

Engagement with the User Community

Matt Heine

BCS subMIT Review

May 22, 2025





Outline



Annual SubMIT Workshop



Users Group Meetings



Classroom Usage



User-Run Workshops using SubMIT



Customizations / User Requests



Current Limitations & Open Challenges

Annual SubMIT Workshop

Project Team Contributions



subMIT

Getting physics things done at MIT

Workshop on Basic Computing Services in the Physics Department - subMIT

📅 Thursday Jan 30, 2025, 9:00 AM → 5:00 PM America/New_York

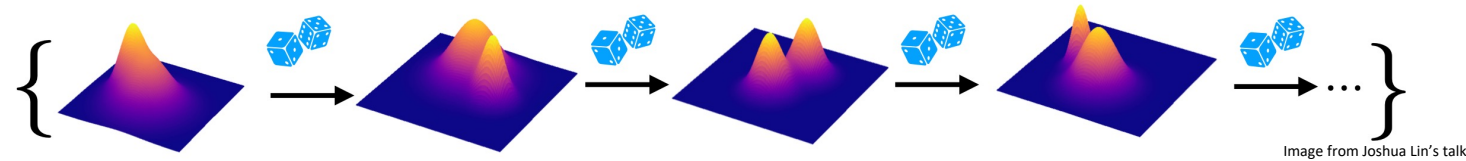
📍 26-414 (Kolker room)

Project Team Information / Tutorial Sessions

- System Overview
- Resources & How to Access Them
- Software Management
 - Environment / Package managers (Conda)
 - Containers (Docker/Podman, Singularity/Apptainer)
- Batch Jobs

2 NEW Tutorials Highly-Requested by Users

Annual Workshop: User-Contributed Talks



Title	Speaker (User)	User Position	
Emulating the atomic nucleus	Antoine Belley	Postdoc	LNS, IAIFI
Symbolic Learning Nuclear Relations	Jose Miguel Munoz Arias	Graduate student	LNS, IAIFI
The needle in a haystack problem of homology-directed DNA repair	Henrik Pinholt	Graduate student	Physics of Living Systems
Lattice-QCD 🤝 Software 🤝 Hardware	Joshua Lin	Graduate student	CTP
Soft unclustered energy patterns in CMS using dask	Pietro Lugato	Graduate student	LNS
Topological Chiral Superconductors beyond Fermi liquid pairing	Luke Kim	Graduate student	CMT
A qubit that corrects itself	Max Geier	Postdoc	CMT
subMIT for cosmological uses	Mikhail Ivanov	Professor / Faculty	CTP
Assessing a new analysis with batch jobs	Sumita Ghosh	Postdoc	LNS

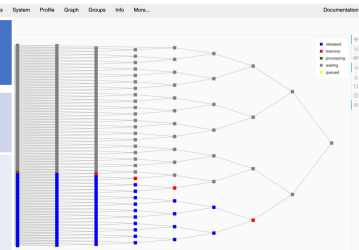


Image from Pietro Lugato's talk

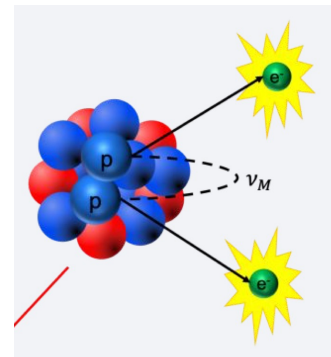


Image from Sumita Ghosh's talk

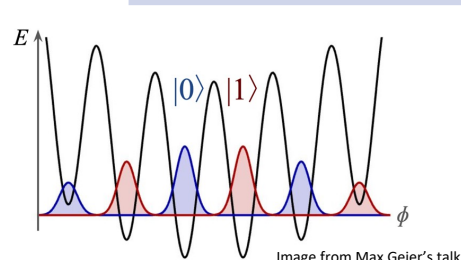
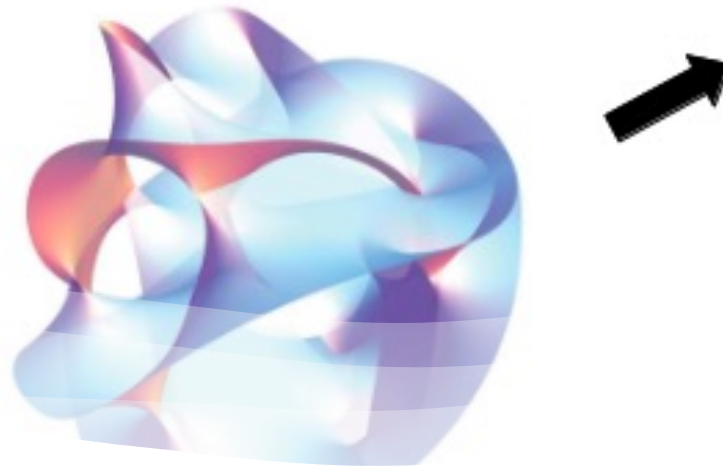


Image from Max Geier's talk

Despite success of last year's user-given tutorials, no users volunteered to present tutorials this year. Although some user talks did heavily highlight tools.

Users Group Meetings



Good active exchange
of feedback/advice
between project team
and user participants
(in both directions)

Also synergistic
discussion between
users at meetings

Project Team
contributed additional
informational session

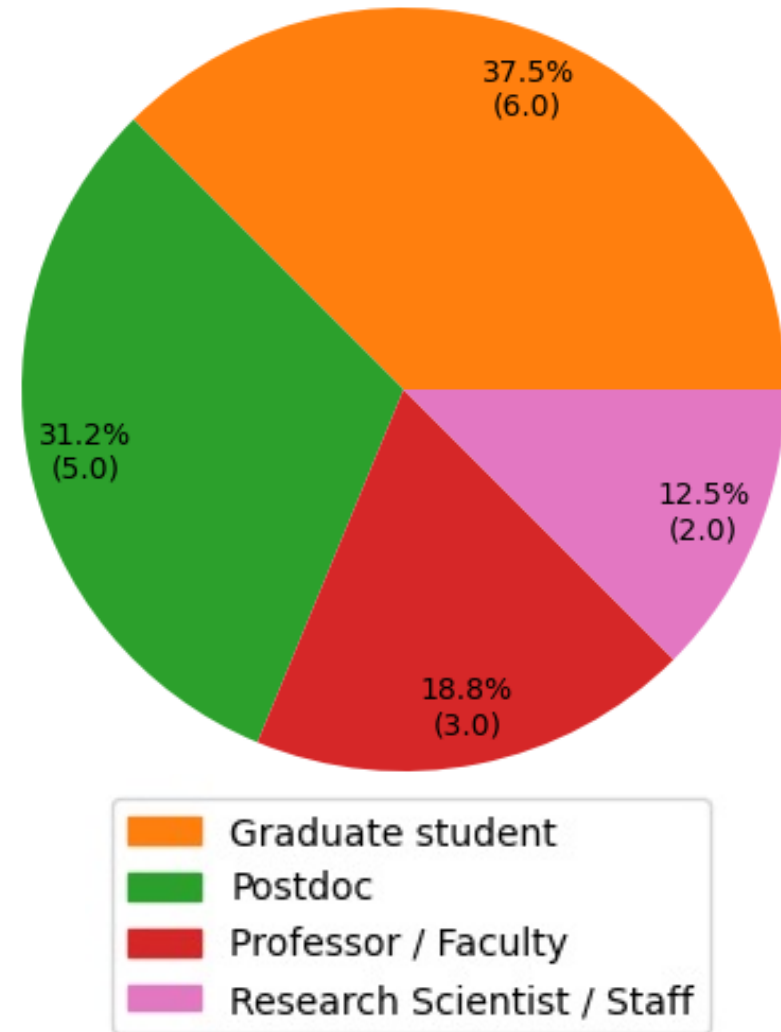
Efforts underway to
increase attendance /
engagement at these

User-Contributed Talks at User Group Meetings

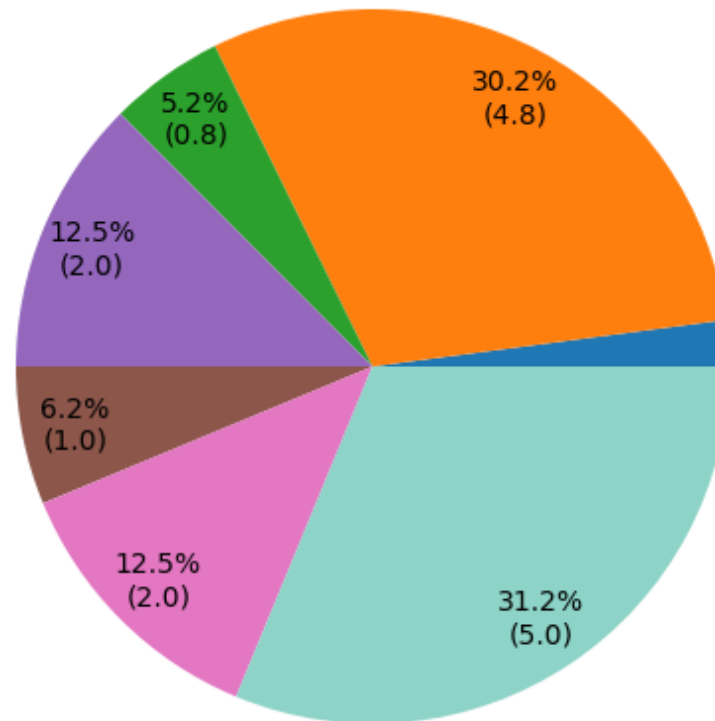
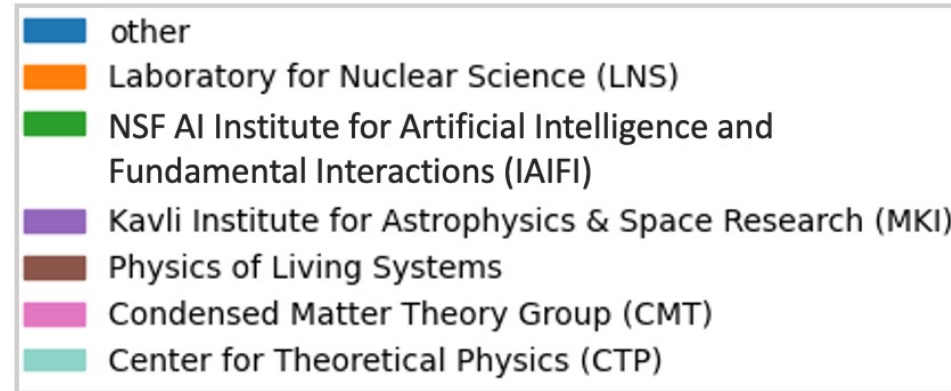
	Title	Speaker (User)	Speaker Position	
July, 2024	Analysis of HERA data with subMIT	Frank Taylor	Professor / Faculty	LNS
October, 2024	W Boson Mass Analysis on subMIT	Josh Bendavid	Research Scientist / Staff	LNS
November, 2024	What I'd like to use submit for	Mortiz Günther	Research Scientist / Staff	MKI
February, 2025	Solving Einstein's Equations on the computer	Josu Aurrekoetxea	Postdoc	CTP
March, 2025	Computing Calabi-Yau Manifolds	Richard Nally	Postdoc	CTP
April, 2025	Modeling gravitational waves using highly distributed parallel computing	Scott Hughes	Professor / Faculty	MKI
June, 2025	Simulating heavy quarks in quark gluon plasma	Jean Du Plessis	Graduate student	CTP

Which Users Gave Talks at SubMIT Events?

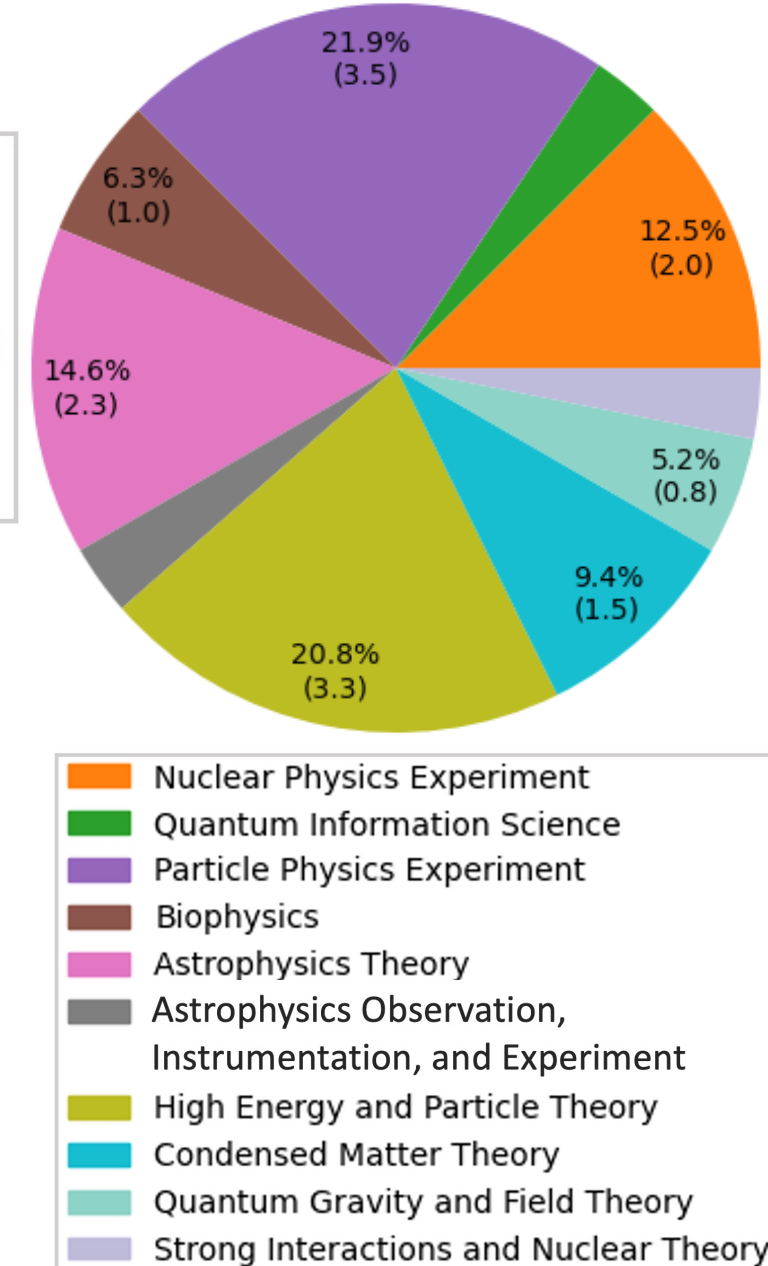
Speaker's Title



Speaker's Lab/Center



Speaker's Research Area



Example Workflows & Tools Presented / Discussed by Users



Markov Chain Monte Carlo



Machine Learning



User-written code / scripts



Singularity Containers (reproducible, portable environments)



Dask (automatic parallelization)



MPI (calculation steps $\sim 1s$ each, 10^6 steps)



Slurm batch jobs / job arrays



Homogenous, simultaneous tasks vs Heterogeneous, few queued (asynchronous) tasks


User-Run Workshops & Classroom Use

Research in Experimental Particle Physics in January

Jan 6, 2025, 9:00 AM → Jan 31, 2025, 9:35 PM America/New_York

24-506 (MIT, Building 24)


Description During January we give students (undergraduates and graduate students alike) the opportunity to learn about the Future Circular Collider project [1] at CERN and perform a small research project that can be adjusted to the skill level of the individual.

 **PPC**

FCC Workshop
January 2025

ABOUT ▾ PPC NEWS PROJECTS ▾ 🔍

From CERN to MIT for the FCC Workshop

 JAN EYSERMANS / JANUARY 26, 2025 / EVENTS, FCC, FUTURE, NEWS

FCC Tutorial
April 2025



Gaia Hackathon

Astrophysics Hackathon

January 27-30, 2025.

All Welcome. Please register [HERE](#)

Run by Prof. Lina Necib (MKI)
Audience: Undergraduates
Requested shared disk space
for datasets (~few TB)
<https://gaiadr3hack.mit.edu>



Classroom Usage



Known classroom (course) usage this year

- 8.01 A2rchi usage
- ESG 8.022 (Jupyter notebook; periodic small reservations)
- 8.13/8.14 Jr Lab
- 8.006 (optional alternative to local development environment)
- Educators may use w/out notifying us

Lighter known classroom usage than last year

- Introductory Physics 8.02 did not use Jupyter Notebooks (consequently nor SubMIT) this academic year
- Less load on the system
- No large reservations

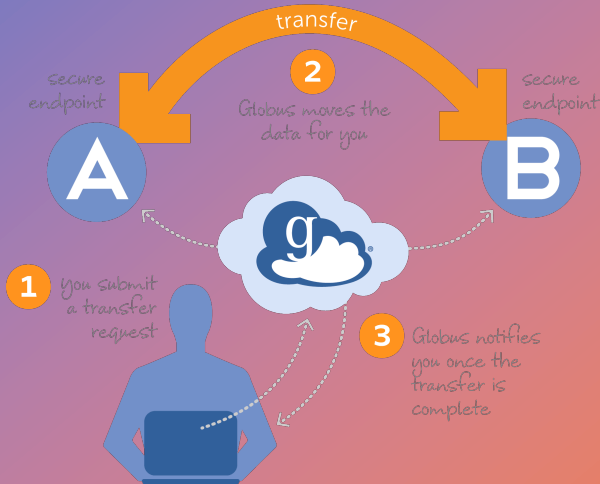
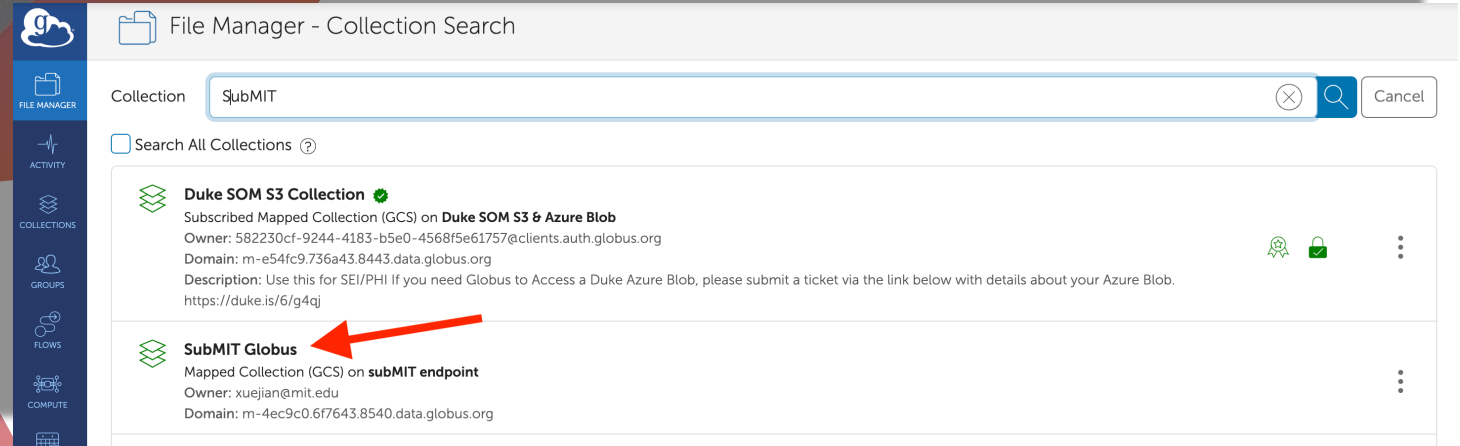
Customizations & User Requests



Image from <https://www.open-mpi.org/doc/5/11/2025>



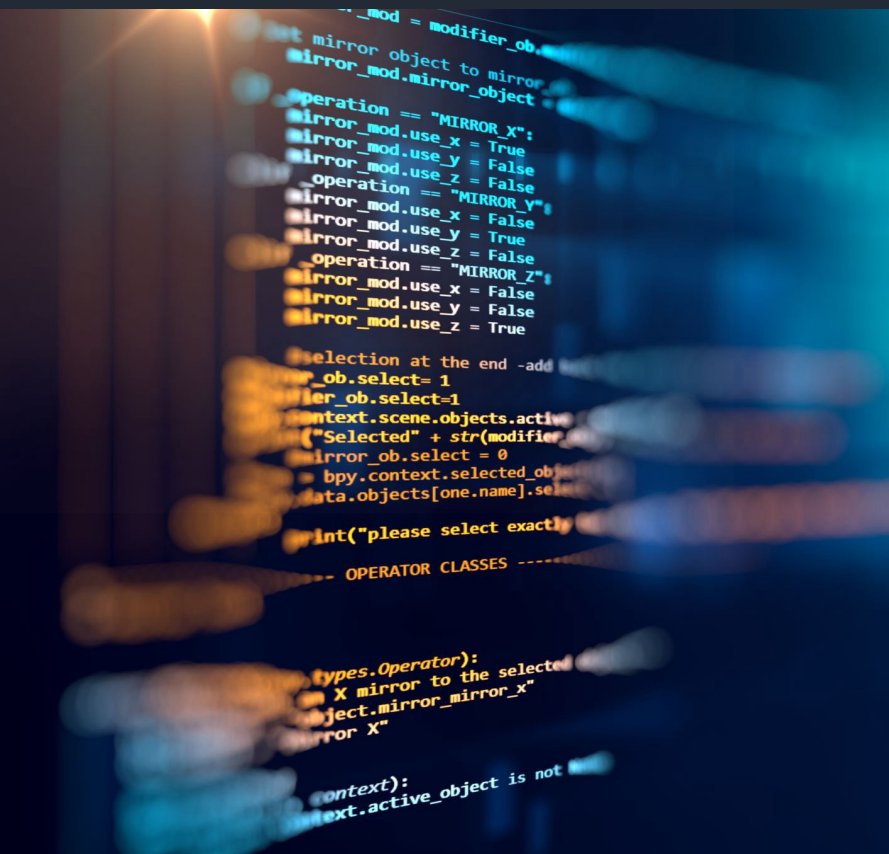
- Central MPI installed upon user requests
 - module system used to manage environment
 - uses RoCE (RDMA over Converged Ethernet) for faster communication



- Globus Endpoint established on SubMIT (user requests)
 - Requested by many users
 - Easy fault-tolerant data transfers
 - User connection/interaction not required during transfer
 - Web Interface

... continued on next slide ...

... Customizations & User Requests (continued)



Group / Project Websites – users requested these requirements:

- Not hosted by an individual user → seamless transitions when researchers move on from groups
- Provides access to data hosted on SubMIT
- Granular data access: some public, some password-protected
- Ability to update content via custom script (no time/labor of dragging & dropping via GUI)
- Easily modified by all group members
- (supplements user websites, which SubMIT always provided)



Priority Access for Purchasing Groups

- Balance availability & sharing
- Investigated several general policies
- Choice: Currently done manually (case-by-case basis) to optimally adapt to changing usage patterns
- E.g. LHCb required priority access to submit81, so it was temporarily removed from SLURM queue
- Direct ssh (non-SLURM) access to purchased machines blocked for non-group members

Current Limitations & Open Challenges

Pilot Customizations

(require additional development before deploying to entire user base)

- Dropbox-like syncing of files (other tasks prioritized over scaling this to all users)

Was unable to accommodate

- Request to accommodate Windows software.
 - SubMIT servers run AlmaLinux 9. Workarounds to support Windows on AlmaLinux (i.e. virtual machines) are too resource-intensive for the shared SubMIT resources.
 - Windows system admin would require significant additional time & specialized Windows expertise (e.g. security, etc.), so SubMIT does not support Microsoft Windows
 - Most likely most efficient for user to run on their own workstation or attempt to use WineHQ on SubMIT
 - Asked user to let us know if this is not a sufficient option

Open Challenges

- Balancing restrictions vs ensuring availability of resources (for purchasing or non-purchasing users)
- Turnaround Time / Availability vs Throughput
- Working to increase attendance of User Meetings
 - Some users thought these meetings were only for User Group Representatives
 - Some potential new users overlook announcements since they are not (yet) subMIT users
 - Great discussions between participants; would be great to expand this more
 - Efforts to reengage User Group Representatives