Novel Approaches to HPC User Engagement

Clair Barrass, David Henty
EPCC
The University of Edinburgh
Edinburgh, UK
e-mail: (c.barrass|d.henty) at epcc.ed.ac.uk

***Abstract*— EPCC operates the UK National HPC service ARCHER, a Cray XC30 with a diverse user community. A key challenge to any HPC provider is growing the user base, making new users aware of the potential benefits of the service and ensuring a low barrier to entry. To this end, we have explored a number of approaches to user engagement that are novel within the context of UK HPC:**

- public outreach activities and promoting the benefits of HPC via a network of ARCHER Champions;
- giving easy access to ARCHER through an online driving test;
- easy-access training including supporting novice users with screencast videos and the Supercomputing MOOC;
- regular long-term training impact surveys to quantify the benefits of the training programme.

Here we choose to focus on those aspects of user support and engagement that we believe are novel and have proven particularly successful, rather than the more standard activities common to most HPC services.

Keywords-user; support; training; outreach; engagement; hpc;

I. INTRODUCTION

We split our discussion of novel approaches to HPC user engagement into the four areas outlined in the abstract. Although funded by a variety of mechanisms, these are all related to the UK national supercomputer service ARCHER www.archer.ac.uk**,** a Cray XC30 system managed by EPSRC on behalf of the UK Research Councils and housed, operated and supported by EPCC.

A. ARCHER Outreach

As part of our public outreach activities, we have developed both physical and virtual HPC demonstrators to promote the understanding of the architecture and applications of modern supercomputers. Wee ARCHIE is a portable supercomputer built using Raspberry Pis for use at schools and science fairs, and we have ported a number of parallel applications which illustrate key HPC concepts. We have also created a web-based "Design Your Own SuperComputer" app to be accessed online, on phones or on tablets to learn about HPC whilst playing a game.

Although networks of local HPC expertise have been developed in other services, for example the XSEDE Campus Champions, this had not been done before in the UK. We will describe our approach to building up a self-sustaining network of ARCHER Champions who can promote the benefits of HPC and be a local source of expertise to enable new users to transition from regional to national HPC services. We will explain how we have kick-started the process with a series of workshops located around the UK, and how we are collaborating with other initiatives such as ARCHER public outreach and the new Research Software Engineer programme.

We run an annual Image Competition and invite submissions from the ARCHER user community of images which best demonstrate "ARCHER Enabling Research". This provides a wealth of material for future outreach and engagement activities; for example, the images are collated into a printed calendar.

B. ARCHER Driving Test

One of the main barriers to entry for new HPC users is simply getting access to resources. Obtaining an account usually entails submitting some form of project proposal, and a new user is unlikely to be willing to go to this effort if they do not even know if HPC can be useful to them. In February 2015 we developed the ARCHER Driving Test which can give new users access to 80,000 core-hours over 12 months to do exploratory work on ARCHER. The driving test comprises 20 multiple-choice questions chosen randomly from a bank of over 100; questions are categorised to ensure an even spread across a range of topics. The aim is to ensure that users have sufficient knowledge of HPC in general, and ARCHER in particular, that they can make productive use of their CPU allocation. Since its launch, over 177 users have gained access to ARCHER by passing the driving test. We present an analysis of the way they use the service, and compare to the behaviour of more traditional users.

C. Easy-Access Training

Although we provide a full range of technical HPC training courses, new users have some particular needs. First, it may not be obvious which courses are relevant and in what order they should be taken. We will explain how we have categorised our training courses and provided suggested training pathways to help users to easily navigate the programme. Second, certain aspects of the system may be completely new to them. For example, all ARCHER account management is done via our online SAFE system and although SAFE is very sophisticated and mature, a new user

is unlikely to have used anything similar before; we do not want something simple like requesting a new password to be a barrier to new users. We describe how we have supplemented our traditional documentation with over 30 short screencasts explaining how to perform key functions. We also describe the “Supercomputing” MOOC, launched in March 2017 under FutureLearn, which was developed as part of the PRACE [1] project.

D. Training Impact

Although it is fairly straightforward to assess the immediate impact of an individual course by having attendees fill in online questionnaires shortly after the event, measuring the longer-term impact of an entire training programme is much harder. For example, did it result in enhanced career progression? This type of feedback, corresponding to levels 3 and 4 of Kirkpatrick's "four level" model [2] of training course valuation, is clearly extremely important but is also recognised as being very difficult to obtain. Others have considered how these levels of feedback can be gathered in the context of HPC training, for example as presented in the talk "Key Performance Indicators and Learning Analytics" [3]. We have so far completed four such surveys which aim to assess benefits that occur on timescales of 6 months or more, the most recent being completed in early 2017. We explain how we have tried to maximise the survey completion rate and our analysis of the results.

II. OUTREACH

A. Wee ARCHIE

One of the issues with explaining the concepts behind parallel computing to the general public is that it is not possible to show them real life super computers - only a tiny number get the chance to visit a machine room. Although a modern laptop is a parallel system, the underlying hardware and software in modern devices is so deeply hidden that the parallelism is hard to demonstrate.

Having a real life portable parallel system is clearly of great value in a hands-on demonstration. The initial work in this area was done by The Oak Ridge Leadership Computing Facility who realised that low cost, low power Raspberry Pi boards made this feasible and created Tiny Titan [4]. They also developed a number of fun engaging demonstration programs based on real applications and provided documentation on the hardware and software for others to follow. Inspired by this we decided to build our own extended version using the newer multi-core version of the Raspberry Pi, designed to explain the architecture of ARCHER - hence the name 'Wee ARCHIE'.

Figure 1. Wee ARCHIE: LED lights display node activity

Wee ARCHIE is constructed from 18 Raspberry Pi 2’s, a network switch, a power supply unit and Ethernet cables. It is a mini supercomputer but robust and easy to transport or ship in a custom-built case [5], created specifically to be taken to outreach events where it demonstrates the component parts of a full-scale supercomputer in ways easy for a lay audience, including school children, to understand, get their hands on and experiment with.

Wee ARCHIE is used to illustrate some of the central concepts of supercomputing and provide examples of how they relate to real-world activities.

These include an airplane CFD demo to model airflow over a wing, computing all sorts of parameters that would be crucial for flight. Participants can then view their wing on a passenger aircraft, see how the easily the aircraft will take off (or not!), and the overall range of the plane on a single tank of fuel. By tweaking the wing parameters and re-simulating on Wee ARCHIE they can optimise the wing for specific aspects – but it can be a trade-off between take-off performance and range.

Other demos include Dinosaur racing and a scaled-down molecular dynamics code running a simulation of pheromones in mouse urine. Although these demos are chosen to be fun and interesting, crucially they are all based on real simulations that run on ARCHER. This has proven invaluable in appealing to a wide audience: young children can play the game while their parents or elder siblings learn about the scientific research applications.

Figure 2. Wee ARCHIE on the ARCHER booth at the Big Bang Fair, Birmingham UK March 2017

Figure 3. CFD simulation running on Wee ARCHIE

In addition to the tech-based activities we also provide some family-friendly, hands-on challenges to demonstrate basic concepts of parallelism. Particularly successful is the bean-bag sorting challenge where teams are timed sorting 'serially' and then in parallel. A leaderboard competition adds to the fun and provides an additional incentive.

At the "Big Bang Fair" [6] held in Birmingham on 14-17 March 2017, the bean bag sort challenge ran 414 times with 1202 participants actively involved. The 'Build your own Supercomputer' app was played by 628 people and 1036 wing designs were made and tested on Wee Archie. In addition many people were watching these activities and interacting with the staff on the booth, so we estimate that in total we reached over 5000 visitors.

B. Build your own supercomputer

This project was initially developed as part of the PRACE Summer of HPC 2015 [7].

Although Wee ARCHIE is designed to be portable, it still requires a physical venue such as a Science Fair or school classroom, transportation and support from a member of the EPCC team. We wanted to develop a hands-on simulation which could be accessed by anyone, anywhere, anytime, online. After a lot of thought and discussion we decided a game-style web page/app was the best approach.

In this game you are in charge of a supercomputer centre and you have to run lots of jobs on your machines, using your budget wisely to buy the best hardware for the job. The main idea behind the game is to create a design of a supercomputer, balancing its components against budget and power efficiency.

Figure 4. Build your own Supercomputer app

Your design must handle a wide range of incoming user jobs, and the completion of these provides money which can be further invested. As you progress through the levels the jobs become more complex (and lucrative) and additional components are available to be included within the machine. Besides passing through the levels, you can also obtain badges that are awarded for specific achievements such as a green supercomputer, profitable machine and the overall number of jobs run.

Although the game can be played just for fun, it is designed to illustrate many fundamental concepts and constraints of a real system. These include power efficiency, terminology such as 'node', 'accelerator' and 'scalability', and the fact that different simulations have varying optimal hardware configurations.

As well as the full game [8], the app can be run in a kiosk mode suitable for e.g. a booth at a science fair where each participant's time is limited to a few minutes. In this situation we also maintain a competition leaderboard.

C. ARCHER Champions

ARCHER Champions [9] launched in March 2016 as an initiative to build a self-sustaining, self-supporting user community and resource.

A champion can be "an individual within a regional or national centre (or institution with equivalent resources) who is engaged in: advising users on the access of appropriate computational resources: and in supporting these users to access and exploit these resources appropriately."

The aim is to provide an informal network of contacts around the UK who are familiar with all aspects of the ARCHER service. We feel that, particularly for a new user, being able to talk to someone in their own institution is a much better way to explore the possibilities of using the national supercomputer than making more formal enquiries via the ARCHER helpdesk.

An ARCHER champion will know about the routes to access and application procedures, making it easier for those making their first foray onto a national service. They may

also have personal experience of running on ARCHER, so that new or prospective users have someone local they can talk to about whether ARCHER is a suitable resource for their work, and who can provide initial guidance on the application process and getting started on ARCHER. A local contact is also best placed to advise on what resources are available locally, meaning that users are more likely to apply to the most appropriate service at the outset.

The first meeting was held in Edinburgh in May 2016, with follow up meetings held every six months at different locations around the UK: September 2016 at the University of Oxford, and February 2017 at the University of Leeds. If possible we try and co-locate with other events: the Oxford meeting preceded an ARCHER training course on Scientific Python, and the Leeds meeting succeeded a meeting of the UK HPC Special Interest Group.

The meetings attracted 23, 25 and 34 attendees from all over the UK, with a wide range of expertise relating to HPC and ARCHER. At each meeting a number of ARCHER staff have presented on topics that are of particular interest to a Champion when providing advice to a prospective new user. Topics have included:

- ARCHER in the UK National HPC Infrastructure
- Migrating to ARCHER: Common Issues Encountered
- The ARCHER Training Programme
- SAFE: the ARCHER web-based user administration system
- The ARCHER public outreach programmme and Wee ARCHIE
- The eCSE Programme for funding HPC software development projects
- ARCHER Support Structures

The Champions themselves have increasingly been able to contribute to the meetings and deliver presentations and talks on their own work and experiences.

1) Tier-2 Centres

ARCHER is the UK's top "Tier-1" service and, as explained above, one of the key aims of the Champions network is to better link ARCHER to smaller local and regional services. The UK has recently funded a network of Tier-2 systems [10] so it is important that ARCHER Champions are closely linked to these new services.

The Tier-2 centres are establishing their own Champions network, and the leader of this activity gave a talk at the most recent ARCHER champions meeting. The whole Tier-2 initiative is very timely and, through support structures such as the Champions networks, we aim to ensure users at all levels across the UK can easily access advice on HPC and that best practice is widely shared.

It is important to realise that there is a two-way flow of information. Through the Champions networks, the Tier-1 ARCHER support team is also learning about what other options are available around the UK. This means that we are better able to appropriately direct users to local services when their needs are not best met by National facilities.

2) Research Software Engineers

Over the past few years the UK has started to recognise the key role that software development plays in scientific research. This has lead to the widespread adoption of the job description Research Software Engineer (RSE), and there is an active community based around the UK RSE Association [11]. Recently the UK Research Councils have funded specific RSE Fellowships [12] for up to 5 years to promote best practice in computational science to a wide range of communities.

It is important that the Champions network links to the RSE community, and from the outset we envisaged that a champion could be "A research software engineer located within a consortium or equivalent group of HPC users. They would act as a champion within their consortium - providing advice on support mechanisms, obtaining training in appropriate HPC technologies and HPC career development skills and engaging in outreach activities." All the existing RSE Fellows were personally invited to attend the initial ARCHER Champions meeting.

Although ARCHER champions is still relatively new, we are already building good links with the RSE community. For example, the RSE page at the University of Sheffield [13] has a section on HPC where they say they can:

"... help you with software installation, training and advice on how to make maximum use out of these HPC resources that are free to all members of The University of Sheffield. We work with the UK National Supercomputing service, ARCHER, as ARCHER Champions. Our role is to advise users on the access of appropriate computational resources: and in supporting these users to access and exploit these resources appropriately."

This is exactly the role we envisaged for the ARCHER Champions and we hope that the activities in Sheffield will be replicated widely across the UK.

D. Image Competition and Calendar

For the last three years we have run an ARCHER Image Competition [14] where ARCHER users are invited to submit images based on their work on ARCHER which demonstrate "ARCHER Enabling Research".

Entrants may submit up to five individual images along with a caption and brief description of the image and the work which prompted it, putting the use of ARCHER into context, and the competition is judged by a panel comprising members of ARCHER CSE team, ARCHER users and other invited contributors.

The winning entry is awarded a cash prize of £250 and the winner is also invited to write up a Case Study of their work.

The competition has been extremely popular, with over 30 entries each year, all of a very high standard. All the entries are displayed in galleries on our ARCHER Web site [15] together with their captions.

The top twelve entries each year, as selected by the judging panel, have been published as a calendar [16] for the past three year. In 2016, 600 copies were distributed to a wide variety of destinations including ARCHER PIs and users, University personnel both here in Edinburgh and

around the UK and Europe, science centres, schools and many other individuals and organisations who have an interest in ARCHER and UK HPC. Each month displays a different image along with the description of how it demonstrates ARCHER Enabling Research.

Figure 5. ARCHER Calendar 2017

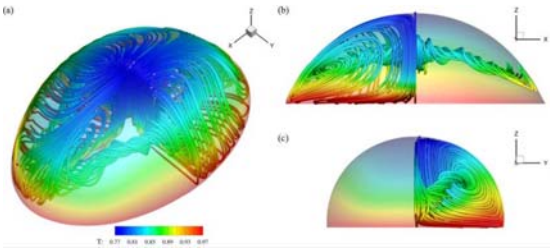

Figure 6. Winning image 2014 : Dr Pedro J. Sáenz, University of Edinburgh Institute for Materials and Processes.

Figure 7. Winning image 2015 : Mr Ivan Langella, University of Cambridge, Department of Engineering.

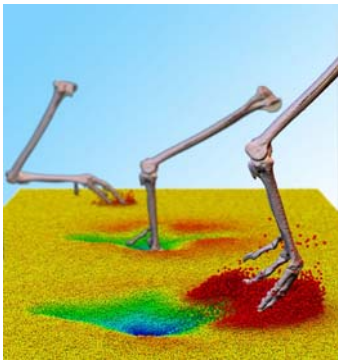

Figure 8. Winning image 2016 : Dr Peter Falkingham, Natural Sciences & Psychology, Liverpool John Moores University.

By submitting images to the competition, entrants agree to the ARCHER Service having rights to use the images in promotional materials. This has proved invaluable in providing a library of attractive images with accompanying text describing the underlying science. Work is always credited to the authors and their institution so they also benefit from raised profile and impact, something which is increasingly important in obtaining funding in the UK.

III. ARCHER Driving Test

The ARCHER driving test [17] was created and launched in February 2015.

We had become aware that novice users were attending our Hands-on Introduction to HPC training course, where they were given access to an ARCHER guest user account, but once the course ended they no longer had any access to ARCHER, or possibly any HPC system at all, in order to practice what they had just learned or try out ideas on how to apply their new knowledge to their own research areas.

Whilst requiring new users to pass a 'test' as a gateway to providing low-barrier access may seem counter-intuitive, in actual fact the test is designed to be extremely easy to pass for anyone who has only the most basic knowledge about HPC. It is comparable to requiring a teenager to have attended a couple of hours of Drivers Ed prior to handing them to the keys to a vehicle, as opposed to the practice and assessment involved in obtaining a commercial pilot's license. Almost 90% (259/292) of people attempting the Driving Test have achieved a pass, though some of them have taken two or more attempts.

We also ensured that all the introductory materials which would enable ARCHER Driving Test candidates to prepare to pass were readily available online.

We video-recorded a run of the Hands-on Introduction to HPC training course and made sure a full set of the course materials, including the PowerPoint presentations, the practical task handouts and resource files were all available on our on-line training web page, with a link to the recordings on our ARCHER YouTube channel [18]. This ensured that anyone could access the materials to either walk through independently or to act as a refresher for the face-to-face training.

We then developed a bank of questions, based on the topics covered in the course materials, using the on-line quiz tool Quia [19]. The Quia user licence is only $49 per year. The questions are sub-divided into different topic areas:

- Compiling
- Hardware
- I/O
- PBS job scheduler
- Programming
- Running jobs

From this bank of over 100 questions, a quiz comprising 20 questions is presented to each site visitor which will pick a range of questions from each of these topics according to parameters set within the quiz e.g.

- 2 x Hardware questions
- 3 x I/O questions

- 4 x Programming questions
- 3 x Compiling questions
- 1 x PBS question
- 3 x Running jobs questions
- 4 x random questions from those not yet used.

In addition to the assessment questions, the quiz also gathers background information about the person including gender and career stage.

We request feedback on how the visitor found the quiz, comments on whether they found anything about the quiz or the online training materials unclear, and we ask them to submit a suggested question that could be used in future.

This background information has been very useful in helping us understand the various types of visitors who have made use of the Driving Test.

As at the end of February 2017, with the Driving Test having been in operation for two years, we have invited 223 to sign up for an account on ARCHER having passed the test, and 165 people have done so. They have run over 15,000 jobs on ARCHER using almost 80,000 kAUs (kilo Allocation Units) = 5.3 million core-hours.

We have also seen that certain projects and PIs and Supervisors have been actively encouraging their students to undertake the training and the driving test to enable them to get hands-on experience before working on their project. This should help to ensure that novice users do not make potentially costly mistakes such as running jobs with settings that burn excess machine time or making errors in their submissions which cause jobs to fail.

We also recognised that not all ARCHER users are writing their own codes; many user communities make extensive use of pre-installed packages and the initial ARCHER Driving Test was particularly tricky for such users since it always included a number of questions on programming and compiling. We therefore also developed a bank of questions which focus on using library modules, and launched a second version of the Driving Test specifically for non-programmers.

This version asks the same background questions but with a distribution of

- 3 x Hardware
- 4 x I/O
- 2 x PBS job scheduler
- 3 x Running jobs
- 8 x Non-programmer module use questions

When in 2016 the Knights Landing (KNL) hardware was added to ARCHER we added a third version of the Driving Test specifically on the KNL training materials which were added to our documentation. Whilst access to the KNL was granted automatically to any existing ARCHER user, this allowed non-users to gain access to the KNL specifically.

The background information collected since 2015 has been very useful in helping us understand the various types of users who have made use of the Driving Test.

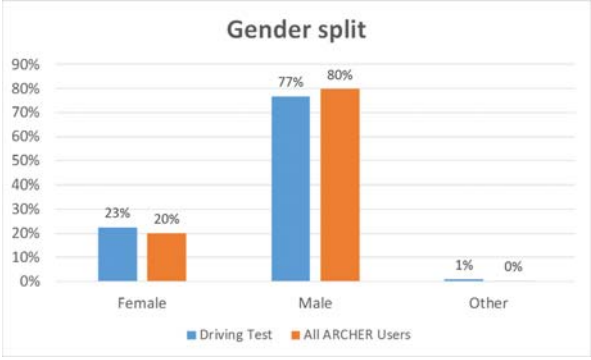

Figure 9. Gender of users

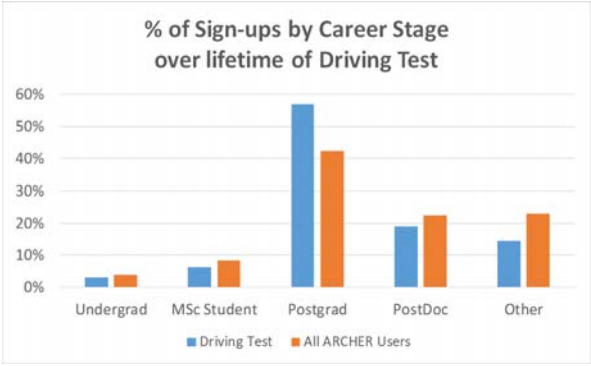

Figure 10. Career Stage of users over lifetime of Driving Test

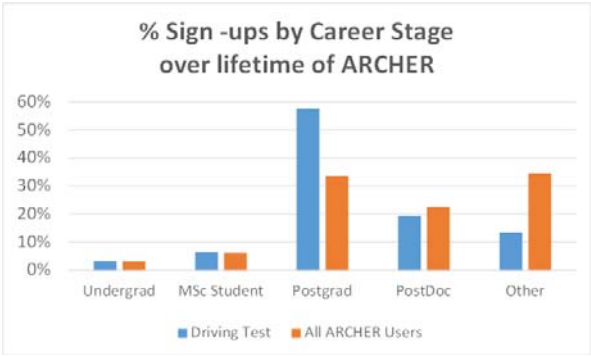

Figure 11. Career stage of users over whole lifetime of ARCHER service

These Figures show some interesting trends. First, the gender split of Driving Test users is the same as the overall ARCHER user base. Although still heavily skewed towards male users, this does demonstrate that the Driving Test does not favour either gender any more than does the entire ARCHER service.

When looking at career stage, we see that the Driving Test favours early-stage researchers: over the lifetime of the Test (2015-2017), the percentage of PhD students gaining access via the Test exceeds the overall ARCHER percentage (Figure 10.). The effect is even more dramatic if we compare to all ARCHER users (Figure 11.), but this is not really a fair comparison as the early users of ARCHER, who signed up well before the Test began, are typically dominated by more experienced users.

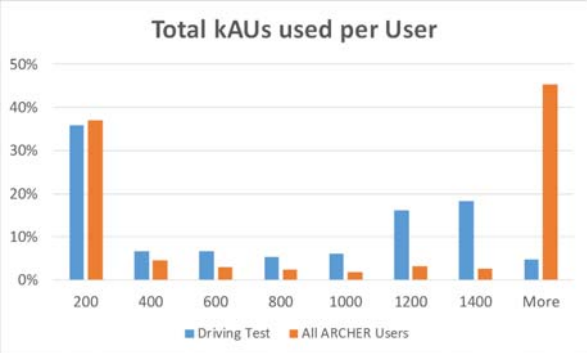

Figure 12. Distribution of usage (binned into intervals of 200 kAUs)

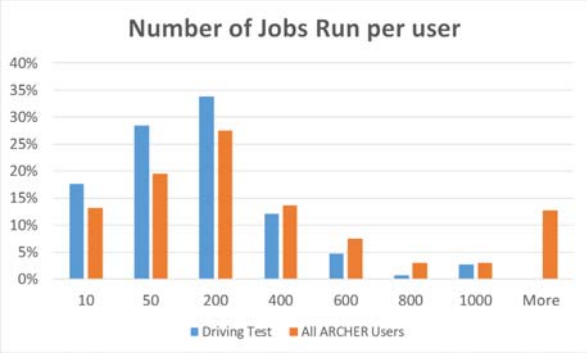

Figure 13. Distribution of number of jobs run (variable bin size)

Overall, Driving Test users use around 50% of their allocation. The distribution of time used shows peaks at small usage and large usage**.** This is presumably because one class of users is investigating ARCHER but not running production jobs, whereas the other class is doing production runs or looking at scaling to large process counts. The number of jobs run per user is remarkably similar to the overall ARCHER user base which is surprising given that Driving Test users have relatively limited time. The only real difference is that they never run more than 1000 jobs which is not surprising given their small allocation.

Note that a significant number of users have managed to use more than their allocation of 1200 kAUs, which happens when their final job uses more time than they have left in their budget. This is possible because the check done on ARCHER is only that they have sufficient budget to cover the job's requested execution time at job launch, and there is some tolerance as the majority of jobs will not run for the entirety of their requested wall time limit. A job whose wall time does approach the requested time can therefore overuse the budget by this tolerance factor.

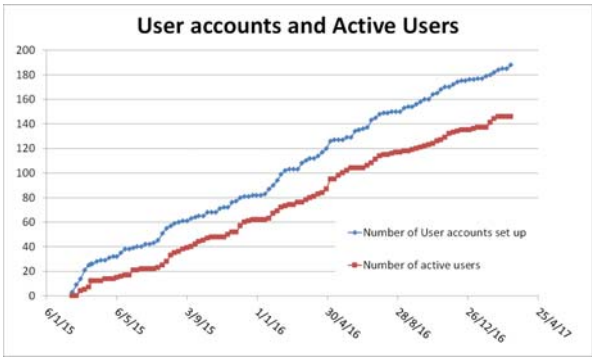

Figure 14. Sign ups vs actual use of accounts

A few users sign up but never make use of their account but a high proportion of around 75% does make good use of it. In the very early days of the Test there was a larger gap between accounts and active users. This effect is because, to try out the Test in a real situation, we asked everyone at a training course of about 30 people to complete the Test. Many of them either already had ARCHER accounts, or were using other HPC facilities, so did not really need ARCHER access.

IV. Training

The ARCHER Training Program is designed to introduce and support all HPC users at all stages of their HPC use, from initial introduction to 'What is HPC and how might it be useful' through to advanced courses for experienced niche users.

The program has been developed over a number of years and evolved in response to demand from users (and potential users). There is a close link with our University of Edinburgh teaching, such as our campus-based one-year taught Masters programmes and our accredited online training courses in HPC and Data Science [20].

A. Face-to-face

The core component of our training is traditional, face-to-face classroom-based teaching in UK Academic institutions throughout the UK. We typically deliver 65 days face-to-face training per year with around three quarters of this taking place outside of Edinburgh. Courses tend to last 2-3 days and are led by one or more members of EPCC, often supported by specialist tutors from the UK or abroad for our more advanced courses. As a PRACE Advanced Training Centre, we have access to additional funding to support the more advanced courses. Some of our courses are delivered as part of an institution's larger training syllabus; for example, we deliver two 3-day courses at the University of Southampton which comprise part of their taught postgraduate programme, but these sessions are also open for anyone to sign up for and attend as part of a more general HPC audience. All the materials are available in a comprehensive online archive [21].

Attending a face-to-face course is a significant investment in both time and money and it is therefore essential that sufficient guidance is available to ensure

people are signing up to courses of the right level and content. Signing up for a course which is too advanced leads to frustration and disappointment. To address this we did three things:

1. Created a single page with short course descriptions for our entire syllabus [22].
2. Classified all the courses by level as 'Introductory', 'Intermediate' or 'Advanced'.
3. Provided guidance about the pre-requisites and suggested paths through the courses on offer. These are linked back to the course descriptions.

This enables people to review the materials for pre-requisite courses, either to bring themselves up to a suitable level or to reappraise which course may be most appropriate for their needs.

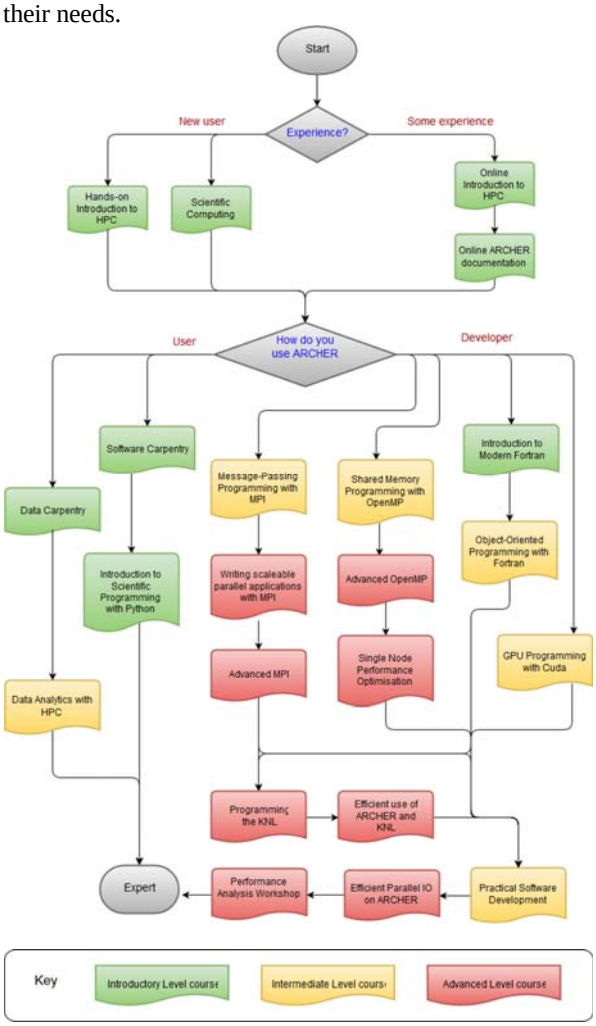

Figure 15. Course classification by level and suggested training pathways

B. *Online training*

In addition to the face-to-face teaching, we provide a wide range of online courses. Some of these are recordings of a face-to-face course which is then available via our YouTube channel [18] for attendees to use to re-cap, or for those unable to attend to cover the material, or to get a feel for the nature of the course before attending a future face-to-face run of the course. As well as the recording, all the training materials are made available online including the PowerPoint presentation slides, the handouts and practical worksheets.

Our online documentation of the ARCHER service [23] [24] also includes both web pages and supporting 'how to' videos. Many of our users are novices in HPC and often also non-native English language speakers. By creating short narrated video walk-throughs of many of the common tasks, users are able to read through the steps and then follow along with the video which guides them through menu choices, button clicks and option selection. The topics cover getting an account set up on the ARCHER online administration system, SAFE, requesting a user account on ARCHER, managing a project as a PI, setting up an SSH client on your machine, logging in using SSH keys, how to access help from the command line, and how to run interactive jobs on ARCHER.

Videos are very short and focused - many are under one minute in duration, demonstrating only one topic - so users do not have to skip back and forth through a long session to find what they are looking for. These have had over 300 views to date.

We also have recordings of the delivery of some of our core training courses such as the "Hands-on introduction to HPC" and "Message-passing programming with MPI" [25]. These longer videos cover one lecture each, gathered together into playlists to cover the entire course. All the supporting materials including PowerPoint slides, handout sheets, practical exercise tasks and supporting files are all made available [26] with these videos via our training website so anyone can download them and work through the entire course at their own pace. The exercises are designed to be able to run on other HPC systems than ARCHER with minimal modification, and links are provided to the Driving Test so new users may complete the test and use the time awarded to complete e.g. the MPI course on ARCHER.

Some of our training is not only video captured but live-streamed to allow unlimited audience. Our "Scientific Computing" course is delivered over four consecutive Wednesday afternoons with local students attending in-person in a University of Edinburgh lecture theatre. The lectures are available to watch in real-time with the option to interact with and ask questions of the presenter via a web-chat. The recording are then made available for anyone to watch or re-watch, along with the practical exercises and materials [27].

We also host online Webinars at least once per month where one or more speakers, from EPCC or external, present a short talk on a relevant topic and then invite questions and discussion on that topic or on ARCHER in general. These are consistently popular, attracting audiences of 15-30 depending on the topic discussed. These webinars are also recorded and made available on YouTube along with the presentation slides and a transcript of questions asked and answered [28].

As part of the PRACE project, EPCC, supported by staff based at SURFsara [29], has led the development of an introductory 5-week MOOC called "Supercomputing" which ran for the first time on FutureLearn from 6th March 2017 [30]. At the time of writing this is still ongoing. It attracted more than 3,000 registrations with more than half starting the course and over a thousand of these actively engaging at the outset. It is well recognised that, for free online courses such as this, attendance will tend to drop off significantly as the course progresses. However, we expect approximately 10% of those starting to complete the full course which would be in excess of 150 people.

There is significant overlap between the materials used in developing the MOOC and those developed for our existing introductory courses and the outreach activities including Wee ARCHIE. Although developing the MOOC was a significant undertaking requiring many person-months of effort, once up and running the learners are largely self-supporting and as educators we only need to provide occasional support contributions, perhaps 3-4 short responses and comments per day. As a result we expect that future runs will be easy to support and therefore provide a very cost-effective way of introducing new people to the subject.

V. Training Impact

We endeavour to get feedback on our training at every stage of the training planning, delivery and post delivery process and also to see how the attendees have benefitted from the training in the longer term: the impact in terms of improved new work and opportunities in their research, successful progression to new machines, new programming languages or techniques, successful funding.

This type of longer term feedback, corresponding to levels 3 and 4 of Kirkpatrick's "four level" model [1] of training course valuation, is clearly extremely important but is also recognised as being very difficult to obtain.

We have worked hard to achieve a high return rate for our user feedback.

During face-to-face training, especially on introductory courses such as “Software Carpentry”, we ask for quick feedback on how the attendees feel the course is going e.g. in terms of depth of coverage of material, pace of delivery via hands up or traffic-lighting post-its where they stick a green post it on their laptop if all is well, and a red if they are stuck or struggling to keep up.

Figure 16. Using PostIt notes to ‘traffic-light’ student progress (Level 1 of the Kirkpatrick model)

We have also made use of mini-quizzes, where two or three quick multi-choice quiz questions are presented on the topics just covered to gauge how well they have grasped it and reveal any gaps or misunderstandings. We have used the Socrative tool [31] for this which is particularly easy to use – it can be open in a browser window or a smart phone app whilst the course is ongoing. The responses are instant and can be anonymous, which encourages full participation.

For our face-to-face training we ask all attendees to complete an on-line feedback form which looks at their overall rating of the course on a scale of Excellent - Very Bad, and for the pre-course information provide, the registration process, the training venue and the overall organisation.

We also ask about the relevance of the topics, the level of challenge of the materials, the lecture presentation and pace, quality of the teaching materials and hands-on activities.

We ask for comments on what they liked most and least and suggestions for improvements and whether they would recommend the course to colleagues.

We obtained a response rate of over 50% in each of the last two years with an overall average rating of 4.3 / 5.0

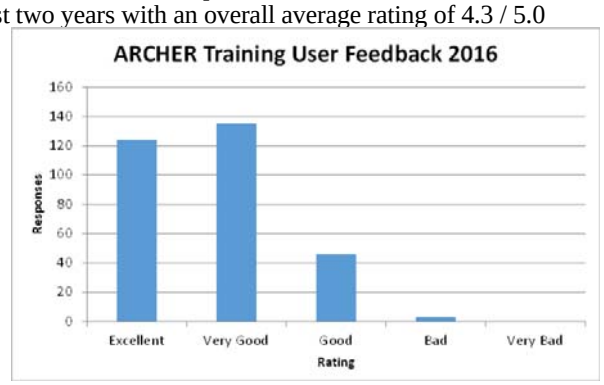

Figure 17. Post-course survey feedback (Level 2 in the Kirkpatrick model)

A. *Longer-term Training Impact Survey*

In addition to seeking feedback on each course, we are also keen to find out what impact the training programme has on attendees behaviour and career progression. The

ARCHER training progamme is approved by an external panel of training experts. We worked with the panel to develop a survey which would investigate how the training programme as a whole had benefitted the users. We have so far completed four such surveys which aim to assess benefits that occur on timescales of several months or more, the most recent being issued in late 2016 [32].

For the follow-up survey, we send an email to all of the people who have attended our training courses between 3 and 9 months previously (excepting those who have expressly asked not to be re-contacted). We follow up with a reminder email a couple of weeks later.

As would be expected, the response rate is significantly lower than the post-course surveys, with 52 responses to the most recent survey, a response rate of 10%.

We have tried various methods to increase this response rate: presently we pledge to make a £1 donation to charity (currently "Children in Need") for each response we receive. It is hard to quantify how effective this inducement has been given limited statistics but at a modest cost of only £52 there appear to be no downsides. We also tried re-inviting anyone who had ever attended an ARCHER training course (i.e. with no 9 month cutoff) but this did not generate any significant increase in response rates.

B. ARCHER Follow-up Survey responses December 2016

Here we present some details from the most recent survey. Sixteen different courses (some of which may have run two or more times in the period) were run in the period. The questions asked included:

Name and Career Stage
Which courses have you attended?
Course Impact - all questions rated on a Strongly Agree-Strongly Disagree 6 point scale:
(a) I have applied the training in my research work
(b) It has resulted in productivity improvements in my work
(c) It has resulted in performance improvements in my codes
(d) It has helped in submission of grant applications (for ARCHER or other systems)
(e) It has enabled me to port new codes to HPC systems
(f) It has enabled me to make use of new HPC packages
(g) It has increased my overall understanding of HPC
(h) I have passed on aspects of the training (formally or informally) to colleagues
(i) The impact has been on my use of: ARCHER/Other HPC Systems/Neither/Other
Are there any courses or topics that you would like us to offer in the future which we do not already offer?
Any other comments or suggestions about the ARCHER Training program?

The questions also allowed a response of 'N/A' and additional comments.

Some of the key question responses are included here.

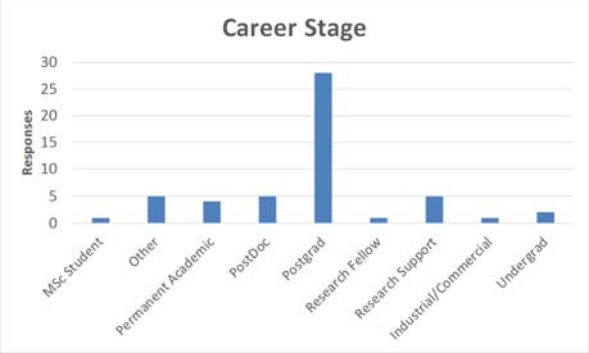

Figure 18. Career stage of survey responders

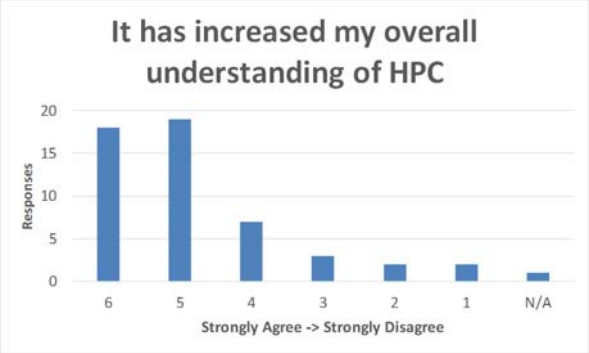

Figure 19. Responses to a typical survey question on 6 point scale. (Level 3 in the Kirkpatrick model)

Figure 20. Longer term impact question in survey (Level 4 in the Kirkpatrick model)

The response rate is dominated by postgraduate students which reflects the demographic of the course attendees. This is a key group to target as we would hope that training would help them in their careers. The responses shown in Figure 19. are encouraging in that they show that the courses generally build user confidence. Attendees have also reported passing on what they have learned on our courses to colleagues. 70% have said that the training also helped them on systems other than ARCHER.

However, it is harder to gauge the quantitative impact this has in the longer term. Figure 20. is less clear cut than Figure 19. but the trend is still generally positive.

We also use the follow up survey as an opportunity to invite attendees to suggest new training courses which they would like us to offer. As a result of the most recent survey responses we are investigating the options for a new course in C++ for HPC.

Despite the acknowledged difficulty achieving high response rates and in quantifying longer term impact, we feel that the survey is still very informative.

VI. Conclusions

We have described how EPCC addresses one of the key challenges to any HPC provider: growing the user base. This requires making new users aware of the potential benefits of the service and ensuring a low barrier to entry. We have outlined how we approach this on three different timescales:

- **short term**: giving new users easy access to the ARCHER supercomputer through the Driving Test, and helping them undertake common administrative tasks via a series of short "How-To" videos.
- **medium term**: making our standard HPC training courses more accessible by placing them in context inside a larger training syllabus, making some courses available entirely online, developing a support network of "ARCHER Champions" and engaging existing users through the ARCHER Image Competition.
- **long term**: aiming to engage the next generation of computational scientists through a range of innovative public outreach activities and a MOOC, and trying to assess the long-term impact of our training programmes with follow-up surveys of past attendees.

Common to all these activities is the requirement for substantial upfront investment of effort. Developing a bank of questions for the driving test, designing and building Wee ARCHIE, devising practical games that illustrate key HPC concepts, creating a MOOC or designing a follow-up questionnaire all take a lot of time.

This is exactly why we are so keen to share experiences with other HPC providers: given the investment required, it is extremely valuable to know what types of activities work well in practice.

However, all of these approaches to user engagement become more cost-effective each time they are run. This is clearly the case for online activities: we expect the overhead of delivering the second run of our Supercomputing MOOC to be at least an order of magnitude less than the first run. However, it is also true of hands-on activities such as the demonstrations we run on Wee ARCHIE and the bean-bag sorting game. Once developed and run a few times by experienced staff at Science Festivals, they can be scripted and handed out to anyone with enough enthusiasm even if they only have a basic knowledge of HPC. Not only is the preparation time minimised, but the delivery load can be shared across a very wide range of staff.

All these activities complement rather than replace the more traditional methods, and we still run face-to-face courses in, for example, Advanced MPI Techniques. However, such advanced courses only target a subset of existing users and do nothing to expand the user base. In a time of increased financial pressures it is essential that HPC services engage users from new applications areas. It is equally important that the general public, who after all support HPC through their taxes, are made aware of the benefits of HPC so that we have public support whenever funding is at risk. Although training users to scale their parallel codes to many thousands of cores is still required, we need to ensure that we continue to have the large HPC systems on which to run them.

Acknowledgment

We would like to acknowledge the work of many of the ARCHER Outreach team www.epcc.ed.ac.uk/discover-and-learn/about/about-us/. The authors acknowledge funding from the ARCHER project www.archer.ac.uk and from PRACE via the EU's Horizon 2020 Research and Innovation Programme (2014-2020) under grant agreements 653838 and 730913.

References

[1] PRACE Research Infrastructure www.prace-ri.eu/

[2] The Kirkpatrick Model, http://www.kirkpatrickpartners.com/OurPhilosophy/TheKirkpatrickModel/tabid/302/Default.aspx

[3] "Key Performance Indicators and Learning Analytics" at PRACE Training Workshop, 23/06/16, ISC2016, Frankfurt https://events.prace-ri.eu/event/498/material/slides/.

[4] Tiny Titan, Oak Ridge Leadership Computing Facility https://tinytitan.github.io/

[5] Edinburgh FUSION team assisted with the design and build of the case for Wee ARCHIE http://www.scifun.ed.ac.uk/fusion/.

[6] The Big Bang Fair 2017 https://www.thebigbangfair.co.uk/

[7] PRACE Summer of HPC 2015 Project reports https://summerofhpc.prace-ri.eu/wp-content/uploads/2016/04/SoHPC2015_final_editorial.pdf

[8] Build Your Own SuperComputer game https://archer-ebapps.epcc.ed.ac.uk/challenge/index.html

[9] ARCHER Champions www.archer.ac.uk/community/champions/.

[10] Tier-2 High Performance Computing Centres https://www.epsrc.ac.uk/research/facilities/hpc/tier2/

[11] UK RSE Association www.rse.ac.uk

[12] RSE Fellowships https://www.epsrc.ac.uk/funding/calls/rsefellowships/.

[13] RSE page at the University of Sheffield http://rse.shef.ac.uk/activities/

[14] ARCHER Image Competition http://www.archer.ac.uk/about-archer/news-events/events/image-comp/

[15] ARCHER Image Competition gallery http://www.archer.ac.uk/about-archer/news-events/events/image-comp/gallery-2016/

[16] ARCHER calendar http://www.archer.ac.uk/about-archer/news-events/events/ARCHER-calendar/index.php

[17] ARCHER Driving Test http://www.archer.ac.uk/training/online/driving_test.php

[18] ARCHER YouTube channel https://www.youtube.com/channel/UCH28DWRmpadD-hEECymaxdA/playlists

[19] Quia quiz tool www.quia.com

[20] EPCC Training https://www.epcc.ed.ac.uk/education-training

[21] ARCHER Training archive online archive http://www.archer.ac.uk/training/past_courses.php

[22] ARCHER Training courses http://www.archer.ac.uk/training/courses/

[23] ARCHER SAFE User guide http://www.archer.ac.uk/documentation/safe-guide/

[24] ARCHER User guide http://www.archer.ac.uk/documentation/user-guide/

[25] Message-passing programming with MPI lectures https://www.youtube.com/watch?v=_h55hwpLwoE&list=PL1b57Q937PoslcIozHf7UMIkfMHQQLVQe

[26] Message-passing programming with MPI supporting materials http://archer.ac.uk/training/course-material/2014/07/MPI_Edi/

[27] Scientific Computing http://archer.ac.uk/training/courses/2015/10/SciComp/index.php

[28] ARCHER Webinars http://archer.ac.uk/training/virtual/

[29] SURFsara https://www.surf.nl/en/about-surf/subsidiaries/surfsara/

[30] Supercomputing MOOC https://www.futurelearn.com/courses/supercomputing/

[31] Socrative quiz tool https://www.socrative.com/

[32] ARCHER Training Follow-up Survey 2016 https://docs.google.com/forms/d/1oLIPcBiwuwr4Nk0soVpcTi-w3ndXnWJvp1sIAaziA0U