The Phantom of the Future

Terrence Tricco
Memorial University of Newfoundland

How did we get here?

2010 – Initial Phantom

JOURNAL ARTICLE

On the diffusive propagation of warps in thin accretion discs FREE

Giuseppe Lodato ✉, Daniel J. Price

Monthly Notices of the Royal Astronomical Society, Volume 405, Issue 2, June 2010, Pages 1212–1226, https://doi.org/10.1111/j.1365-2966.2010.16526.x

Published: 11 June 2010 **Article history**

PDF Split View Cite Per

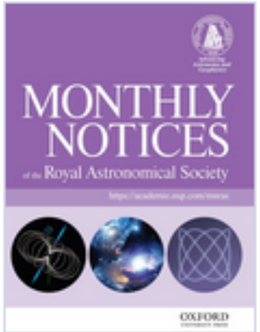

JOURNAL ARTICLE

A comparison between grid and particle methods on the statistics of driven, supersonic, isothermal turbulence FREE

Daniel J. Price ✉, Christoph Federrath

Monthly Notices of the Royal Astronomical Society, Volume 406, Issue 3, August 2010, Pages 1659–1674, https://doi.org/10.1111/j.1365-2966.2010.16810.x

Published: 03 August 2010 **Article history** ▾

PDF Split View Cite Permissions Share ▾

2010 – Initial Phantom

- *"written especially for studying non-self-gravitating problems"*

- Only hydrodynamics! No gravity, no other physics.
- No tree! Neighbour finding using fixed grid and linked lists.
- Density and force in a single loop!

2018 – Phantom Public Release

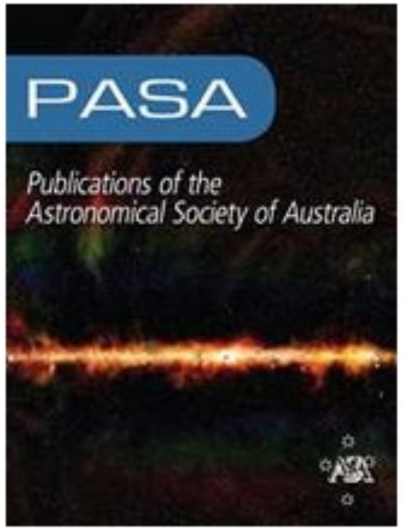

Publications of the Astronomical Society of Australia

PHANTOM: A Smoothed Particle Hydrodynamics and Magnetohydrodynamics Code for Astrophysics

Published online by Cambridge University Press: **25 September 2018**

Daniel J. Price, James Wurster, Terrence S. Tricco, Chris Nixon, Stéven Toupin, Alex Pettitt, Conrad Chan, Daniel Mentiplay, Guillaume Laibe, Simon Glover, Clare Dobbs, Rebecca Nealon, David Liptai, Hauke Worpel, Clément Bonnerot, Giovanni Dipierro, Giulia Ballabio, Enrico Ragusa, Christoph Federrath, Roberto Iaconi, Thomas Reichardt, Duncan Forgan, Mark Hutchison, Thomas Constantino, Ben Ayliffe, Kieran Hirsh and Giuseppe Lodato

Show author

<u>2018</u> – Phantom Public Release

- “*initial release of Phantom has been developed with a focus on stellar, planetary, and Galactic astrophysics, as well as accretion discs.*”
- “*MHD, dust-gas mixtures, self-gravity, and a range of other physics.*”

- Phantom had been used in ~15-20 publications by this point.
- Tree-based neighbour finding!
- Density and force in separate loops!
- Range of initial setups and phantomanalysis.

2018 – First Phantom Workshops

1st Phantom Users Workshop – Melbourne

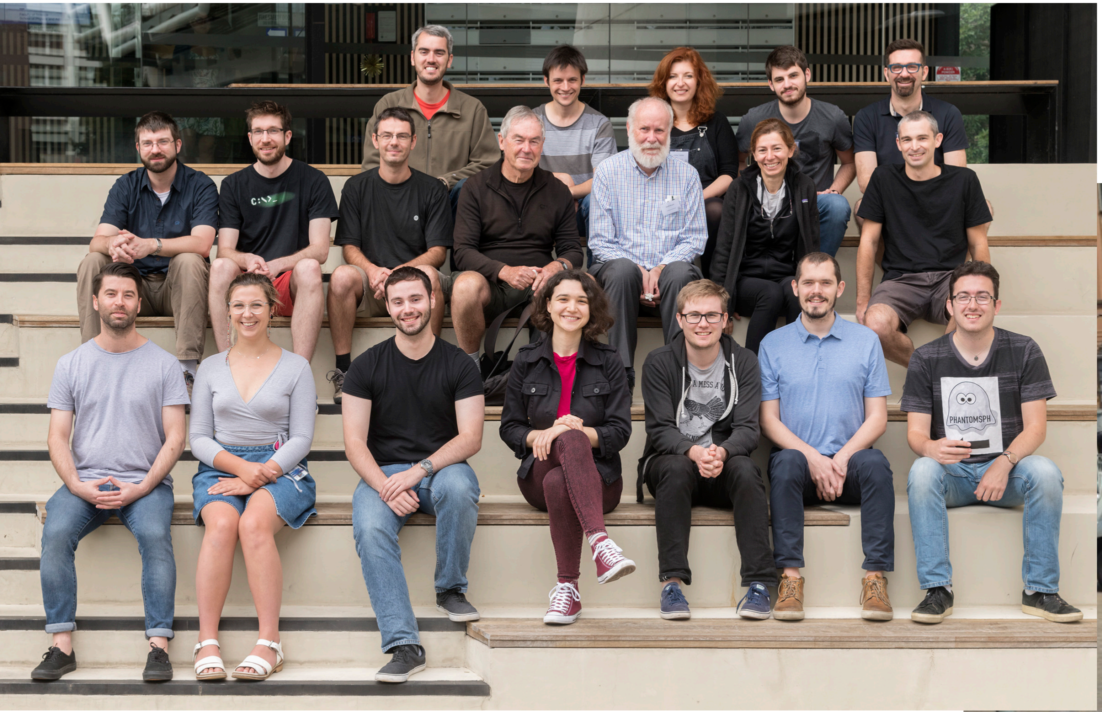

1st Phantom European Users Workshop – Milan

2024 – Phantom

- Phantom used for dozens of research applications.
- Code paper has hundreds of citations.

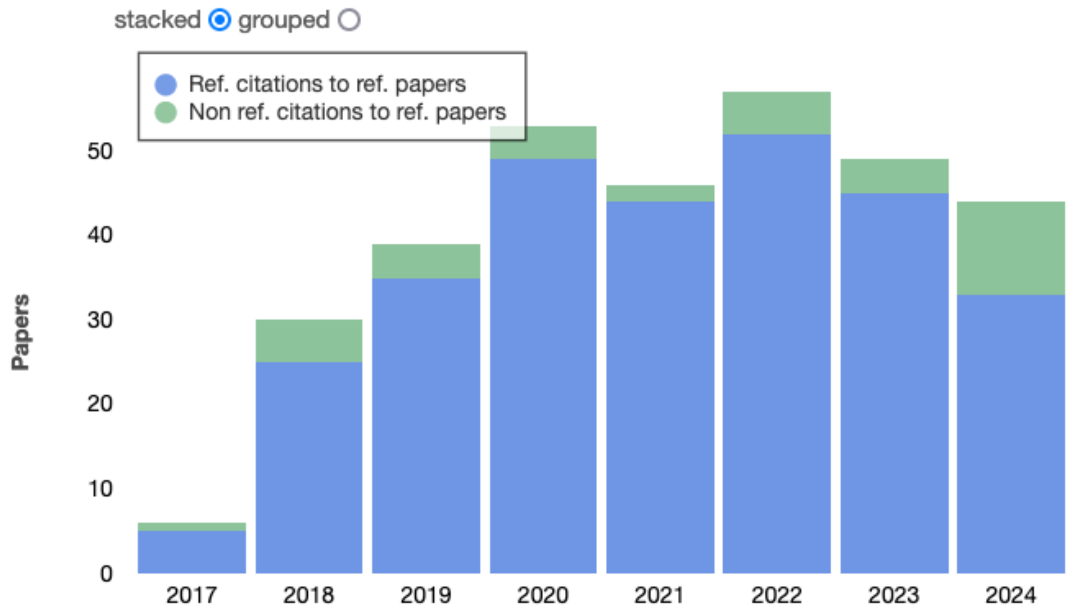

2024 – Phantom

- 100+ initial setups available.

SETUP=
BHL
adiabaticdisc
alfven
asteroidwind
balsarakim
binary
binarydiscMFlow
blob
bondi
cluster
common
converging
default
disc
dustsettle
dustybox
dustydisc
dustyisosgdisc
dustysedov
dustysgdisc
dustyshock
dustystar
dustywave
empty
evrard
exoALMA
firehose
flrw
flrwpspec
galaxies
galcen
galdisc
galdiscmhd
gr_testparticles
grbondi
grbondi-inject
grdisc
growingdisc
growthtomulti
grstar
grtde
gwdisc
hierarchical
ismwind
isosgdisc
isowind
jadvect
jet
jetdusty
jetnimhd
kh
lightcurvedisc
mhdblast
mhdrotor
mhdshock
mhdsine
mhdvortex
mhdwave
neutronstar
nimhdshock
nshwdisc
nsmerger
orstang
planetatm
planetdisc
polytrope
prtest
quebec
raddisc
radiativebox
radiotde
radshock
radstar
radwind
sedov
sgdisc
shock
solarsystem
sphereinbox
srblast
srpolytrope
srshock
star
taylorgreen
tde
test
test2
testcyl
testdust
testgr
testgrav
testgrowth
testkd
testlum
testnimhd
testparticles
tokamak
torus
turb
turbdrive
wave
wavedamp
wd
wddisc
wind
windtunnel

2024 – Phantom

- Included physics has been expanded (GR, radiation, etc)
- CI/CD pipelines with automated Github actions.
- Documentation hosted on read the docs.
- Pull request templates.

2024 – Phantom

1st North American Phantom Users Workshop

Our 6th Phantom Users Workshop!

Where do we go next?

Phantom Userbase

Has user adoption peaked?

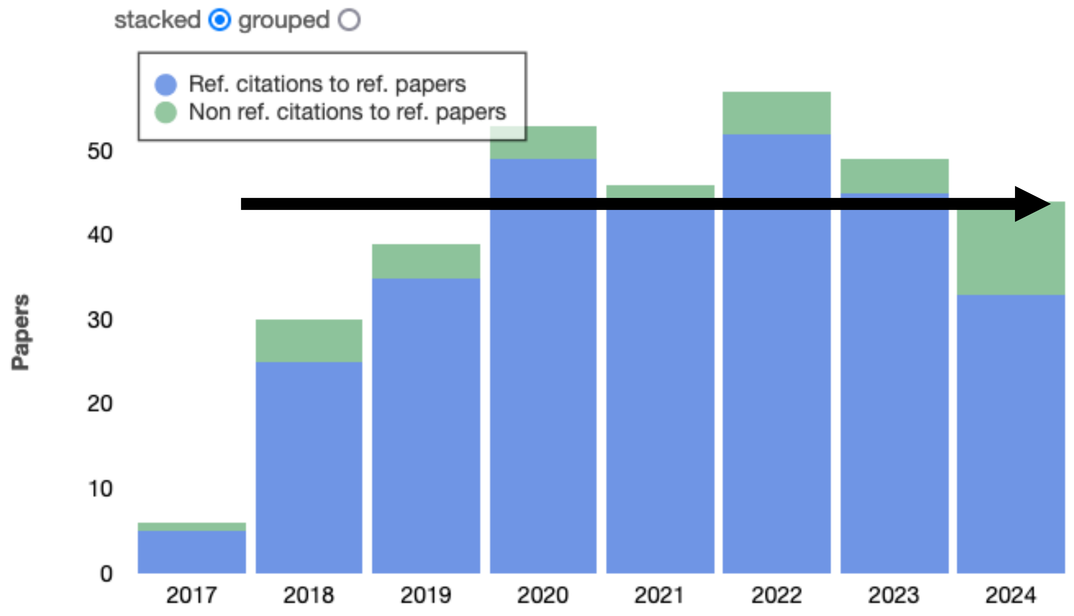

Phantom Userbase

Has user adoption peaked?

If so, **why?**

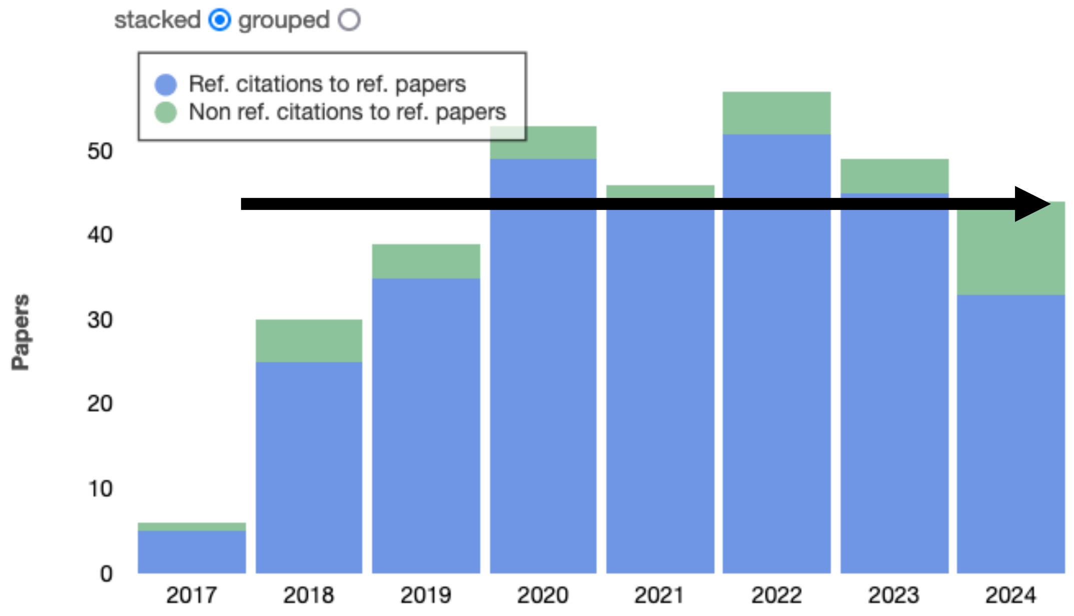

Phantom Scaling

75% of papers that used Phantom used **2.5 million particles** or fewer.

75% of papers used **2 million particles** or fewer.

2018: 18 papers
2024: 27 papers

Top 500 Supercomputers

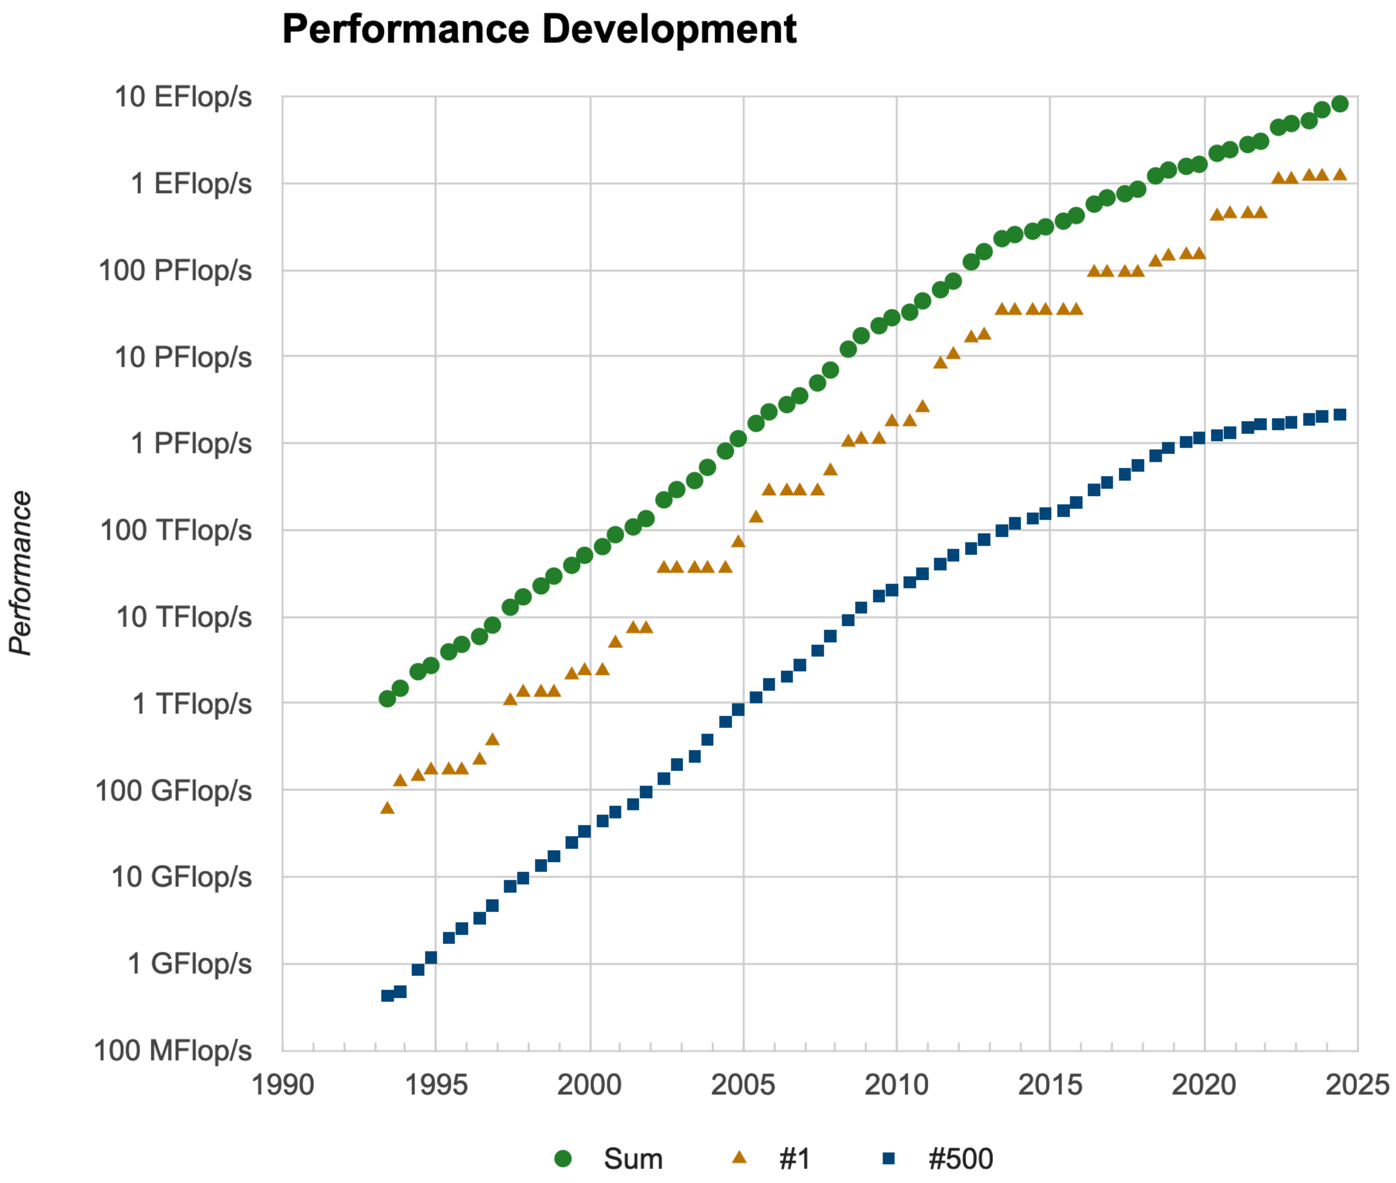

Phantom Scaling

Here's the problem – Phantom only uses OpenMP parallelization.

* MPI implementation available, but is not robust?

Top 500 Supercomputers

- Computing power is exponentially increasing (~doubling every 2-3 years).

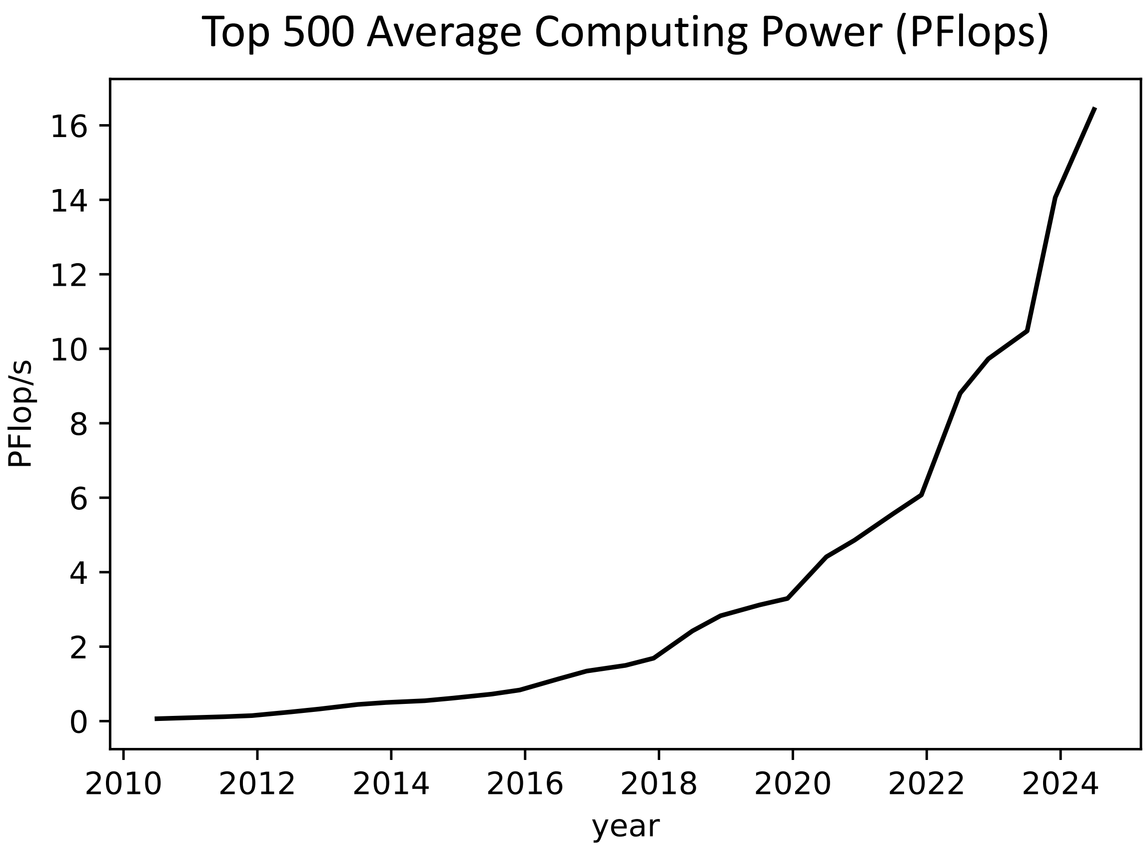

Top 500 Supercomputers

- Cores per node increasing much more slowly.

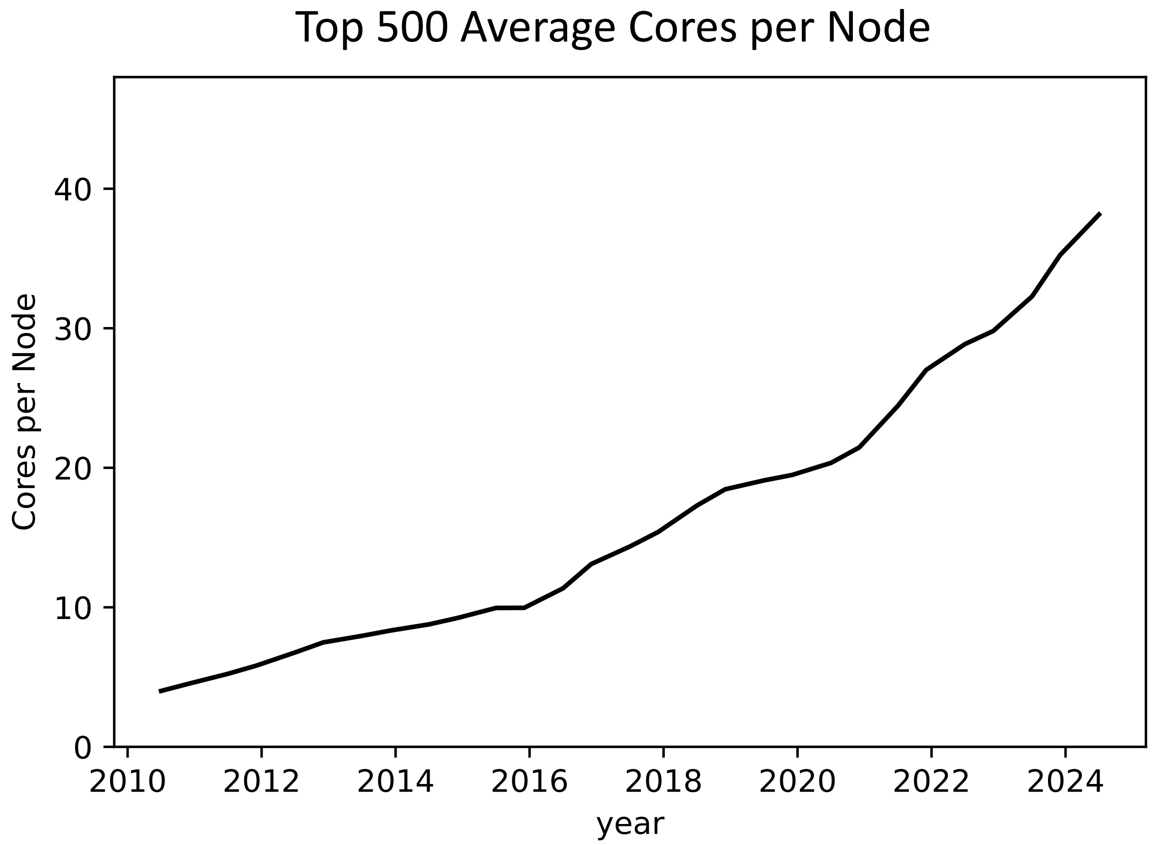

Top 500 Supercomputers

2018 average supercomputer:

- 2.4 Pflops
- 17 cores/node

2024 average supercomputer:

- 16 Pflops (8x increase)
- 38 cores/node (2x increase)

The main driver of increasing compute power is **not** increasing cores/node!

Phantom Scaling

Problem #1: Clusters continue to increase in node count.

Phantom Scaling

Problem #1: Clusters continue to increase in node count.

Problem #2: Clusters continue to rely on larger complements of GPUs.

Top 500 Supercomputers with GPUs

- Heterogenous architectures – 40% of clusters in 2024 have a complement of GPUs.

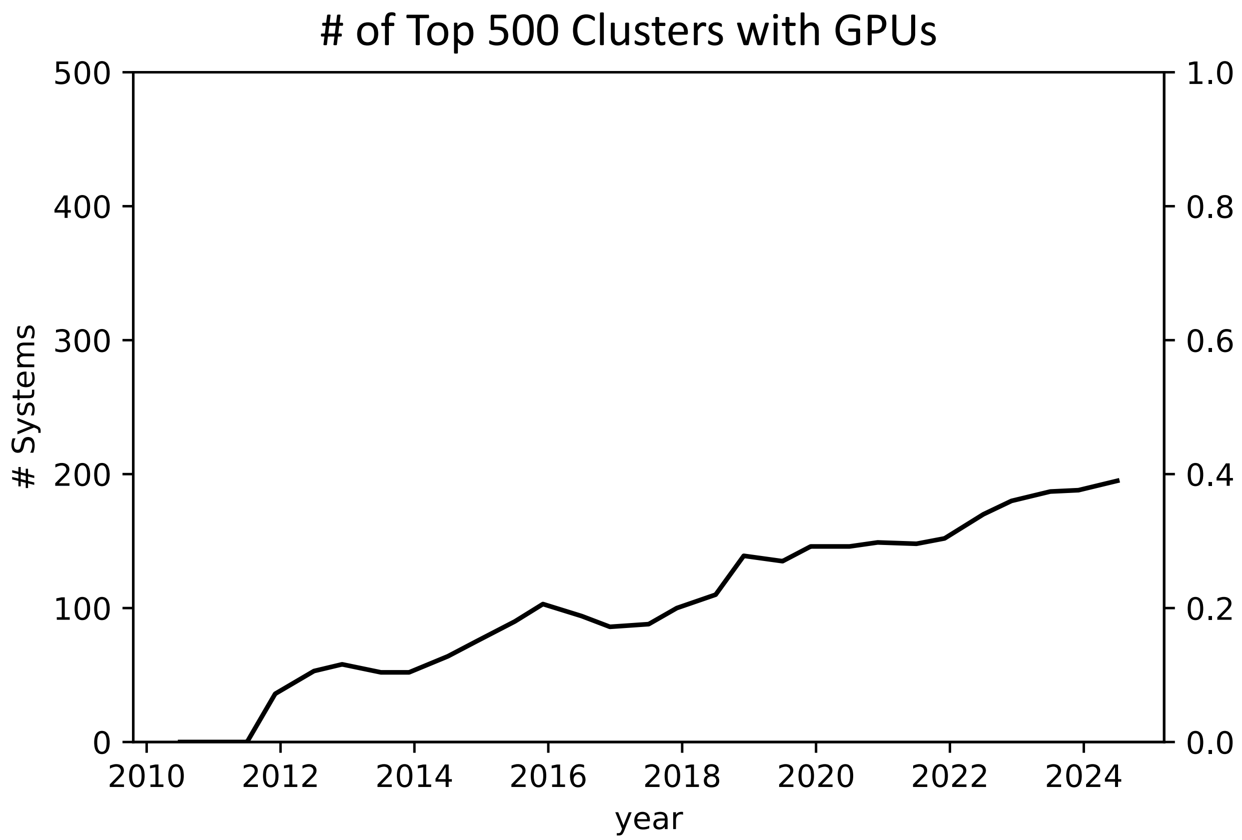

Phantom Scaling

Phantom needs to move beyond just OpenMP parallelization.

- Cannot rely on scaling inherently with increasing core counts.
- Will need MPI or other solutions to scale onto many node clusters.
- Will need GPU capability to fully utilize available hardware.

MPI / CUDA / etc

- Implementing MPI, CUDA or another solution is a big task.
- Not just from the refactoring required of the code… (which is huge)
- But also how to create implementations that handle all the various physics?
- Individual timestepping?
  - 100 particles evolving on the smallest timestep won't scale to 10,000 cores.

2030 – Phantom of the Future

What does Phantom look like in 5 years time?

2030 – Phantom of the Future

1. Development cycle (CI/CD)

2. Initial conditions / Initial setups

3. Output analysis

4. Governance structure

2030 – Phantom of the Future

1. Development cycle (CI/CD).
   - Github actions, unit tests, official release cycles.

2. Initial conditions / Initial setups
   - No more phantomsetup. Python-based initial setups?

3. Output analysis
   - No more phantomanalysis. Python-based analysis (Sarracen!)

4. Governance structure
   - Longevity and management of the codebase.

Input / Output to Phantom

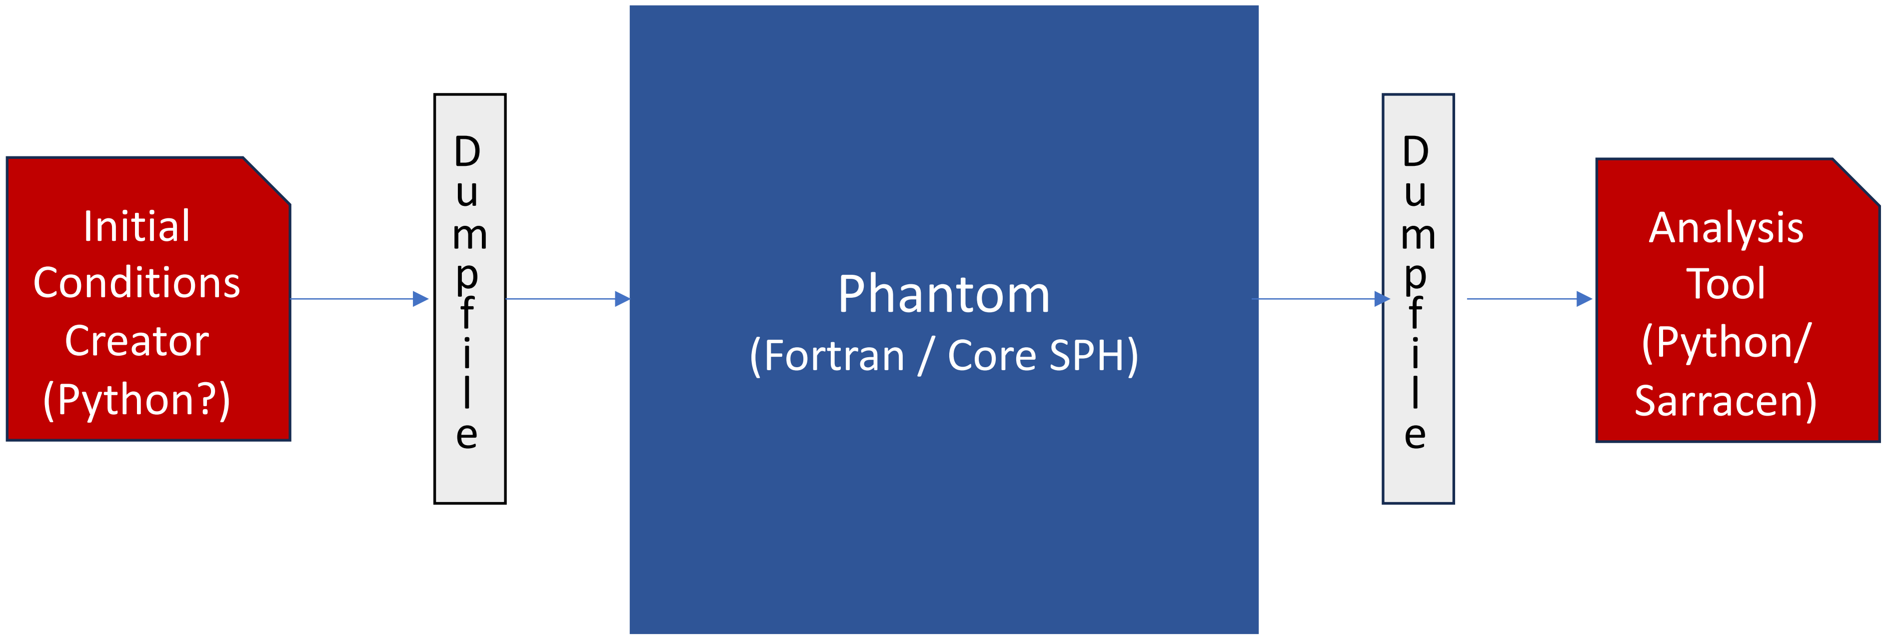

- No more phantomsetup, phantomanalysis, phantommoddump.
- Phantom is a pure SPH engine:
  - Accepts input dumpfile.
  - Crunches the SPH.
  - Returns output dumpfiles.

Input / Output to Phantom

- Fortran is a painpoint.
- (Daniel once called it the language of the gods; I disagree!)

- Fortran is efficient – good for scientific simulations.
- But cumbersome to work with – bad for tasks that don't require performance (e.g., interactivity).

- Non-simulation components of the code should move out of Fortran (preferably Python?).

Machine Learning

- Growing presence of ML and AI will inevitably intersect with traditional numerical simulations.

- Already are seeing a growing use of PINNs (physics informed neural networks) to generally solve physics equations.

- Will traditional simulations be needed if AI can do the job faster and more accurately?

Machine Learning

Will AI replace SPH?

Machine Learning

- It remains to be shown that AI can outperform traditional simulations.
- Computational cost vs accuracy?

- I believe that ML will provide a complementary solution to traditional astrophysics simulations.

- SPH and grid-based codes are able to co-exist.
- Multiple distinct solvers that agree on a solution is a positive thing!

Machine Learning + SPH

- Some aspects of SPH simulations that may be improved by ML:
  - AI Kernels (better accuracy for cheap).
  - Shock capturing (artificial viscosity / replace switches).
  - Sub-grid physics (any sort of heuristic algorithm).
  - Data / workload decomposition for parallel optimization.
  - …

Summary

- Phantom is a great code with a great community.
  - Regular users workshops. Used for many research projects.
  - Respected as a trustworthy code that is easy to use.

- Biggest shortcoming of Phantom is its scalability.
- The OpenMP implementation scales extremely well…
- But relying just on OpenMP isn't viable in the long term.

Phantom of the Future

1. Phantom's Fortran base is only for simulation.
2. SPH base code scales onto modern cluster hardwares.
3. Multi-GPU capability.
4. Initial conditions creation and output analysis handled in Python.
5. Documentation rewrite with new user guides, developer guides, feature implementations, etc.
6. Leveraging ML / AI to create future SPH algorithms.