Software Sustainability through Community Building

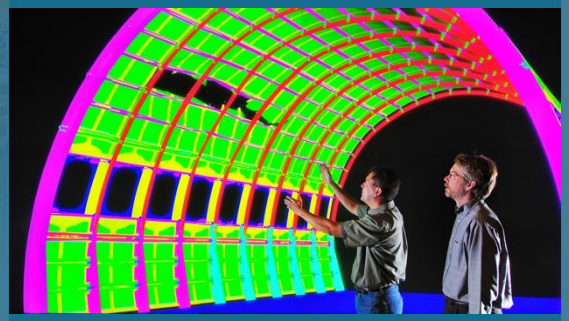

kokkos

Christian Trott

Sandia National Laboratories, Center for Computing Research

SAND2025-03093C

Sandia National Laboratories

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

SAND2025-03093C

HPC Community 25 Years Ago

- ”one-off” codes:
  - each grad student writes their own
  - New problem == new app
- Library reuse limited
  - Little besides MPI, BLAS and LAPACK
- Limited software engineering know-how
  - Unit testing basically unheard of
  - Avantgarde was to use version control (cvs etc.) – git didn’t exist yet
- Software publishing as tar balls on websites

Typical HPC application developer team

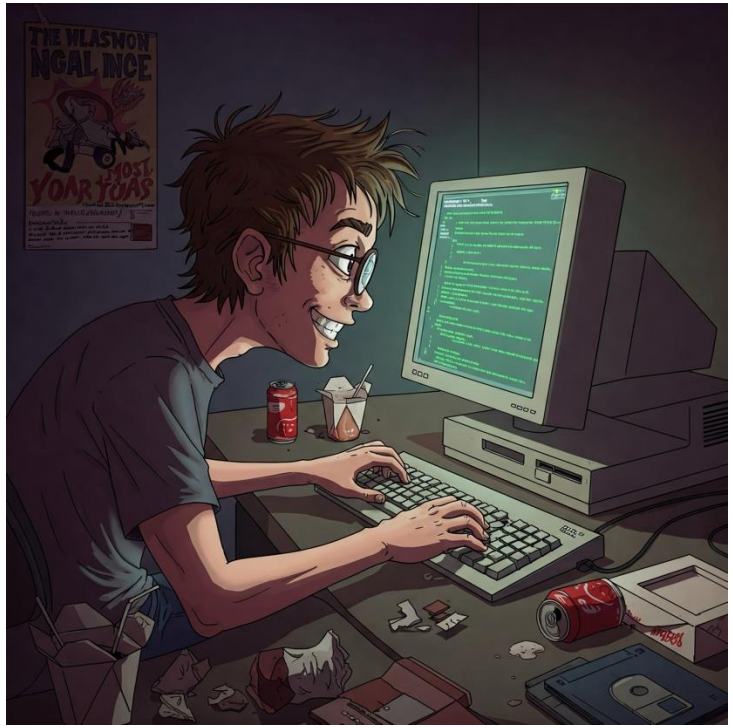

Me working on a Fortran MD code

The need for Change

We got more ambitious:

- More compute capability enables more physics
- Replacing more and more RL experiments with modelling

Ca 2000 CFD

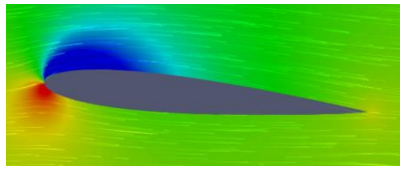

Ca 2025 CFD

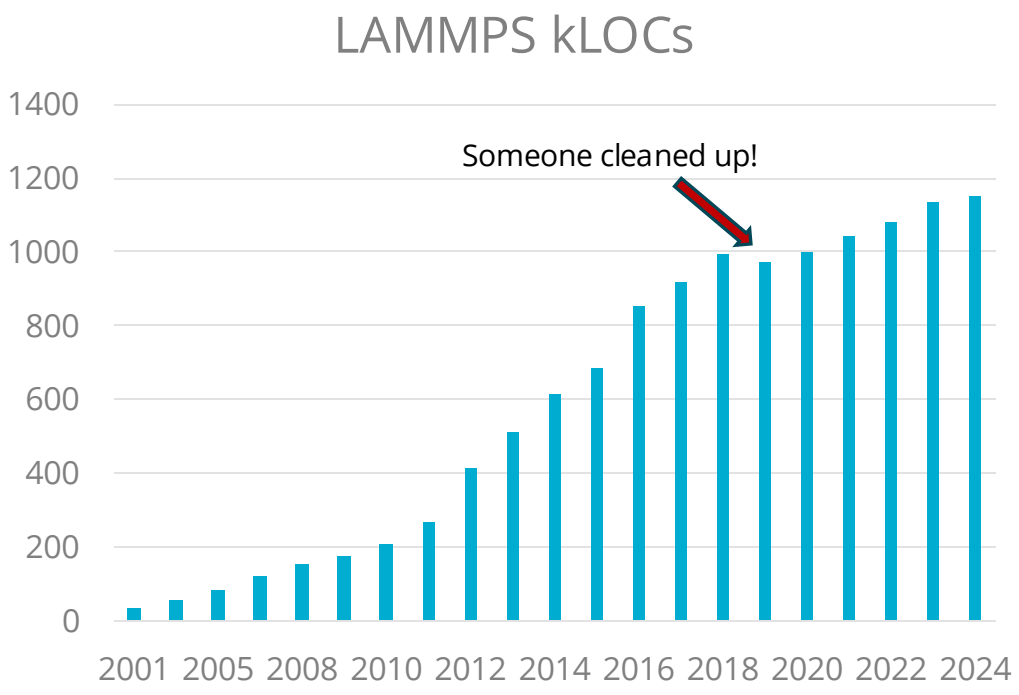

More Physics == More Code

- Complexity of software has increased dramatically
  - Grows super linear with amount of captured phenomena
- Treating corner cases correct and doing less physics simplifications/approximations are drivers
- Also increases demands on solver algorithms due to scalability issues

And then the hardware landscape got more complex

TOP 500 Supercomputer

CPUs:

OpenMP

Threading
Building
Blocks

std::thread

Accelerators:

NVIDIA CUDA

AMD ROCm

oneAPI

OpenMP

OpenACC
Directives for Accelerators

Hardware Architectures and the Means To Program Them Have Diversified

You can’t go it alone

Not enough time/resources to reinvent everything

- 10,000 lines production code per year per dev
  - Ok maybe 30,000 lines of hacky research code
- Something like Kokkos is 30 years worth of dev time

HPC is niche: More stakeholders == more attention

- Help from vendors is easier to get if you have critical mass
- We are already small enough compared to industry

More users == more feedback

- Attracting more users from more institutions tests more
- Makes software better for your own people

Impediments to collaboration

Not invented here Syndrom

- Trust issue: will it still be supported?
- Credit issue: will I be recognized for contributing vs. doing something new?

Disaggregate Funding Justification

- Justifying money to do only part of something is harder
  - How is the risk picture of this?
- Reporting is more complicated if everything is a team effort
  - Are sponsors ok with “partial” claim?
- How to pay for common needs (meeting, CI hardware)

Need a Remote Work ethic

- For a tight collaboration constant communication essential
- Avoid devs from other institutions feeling second class

Kokkos Growth Story

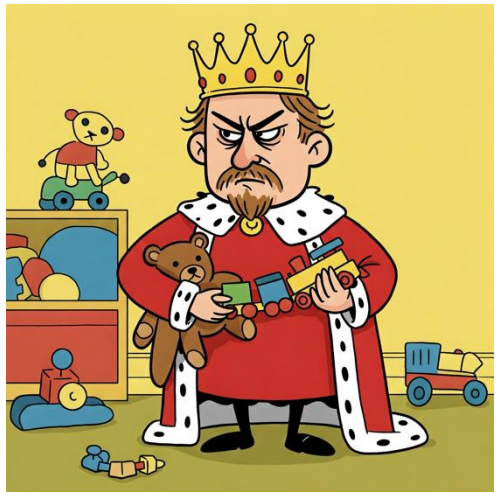

Open Source vs Open Community

- Recognized early that capturing market & mind share we need open community
- Not enough to make code open: need to give up control
- **Very large** initial investment in training events etc.
- Pandemic helped integrate the team!

The Role of "Expansion Funding"

- Originally only funded by Sandia
- ECP initially enabled more interactions
- Explicit developer funding for other labs in 2019
  - Established dev team nucleus at other lab that was then able to secure follow up funding
- Latest example: CEA CExA => explicit funding to build Kokkos team at CEA

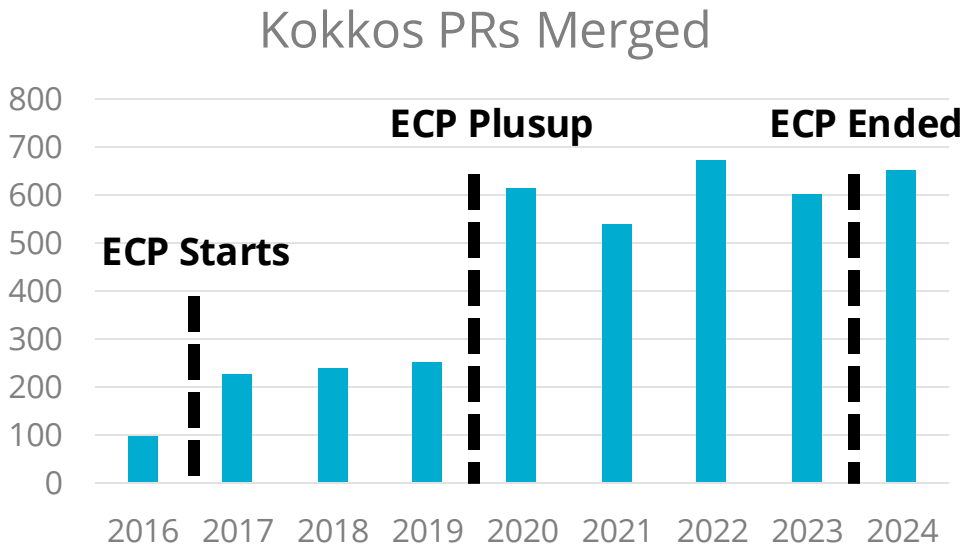

Joining the LF High Performance Software Foundation

Kokkos is now a Linux Foundation project in the High Performance Software Foundation!

The High Performance Software Foundation

Launched at ISC in May 2024!

Sponsors:

- *US Labs:* SNL, LLNL, ORNL, LANL and ANL
- *Industry:* HPE, Amazon (AWS), NVIDIA, AMD, Intel, Kitware, ARM
- *Others:* CEA

Scope:

- Help sustain critical software technologies for HPC and related areas

What does it do:

- Guarantees Open Governances (e.g. github.com/kokkos is owned by LF)
- Help organize and pay for meetings and workshops
- Help finance and organize project infrastructure (e.g. Slack)
- Work with other organizations such as OpenSSF, LFEdge for overlapping concerns

How HPSF helps with multi institutional collaboration

- **Optics:** Participating means being “Maintainer of LF Project” vs. “Contributor to SNL Project”
  - Reality: clear rules for governance
  - Reality: common resources owned by neutral party
- **Resource Pooling:** Fees are easier than contracts
  - One way agreements with Linux Foundation
  - Pooled resources can be used for common needs
    - Conference/Meetings
    - Testing Resources
    - Website/Outreach

- **Credibility:** Easier to convince non-DOE people that this is serious
  - It takes some effort to become LF project – not just random person pushing something to github
  - HPSF has actually paid staff from LF
    - Does a lot of the onboarding management
    - Helps with outreach and maintaining connections
  - LF helps to professionalize web presence

A Pitch: Join us for HPSF Conference 2025!

- **Chicago May 5th - 8th**
- **Early registration by March 21st**
- **Over 130 speakers**

HPSF
HIGH PERFORMANCE SOFTWARE FOUNDATION

https://hpsf.io

General Sessions

May 5th

- HPSF Overview
- Project Updates
- Project Updates
- Panel: Status & Trends in HPC Sys.

May 6th

- Performance & Usability
- Panel: Processor Trends
- Working Group Breakouts
- Community BoF

Project Days

May 7th

- Kokkos
- Spack
- Charliecloud
- AMReX
- Apptainer
- Trilinos

May 8th

- Kokkos
- Spack
- Charliecloud