A Large-Scale Study of MPI Usage in Open-Source HPC Applications

Ignacio Laguna, Ryan Marshall, Kathryn Mohror,
Martin Ruefenacht, Anthony Skjellum, Nawrin Sultana

THE UNIVERSITY OF TENNESSEE CHATTANOOGA
Lawrence Livermore National Laboratory

Message Passing Interface

- The de facto internode communication framework

- MPI can be divided into feature sets
  - Point-to-Point
  - Collectives
  - etc…

- Critical to application performance at scale

Why survey MPI usage?

- Implementers
  - Applications
  - Libraries
  - Tools
- Procurement
- Vendors
- MPI Forum
- But there have been surveys for decades...

Why another survey!?

- Narrow set of applications
- Only single site
- Specific to a project

- Ultimate goal was a large population sample
  - 110 open-source HPC applications
  - We want all possible
- Automation

The Analysis

The magic of analysis

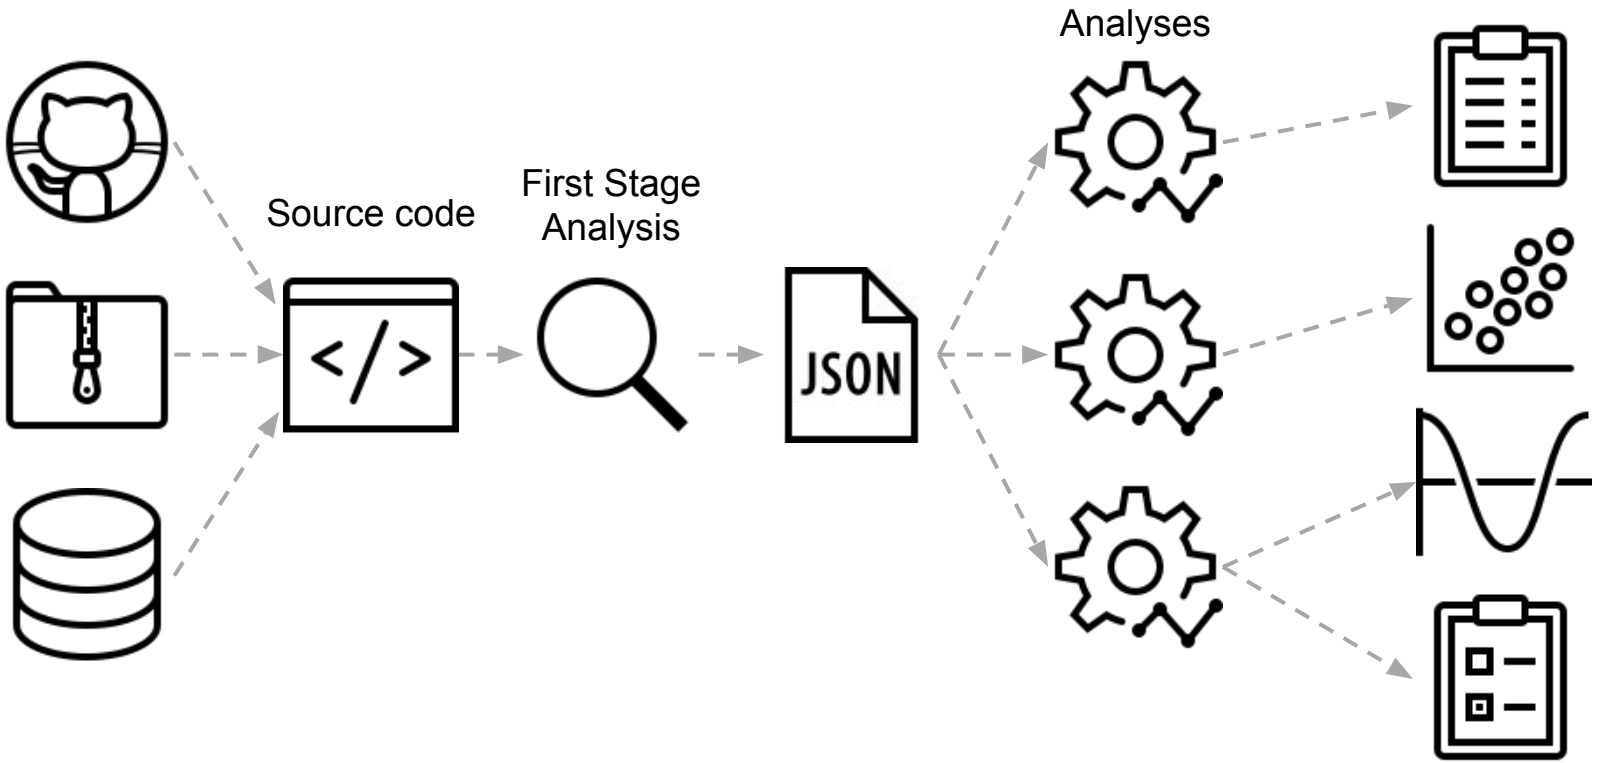

Why static analysis?

	Static	Dynamic
Dependencies	No	Yes
Compiling	No	Yes
Variability	No	Yes
Complexity	Various	Various
Coverage	Difficult	Easier
Cost	Cheap	Expensive

Input of the Analysis

- We searched manually
  - No database of HPC applications

- For secondary analysis we needed
  - Source code classification
  - Date of first and latest release
    - Dates are difficult especially the first

Analysis Limitations

- Unusual usage of MPI is not detected
- Counting call sites, not actual call count
- We use pattern matching, instead of more sophisticated methods

Some Key Insights

Surveyed software

ADCIRC
CGM
CloverLeaf
E3SM
GTC-P
HPCG
LAMMPS
MLSL
Neckbone
PICSAR
Qmcpack
SWFFT
Trilinos
amgcl
feltor
miniAero

ADIOS
CLAMR
CoMD
EigenExa
Gadget
HPGMG
LULESH
MOAB
Nek5000
Paradis
Quantum_Espresso
Samrai
Tycho2
arbor
ffm_mpi
miniFE

AMG2013
Cactus
CoSP2
Elemental
Gamess
HYPRE
Laghos
MUMPS
Nyx
PlasmaApp
SCOREC
ScaLAPACK
VMD
ascent
horovod
miniMD

ASPA
Caffe-MPI
Code_Saturne
Enzo
Genesis
HemeLB
MACSio
MeshKit
OpenFOAM
PrincetonCBEMDMPI
SNbone
Scale
VPFFT
cp2k
meep
miniSMAC2D

BioEM
Chatterbug
Comb
ExaMiniMD
Gizmo
Kratos
MCB
NAMD
OpenMD
ProfugusMC
SPPARKS
SimpleMOC
WRF
elmerfem
metag_partitioning
miniTri
phyml

CFDEMCoupling
Chombo
CoreNeuron
FLASH
Gromacs
Kripke
MITgem
NWChem
PENNANT
QBox
SW4lite
TCLB
abinit
faunus
miniAMR
miniVite
pism

yambo
thornado-mini
primecount
pypar
rocALUTION
siesta
souffle
splatt
miniXyce
mxx
nest-simulator
ntchem-mini

Only a small subset of MPI is used

Minor features are rarely used

Persistent and one-sided are not used much

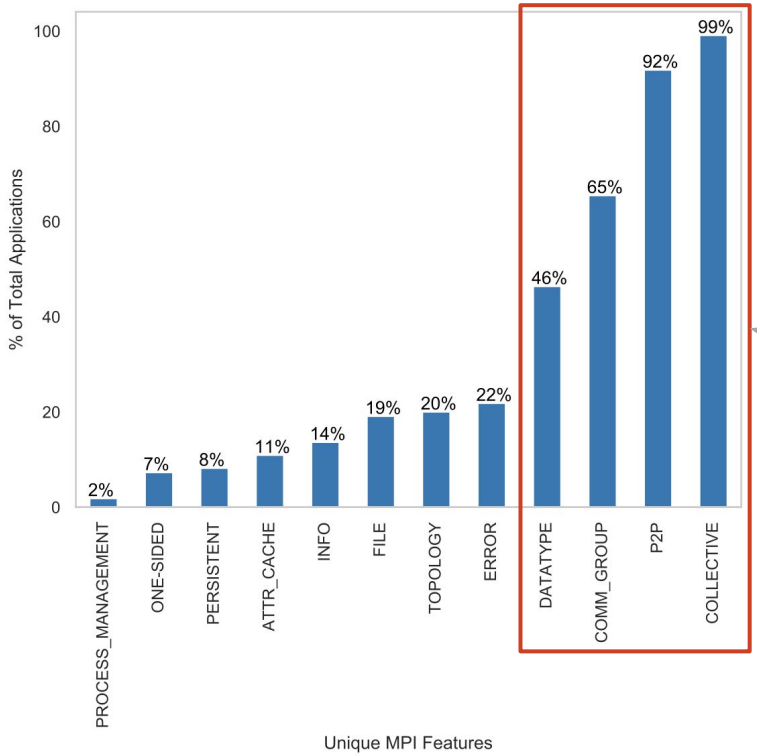

Core set is used

Advanced features are used sparingly

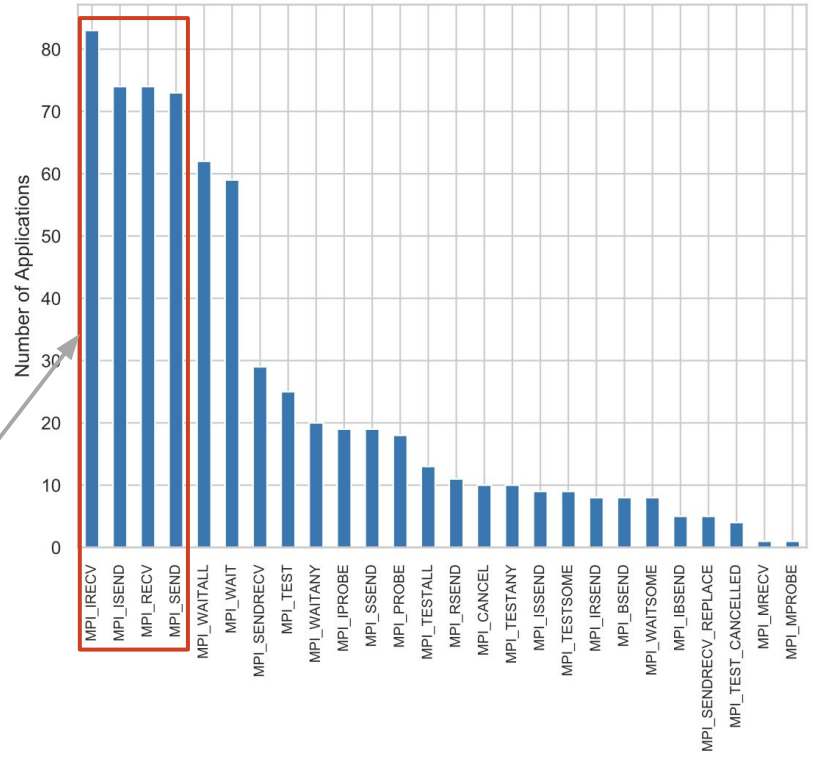

"Simple" functions are used commonly

67% use blocking send/recv

Similar effect for collective operations

MPI 1.0 is sufficient for most applications

80% of applications use are covered by MPI 2

12 applications use MPI 2 for MPI_Init_thread

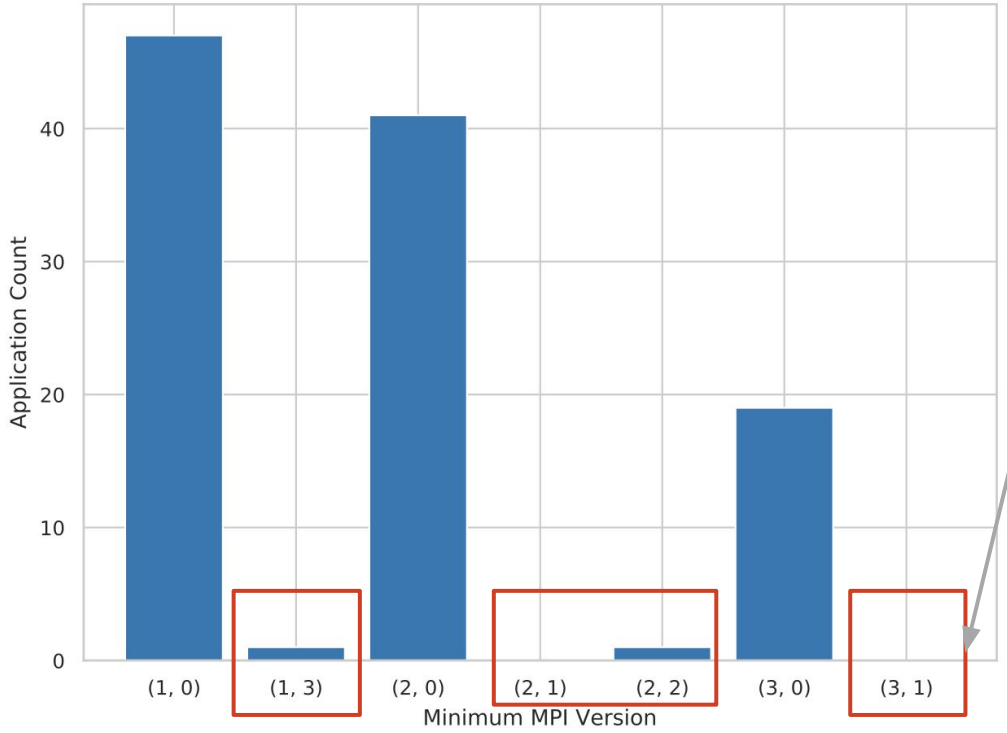

Minor MPI versions are essentially useless

OpenMP is the most used X in MPI+X

Two thirds use a mixture of programming models

Nobody likes pthreads

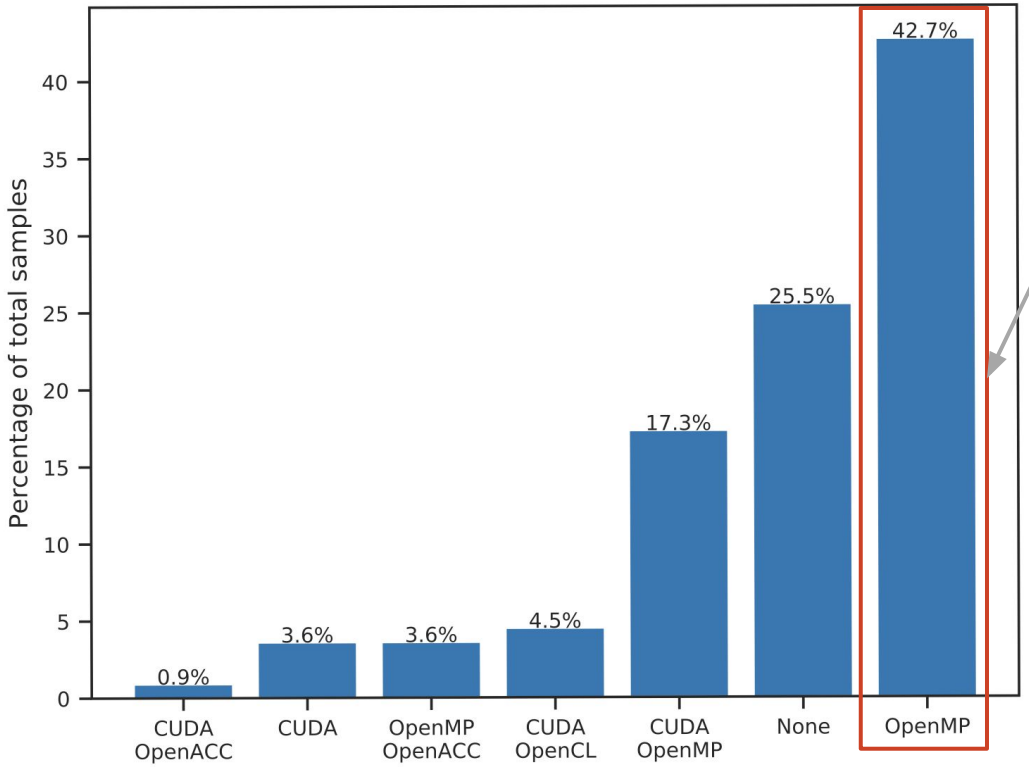

OpenMP is used most

C++ is the most used language to program MPI

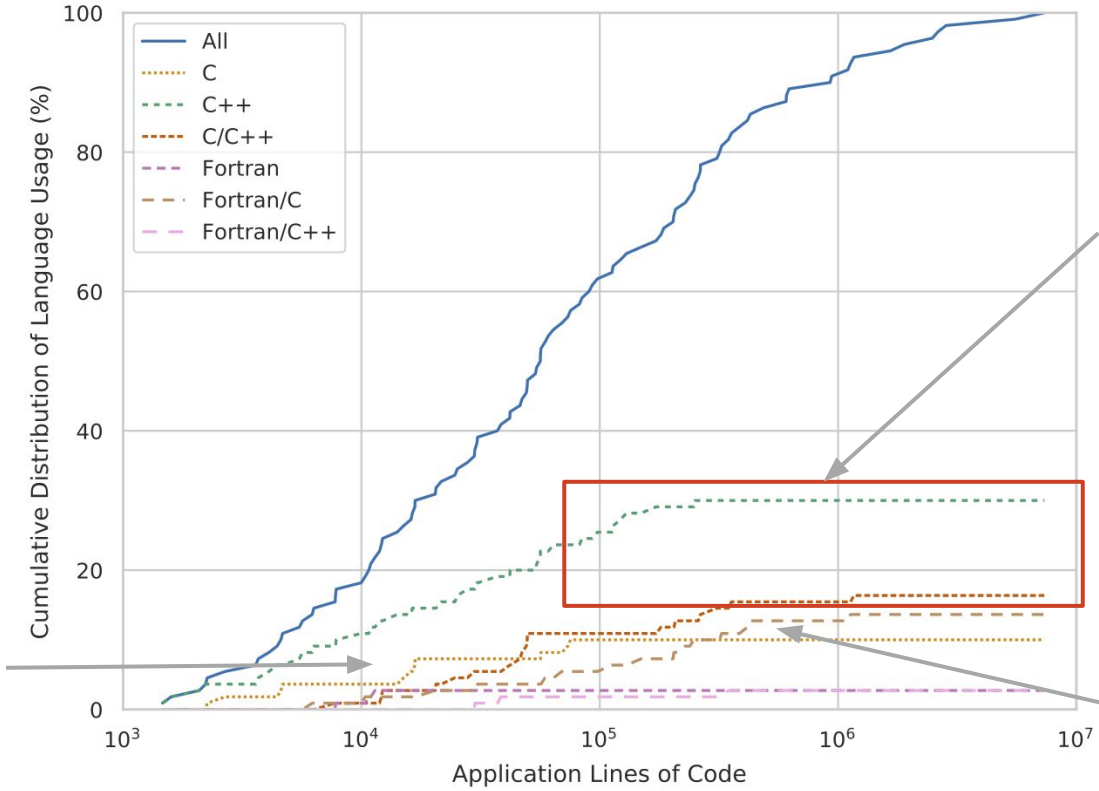

C++ is most used

Pure C is used for small or medium sized codes

Fortran is used for the largest codes

Disclaimer

- Left out obvious (in retrospect) applications & libraries
  - PETSc
  - Probably more

- Overlapping repositories
  - Charm++ included in NAMD

The Next Step

Ongoing magic!

- Improve our communal knowledge of how HPC exists
- Address the limitations of our study
- Greater number of source codes
  - Provide centralized database
- Enhance current analyses
- Expand the number of analyses
- Further automation

Submit your application for inclusion!

Applications!

Libraries!

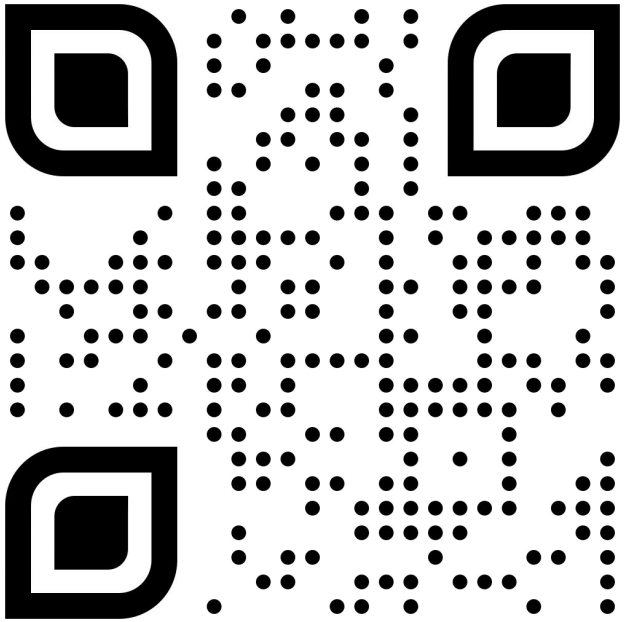

Tools!

Anything with HPC!

https://mpisurvey.org

Submit a New Program

Submit a new MPI program here.

Your email:	
A short description of the application:	
Code URL (e.g., github repo):	
Latest version:	
Date of first release:	01/11/2019
Date of latest release:	01/11/2019
Send	

Summary

- Only a small subset of MPI is used
- Advanced features are used sparingly
- MPI 1.0 is sufficient for most applications
- OpenMP is the most used X in MPI+X
- C++ is the most used language to program MPI

Summary

- Many insights from a large sample
- Automated, open-source tools for the analysis
- Ongoing project to enhance everyone's knowledge

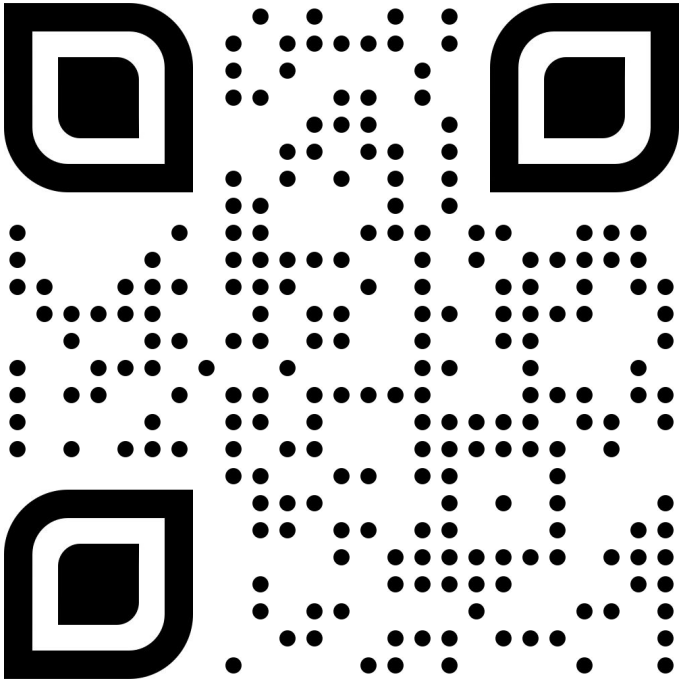

https://mpisurvey.org