



LArTPC multi-threading and acceleration workshop

March 2 – 3, 2023
Fermilab

Note: sessions will be recorded!

Related workshops coming up

- [AI in Production Mini-Workshop, March 7](#)
 - “Principal goal is to bring stakeholders and service/facility providers together and understand what AI for production/operation workflows will look like in the near term for both training and inference. In particular, stakeholders should discuss software and hardware needs. CSAID will use the product of this mini-workshop to plan resource acquisition, deployment, and developer effort”
- [Workshop on Neutrino Event Generators, Mar 15–17](#)
 - “Review progress and assess requirements for common simulation software tools needed by the neutrino physics community. Discussions will include technical topics, systematic uncertainty quantification, data comparisons, standard output format.”
- Frameworks workshop
 - Currently in the planning phase.

Why a multi-threading and acceleration workshop?

1. Resource optimization and throughput bottlenecks on existing resources
 - All LArTPC neutrino experiments at the lab report significant fraction of jobs running on more than a single grid slot due to memory consumption
 - Many LArTPC computing problems are parallelizable and would benefit from various types of acceleration
2. HPC
 - Funding agencies pushing lab / experiments to use more HPC
 - Many experiments / groups have experience with this already
 - Multi-threading / optimizing for GPU also help with this transition, or are already part of it
3. Uniformity of LArTPC technology
 - LArTPCs are well-suited for direct sharing of code, techniques, technologies

Why a multi-threading and acceleration workshop?

1. Resource optimization and throughput bottlenecks on existing resources
 - All LArTPC neutrino experiments at the lab report significant fraction of jobs running on more than a single grid slot due to memory consumption
 - Many LArTPC computing problems are parallelizable and would benefit from various types of acceleration
2. HPC
 - Funding agencies pushing lab / experiments to use more HPC
 - Many experiments / groups have experience with this already
 - Multi-threading / optimizing for GPU also help with this transition, or are already part of it
3. Uniformity of LArTPC technology
 - LArTPCs are well-suited for direct sharing of code, techniques, technologies

Why a multi-threading and acceleration workshop?

1. Resource optimization and throughput bottlenecks on existing resources
 - All LArTPC neutrino experiments at the lab report significant fraction of jobs running on more than a single grid slot due to memory consumption
 - Many LArTPC computing problems are parallelizable and would benefit from various types of acceleration
2. HPC
 - Funding agencies pushing lab / experiments to use more HPC
 - Many experiments / groups have experience with this already
 - Multi-threading / optimizing for GPU also help with this transition, or are already part of it
3. Uniformity of LArTPC technology
 - LArTPCs are well-suited for direct sharing of code, techniques, technologies

Workshop goals

- To learn the multi-threading and acceleration capabilities of frameworks and common toolkits used by LArTPC experiments;
- To share experiences across experiments about existing resource utilization and throughput problems that lend themselves to multi-threaded or acceleration solutions;
- To explore how multi-threading and acceleration is being used to address these problems and open avenues to the use HPC resources more broadly;
- To discuss the results of applying these techniques and capabilities

Organization of the program

FCC1W

Common tools and support

09:00	Introductions	09:00 - 09:15
	Multi-threading support in LArSoft	<i>Kyle Knoepfel</i> 09:15 - 09:45
	Handling external libraries with conflicting thread pools	<i>Christopher Jones</i> 09:45 - 10:05
10:00	Parallelization in LArSoft reconstruction - SciDAC4 developments	<i>Giuseppe Cerati</i> 10:05 - 10:35

Simulation tools

	Break	10:35 - 11:00
11:00	Geant4 multi-threading and tasking	<i>Soon Yung Jun</i> 11:00 - 11:35
	GEant4/CaTS/Optics: optical photon propagation on a GPU	<i>Hans-Joachim Wenzel</i> 11:35 - 12:10

Open working time: 13:30 – 17:00
(per request)

Zoom will remain open during this time

12:00	Lunch break	12:10 - 13:30
13:00	Feynmann Computing Center, FCC1W: Open working time	

Organization of the program

WH8X “Hornet’s Nest”

Note 8:30 am start time!

Experiment tools and experience I

Experiment tools and experience II

08:00

Wire Cell Toolkit + art multi-threading strategies

Brett Viren

08:30 - 09:00

09:00

Wire Cell Toolkit + art multi-threading usage

Haiwang Yu

09:00 - 09:30

Pandora

Ryan Cross et al.

09:30 - 10:00

10:00

ICARUS multi-threading production workflow

Tracy Usher et al.

10:00 - 10:20

Break

10:20 - 10:50

11:00

DUNE event serialization, low level data processing and production

Thomas Junk

10:50 - 11:20

Multi-processing for ND-LAr in larnd-sim and ndlar_flow

Matt Kramer

11:20 - 11:50

12:00

Applications of GPUs to DL inferencing

Michael H L Wang

11:50 - 12:20

Ground rules for the sessions

- Please be courteous and respectful at all times.
 - Workshop is subject to the [Fermilab Community Standards](#)
- This is a hybrid meeting, so:
 - Please raise hands on zoom even if in the room
 - For direct responses to the current speaker, put “DR” in the zoom chat
- Speakers in the room, please use the pointer on your computer so virtual audience can see it
- If you are remote and cannot hear something being said in the room, please say something
- Please try to keep to time.

Acknowledgements

LArSoft Project

- Katherine Lato

Events Office

- Joy Pomello
- Melody Saperston
- Aaron Zinsmeister

CSAID Administrators

- Rachel Hurd
- Carla Lloyd

The end