The ASTAC Supercomputer Time Allocation Scheme

Bernhard Mueller (Monash)

ANITA Green Computing
Summer School
3 June 2021

Astronomy Australia Ltd.

Your Computing Needs – Simplified Picture

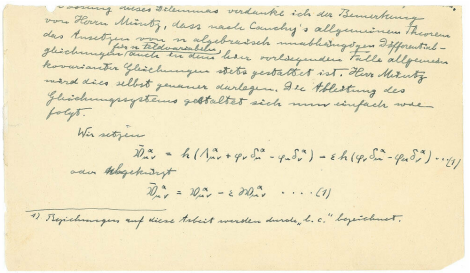

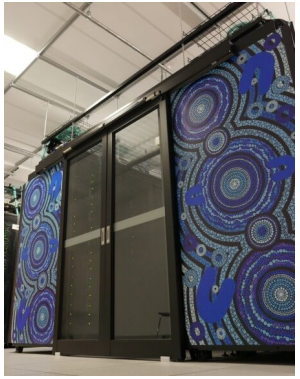

Core hours	Example	Where to get
1-10	• 1D stellar evolution model • Statistics with modest amounts of data	Desktop PCs & laptops Provided by university
1-10^4	• Small multi-D simulations • Bundles of hundreds of simple job	Workstations, local clusters owned by group/university Discretionary use or internal application
>>10^4	• Big multi-D simulations • Big data reduction tasks (e.g. TBs/PBs of radio data)	**Supercomputers** **Peer-reviewed projects (unless you're very lucky)**

High-performance Computing (HPC) Schemes in Australia

NCMAS:

- Annual calls
- Projects of ~100k-20M core-h
- Competitive (3x oversubscribed)

ALCG (since 2020):

- 4-5 projects a year, >15M core-h each

Specialised or ad hoc schemes:

- Early Adaptor schemes, etc.

Significant hurdles

AAL schemes for the general astronomy community:

- Discretionary use of OzStar computer at Swinburne
- ASTAC:
  - two calls per year
  - a few million core-h in total
- ASTAC Large Scheme (since 2020):
  - one call per year
  - Several million core-h per project

Purpose & Philosophy of ASTAC

- Lean application process
- Review by panel of astronomers
- PhD students can be lead-CI!
- Opportunity to build HPC track record
- Panel and AAL as advocates for HPC capabilities of Australian Astronomy – we want projects to succeed

The Major Workhorses

OzStar (Swinburne)

Also OpenStack Virtual Machines on NeCTAR cloud available in recent call

Gadi (NCI National Facility, Canberra)

Application Process – Biannual Calls

- Calls typically in May/November
- Required information
  - Proposal summary & science case (2-3 pages)
  - Track record of chief investigators (publications, awards, honours, prizes, research contributions and experience,...)
  - Technical justification (computer time, storage, code performance, computational paradigm)
  - Past project report if applicable
- Usually easier to get a few hundred thousand core hours than from NCMAS
- Format of Large Calls is similar but with strong emphasis on scalability & HPC experience

https://tac.adacs.org.au/

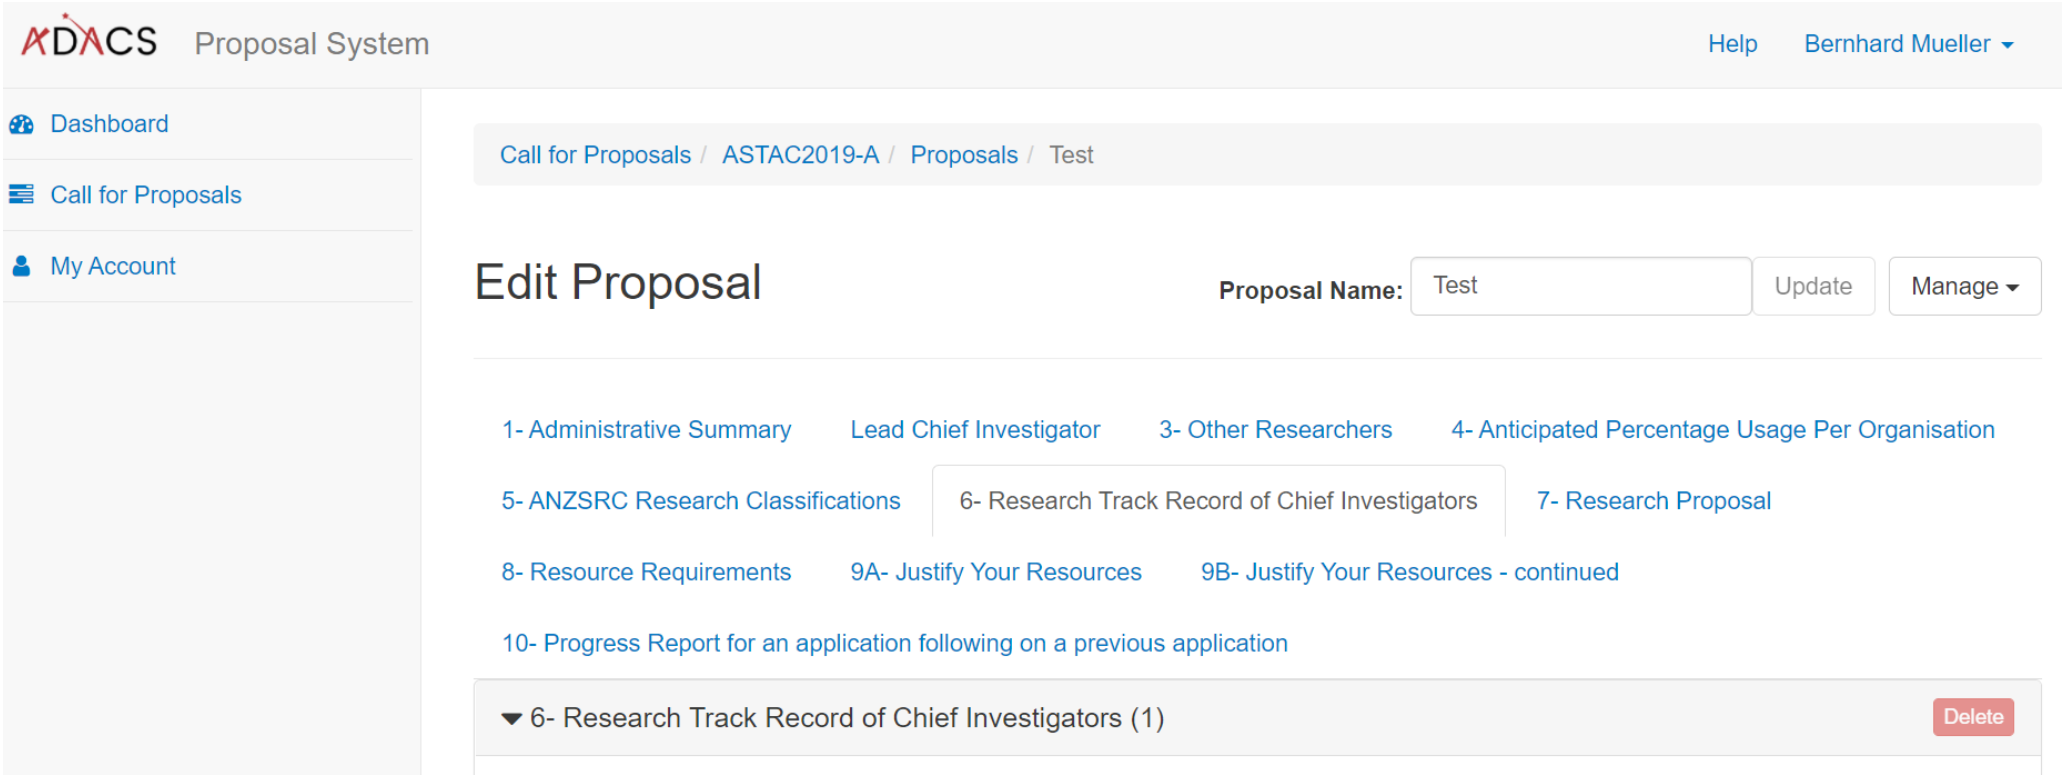

Science Case

- Should cover “aims, significance, impact and innovation”:
  - Context & motivation
  - Problem to be solved
  - Methodology – why appropriate and “best”
  - Work schedule & role of investigators/researchers
- Big picture important, but questions, goals & work plan should be **concrete**
- Contingency plan often a good idea – could you use a partial allocation?
- Can indicate next steps for the future (but be clear what’s part of the project

Technical Justification

- **Much more important than, e.g., in ARC funding schemes!**
- Breakdown of computer time and storage requirements of tasks
- Can use estimate, but should be good and transparent
- Clearly demonstrate parallel performance or suitability for specialised hardware (GPUs – goes down very well!)
- Ideally: perform tests on target machine
- Helpful to indicate suitability of alternative machines
- Before and after proposal: **demonstrate good usage of allocated resources**
- **Show your project can work within the next few months**

Parallel Scaling – Use table or plot

# of cores (Gadi)	wall clock time [s] (Gadi)	parallel efficiency
128	230.53	100%
256	116.6	98.9%
512	61.07	94.3%
1024	34.22	84.2%
2048	19.8	72.7%

Efficiency = (N0 x time on N0 cores) / (N x time on N cores)

- **Strong scaling** most interesting: How much speed-up do you get for a **given task size**?
- Weak scaling: Scaling for fixed problem size per core

Track Record

- Raw stats aren’t everything
- Senior CIs should demonstrate leadership & productivity by standards of field
- No harm in PhD student going in as lead-CI!
- Demonstrate technical competence:
  - Past HPC usage
  - Active use and perhaps development of codes
  - HPC summer school attendance, etc.

Selection Process (may vary slightly)

- 6 panel members (different institutions and fields), assisted by E. Ali
- At least 3 detailed assessments + numerical grades for each proposal
- 10-20min of panel discussion for each proposal to decide fundable/not fundable
- Discussion of allocations in case of oversubscription
- Applicants get a decision summary
  - Explain allocation
  - Provide feedback and advice

Why should I do this – Some ECR Success Stories

- T. Nordlander (postdoc, ANU): Two ASTAC Large projects in 2020 and 2021
- J. Powell (Swinburne): Several ASTAC projects since 2018 – then ARC DECRA 2019
- Demonstrate ability to manage complex projects responsibly

Value of Computer Time

- Current value for academic usage at NCI National Facility: 1 core-h = 2 service units = $0.08
- So a medium ASTAC project (500,000 service units) is worth $20,000
- A large NCI project (5,000,000 service units) is worth $200,000

How do I prepare?

- Get account on OzStar or NCI (NCI Start-up Scheme) to check that code runs
- **No need to be a computer buff**: running batch jobs to take some timings is usually sufficient for decent technical justification
- Consult with senior colleagues
- Contact the OzStar/NCI service desks in case of unexpected behaviour – it happens to everyone!
- If code doesn’t run well: Consider ADACS support (over to Greg...)