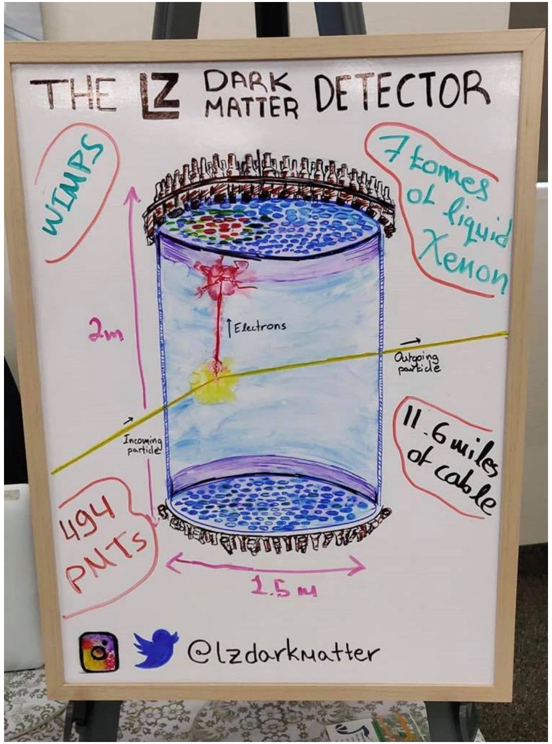

Software management, deployment, etc.

- Software management and deployment:
  - Code repository is hosted on **gitlab.com**
  - Ample use of gitlab's CI/CD + **Scikit-HEP**
  - Code is deployed to a **CVMFS** partition
- Web services, DataBases, etc.:
  - Leverage **SPIN**, which is a container-based platform at NERSC
  - Services hosted on SPIN include: Data transfer service, Offline event viewer and data quality monitor, Job submission engine, Databases and their mirrors and interfaces, etc.

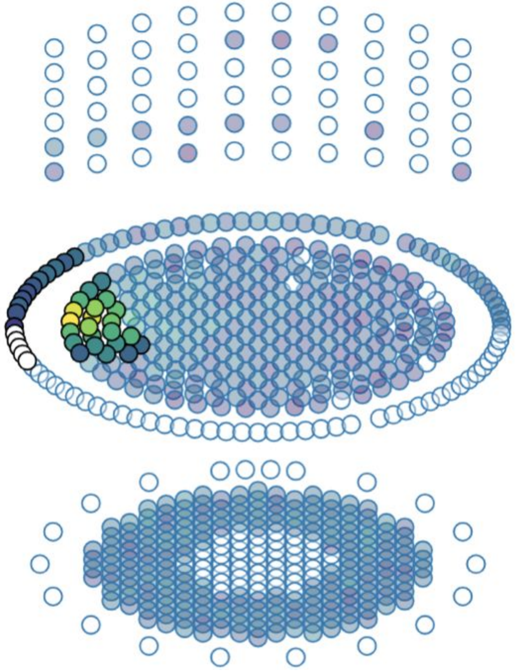

A non-exhaustive list of community tools used in LZ

- Simulators/event generators: **Geant4**, **NEST**, Opticks, SOURCES4A, DICEBOX, G4EMLivermorePhysics, G4HadronPhysicsQGSP, etc.
- Event processing/reconstruction/analysis frameworks: **Gaudi**, **ROOT**, Boost, CLHEP, ANTS2, FFTW, QT5, AIDA, TBB, python/jupyter
- Machine Learning tools: ONNX Runtime, scikit-learn, TensorFlow, etc.
- Miscellaneous tools: gitlab.com, **CVMFS**, Shifter, Singularity, Valgrind, Coverity, SPADE, Scikit-HEP, mongoDB, mysql, Dirac, SLURM, RabbitMQ, etc.

Unique ML challenges in DM detection

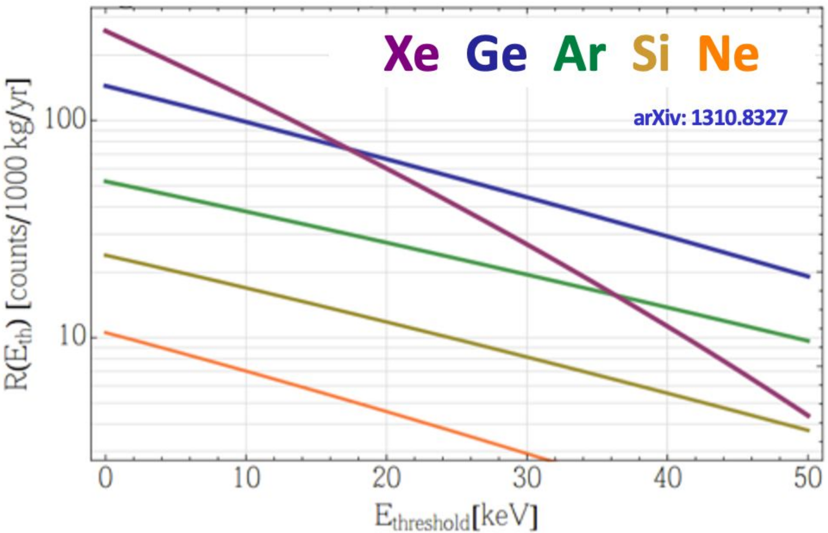

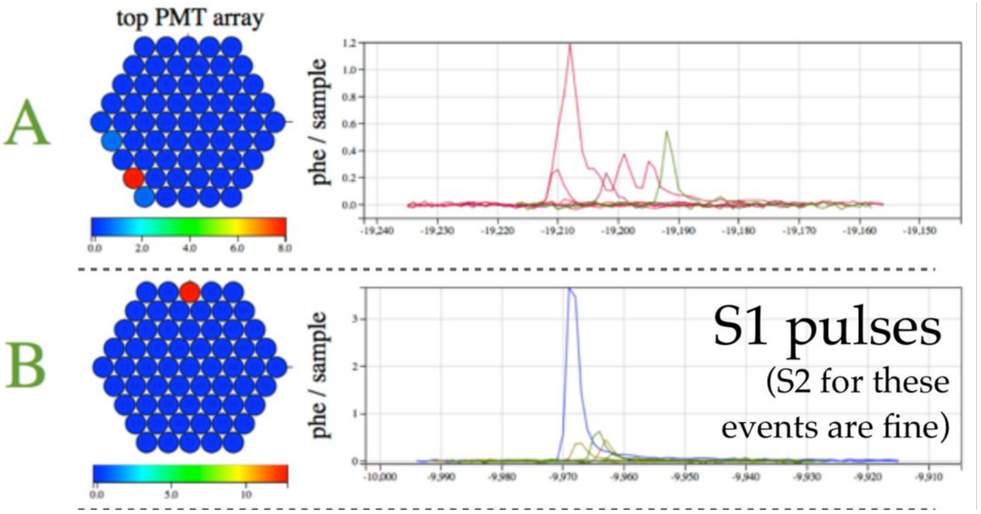

- Sensitivity for DM experiments is dominated by events at threshold, where the information is exceedingly sparse: how do we appropriately train our models?
- (Mis-?)perception of ML as a "black box": challenges of interpretability and uncertainty quantification, which are common to all physics applications

Coping with rare, unmodeled backgrounds

- Most direct detection experiments adopt some form of bias mitigation (blinding, salting)
  - ALL of them found residual background events once they opened the “search box”
  - These were rare & unmodeled backgrounds (we discovered them after opening the box)
- Can we use anomaly detection for this? (LOI)
  - Similar application to intrusion detection in IT
  - Challenge: we still may need to feed the whole 5PB dataset to the network at some point

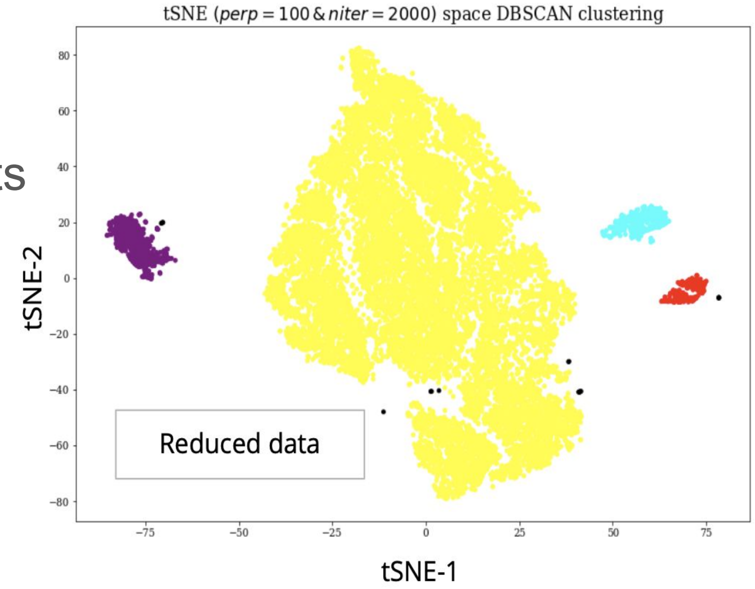

Clustered events in t-SNE space
Unexpected populations such as pulse finder issues are evident in this space but not in common analysis space

Conclusions and Outlook

- LZ and other dark matter experiments are rapidly ramping up on large-scale computing needs:
  - LZ will write ~1 PB of raw data a year to disk
  - Extreme “needle in a haystack” challenges
  - Plus: rapid evolution of computing paradigms
  - Computing support was historically modest
- Needs and opportunities for the next decade:
  - Support for HEP “community tools” is crucial
  - Ongoing efforts to share tools and resources
  - Many creative avenues in simulations and ML (e.g. modeling/rejection of rare backgrounds)