



Benefits of the ARCHER eCSE Programme

V1.0

Lorna Smith, Alan Simpson, Chris Johnson, Xu Guo,
Neelofer Banglawala



1. Executive Summary

The eCSE programme has allocated funding to the UK computational science community through a series of funding calls over a period of 5 years. The goal throughout has been to deliver a funding programme that is fair, transparent, objective and consistent. The projects funded through this programme were selected to contribute to the following broad aims:

- To enhance the quality, quantity and range of science produced on the ARCHER service through improved software;
- To develop the computational science skills base, and provide expert assistance embedded within research communities, across the UK;
- To provide an enhanced and sustainable set of HPC software for UK science.

The eCSE programme is a significant source of funding for the Research Software Engineering community and all UK Higher Education Institutions are able to apply for funding. This document provides more detail on the programme, looking at how the funding has been spent and examining the various benefits realised from the programme.

2. Benefits

A series of benefits have been identified. While realisation of these is an on-going process (as projects are still active), these will ultimately provide confidence that the overall goals/aims of the programme have been met. The benefits identified are:

- A high quality, fair and objective eCSE selection process, delivering maximum value to the community;
- Increased science productivity – financial saving reinvested to allow scientists to achieve more science from the same resource allocation;
- Increased novelty and range of science on the system, both traditional and new;
- Enhanced computational science skills base across the UK.

In this document we consider each benefit in turn, providing evidence of associated outputs and metrics.

A high quality, fair and objective eCSE selection process, delivering maximum value to the community

The goal over the past 5 years has been to deliver a funding programme that is fair, transparent, objective and consistent. A whole range of procedures have been implemented to achieve this. However the following were key:

Regular calls

The aim was to hold a series of regular calls across the service. This allowed applicants to plan their submissions and prepare appropriately. This also makes a positive equality impact, reducing barriers to those on maternity and paternity leave by providing frequent and regular calls. We have held 12 calls, three a year, with dates published well in advance on the web site. To date we had 201 proposals, with 90 funded projects giving an average success rate of 45% across all the calls.

Not for profit, FEC costing model and transparency in costs

Contractually, we were required to deliver 14 FTEs per year and this was delivered. However EPCC submitted a not-for-profit bid and committed to using all funds for the programme. Having delivered 14 FTEs, remaining funds were utilised to fund an additional call. At this call we funded an additional 40 project months.

We utilised an FEC funding model to make the programme as accessible as possible to research groups across the UK. Utilising this standard and approved model for research projects reduced the complexity and barriers to submission while ensuring that projects were funded at an appropriate level, but without profit.

Figure 1 shows the number of institutions with active eCSE projects during any six-month period of the service. As you can see, after an initial ramp up, this ranged from 20 to 30 institutions, a significant number. In total, 45 different institutions were involved in the eCSE programme, demonstrating the success of the FEC model. More details on the institutions and their locations are provided later in this document.

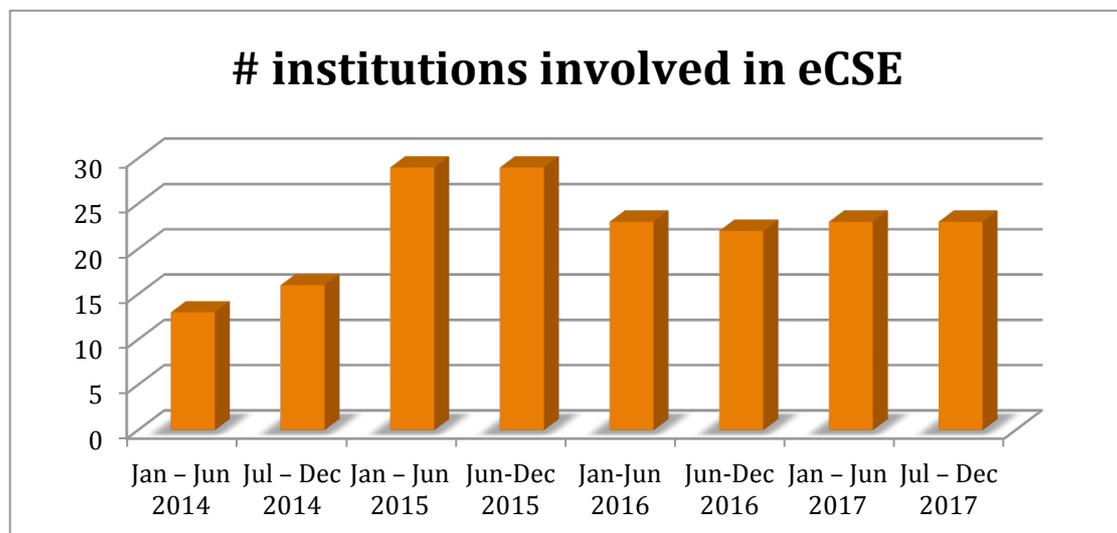


Figure 1: Number of institutions involved in an active eCSE across a 6-month period.

Planned spending profile

From the start we have tracked and planned for a managed number of active projects across the service – the aim was to ramp up and fund slightly more projects near the start of the service, to maximise the utilisation time post-project. Figure 2 shows the spend profile across the service. We also looked to fund a similar number of projects at every call and provide a similar level of competition. We published data on successful projects to ensure transparency.

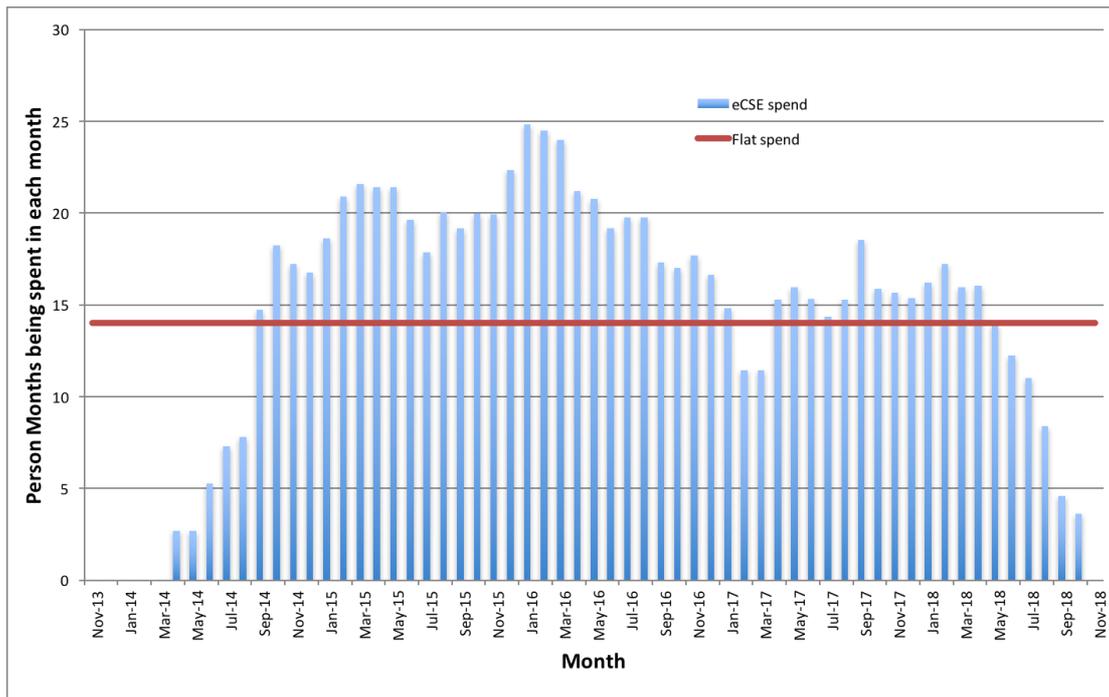


Figure 2: eCSE spend profile – number of project months being spent each month.

Independent panel members

The programme had a set of independent panel members from a broad spectrum of science areas. EPSRC and NERC representatives were present at every panel meeting. Clear guidance was provided to panel members on selection criteria and a robust conflict of interest procedure was applied.

Increased science productivity – financial saving reinvested to allow scientists to achieve more science from the same resource allocation

Many eCSE projects generate significant financial savings, which in turn are reinvested to allow scientists to achieve greater and enhanced science from the same amount of resource.

Often eCSE projects result in improvements in the performance of the software or the addition of new functionality to the code that, in turn, allows a more rapid time to solution. By measuring CPU utilisation on ARCHER after the project, these improvements can be used to determine a quantifiable measure of the amount of saved CPU resource and therefore financial saving. This process was described in [1] and utilises data provided in eCSE project final reports.

Amalgamating these improvements across all projects provides an estimate of the overall financial savings generated by the eCSE programme. Figure 3 shows the measured financial saving for each 6-month period of the service. From this figure it is clear that the eCSE programme provides significant financial benefit. This benefit will only increase as time goes on and the codes continue to be utilised by the community.

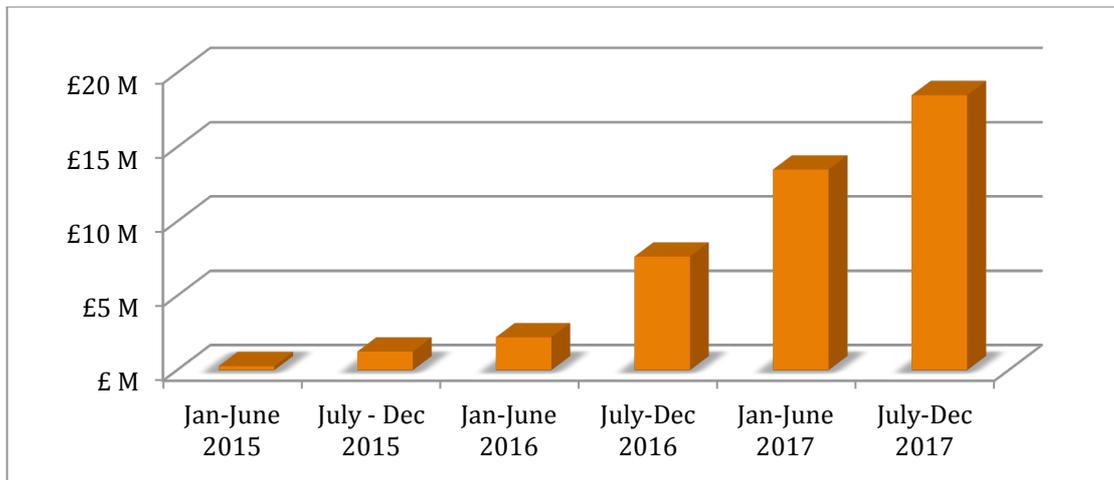


Figure 3: Overall cumulative financial saving from the eCSE programme on a 6-monthly basis.

The overall cost of the eCSE programme was around £6M and the reported benefits are more than £18M, already representing more than 3 times return on investment. These figures are significant, but do represent the cost of the science had the codes not undergone optimisation. The data is based on final reports, so for example, this assumes the benefit is to all users of a code unless noted in the final report. Where a range of improvements (speedups) are reported, the most conservative estimate has been utilised. Not all projects are complete, the data presented is for 58 of the 90 projects. As more projects complete and more CPU time is utilised by the optimised codes, the financial benefits will continue to increase.

This financial saving figure is compelling, but needs to be placed in context. In real terms, the eCSEs ensure existing science can be done within a smaller CPU budget, freeing up resources for additional science – either by the same group or by new research groups. For example, an eCSE carried out by the University of Hull on a bone simulation code resulted in immediate financial savings, but the real benefit was to allow fine details of bone models to be studied which has impact for implant design and research in to musculoskeletal conditions [2]. Scientific highlights are therefore produced by each project on completion and these are published on the ARCHER web site. These showcase the science and impact generated from the eCSE programme. To date, 52 have been published.

Increased range of science on the system, both traditional and new

It is important that the eCSE programme benefits the full science community on ARCHER, rather than specific science areas. To measure this, we have tracked the different funding areas for each project and have looked at the 40 most used codes on the system and assessed how many of these have benefitted from eCSE support. Combined, these metrics give a strong indication that the whole ARCHER community has benefitted from this support.

In addition to funding proposals from established ARCHER communities, proposals from scientific areas not already exploiting ARCHER were also encouraged so as to help grow the user base for ARCHER and future national HPC systems. Since the 4th eCSE call, we actively encouraged proposals from these "New Communities," and 10 such proposals were funded.

Figure 4 shows the breakdown of science areas for active projects during each six-month period of the service. The breakdown is given in staff months, not number of projects. This shows that the major science areas are all well represented, with the ocean and modelling component similar in proportion to the NERC allocation on ARCHER.

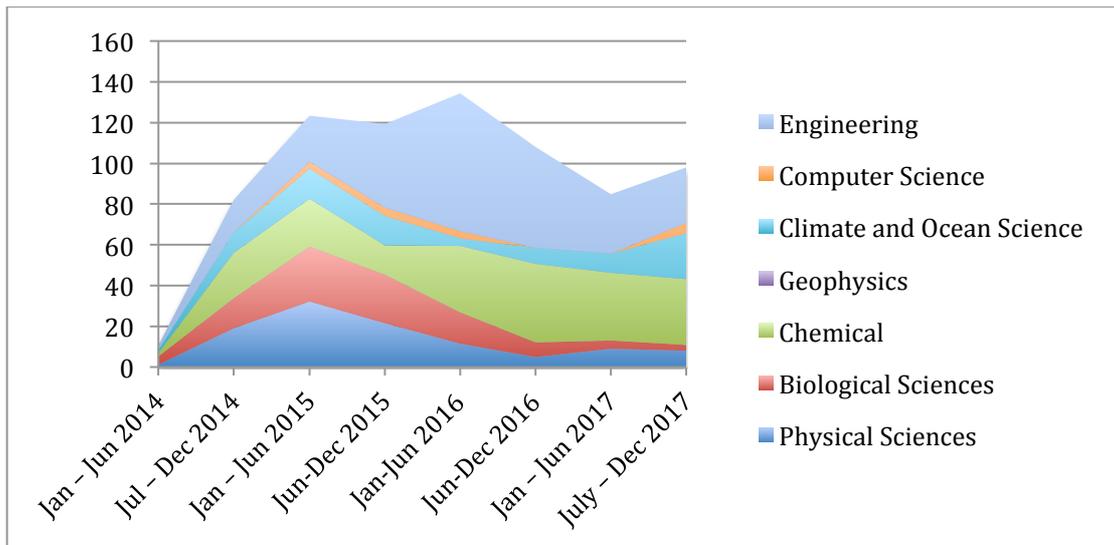


Figure 4: Breakdown in major science areas funded by the eCSE programme (in project months).

Looking at the number of codes on ARCHER that have benefitted from eCSE support gives a good indication of the breadth of codes and communities that have benefitted from this support. Figure 5 shows the percentage of codes that have benefitted from eCSE support, based on the 40 most heavily used codes. In the last 6-month period, over 40% of the top 40 codes had benefitted from some form of eCSE support.

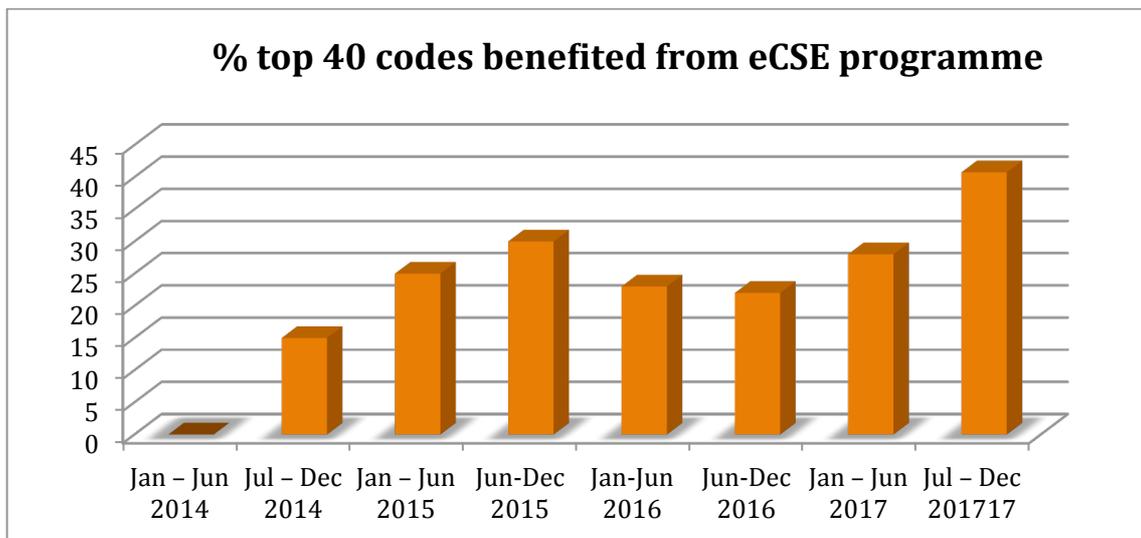


Figure 5: Percentage of major codes (top 40 codes) that have benefitted from eCSE support.

Benefit: Skilled embedded workforce across UK

A key aim of the eCSE programme is to fund software engineering projects across the UK, embedding staff within computational science groups and with key ARCHER users.

By ensuring this geographical distribution, appropriate knowledge and skills are available to the relevant communities, not simply hosted in a few centres. This also develops a strong set of UK research software engineers, with appropriate skills in software development, High Performance Computing and computational science. This in turn helps drive productivity on the ARCHER service and provide the UK HPC community with a competitive advantage over other countries. Ultimately

this enhances the quality and quantity of science output on the UK National HPC services (i.e., ARCHER and beyond).

The eCSE programme provides dedicated funding for software development, allowing those employed on the projects to focus and enhance their software development skills. Individuals also have access to the ARCHER service's training resources, online courses and virtual tutorials, all of which enhance overall skills.

Figure 6 shows the geographical distribution of eCSE institutions. Stronger colours represent larger numbers of projects. As is clear from the figure, the eCSE programme has funded a geographically dispersed group of projects, ensuring knowledge and skill have been developed across a wide range of communities.

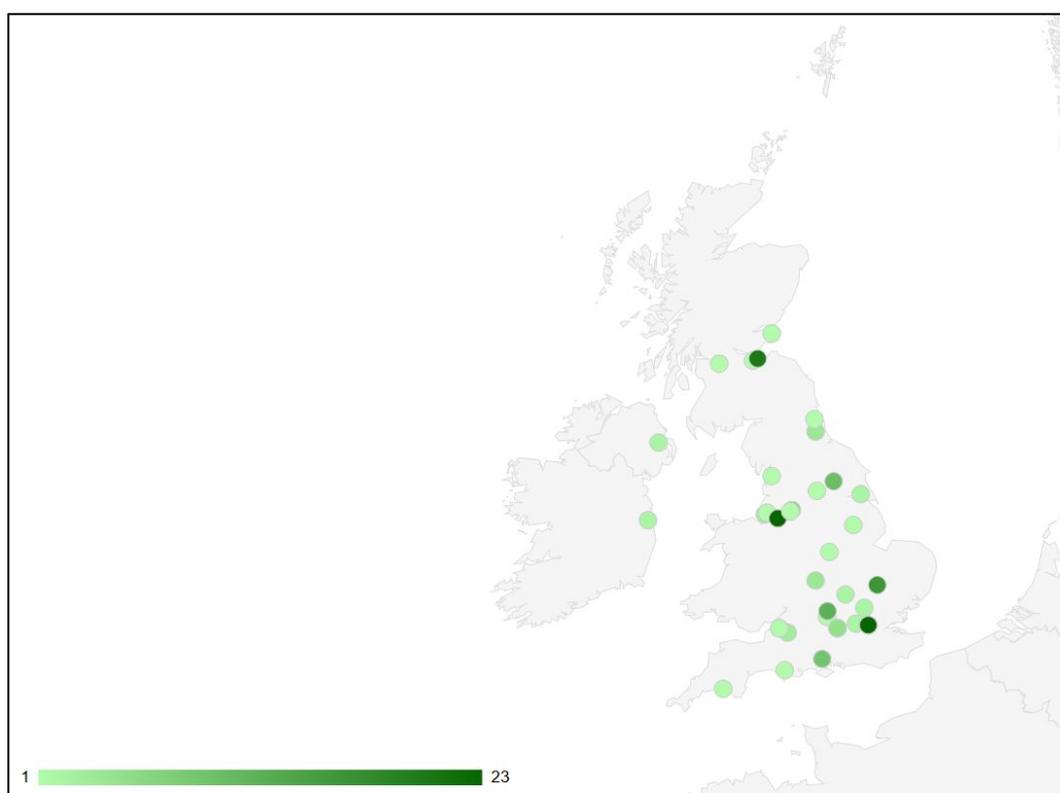


Figure 6: Geographical distribution of institutes involved in the eCSE programme.

Conclusion

The eCSE programme has allocated funding to the UK computational science community through a series of funding calls over a period of 5 years. While some projects are still running, it is clear that the programme has already:

- provided a consistent, fair and not-for-profit funding programme, a programme that has funded significantly more projects than contractual requirements;
- funded projects across the UK and across science areas, embedded within research communities;
- enhanced the skills base of the UK computational community, again across the UK;
- generated considerable financial benefits, already representing a three times return on investment. This manifests itself in more science from the same resource.

As the codes continue to be used over the following years we anticipate a wide range of high quality science will be performed and published in scientific journals. Currently, a large number of scientific highlights have been produced, presenting the anticipated benefits of the code development.

References

[1] "eCSE Benefits Realisation", L.Smith, C. Johnson, X. Guo,2016

[2] "Understanding how bones develop and respond to disease and the use of implants", see:
http://www.archer.ac.uk/casestudies/ARCHER_casestudy_VOXFE.pdf