ADACS + Phantom

Phantom Users Workshop 2023

Conrad Chan, David Liptai

ADACS

ASTRONOMY DATA AND COMPUTING SERVICES

ADACS + Phantom

Phantom Users Workshop 2023

- Purpose of this talk: Bring developers up to speed on ADACS' involvement in Phantom
- Testing via modern CI/CD
- MPI + OpenMP Parallelism: a deep dive into the latest developments

Why Phantom and ADACS are a great match for each other

Who are ADACS?

(in the context of Phantom)

- We provide software development services to astrophysicists like you https://adacs.org.au/software-services/
- Our team is comprised of:
  - Former astrophysicists - we understand the science and scientific workflows!
  - Software developers - promoting the best development practices

Why Phantom and ADACS are a great match for each other

- Let the PhD students and researchers focus on their science problems, while we worry about code problems and performance optimisations
- ADACS developers are also system administrators at OzSTAR - this puts us in a unique position to best optimise and test Phantom on HPC hardware
- We love working on Phantom because:
  - The developers understand the importance of testing in a scientific workflow
  - It uses interesting algorithms to solve interesting problems

Contributions to Phantom

GitHub Pull Requests

- #242 - MPI support for GR
- #247 - OpenMP critical section optimisation
- #259 - Optimisations for running MPI with a single task
- #262 - MPI and OopenMP memory allocation optimisations
- #310 - OpenMP-parallel MPI send (*in progress!*)
- #316 - Particle balancing optimisation
- #217, #220, #322 - MPI testing
- #138, #223, #224, #229, #317 - CI/CD and pre-commit support
- #132, #153, #157, #158, #160, #164, #165, #206, #209, #212, #222, #229, #232, #234, #235, #237, #243, #256, #265, #269, #320 - Bug fixes
- #146, #251, #254, #258, #261, #270 - Code refactoring

Test as You Fly, Fly as You Test

The process of end-to-end system verification may be compromised when it is not consistent with the mission profile

NASA / JPL: http://www.jpl.nasa.gov/pictures/solar/mplartist.html

NASA / JPL: https://llis.nasa.gov/lesson/1196

Most scientific codes and scripts

Early Phantom

Current Phantom

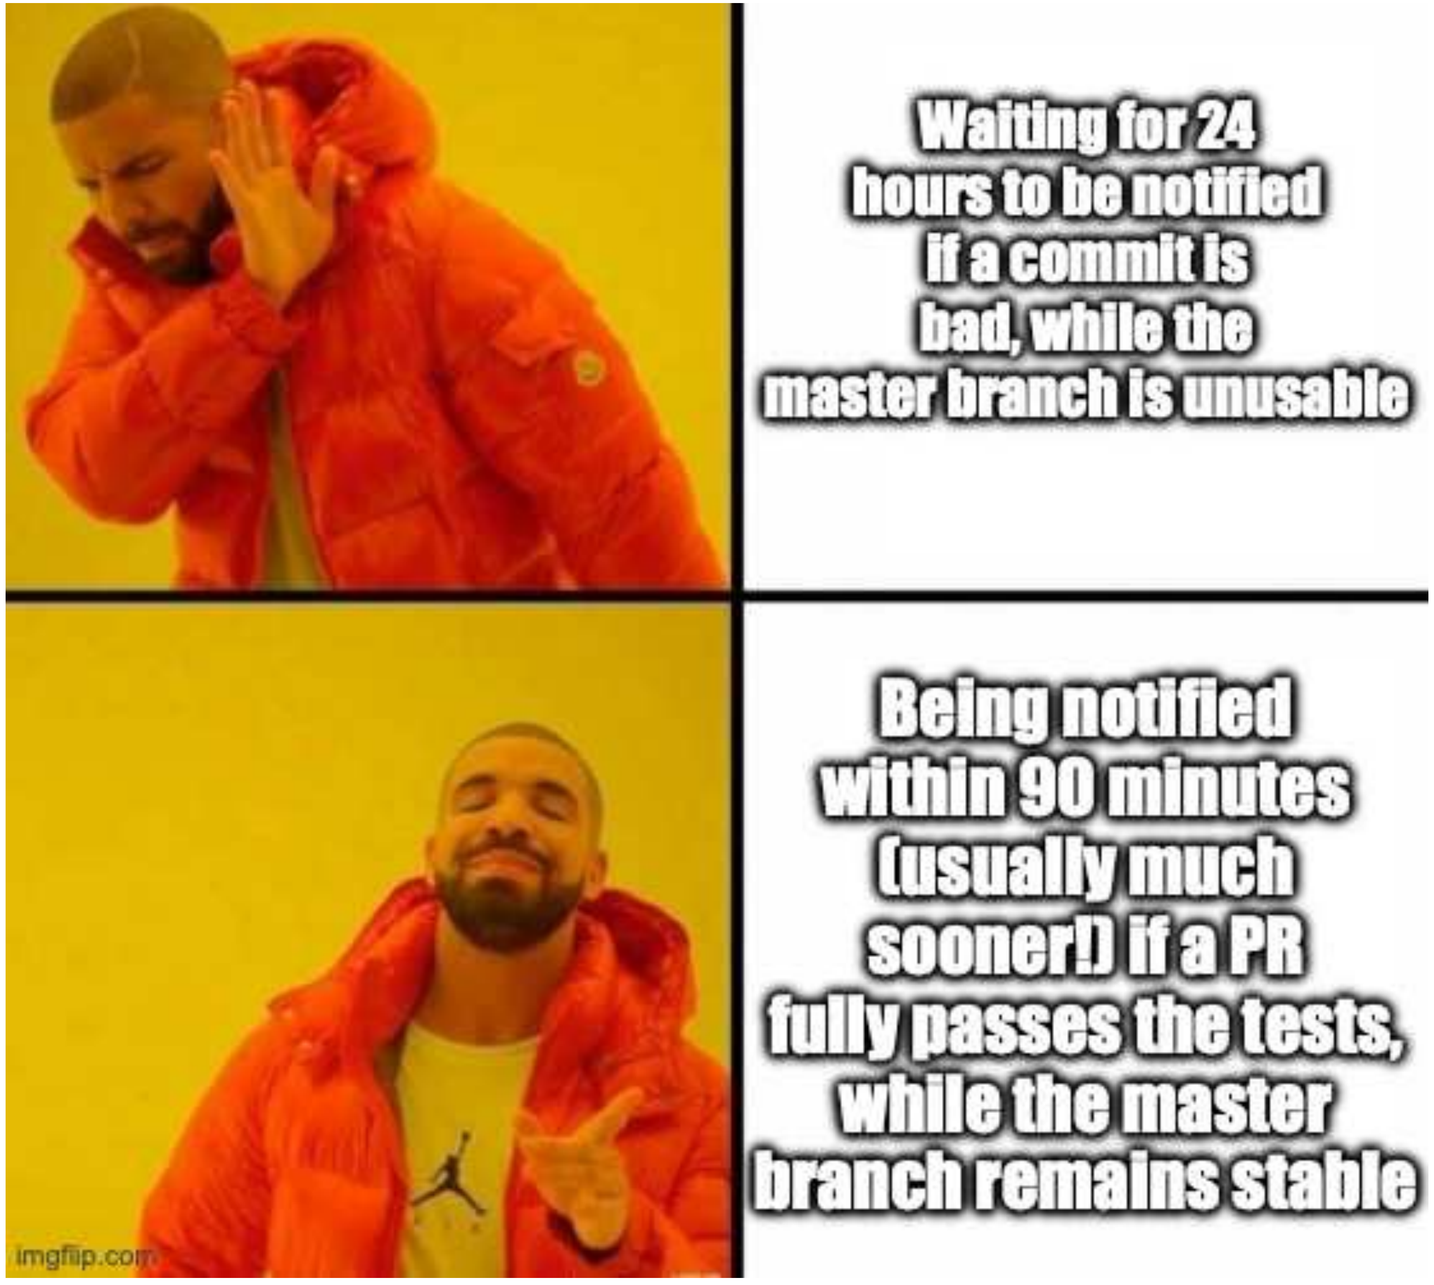
Waiting for 24 hours to be notified if a commit is bad, while the master branch is unusable
Being notified within 90 minutes (usually much sooner!) if a PR fully passes the tests, while the master branch remains stable
imgflip.com

Rigorous testing provides a framework to overhaul the core algorithms

MPI + OpenMP

Multiple Levels of Parallelism

- **Goal**
  Use MPI to scale beyond more than one node while preserving the performance benefits of OpenMP within nodes

- **What this looks like**
  Each node runs 1 MPI task, with each task running multiple (e.g. 32) OpenMP threads

- **Challenge**
  How can each thread send MPI messages under the same task?

MPI + OpenMP

Multiple Levels of Parallelism

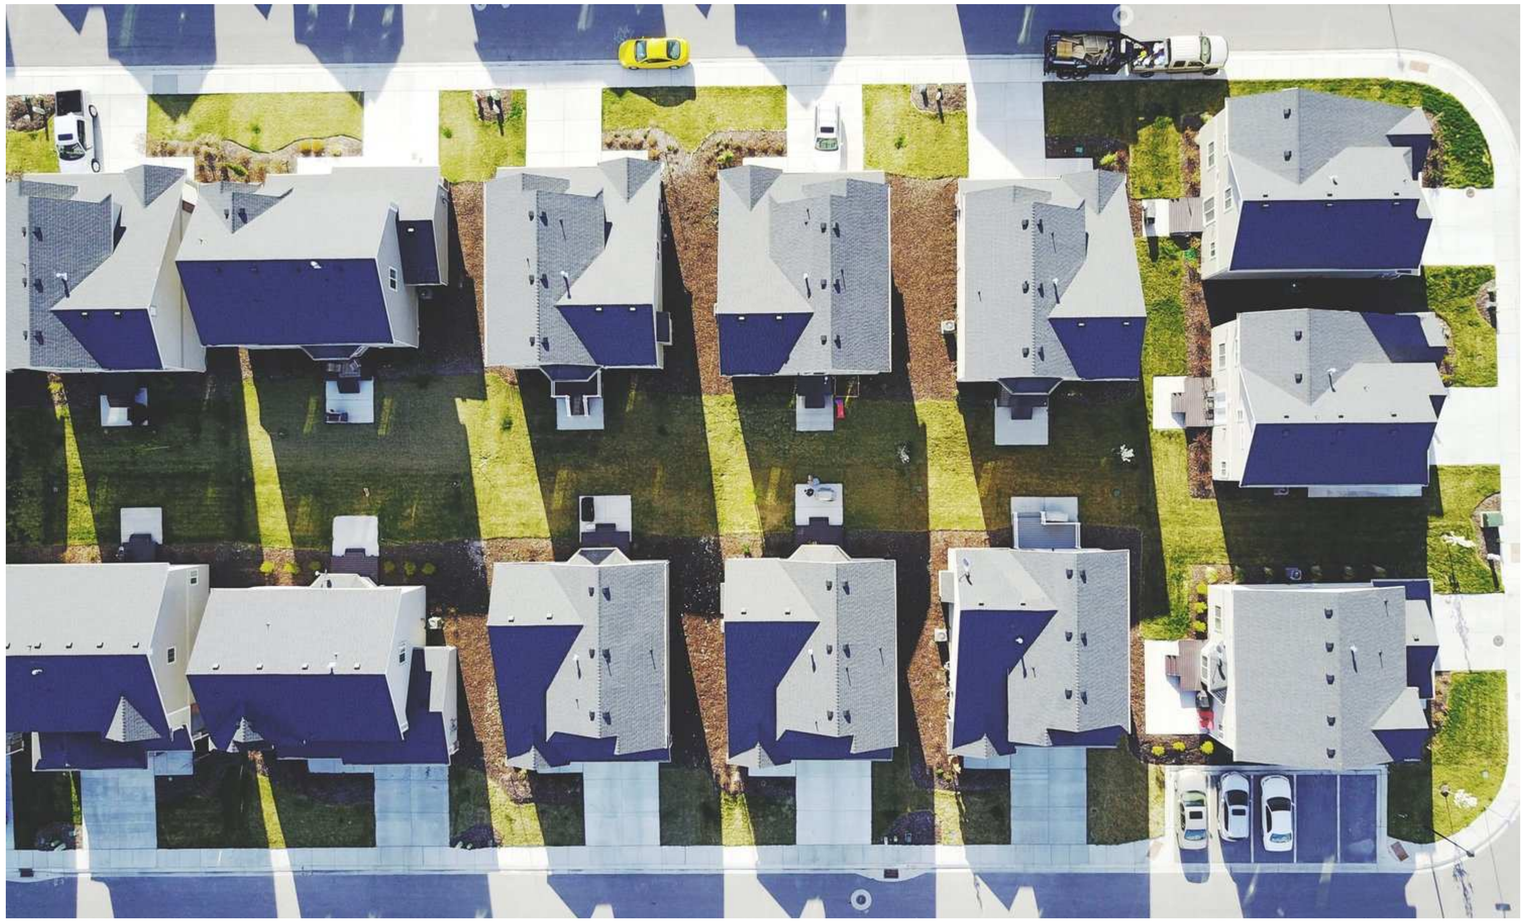

MPI + OpenMP

How can each thread send MPI messages under the same task?

Currently (serial communication)

- Each thread waits its turn until the communication buffer is free to be used
- Relies on OpenMP critical sections, which don’t scale well with large numbers of threads

Proposed solution (parallel communication)

- Each thread has its own communication buffer and can make MPI calls independently
- MPI_THREAD_MULTIPLE makes this possible - however, it may incur performance overheads and is not supported everywhere
- Work in progress! (PR #310)

Roadmap to better parallelism

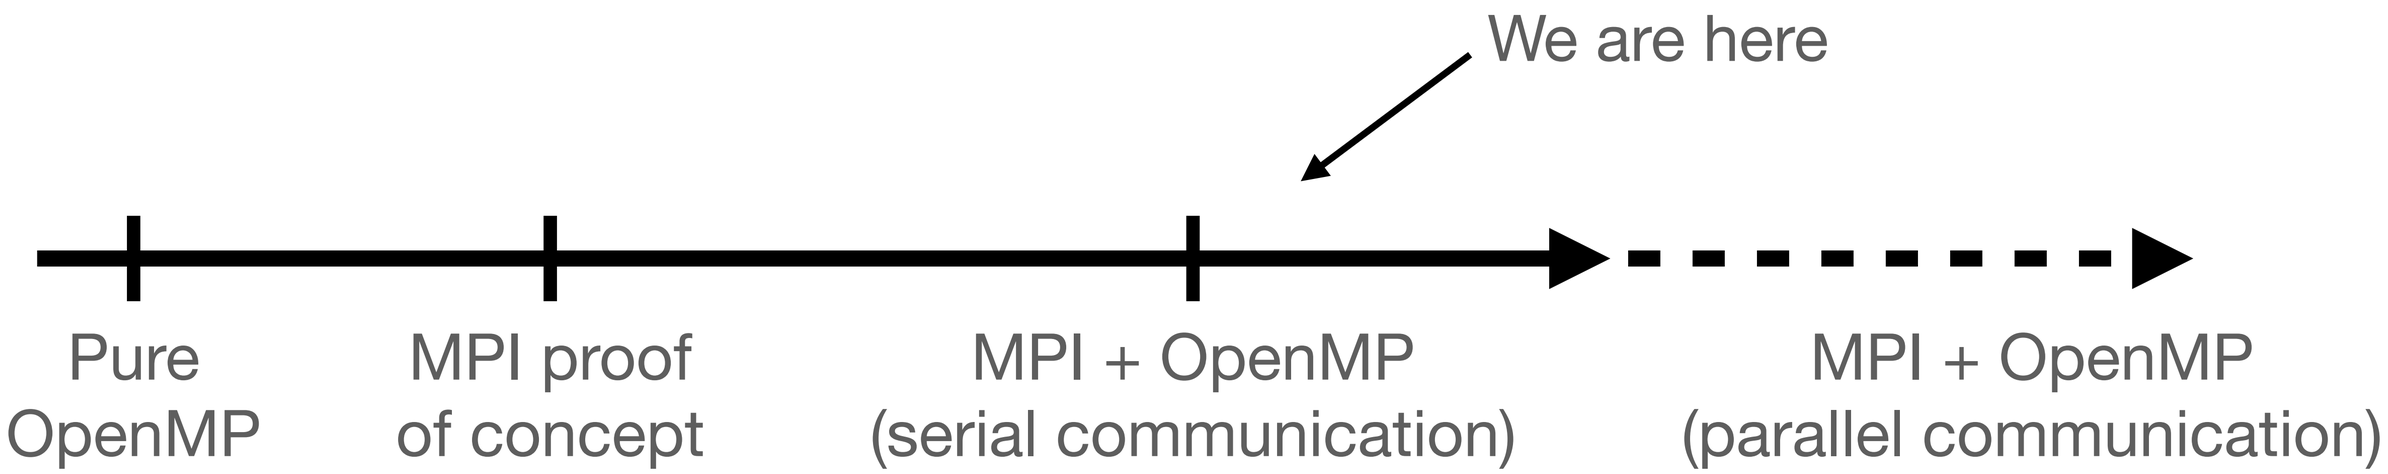

If…

- Your work would benefit from parallelism in Phantom
- You have ideas for how to further improve performance in Phantom
- Or you’d like to learn more about ADACS’ software support program

We’d love to hear from you!
Come and chat with us during the hack sessions