

GPU: embedded CSE (eCSE)

Chris Johnson, Lorna Smith, EPCC

8th February 2024

<https://www.archer2.ac.uk/ecse/calls/>



Reusing this material



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.

<https://creativecommons.org/licenses/by-nc-sa/4.0/>

This means you are free to copy and redistribute the material and adapt and build on the material under the following terms: You must give appropriate credit, provide a link to the license and indicate if changes were made. If you adapt or build on the material you must distribute your work under the same license as the original.

Note that this presentation contains images owned by others. Please seek their permission before reusing these images.

Partners



**UK Research
and Innovation**



**THE UNIVERSITY
of EDINBURGH**



Overview

- What is the eCSE funding programme?
- Who is eCSE funding for?
- Programme logistics
- Examples of eCSE work
- Applying for a project – walk through
- Summary
- Questions

What is the eCSE funding programme?



- The Embedded CSE (eCSE) programme funds software development for codes to run on GPU architectures
- Allows software to be developed in a sustainable manner to improve research on GPU architectures
- Facilitate efficient use of GPU resources through enhanced code performance/functionality
- For developing software that facilitates research targeted at UKRI's digital research infrastructure
 - e.g. future Exascale services, UK national AI services, national Tier-2 HPC services
- Allows
 - employment of a Research Software Engineer (RSE) or PDRA
 - software development leading to research **anywhere within the UKRI's remit**

Who is eCSE funding for?

- Two ways to be directly involved
- As a Co-I / PI
 - e.g. a researcher who works with a code which runs (or potentially runs) on GPUs and requires an improvement
 - May not have the skills/availability to do the technical work
 - May/may not have expertise within own group
 - Someone else within group or elsewhere doing technical work
- As a technical member of staff
 - RSE or PDRA has, or is gaining, software development expertise within HPC
 - PhD student with similar role coming to the end of studies
- The panel will assess the teams ability to deliver on the project
 - e.g. staff skills, support from PI/Co-Is, training plan

GPU eCSE Programme administration



- Programme is administered by EPCC through ARCHER2, the UK National HPC Service
- However, applications are not limited to projects targeting the GPU component of ARCHER2
- Projects may utilise other GPU resources within the UK digital research infrastructure.
- Successful applications can get access to the ARCHER2 GPU cluster – or should arrange suitable access to other GPU resources
- Call is all about software development and not production runs
 - Other routes available to gain CPU time e.g. Access to HPC Calls

eCSE Programme Funding



- Funding for can be requested for staff located at
 - institution of the PI
 - staff from the ARCHER2 centralised CSE support team
 - third party institution(s)
 - any combination of the above
- A project funds 1 to 36 person months in total across all technical staff
 - Flexible: Effort can be spread over a longer period or employ multiple people simultaneously
 - Max duration 24 calendar months

eCSE Programme Funding (contd.)



- We run a regular series of calls for proposals, with three calls per year
- Support and encourage early career researchers
 - Many RSEs are early career
 - Good opportunity to develop software development skills – great for CV
- Observers
 - Early career researchers invited to on-line panel meetings as observers
 - Usually have 3-4 observers at each panel meeting
 - Good opportunity to see what reviewers are looking for
 - Call open alongside GPU eCSE call – deadline 27th Feb
- Proposals are selected by an independent diverse panel of experts

Example eCSE work

- Transforming CPU code -> GPU code
- Make improvements to existing GPU codes
 - portability to allow efficient running on wider range of architectures
 - performance
 - scalability
 - sustainability and maintainability
 - Implementing new algorithmics within code
- Improvements to GPU-based software to allow new research
- Adding new functionalities to existing GPU-based software
- Enhancing existing GPU software for use in new research areas/workflows

Applying for an eCSE Project



Steps for applying to eCSE call



- First steps
 - Start proposal in SAFE
 - Invite other team members
 - Start getting costing information
 - Costing not needed for ARCHER2 CSE team (we can get these internally)
 - Check financial arrangements
 - in particular if RSE not at institution of PI or part of ARCHER2 CSE team then money is paid to PI's institution who must arrange to pass this on
 - Look through proposal forms to fill in (both on-line form and Word document)
 - Check through Terms document with your department
 - Read the guidance document!
 - Start requesting letters of support (optional)
- Next steps
 - Complete main proposal form
 - Complete the on-line form – linking in completed proposal form, costing docs, CVs, letter(s) of support, etc.
- On-line form must then be completed with
 - All project members having accepted invitation
 - All relevant documents attached
 - Submitted by **16:00 19th March 2024**

Documents required

SAFE form



Costing document(s)
(required for each RSE)

Title of eCSE Project. A. N. Other 01/09/2022 - 31/08/2023 (12 months). Full time % : 100%.

Cost Category	Total 100%	Total 80%
Directly Incurred		
Staff	£50,000	£40,000
Travel	£0	£0
Equipment	£0	£0
Other costs	£0	£0
	£50,000	£40,000
Directly allocated		
Investigators	£0	£0
Estates	£16,000	£12,800
Other	£0	£0
	£16,000	£12,800
Indirect	£56,000	£44,800
Total	£122,000	£97,600

SAFE form screenshot showing project details:

- Project Title: GPU test proposal
- Subject Area: STFC (Science and Technology Facilities Council)
- Project Summary: This is all about my project
- People associated with the application:

Name	Email	Role	Institution	Contribution and Background
Dr Chris A. Johnson	chrisj@epcc.ed.ac.uk	Member	University of Edinburgh	
- Project Details: Current status of this proposal: Pending

CV (required for each RSE)

CV template structure:

- PERSONAL DETAILS**
 - Name: John Smith
 - Address: Edinburgh, UK
 - Phone Number: +44 123456789
- EDUCATION**
 - BSc Physics 2010-2013 University of Manchester
 - MSc More Physics 2013-2014 University of Leeds
 - PhD Even more Physics 2014-2017 University of Edinburgh
- PUBLICATIONS**
 - [1] A. N. Other & J. Smith First publication listed here, Journal of Physics 94 148-165 (2019)
 - [2] A. N. Other & J. Smith Another publication listed here, Nature 37 256-267 (2019)
 - [3] J. Smith First publication listed here, Further Journal of Physics 32 122-134 (2020)
- REFERENCES**
 - My PhD supervisor: University of Edinburgh, Physics, etc.

Letter of Support (optional)

Letter of Support template structure:

- Header: THE UNIVERSITY OF EDINBURGH, EPCC, The University of Edinburgh, The Royal Centre, 47 Robertson, Edinburgh, EH8 9BT. Tel: +44 (0) 131 650 5533, Fax: +44 (0) 131 650 5535.
- Date: Thursday, 01 February 2024
- Recipient: Dear Professor J. Smith,
- Body: This letter confirms our support of your eCSE GPU project "Improving the scalability of BioGPUCode". At present our research is limited due to the relatively poor scaling of the code meaning we are at present unable to run really big and exciting simulations. The work to be carried out in your upcoming project will enable us to, etc., etc. More text.
- Text: More text giving more detail about what we will use the code for. More text and more text. Lots more text. Some more text, etc. Maybe a bit more about what we do and why it's really important.
- Text: Even more text. More stuff. More info. More, more, lots more.
- Text: Finally, more to sign off with. A bit more.
- Closing: Yours sincerely,
- Signature line: |
- Footer: Letter sender, Job Title

Proposal Document (required)

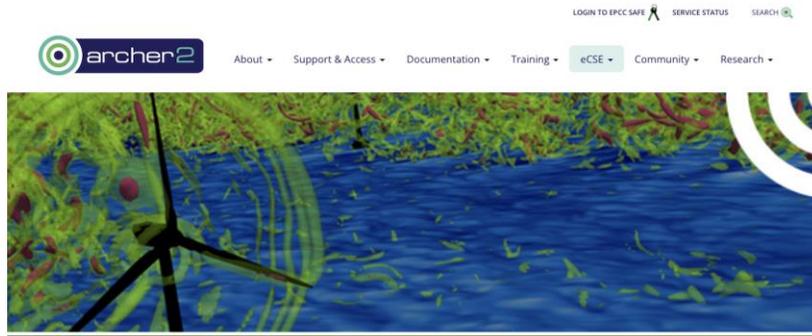
Proposal Document template structure:

- Header: epcc, THE UNIVERSITY OF EDINBURGH, archer2, UK Research and Innovation
- Title: Proposal for a GPU eCSE Application
- Text: Applicants for a GPU eCSE application should use this template. This GPU eCSE call will close at 16:00 on Tuesday 19 March 2024. No applications will be accepted after this time.
- Please note:
 - There is a hard page limit for each section which is shown at the start of each section.
 - The font size should be no smaller than 11pt and the margins should be no smaller than 2.5cm.
- Text: Please upload your proposal when you fill in your eCSE application online form via the eCSE Funding Calls section within the SAFE: https://safe.epcc.ed.ac.uk/
- Text: Please read the Applications Guidance (linked to the page https://www.archer2.ac.uk/eCSE/calls/) for further instructions on how to complete your eCSE online application.
- Text: If you have any queries or require assistance regarding your application, please contact the ARCHER2 service desk: support@archer2.ac.uk
- Section 1: Project Title (as given in on-line SAFE form)
- Section 2: PI Name and Institution
- Section 3: Description of code(s) (max. 2 page)

Beginning an eCSE Application

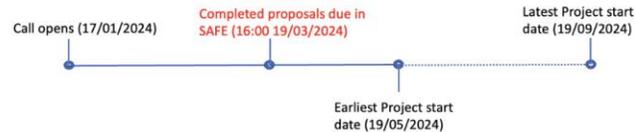
- Visit call page <https://www.archer2.ac.uk/ecse/calls/>
 - links to templates, guidance and terms
- Start in the SAFE
 - <https://safe.epcc.ed.ac.uk/>
- Enter Project Title, subject area and research area (remit)

GPU eCSE Calls on-line



GPU Embedded CSE (eCSE) support Call

Details of open calls



The first GPU eCSE call and Early Career eCSE Panel Observers call

GPU eCSE support provides funding to the UKRI research community to develop software in a sustainable manner to run on GPU-based architectures. We invite proposals to develop software that facilitates research targeted at UKRI's digital research infrastructure e.g. future Exascale supercomputing services, UK national AI services, national Tier-2 HPC services. The call is open to research across all of UKRI's remit.

The GPU eCSE call (GPU-eCSE01) opened on the 17 Jan 2024. The final deadline for proposal submission is 16:00 on 19 Mar 2024 via the ARCHER2 SAFE. Please note that for this call there is no separate technical evaluation stage or separate technical evaluation form. The early career eCSE Panel Observers opened alongside the eCSE call and the deadline for applications for this call is 16:00 on 27 Feb 2024.

GPU eCSE call

Funding will enable the employment of an RSE (a Research Software Engineer, PDRA or equivalent) to work specifically on the relevant software to enable new features or improve the performance of the software to be run on GPU-based architectures. Funding can be requested for up to 36 person months of effort per project for a duration of up to 2 calendar years. There is flexibility in the way the effort is requested. For example, a project may have more than one person working on it or the effort could be spread over a greater number of calendar months, i.e. a staff member could work on the project at less than 100% alongside other commitments or it could be worked on by a staff member who doesn't work full time. Please just get in contact with the ARCHER2 Service Desk if you wish to discuss possible scenarios. Projects are expected to start between 2-6 months after a call closes.

The GPU eCSE programme is administered by EPCC through ARCHER2, the UK National HPC Service. However, applications are not limited to projects targeting the GPU component of ARCHER2. Projects may utilise other GPU resources within the UK digital research infrastructure.

The text below gives an overview of the call but all applicants should read the [guidance for submitting a proposal](#) to this call in full.

Early Career eCSE Panel Observers call

One of the aims of the eCSE programme is to provide support for researchers in the early stages of their careers and we will be offering a small number of early career researchers the opportunity to attend the eCSE panel meeting as observers. Please see the call information [here](#). Please see the [recent blog article](#) describing the experience of attending a panel meeting from a similar programme.

If your proposal is selected for funding, you will be provided with an award letter and will be asked to agree to the eCSE Terms.

Please view the Terms in advance of submitting the proposal.

Proposal templates



- The eCSE form has an online form to which documents should be attached before submission
- MS Word templates available for proposal document
- Proposal template
 1. Project Title
 2. PI Name and Institution
 3. Description of code(s)
 4. Project Objectives and success metrics
 5. Project Overview, Technical Information and Workplan
 6. Sustainability, maintenance, validation and availability of code(s)
 7. Impact and Benefits

Description of codes(s)



Up to 2 pages

- Tell us about the code! What is it for? what does it do? Who uses it?
- How big is it? How is it structured (how complex)?
- What language(s) does it use and any dependencies (libs, compilers, other software)?
- What is the code's status in terms of usage and performance?
- Performance data may include scaling (if multi-GPU), profiling data, or comparison of GPU code with CPU code
- Can include links but include summary in text

Project Objectives and success metrics



Up to 1 page

- Number of objectives and associated success metrics expected
 - 1-2 objectives for a proposal requesting <6 PMs of effort
 - 5-6 objectives for a proposal with the maximum requested effort of 36 PMs
 - This is just a guide!
 - Expect each objective to have an associated success metric
- Success metrics should be specific and measurable. E.g. the following sounds reasonable
 - *“We will significantly improve the scaling of the GPU code compared with the existing code on a set of benchmark codes”*
- but isn't really specific or measurable
 - Which codes? How many? What improvements are expected? For what problem sizes?

Project Overview, Technical Information and Workplan



Up to 2 pages

- This is where you describe what you will do in the project
- Describe from a technical viewpoint
 - But should be understandable to someone who doesn't know the code specifically
- Should include
 - associated timescales
 - Gantt chart (or equivalent) of the workplan with milestones and effort broken down by RSE (if > 1 RSE)
 - Risk – with likelihood and possible mitigation

Sustainability, maintenance, validation and availability of code(s)



Up to 1 page

- Should describe how code(s) will be
 - made available to users
 - maintained
- State what coding standards will be used, etc.
- Should include any information on the licensing model
- Describe both for users
 - wishing to run the code(s) and
 - wanting access to the underlying source code(s)

Impacts and Benefits



Up to 1 page

- Describe potential impact of the science enabled by the project
- How will impact be achieved?
 - Some impacts may be realised after the end of the project
- Consider
 - both the computational and scientific benefits
 - who will benefit

Applying for an eCSE Project



- Final deadline for full proposals submitted to the SAFE is then **16:00 19th March 2024**
 - Late proposals will not be considered!
- Previous eCSE calls had a separate technical evaluation, this is now part of the main proposal
- Any project team members must be invited to a proposal and have accepted the invitation via the SAFE before final submission
 - Otherwise they will not be considered part of the team
 - Allowable roles are PI, Co-I, technical staff, Contact, Member
 - Technical staff add their own CVs
 - **All** other team members must provide track record and describe what they will bring to project
 - When inviting people, please use the email address they use in the SAFE
 - Note that you **cannot submit without technical staff** (or unnamed staff selected)

Costings



- Please provide 100% FEC value for each member of technical staff
 - We pay 80% FEC for funded proposal
 - Funding only includes technical staff
 - Include direct and indirect costs but not additional costs (e.g. PI/Co-I costs)
 - Provide breakdown of costs as a PDF (Worktribe or equivalent)
 - No need to provide costings for ARCHER2 CSE staff
- Travel is given separately from above
 - Funds are limited – see guidance
 - Can include project members traveling with UK to meet each other
 - Additionally each RSE can be funded to attend one national/international conference
- If your institution does not use FEC please contact us in advance
- We pay the PI's institution via invoice
 - Any payment to third parties is paid via the PI's institution
- Recommend you use named staff where possible
 - Ability of technical staff to complete the work is assessed
 - Considered together with PI/Co-I expertise

eCSE Costing document

- Example costing for 12-month project
- Made up numbers
 - This is not a real document!
- Shows the information we need

Enter **100% total** costing in proposal form Form then calculates 80%

Title of eCSE Project. A. N. Other 01/09/2022 - 31/08/2023 (12 months). Full time % : 100%.

Cost Category	Total 100%		Total 80%
Directly Incurred			
Staff	£50,000	X	£40,000
Travel	£0		£0
Equipment	£0		£0
Other costs	£0		£0
	£50,000		£40,000
Directly allocated			
Investigators	£0		£0
Estates	£16,000		£12,800
Other	£0		£0
	£16,000		£12,800
Indirect	£56,000		£44,800
Total	£122,000		£97,600

eCSE – Submitting a proposal



[Your details](#)

[Service information](#)

[Projects](#)

[eCSE Funding](#)

[Login accounts](#)

[Help and Support](#)

[<< Back to the Main Page](#)

New Funding Application

This application form is to be completed by the applicants for the eCSE call. If any missing information or deficiencies are found, the applicants may be contacted to supply the missing information.



To create a new application all that is required is the project title (which can be changed once the application is created) and to select the Funding Call you are creating the application for. Once you've created the application you will be able invite collaborators and technical staff to collaborate on the application.

Create

[<< Back to the Main Page](#)

eCSE – things to watch out for



- Doesn't look technically feasible (either within the timescale or at all)
- Poor estimate of effort - may be over or under ambitious for effort requested
- Sometimes just missing the basic “why”, “how”
- Proposal may technically desirable but missing the “why”
 - E.g. project has clear plan for reducing time taken for a code by large percentage...
 - ...but why does that matter? Who uses code? For what? Why is the code important?
- Staffing not suitable
 - Missing needed expertise with code or GPU skills – or a plan to obtain these skills
- Missing suitable analysis of risk
 - E.g. how will WP4 work if WP3 isn't successful?
- Success metrics not quantified
- Letters of support not specific
- Do proof-read (or get a colleague to)!
- Sometimes nothing specifically wrong but details just not well enough explained

Summary



- The first GPU eCSE call is open – see time timetable for future calls
 - Expect further calls over the year
- Final deadline for full proposals: **16:00 19th March 2024** -> SAFE
- Early Career Observers call open alongside (deadline 16:00 27th Feb 2024) -> ARCHER2 service desk
- Details of calls
 - <https://www.archer2.ac.uk/ecse/>
 - Any questions please ask here or via the ARCHER2 service desk support@archer2.ac.uk



Any Questions?

www.archer2.ac.uk

