ADACS

ASTRONOMY DATA AND COMPUTING SERVICES

NCRIS
National Research Infrastructure for Australia
An Australian Government Initiative

Astronomy Australia Ltd.

SWINBURNE
SWINBURNE UNIVERSITY OF TECHNOLOGY

Curtin University

PAWSEY
supercomputing centre

ADACS Introduction

- **Vision:**
  - ***astronomy-focused*** training, support and expertise to maximise scientific return on investments in astronomical data & computing infrastructure
  - 3 service components:
    1. Training
    2. Data & eResearch
    3. Computing access & support
- **Two nodes:**
  - Swinburne University (Melbourne)
  - Curtin University + Pawsey Supercomputing Centre (Perth)
- **Commenced operations March 2017**
- **Funded by Astronomy Australia Limited (AAL) through the astronomy National Collaborative Research Infrastructure Strategy (NCRIS) allocation**

ADACS BoF eResearch Australasia

Delivering Software Solutions to Astronomy Researchers

- Goal: Explore the relationship between software engineer and researcher when delivering professional software services to the academic community

Agenda

- Overview of the ADACS software support program
- Case study from both perspectives
- Discuss:
    - Ingredients for effective delivery?
    - What is best practice?
    - Lessons learned? Mistakes made?
    - How do we ensure long-term success for such programs?
    - Should we consider skill-sharing across like-minded programs?

ADACS

ASTRONOMY DATA AND COMPUTING SERVICES

Software development services:

Overview

Software Development Services

- Merit-based allocation of professional software development
- Program is in its 3rd semester (2/yr)

- Developers with expertise covering:
    - System analysis and design,
    - Scientific computing,
    - High performance computing,
    - Data science,
    - Web development,
    - Large-scale scientific databases,
    - Cloud computing, and
    - Scientific visualisation.

... integrated with suitably skilled astronomers to ensure adequate domain knowledge.

Source: imgglobalinfotech.com

Software Development Services

ADACS
ASTRONOMY DATA AND COMPUTING SERVICES

- Merit-based allocation of professional software development
- Program is in its 3rd semester (2/yr)

- Methodology:
  1. Users respond to calls for Expressions of Interest (EoIs); 1 page description of project
  2. ADACS developer interviews applicant to coax-out detailed technical specifications
  3. Once all interviews are complete, ADACS meets as a team to develop an assessment of required development time and skills required
  4. Users complete detailed application and quote the ADACS assessment for their project
  5. An independent time allocation committee (TAC) selects projects to be supported, reconciling requested and available resources

Source: imgglobalinfotech.com

Software Development Services

ADACS
ASTRONOMY DATA AND COMPUTING SERVICES

- Merit-based allocation of professional software development
- Program is in its 3rd semester (2/yr)

- 5 projects completed:

1. Web front-end to a rotation curve fitting code
2. GPU acceleration of the photon-diffusion physics of a galaxy-formation model
3. Automated data reduction pipeline for for an optical telescope
4. ***GPU acceleration of a binary black hole gravity wave model code****
5. Prototype web app for a machine learning citizen science project (Radio Galaxy Zoo)

Legend

Web application

GPU optimisation

* *We will look at this project in more detail later*

Software Development Services

ADACS
ASTRONOMY DATA AND COMPUTING SERVICES

- Merit-based allocation of professional software development
- Program is in its 3rd semester (2/yr)

- 5 projects underway:
  1. GPU acceleration of a model-emulator/ parameter-estimation code
  2. Web interface for performing gravity wave analysis/parameter-estimation
  3. Radio telescope survey team web app
  4. GPU acceleration of a galaxy clustering analysis code
  5. GPU optimisation of a radio telescope data calibration pipeline

Legend

Web application

GPU optimisation

Software Development Services

ADACS
ASTRONOMY DATA AND COMPUTING SERVICES

- Merit-based allocation of professional software development
- Program is in its 3rd semester (2/yr)

- 5 projects underway:
  1. GPU acceleration of a model-emulator/ parameter-estimation code
  2. Web interface for performing gravity wave analysis/parameter-estimation
  3. Radio telescope survey team web app
  4. GPU acceleration of a galaxy clustering analysis code
  5. GPU optimisation of a radio telescope data calibration pipeline

Legend
Web application
GPU optimisation

All but one selected project has been a web application or a GPU optimisation project

Software Development Services

ADACS
ASTRONOMY DATA AND COMPUTING SERVICES

- Some general issues encountered:
  1. Web applications and GPU optimisations are dominating supported projects
     - Turns-out: astronomers have a lot of good web application ideas/needs
     - UX expertise?
  2. Who owns/is-responsible for code developed?
  3. Ongoing support?
     - Burden grows with time if we offer support
     - On the other hand: don't want to be building tools that fall into states of disrepair
  4. Role of TAC
     - What influence should the TAC have on policy?
  5. Project management
     - Agile development practices, etc.?
     - Astronomers aren't necessarily used to the language, let alone the methodologies
     - They are generally accustomed to a great deal of control over development work

Source: imgglobalinfotech.com

ADACS

ASTRONOMY DATA AND COMPUTING SERVICES

Software development services:

Case study - project background

Some background on gravitational waves

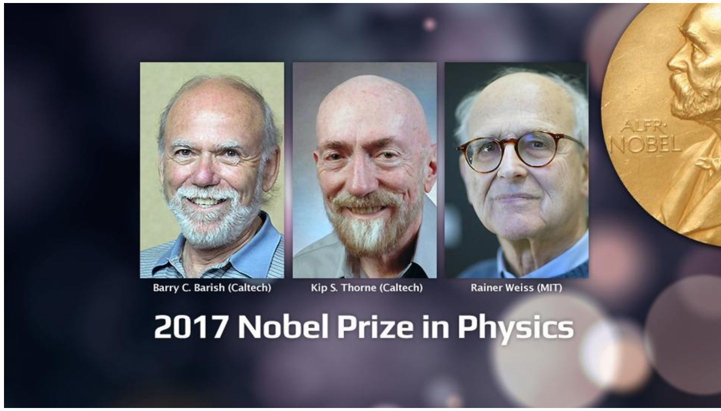

Gravitational wave astronomy:

* Colliding black holes and neutron stars
* Probe extreme gravity
* Ultra-dense nuclear matter

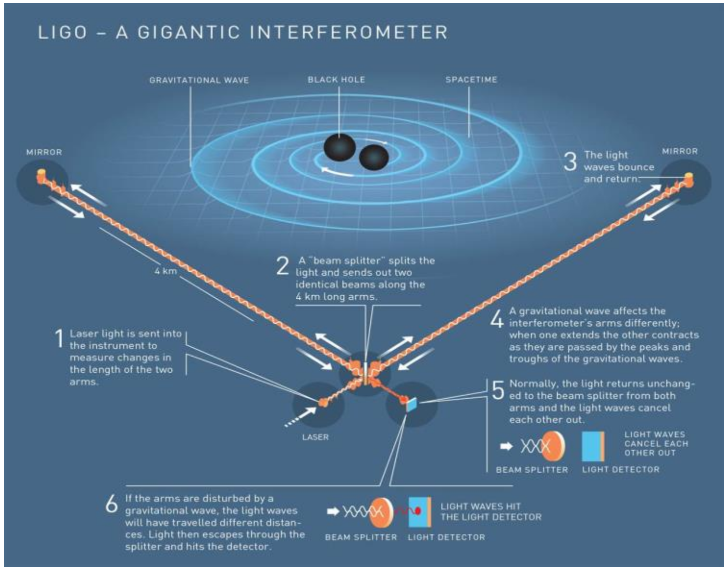

LALSuite development - project background

Gravitational wave astronomy: employs MCMC algorithms to infer astrophysics from gravitational waves
* Compare models of signals to data to infer source properties

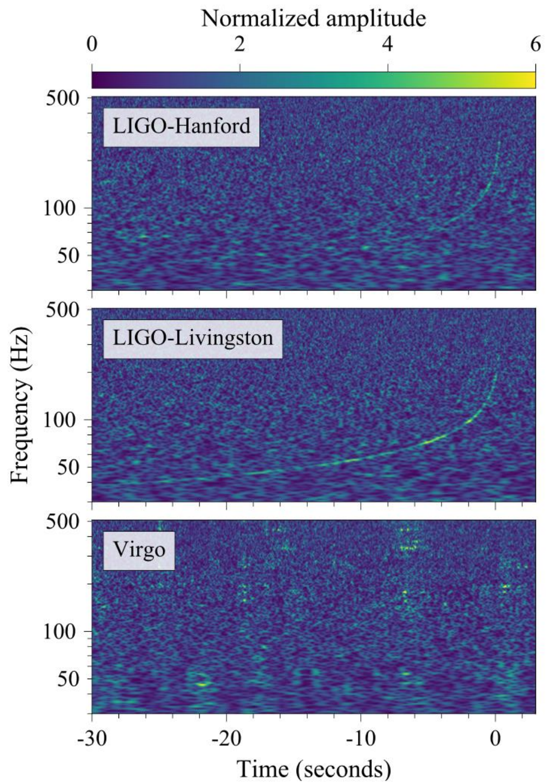

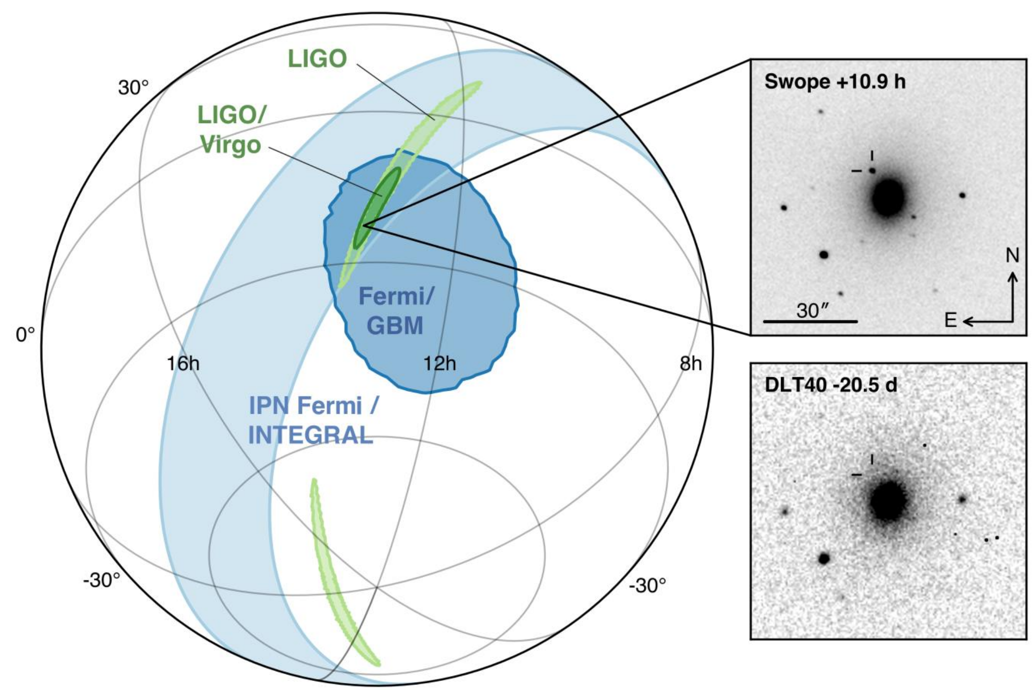

LALSuite development - Overview

- Overview:
  1. LALSuite: Main analysis library of the LIGO Scientific Collaboration
  2. Current analyses take between hours to days
     1. Wall time set to increase by order of magnitude within next few years
     2. Number of detections set to increase by at least an order of magnitude
  3. Need to improve *scalability* of analyses to maximise astrophysics yield

LALSuite development - project background

- GPUS could offer best tradeoff between speed and flexibility
    1. Low risk (already have hardware; OzSTAR cluster)
    2. In house expertise w/ ADACS/OzGRAV
    3. Predictable scaling of codes/analysis
    4. “inner most loop” of analysis is embarrassingly parallel, even if MCMC isn’t easily paralleled

LALSuite development - project background

- Goals:
  1. Create a GPU implementation of gravitational-wave signal models for use in data analysis within LALSuite
  2. Profile GPU implementation
  3. Determine how to proceed, e.g., should we consider large-scale GPU-ifying of LIGO data analysis

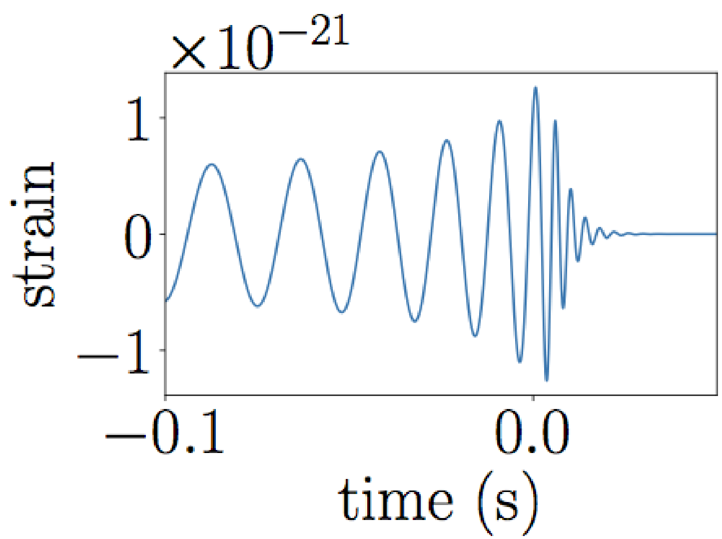

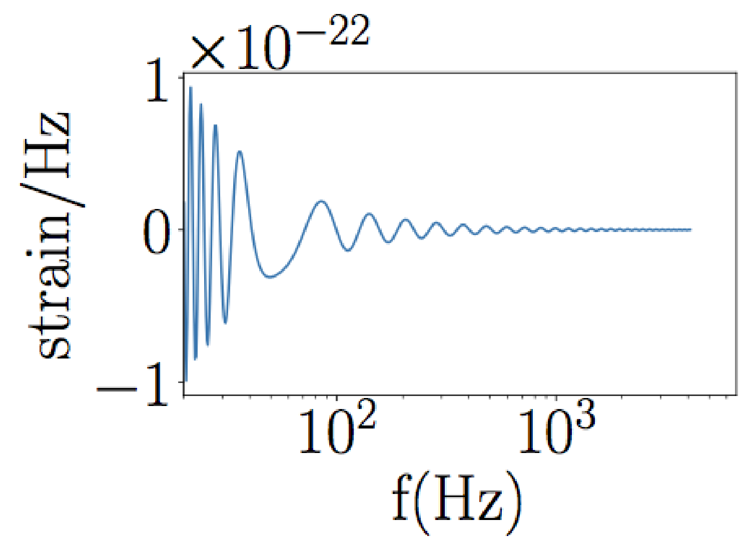

LALSuite development - project background

- Deliverables:
    1. Cuda-compilable version of LALSuite
    2. Cuda implementation of a “workhorse” signal model

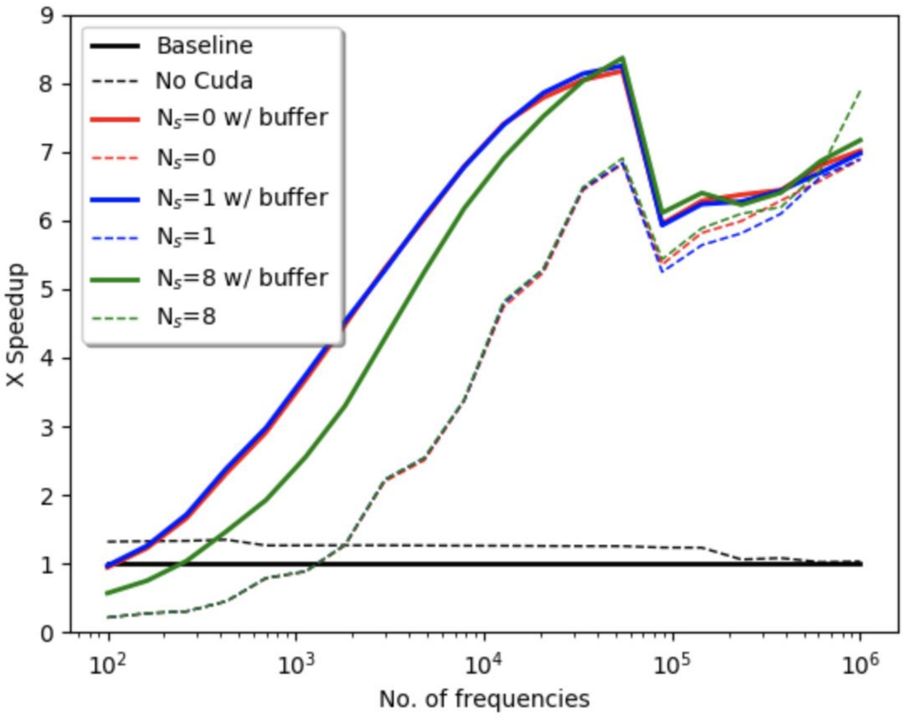

https://adacs-ss18a-rsmith-python.readthedocs.io/en/latest/

LALSuite development - project background

- Takeaways:
  1. Order of magnitude speed up “out of the box”
  2. High latency between CPU and GPU (expected)
  3. Expect a further improvement of a factor of a few by keeping everything on GPU card

LALSuite development - project background

- Moving forward:
    1. Promising path moving forward
        1. Can be used out of the box for upcoming LIGO science run (starting early 2019)
    2. Motivates a full GPU analysis library (future collaboration with ADACS + students)
        1. Want to get the lowest latency —> bypass CPUs!

ADACS

ASTRONOMY DATA AND COMPUTING SERVICES

Software development services:

Case study - developer's perspective

LALSuite development - ADACS perspective

- Some thoughts:
  1. Old and very large codebase (~2M lines)
     - Many design decisions that made integration with CUDA difficult
     - Large community invested in design as-is
     - Need to balance this with need for change
  2. managing expectations
     - Lots of time wrestling with build system

       ⟶ looks like little is going on to user!

Source: imgglobalinfotech.com

ADACS

ASTRONOMY DATA AND COMPUTING SERVICES

Software development services:

Case study - user's perspective

LALSuite development - scientist's perspective

- Some thoughts:
  1. Excellent model for advancing infrastructure :
     - ADACS scientists bridge “expertise gap”
     - More fully utilise advanced computing resources
  2. Streamlines large-scale astronomy
     - The field needs efficient code deployable at scale
     - Makes astronomy more cost effective!

Software Development Services

- Some general issues encountered:
  1. Web applications and GPU optimisations are dominating supported projects
     - Turns-out: astronomers have a lot of good web application ideas/needs
     - UX expertise?
  2. Who owns/is-responsible for code developed?
  3. Ongoing support?
     - Burden grows with time if we offer support
     - On the other hand: don't want to be building tools that fall into states of disrepair
  4. Role of TAC
     - What influence should the TAC have on policy?
  5. Project management
     - Agile development practices, etc.?
     - Astronomers aren't necessarily used to the language, let alone the methodologies
     - They are generally accustomed to a great deal of control over development work

Source: imgglobalinfotech.com

http://ADACS.org.au

@AdacsAus