



UNIVERSITY OF
BIRMINGHAM

Research Software Group

Advanced Research Computing

Annual Report 2022

BEAR
BIRMINGHAM ENVIRONMENT
FOR ACADEMIC RESEARCH



Contents

Research Software Group	3
Advanced Research Computing	5
Research Software Engineers	6
BEAR Software: Advice, Coaching, Coding and Mentoring	7
Case Studies	11
BEAR Applications	17
Training	22
Conferences and Events	25
Contact Us	28

Research Software Group

The Research Software Group (RSG) is part of Advanced Research Computing at the University of Birmingham. The RSG was formed in November 2017, starting with two Research Software Engineers (RSEs): Simon Branford and Andrew Edmondson. We designed the free Advice, Coding, and Coaching model, and began working with researchers with the aim of improving the research software they write and use. BEAR Software's mission is summed up in the words of the Software Sustainability Institute: "better software, better research". [<https://www.software.ac.uk/resources/publications/better-software-better-research>]

Over the next few years the group grew steadily and at the start of 2023 the RSG will have 14 team members.

The RSG work is split into three parts:

1. We support users of BEAR's high performance computing services BlueBEAR, BEAR Cloud, CaStLeS and Baskerville. A significant part of this work is installing, curating and supporting over a thousand applications and libraries.
2. We provide free advice, coaching and coding work, collaborating with researchers from PhD students to senior academics across the whole University.
3. We join research projects or groups for longer-term funded collaborations. This can be via funding dedicated RSEs managed by ARC, or by projects requesting time from our growing Pool of RSEs for funded work.

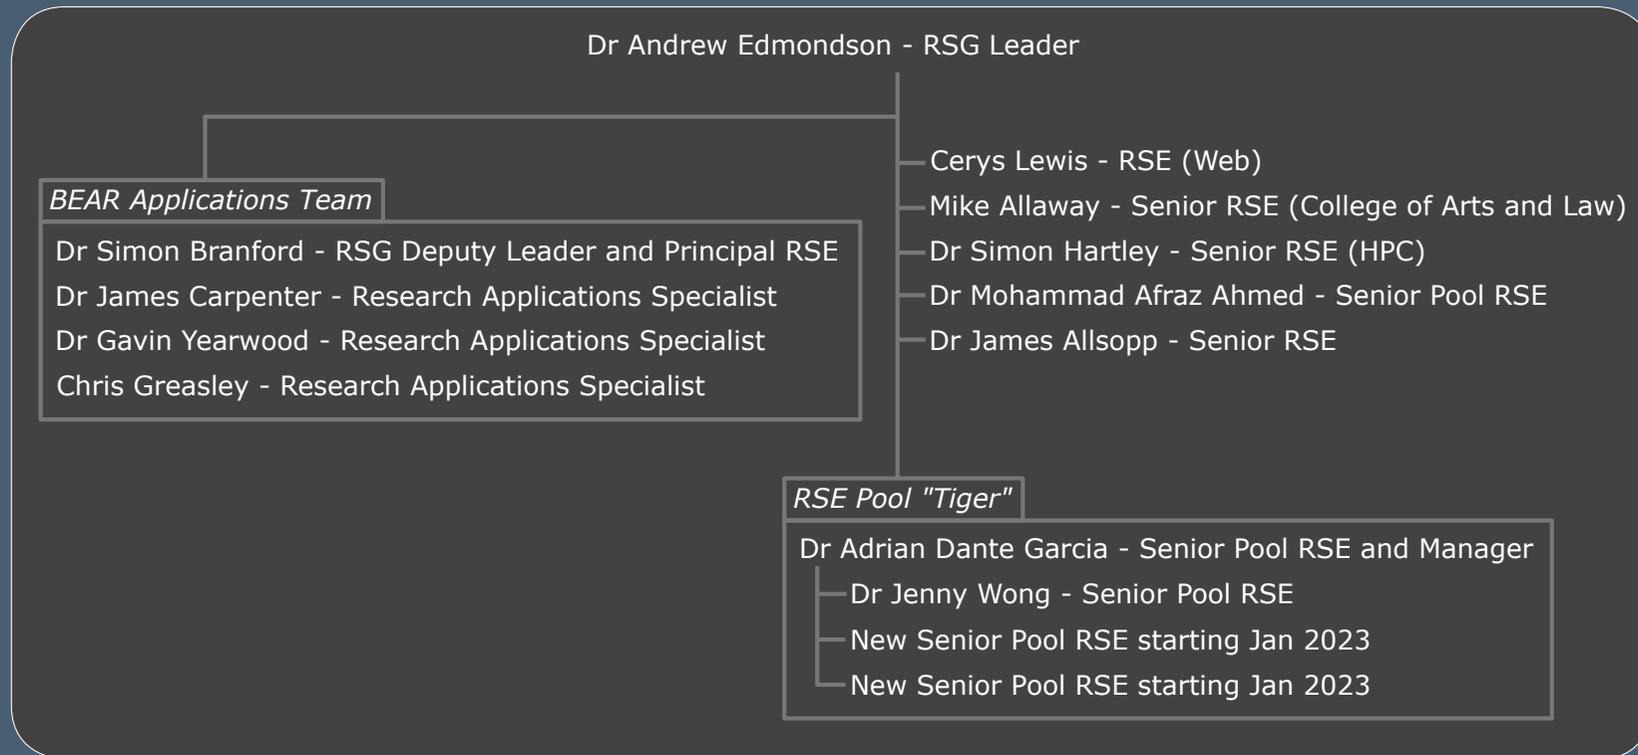
The Research Software Group exists to:

- Enable the University of Birmingham's research community to get the best from their research software
- Provide specialist software engineering advice and support to researchers and research software engineers (RSEs)
- Help to enhance the University's reputation for high quality research
- Help researchers get the most from BEAR services, maximising the return on the University's investment in BEAR



Research Software Group Structure

The RSG has grown substantially this year, with three new starters in 2022 and two more coming in January 2023. We are currently organised as follows:



BEAR

BIRMINGHAM ENVIRONMENT
FOR ACADEMIC RESEARCH

Advanced Research Computing

The Birmingham Environment for Academic Research (BEAR) is a suite of services managed by Advanced Research Computing (ARC). These services are designed to support research and many are free at the point of use. More than 2,800 research projects with over 4,700 staff and students are currently registered to use BEAR services. Since starting to collect the data in 2021, our academics have told us about over 900 publications that resulted in full or in part from using BEAR.

ARC is based in IT Services, led by Carol Sandys and reporting to the Chief Information Officer, Mark Gee. The Team grew rapidly following the 2015 Vice Chancellor's Review of IT Services and in direct response to the increasing research computing needs of our academic community. The need for and commitment to continued expansion of BEAR was reasserted in the most recent Review and in the Digital Strategy with renewed funding for infrastructure as well as support for the recruitment of additional research software engineers.

ARC is comprised of three groups:

- Architecture, Infrastructure and Systems Group
- Research Engagement and Data Group
- Research Software Group

'These groups provide experts to design and operate ARC's award-winning compute, storage and related services and also specialists at engaging with and supporting our wide and varied community of researchers. The report that follows provides an introduction to the work of the Research Software Group and gives brief case studies to illustrate the benefits partnership with the RSG, and ARC, can deliver.' - **Carol Sandys, Head of Advanced Research Computing and Research Engagement**



Research Software Engineers

The term 'Research Software Engineer' (RSE) was coined in March 2012 at the Collaborations Workshop and you can find all about the history of RSEs from the State of the Nation report from 2017: <https://doi.org/10.5281/zenodo.495360>. RSEs use software engineering practices to develop and support software in research applications. RSEs are frequently multi-skilled and adaptable, turning their hands to anything from coaching researchers in data science, optimising high-throughput data pipelines, developing full-stack web applications, or implementing AI models on supercomputers - helping improve the quality of the University's research. All the RSEs at ARC have a wide variety of technical skills and most have been researchers at some point in their careers.

The value of good Research Software Engineering

Research software engineering covers a broad spectrum of applications and good Research Software Engineering principles result in more reproducible and accurate research. Or as the Software Sustainability Institute (SSI) puts it:

'Better software, better research' - **Carole Goble, 2014**, <https://www.software.ac.uk/resources/publications/better-software-better-research>

'Software can be a few lines of code written by a single researcher or a major framework developed over decades by dedicated teams of researchers and software engineers. In academia, 92% of UK researchers use research software and 69% report that it is fundamental to their research with software development a research activity in its own right. The near-ubiquity of software means that it is not possible to disentangle the quality of the software from the quality of the research.' - **UKRI, 'The UK's research and innovation infrastructure: opportunities to grow our capability', 2019**, <https://www.ukri.org/files/infrastructure/the-uks-research-and-innovation-infrastructure-opportunities-to-grow-our-capacity-final-low-res/>, p125

Having an RSE (or two) on your research project is a cost effective way of increasing the quality of your research.

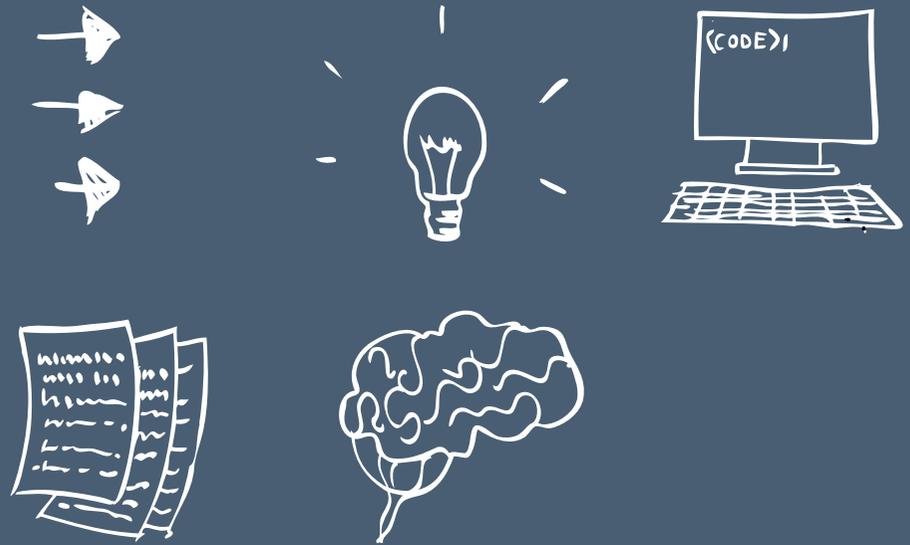
BEAR Software: Advice, Coaching, Coding and Mentoring

The RSG offers advice sessions and coaching, coding or mentoring engagements to all researchers in the University under the banner of 'BEAR Software'.

As is usual for BEAR Services, there is a standard, free service and premium, paid-for options. Our free services include one-off advice sessions, mentoring and short coaching or coding engagements.

We have two models of funded coding or coaching engagement - allowing more substantial time commitments. Standard funded engagements take the form of a percentage of one or more RSEs for a number of months or years. We also offer flexible engagements where the workload can be defined in a more complex manner, with individual work packages defined that may differ from each other while being part of a single overall engagement.

More detail on our free, funded and flexible services can be found later in this report or on our website: <https://www.birmingham.ac.uk/bear-software>



BEAR Software: Free Engagements

The RSG offers advice and mentoring, and short coaching or coding engagements free to all researchers in the University.

Advice: If a researcher either has need for new research software or wishes to improve their existing software, then BEAR Software RSEs can be engaged to talk through the needs of the researcher and offer advice on how best to achieve those needs. BEAR Software RSEs are able to offer advice on the whole life-cycle of research software. Examples might include using version control to improve software quality, or porting to or optimising software on a new compute platform. By following the BEAR Software Advice a researcher should be able to improve the reliability and maintainability of the research software they write/use and this will help the researcher with the reproducibility and robustness of their research.

Mentoring: Researchers who code, or RSEs embedded in research groups, can often be the only such person in their group. We offer the option free to any researcher or RSE who would like to develop a longer-term mentoring relationship with a member of the RSG. Through this mentoring relationship the mentee would develop their technical and non-technical skills as they discuss their work, career aspirations and training needs. Typically mentoring meetings take the form of an informal conversation a few times a year. Space is limited for this service and sometimes there may be a delay before a new mentoring engagement can start.

Coaching: If a researcher, or research group, has the need for specific research software expertise, then an RSE can be requested, free of charge, for multiple half-day coaching sessions (max. 10). During these sessions the RSE would meet the researcher or research group (physically or virtually) to work with them to complete specific agreed objectives. Examples might include learning a new skill, reviewing code or designing and implementing a testing framework. These tasks are typically achieved with a weekly meeting for example doing pair programming or teaching/homework model similar to PhD supervision.

Coding: If a researcher, or research group, has the need for specific research software expertise but doesn't want, or doesn't have the time for, a coaching engagement, then we offer up to ten days of coding time (sometimes spread over several weeks). Examples might include re-implementing an algorithm in a different language, adding parallelisation or CUDA (GPU) support to an application, or developing a research website that requires functionality not available in the Intranet or other central services.

For all these engagements, we write some lightweight paperwork to agree who, what and when. Later free follow-on engagements are available, demand permitting.

To request advice, mentoring, coaching or coding please fill in the form at <https://intranet.birmingham.ac.uk/bear/sd/bear-software-advice> and specify which service you would like.

BEAR Software: Funded Engagements

The RSG has a growing Pool of RSEs, available for funded projects. In many ways the funded engagements are similar to our free engagements, but generally longer or involving more people. Pool RSEs can be included in grant applications or bought out using existing research funding.

The idea is to allow a research group/project to pay for one or more RSEs with particular skills for a period of time. This could be a short project (e.g. two RSEs running a two-week workshop) or a long project (e.g. 50% FTE for 2 years). To achieve this we have a central pool of RSEs that can work on such projects without them leaving the university when the project finishes.

We offer two main types of funded engagement:

- Short projects (e.g. a month or less) can have any percentage FTE of RSEs - e.g. 20% of one RSE and 90% of another for a month - depending on availability.
- Long projects (typically 3 months or more) should be defined as essentially 50% of an RSE, or multiples. E.g. two RSEs at 50% for 2 years. For each 50% FTE we will actually quote 45% of the RSE, and 5% of a senior member of the RSG who will provide code review, support, mentoring and advice to the RSE. Our policy is that no RSE will work more than 45% on a single long project, so for larger requirements we will provide multiple RSEs as needed.

We also offer a funded Flexible RSE Engagement - which is covered in detail in its own section.

It is rare that an RSE engagement will have static requirements specified up front - and much more likely that the researcher and RSE will work together defining and re-defining the goals of the project as time goes on. As such it is valuable to think of RSE engagements as research collaborations, and RSEs as more like researchers than consultants.

For more detail about our funded engagements and for instructions on how to include an RSE in a [Worktribe](https://intranet.birmingham.ac.uk/bear-software/funded) project see <https://intranet.birmingham.ac.uk/bear-software/funded>.

If you are considering including an RSE on a grant application please get in touch with us at an early stage by emailing:

bear-software@contacts.bham.ac.uk

We cannot guarantee to supply grant-funded RSE time if we were not consulted prior to the bid being submitted to the funding body. Nevertheless, if you find yourself in this position do get in touch and we will try to help.



BEAR Software: Flexible Engagements

In addition to our standard Funded Engagements, where one or more RSEs are defined as working a certain percentage of their time on a project for a defined period of time, we also offer a funded Flexible Engagement.

A Flexible RSE agreement is one where a researcher funds some RSE time to be used as required. It provides a framework for individual pieces of work ('work packages'). These work packages can then be defined at any time in the life of the project, to suit the project's needs.

For example, a project might need RSE input for, say, three months at the start of year two to help review and validate software written by the researcher before starting using it in earnest in simulations or experiments. And then again, a year later to help to package up and publish the software for use by third parties. Or RSEs might be needed to run a variety of training courses or coaching sessions periodically throughout a project's life.

With suitable notice and planning the same RSEs could work on several work packages as part of a research project, spread over years. Equally, different combinations of RSEs may be more appropriate to deliver different work packages. Work packages can be all defined up front, or ad hoc throughout the life of the project. Please note that we cannot guarantee that an RSE will be available at short notice for a newly defined work package.

If you are considering including a Flexible RSE Engagement as part of a grant application please get in touch with us at an early stage by emailing:

bear-software@contacts.bham.ac.uk

We cannot guarantee to supply grant-funded RSE time if we were not consulted prior to the bid being submitted to the funding body. Nevertheless, if you find yourself in this position do get in touch and we will try to help.

Case Studies

We have helped researchers across the university via advice sessions this year. Here are some examples.

Advice Session: Working with R on BlueBEAR

We held an advice session with a researcher in the College of Arts and Law to discuss how he and his team could use BlueBEAR to run his R code. The code (broadly based around Monte Carlo simulations) isn't suitable for extensive parallelisation, which is a major benefit of BlueBEAR. However, it does require a lot of RAM and while his local machine has 96GB BlueBEAR can offer over five times this amount. Furthermore, his team are mostly off campus, so the ease of access to BlueBEAR remotely is a major advantage compared to configuring a local machine for remote access by the team. We also advised the researcher of the available in-person and online training courses offered by ARC (for example introductions to Linux and BlueBEAR) and further afield such as from [LinkedIn Learning](#).

Advice Session: Conceptualising Transitivity Networks through Pattern-based Constructions

A researcher was aiming to create a database of approximately 500 records relating to conceptualising transitivity networks through pattern-based constructions, including data about: patterns, constructions, verbs, examples, cross references, and indices. From this data, she aims to produce approximately 100 data visualisations. We discussed the best ways for the project to store its data, comparing the advantages and disadvantages of various formats, such as Microsoft Excel spreadsheets, a relational SQL database, and XML files. We also discussed possible tools for generating the data visualisations. The researcher also required a public facing website to openly share this data and the project's outputs. Following discussions around the options available to her, we have scheduled a free coding engagement in which we will deliver this public website and database.

Advice Session: Python Workflows on Virtual Machines

We talked with a researcher regarding a project using various technologies, including Python, Amazon Bucket and a BEAR Globus endpoint. We advised the researcher how he could utilise BEAR services to enhance his research. First, the researcher uses PyCharm (an IDE tailored towards Python development), however it's not possible to use this on the virtual machine (VM) he uses and so we provided advice about setting up environments so that he could work seamlessly across his local machine and remote VM. This advice included switching to Microsoft Visual Studio Code as his IDE, as this can be run as a server application on the [BEAR Portal](#) and replication of the workflow locally is straight-forward. Secondly, we discussed options around potentially coaching the project group on database schema design. Thirdly, we discussed a project that has been running for two years and uses Digital Ocean to store energy-use data in a PostgreSQL database. Data is no longer available from the energy suppliers and the researcher would like to ensure that the data remains available to researchers by storing it on university storage. We discussed further automation and improvements to his data collection mechanisms to enhance the repeatability of the data cleaning aspects. We discussed possible future coding engagements for the delivery of a database and associated website.

To request an advice session please fill in the form at <https://intranet.birmingham.ac.uk/bear/sd/bear-software-advice>.

Social Bodies

Researcher(s)

Professor Karen Harvey, College of Arts and Law

Technologies

Python, Django, SQL, HTML, CSS, JavaScript



Project Description

“Material Identities, Social Bodies: Embodiment in British Letters c.1680-1820” is a Leverhulme Trust funded project based in the School of History and Cultures. The **Social Bodies** project uses eighteenth-century letters to explore the relationships between the physical body, self and social identity, and experiences of “embodiment”. The project team are transcribing and analysing thousands of letters between correspondents of family and kin, friendship, faith and business.

In a previous coding engagement we designed a relational database and private web interface for the project team to submit and store their research data. This new coding engagement set out to present this data through a public web interface, allowing users to see images of letters along with transcriptions and other related data. Additionally, a new section of the website allows for the general public to attempt their own transcriptions of selected letters from within the database.

This coding engagement continues an ongoing collaboration with the project, following a coaching engagement in 2021 in which we taught the project team how to manage and analyse their research data directly using SQL.

The total sum of this collaboration means that the project team have been able to successfully record, manage, analyse, and share their research data. The contributions of the RSG have been important to the ongoing success of the project.

‘The ongoing work and support from Mike Allaway in the Research Software Group has been essential to the working, completion and success of the project. The digital methods facilitated by the database have transformed the sophistication of our research. The project team have been supported in acquiring the skills to undertake this research ourselves, using SQL. Finally, the public website and database is enabling an exciting public engagement aspect of the project.’ - **Professor Karen Harvey**



Midlands Open Bioimaging

Researchers

Professor Davide Calebiro, Dr Joao Correia and Dr Jeremy Pike - COMPARE Research Group, College of Medical and Dental Sciences

Technologies

LabView

Project Description

The RSG was asked by the Midlands Open Bioimaging team to contribute IT expertise to the state-of-the-art 4D Automated Remote Imaging System being built by the COMPARE Research group at the University of Birmingham's Medical School. This is an innovative system that enables researchers to take 3D movies of samples across multiple tracers, allowing researchers to follow processes occurring at cell boundaries in real-time. Drug discovery is one of the many areas that will benefit from this, enabling scientists to refine their choice of potential targets well in advance of more expensive tests.

The system represents a step change, not just in its capabilities, but in its whole way of working. This project was conceived as a remote instrument from the ground up allowing researchers from across the world to use this from wherever they are. In future, this microscope will work as a completely automated system where taking images will just involve a researcher placing a job in an online queue and being notified when their images have been produced.

Jenny Wong and James Allsopp from the RSG have been helping from the very beginning, providing input on the specification of the hardware and its validation, to working with the University to solve networking problems, acquiring the correct licences for software, and writing code to integrate this system.

The [Midlands Bioimaging project](#) is a collaboration between the Universities of Birmingham, Leicester, Nottingham, and Warwick. Funded by the Biotechnology & Biological Sciences Research Council (BBSRC), part of UKRI.

'The RSG team provided highly professional and organized support to our project, giving a major contribution to its successful delivery. We are looking forward to continue working with the team during the next, exciting phase of our project.' - **Professor Davide Calebiro**

Deploying AIDA on the European Space Agency's SWE Portal

Researcher(s)

Dr Sean Elvidge and Dr David Themens, Space Environment & Radio Engineering (SERENE) Group, College of Engineering & Physical Sciences

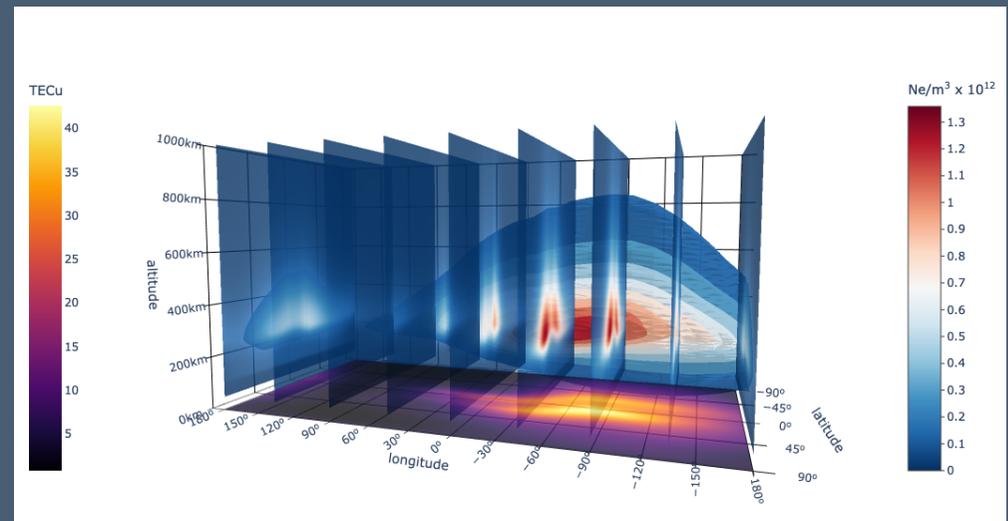
Technologies

Web, Django, Python, HTML, CSS, JavaScript, Celery, SLURM, Docker

Project Description

The ionosphere is a rapidly changing environment with complex statistical properties. In the last 15 years, a wide range of ionospheric data assimilation models have been developed. They employ a number of different background models and assimilation techniques. In all cases, they have been implemented with the aim of making a prediction of the current state of the ionosphere and then updating this prediction with data.

The SERENE group have developed AIDA, which is based on a state-of-the-art data assimilation technique known as a particle filter. This coding project is for the deployment of AIDA on the European Space Agency's (ESA) SWE portal.



3D surface plot of the ionosphere electron density (red/blue colormap) with a 2D contour plot of the vertically integrated electron density (purple/yellow colormap)

[Deploying AIDA on the European Space Agency's SWE Portal - Continued...]

A Django website (<https://spaceweather.bham.ac.uk>) that can submit Slurm jobs to run on the BlueBEAR supercomputer has been developed. The website uses Celery to schedule regular runs of the AIDA model on BlueBEAR. AIDA uses real-time Networked Transport of RTCM via Internet Protocol (NTRIP) data as input, which requires simultaneously downloading data from several thousand GPS stations, using the German Federal Agency for Cartography and Geodesy (BKG) NTRIP Client. These files are then converted to HDF5 format before being used by the AIDA model. All of these processes have been automated and are controlled using Celery. Interactive 3D visualisations of data that is output from the AIDA model are also currently being developed.

Within the SERENE group, improvements are currently being made to the AIDA model. Once these have been incorporated and tested, the model output will be made accessible via the ESA's SWE portal.

In addition to the deployment of AIDA, the RSG have:

- developed all other pages on the spaceweather.bham.ac.uk website;
- installed several applications on BlueBEAR, including TIE-GCM (<https://www.hao.ucar.edu/modeling/tgcm/tie.php>) and CESM (<https://www.cesm.ucar.edu>);
- developed a Django REST Framework API which allows users to make a request to run a model that returns a prediction of the total electron content (TEC) in the ionosphere for a given latitude and longitude.

'The support from RSG for this project has been critical for the success of AIDA. Their expertise has allowed for the development of a robust automated system on the BlueBEAR HPC that will be used as critical infrastructure for clients around the globe, supplying reliable, accurate and approachable ionospheric corrections for navigation, communications and remote sensing technologies. With their help, we have also established a professional client-facing website and web interface that is easy to maintain and will serve us well into the future. The RSG staff on our project have made phenomenal contributions, making a logistically onerous task tractable.' - **Dr David Themens**



Nano-Org Protein Databank

Researcher(s)

Professor Dylan Owen, College of Medical and Dental Sciences

Technologies

Web, Django, VMs, R, SLURM, GPUs/HPC.

Project Description

Dylan's research group in the Institute of Immunology and Immunotherapy have published novel algorithms for analysing point-cloud microscopy images. The aim of the project was to provide a databank where researchers could upload their own images and have them analysed, running the algorithms on BlueBEAR. One exciting use of this project is to allow researchers to search through the databank and have the system identify datasets that most closely resemble their own microscopy images. The website is hosted on a dedicated Virtual Machine that is connected to our Research Data Store and High Performance Computing clusters. This means that researchers from all over the world will be able to access and run Dylan's computationally expensive algorithms, even if their local compute access is poor.

'Working with the RSG has been critical to the success of this project. They have successfully implemented a database for microscope data and the website to access it, and have implemented analysis algorithms to exploit that database. The interactions have always been punctual and professional and the RSG have project managed the work well. I would be happy to continue working with the RSG in the future.'

Dylan Owen

BEAR Applications

In 2022 the Research Software Group handled over 2000 Service Now tickets, covering a broad range of areas including BEAR Project admin requests, application installations (see separate section for further information), licensing queries, HPC-specific queries and other general support tickets.

There is further information on the support we provide to users on Baskerville and Sulis in this report. There is also information on some of the development done for the BEAR and Baskerville Portals.

Installations

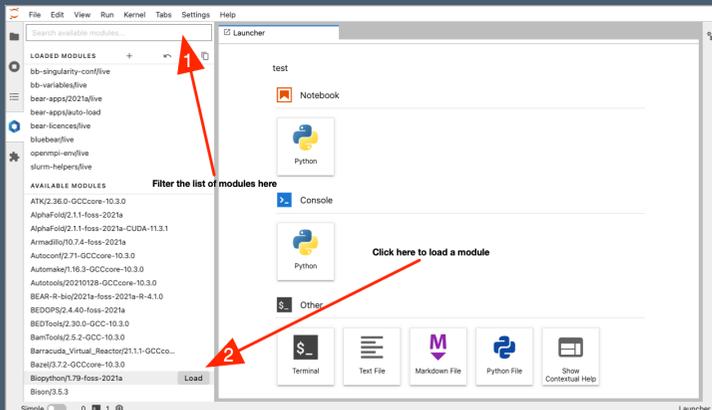
We built **641** applications (counted by name only), or **842** counting installed modules, therefore including versions and multiple toolchains. In total we now have **1391** applications (counted by name) across the BEAR systems.

Application installation requests (recorded via Service Now) during 2022: **208**



BEAR Portal and Baskerville Portal Developments

There have been a number of significant updates to our [BEAR Portal](#) and [Baskerville Portal Interactive Applications](#) over the past year. In general we have shifted towards using containers, offering us greater control over the individual environments that are providing these apps. Another key change is the introduction of the *BlueBEAR GUI* app.



JupyterLab Interactive Application

Using a container to provide the Jupyter server backend means that we are no longer constrained by compatibility issues inherent to using a module. It also allows us to include the *Jupyter Lmod* server extension so that users can add modules to their kernel, e.g. extending Python with *numpy* or *PyTorch*.

As part of this change on BEAR Portal we have also redesigned the submission form which now changes dynamically depending on the selected QoS, for example only revealing GPU options when they're appropriate and also modifying the list of available kernels so that they match the hardware options selected.

On Baskerville we now support Conda environments and there is a checkbox on the submission form which causes the app to search for available Conda envs and then adds them to the list of available kernels in the Jupyter Launcher.

BlueBEAR GUI

The *BlueBEAR GUI* application offers a 'desktop-like' experience, offering greater control of the environment to users who need to work graphically. Based on a lightweight *xfce* interface, the *BlueBEAR GUI* app bundles the bare minimum required, primarily a terminal application from which users can execute their required commands.

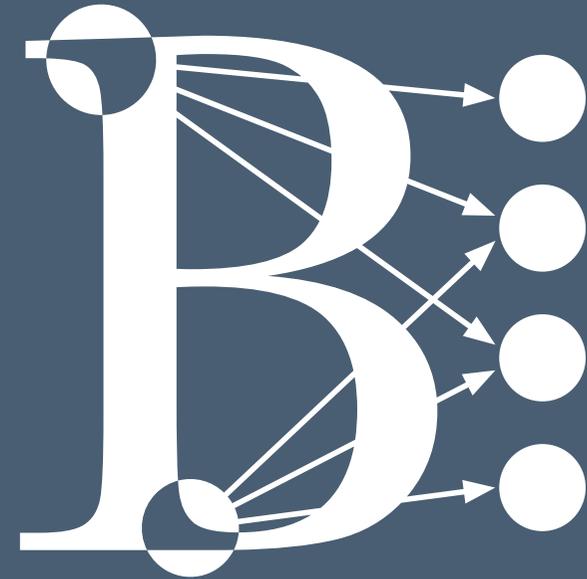
RSE Support for Baskerville

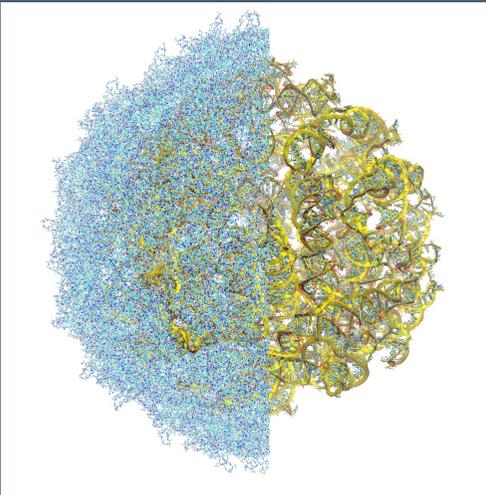
Since Jenny Wong and James Allsopp joined the RSG at the end of last year, half of their time has been spent engaging with Baskerville users, both at the University of Birmingham and other partner institutions.

Part of this has been holding regular meetings with remote partners to find out their concerns and to try and address those. An example of this is the implementation of a Globus system in response to concerns over the length of time it would take to upload data to Baskerville. Other examples are improving our documentation and augmenting that with video tutorials showing how to [log on for the first time](#).

We've also been involved in responding to users' requests to use Baskerville in ways not originally foreseen with the use of the Portal and through the use of [Apptainer](#). Apptainer allows users to bring their containerised applications to Baskerville from either their desktop PC or other HPC provider with little or no reconfiguration. These containers include any customisations the researchers may have made themselves, reducing the workload on Baskerville system administrators, as well as making some otherwise-unfeasible projects possible.

In addition to this, we've been reaching out to users within the University of Birmingham to find out what their requirements are and how they're using Baskerville. For example, we spoke to one user whose [published work](#) used Deep Learning on Baskerville to approximate the computationally-expensive model evaluation step in a Markov-Chain Monte Carlo simulation.





RSE Support for SULIS

Funded by the UKRI [World Class Laboratories fund](#), and the HPC Midlands+ Consortium, [Sulis](#) was launched in 2021 and replaced Athena, the previous HPC Midlands+ Tier 2 cluster. Throughout 2022, the RSG has delivered support for Sulis' Birmingham and Aston users, with both universities funding 50% of an RSE post in the RSG. The dedicated RSE time is led by Dr Simon Hartley (for Aston), Dr Simon Branford (for Birmingham) with support from Chris Greasley.

RSE support has comprised various forms while engaging with researchers at Birmingham and Aston. [Carpentries](#) courses have been delivered to Aston researchers, with Introductions to Git/HPC/Linux covered by Simon Hartley. Various applications, managed through Easybuild, have been built to meet user requirements and a separate application install area for Birmingham's Sulis users has helped ensure timely responses to software requests. RSE support has also extended to helping researchers optimise their workloads, with bespoke support for users of MATLAB, COMSOL, GROMACS and Fortran packages.

Sulis is used by researchers across a broad range of disciplines from Birmingham and Aston and the full scope of Sulis' userbase was highlighted in a showcase event at Warwick. Birmingham and Aston research undertaken on Sulis includes simulations of binary black holes for model development of their gravitational-wave signal, computational structure prediction investigating novel cathode materials for better rechargeable battery performance, and reconstructions of entire models of viruses.

Researcher feedback of our Sulis support has been positive. Dr Dmitry Nerukh (Aston) highlighted the RSG support for managing Aston's Sulis use, particularly the workshops Simon Hartley organised for Aston HPC users, as well as his expertise in assisting with GROMACS work. Professor Andrew Morris (Birmingham) highlighted the importance of being able to access Sulis' additional Tier-2 resource to accelerate his group's research with prompt RSG support for their software allowing them to meet their deadlines.

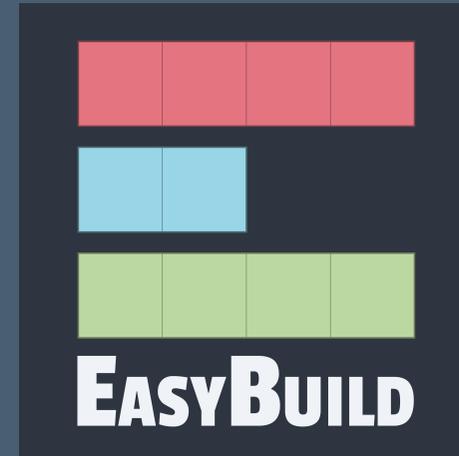
([DOI for image](#))

EasyBuild

The BEAR team continues to use EasyBuild to deploy nearly all of the centrally installed software on both BlueBEAR (<https://intranet.birmingham.ac.uk/bluebear>) and Baskerville (<https://www.baskerville.ac.uk/>), and also a small stack of custom applications provided to University researchers on Sulis (<https://sulis.ac.uk>). EasyBuild allows us to effectively and efficiently install various pieces of software across our computing platforms. This enables us to leverage the collaborative efforts of various worldwide HPC people and we have contributed to the community by providing hundreds of pull requests to the various EasyBuild repositories in 2022.

The Research Software Group's Deputy Leader Simon Branford is an EasyBuild maintainer. He regularly contributes to the project by facilitating with the processing of the community's large number of GitHub Pull Requests, alongside other strategic decision making, which is of huge benefit to the RSG as it ensures that we're fully abreast of any forthcoming updates, known issues etc. Simon also hosted the Seventh EasyBuild User Meeting (<https://easybuild.io/eum22/>) and the EasyBuild Tech Talks (<https://easybuild.io/tech-talks/>).

In the last quarter of 2022, James Carpenter and Simon Branford heavily contributed to the project to port the EasyBuild documentation (<https://docs.easybuild.io/>) from reStructuredText to Markdown and switching to using MkDocs (<https://www.mkdocs.org/>). The knowledge already gained by the BEAR team in using MkDocs for Baskerville Docs (<https://docs.baskerville.ac.uk/>) aided in configuring the EasyBuild MkDocs setup and repository.



Training

This year we have continued to deliver training to members of the University and, in the case of Baskerville users, beyond. A significant number of the Research Software Group are responsible for delivering this content and courses range from the *Introduction to BlueBEAR* half-day course (which introduces University staff and students to using the BlueBEAR HPC system) to Carpentries modules and GPU-focused NVIDIA training. In addition to delivering remote and in-person courses we have started developing asynchronous video-content, covering various aspects of our services and designed to complement our existing offerings.

Baskerville

Documentation on <https://docs.baskerville.ac.uk> was further developed this year to improve accessibility and usage of Baskerville. New items include:

- Updated contact details for support, including details of lifecycle management of Baskerville projects
- A video tutorial of first time access and logging onto Baskerville
- An explanation of how to use the Globus app for data transfer
- Updated guide on using Apptainer for containerisation
- Instructions for using conda environments for JupyterLab on Baskerville Portal
- A cheatsheet for quick access to commonly used commands.

James Allsopp and Jenny Wong qualified as NVIDIA DLI Instructors this year and delivered many Baskerville training events. For more information, please see Training page: NVIDIA and Conferences and Events: Baskerville.



Video content

Here are two examples of the video content that we have developed this year, both available to view on YouTube:

- [BEAR Portal: JupyterLab intro](#)
- [Baskerville Basics: First Time Access](#)



NVIDIA training courses

Funding from the Baskerville project allowed two of our RSEs, Jenny Wong and James Allsopp, to become NVIDIA University Ambassadors allowing them to teach certified **NVIDIA DLI** courses across the University and to our Baskerville partners. These NVIDIA DLI courses are hands-on full-day courses which show participants how to make the most of GPU systems. The two courses we currently teach are:

- Accelerating Data Engineering Pipelines
- Fundamentals of Deep Learning

We've instructed thirteen participants in the Accelerating Data Engineering Pipelines and a further twenty participants in the Fundamentals course, including both PhD students and academic staff. In addition, parts of the Accelerating Data Engineering course have been delivered as workshops at **RSECon22**, and in full to thirty undergraduates for **BEAR Challenge**. With the assistance of NVIDIA, we also delivered a multiple-GPU course for BEAR Challenge.

We plan to extend this offering by including courses using CUDA and multiple GPU Deep Learning in 2023. All of these courses are extended by ourselves to show how the material can be applied using Baskerville.

Carpentries Training

The Carpentries (<https://carpentries.org/>) lessons form an integral part of the training provision offered to University researchers. This training is provided by ARC, coordinated by REDG, and using a mix of instructors and helpers from within the RSG and from the wider research community at the University.

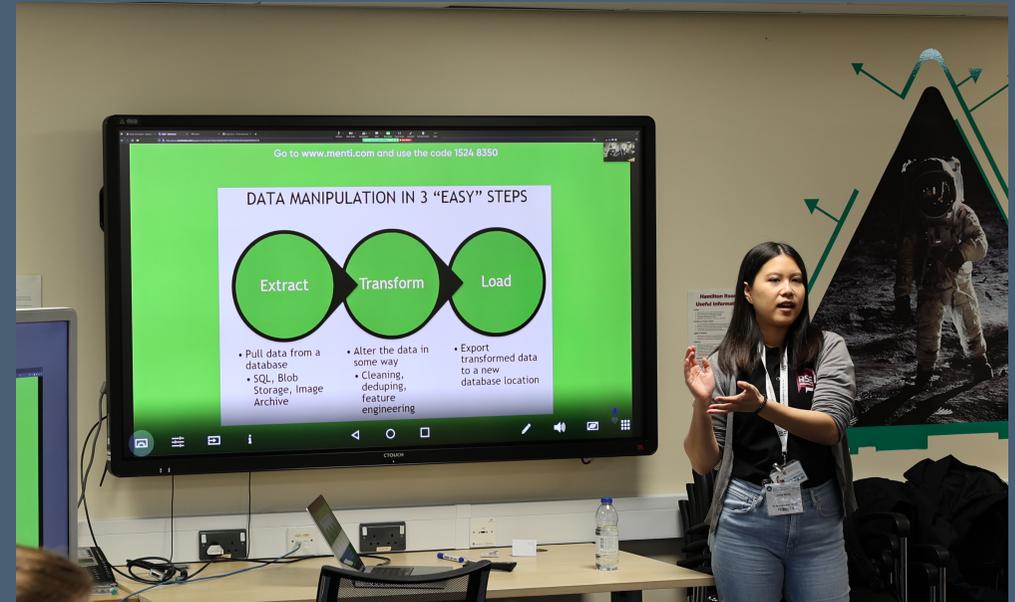
From the RSG, James Allsopp, James Carpenter, Simon Branford, Simon Hartley, Gavin Yearwood, and Jenny Wong are all qualified Carpentries instructors and between them have provided workshops covering Python, Bash, Git, R, and MATLAB. In addition, Afraz Ahmed, Mike Allaway, Adrian Dante Garcia, and Cerys Lewis have helped on Carpentries workshops.



The Rosalind Franklin Institute

As an concrete example, we've taken these courses and built them into a full training package for our partners. In November 2022, Jenny Wong and James Allsopp travelled to the the Rosalind Franklin Institute, situated at the Harwell Science and Innovation Campus, and delivered Baskerville training on-site to over twenty researchers.

We augmented the two NVIDIA courses we teach with an *Introduction to HPC on Baskerville*, where we talked about the architecture of the system and how running jobs is different from just running a program on your desktop. We also gave a practical walkthrough of the [Baskerville Portal](#), which provides access to Baskerville's resources through a web browser. As researchers are increasingly drawn from non-traditional HPC backgrounds, lowering the barrier to entry to HPC through services such as the Portal allows us to benefit a wider cross-section of the research community than would otherwise be possible.



Jenny Wong presenting at the Rosalind Franklin Institute

'The Rosalind Franklin Institute were glad to have the Baskerville RSE team travel down to deliver two days of NVIDIA Deep Learning and Data Science courses. Participants spanning a wide range of skillsets and programming experience were able to benefit from different aspects of the training and the opportunity to speak directly to the Baskerville RSE team and ask questions was invaluable. We look forward to hosting the Baskerville team again as they continue to develop new training courses.' - **Joss Whittle**

Conferences and Events

Members of the Research Software Group have attended, participated in and organised various events and conferences this year.

Ed (Andrew Edmondson) spoke at UCL's 'Kickstart your HPC Journey' event, highlighting the use of software and HPC and how researchers can combine these effectively.

Ed attended the N8 DRI retreat (<https://n8cir.org.uk/dri-retreat>), and sat on panels about *Usability*, *Outreach*, *Novel Technologies* (as chair), *Managing a Portfolio*, *Strategic Thinking and Vision*, and *Developing a Profession*.

Simon Branford presented at the HPC-SIG technical session on the use of ReFrame for testing the software we install on both BlueBEAR and Baskerville.

Ed hosted the RSE Leaders meeting and presented an update on the Research Software Group and the work we do.



Baskerville

The RSG has represented Baskerville, our national EPSRC-funded Tier 2 HPC facility, at many conferences and events this year. This has been key to growing our continued relationships with Baskerville's research partners and user base. Here we report on a selection of conferences and events attended this year.



BEAR Challenge

After a 2 year hiatus, the **BEAR Challenge** made its return to the University of Birmingham. The popular three-day event was held on June 21st-23rd and attended by seven teams of 3-5 undergraduate and taught postgraduate students. The teams tackled an array of machine learning, deep learning, data engineering, parallelisation, and computer cluster design challenges using the GPU resources available on Baskerville. The event was organised by Advanced Research Computing (with significant contributions from all three groups) and industry partners NVIDIA, Lenovo and OCF.

RSECon22 Workshop

RSECon22 took place September 6-8th at Newcastle University, with all 250 tickets sold and a record number of submissions this year. The RSG's submission was one of the six accepted workshops, out of the total 17 applications received. Jenny Wong led a workshop on Accelerating Data Engineering Pipelines, which introduced the basics of using Python packages, such as cuDF and Dask, to scale up and scale out data engineering tasks with multi-GPUs on Baskerville using Jupyter notebooks in a web browser. Collaborative and competitive teamwork was the order of the day, as participants worked in pairs to visualise a dataset with the best GPU performance possible.

Other

We hosted an NVIDIA GPU Hackathon in February where Gavin Yearwood gave a presentation on Baskerville. In May, we hosted a Data Study Group in collaboration with The Alan Turing Institute, in which James Carpenter gave a talk on how to use Baskerville for the challenges. A Baskerville-centric Coding Club in June was delivered by Gavin Yearwood and Simon Hartley to researchers at the University of Birmingham. Presentations about Baskerville were given at the College or School level throughout the year, including Computer Science, Chemistry, and Business School.

For additional information on Baskerville training delivered this year, please see the training pages on Baskerville and NVIDIA.

RSE Midlands

Birmingham hosted the inaugural **RSE Midlands** conference on 8 June 2022, an event to build a Midlands regional RSE community. RSE Midlands is a Regional Group affiliated with the Society of RSE (<https://society-rse.org/community/regional-groups/>). The committee was comprised of individuals from both ARC and RSEs from the University of Warwick and was chaired by Gavin Yearwood. The event had financial support from: NVIDIA, Lenovo, OCF and the **Society of Research Software Engineering**. Together this financial contribution paid for the hotel, food, merchandise and transportation costs for speakers and attendees.

The first half of the event had talks by Claire Wyatt on Regional RSE Groups and by Rachael Ainsworth about The Software Sustainability Institute Community and Events. The first half also had interactive workshops led by Gavin Yearwood to help the attendees define themselves as RSEs and what RSEs do. Fellow ARC committee members Ed (Andrew Edmondson), Adrian Dante Garcia and Jenny Wong shared their experience as RSEs during lightning talks. The second half of the event had talks from NVIDIA (Paul Graham), Intel (Andy Mallinson) and Lenovo (Jim Roche) and it covered many tools available for RSEs to use.

The legacy of this event is:

- plans for another RSE Midlands event in Nottingham with members of ARC (Gavin Yearwood, James Allsopp and Simon Hartley) also joining the 2023 committee.
- a new RSE Midlands coding club (chaired by Gavin Yearwood) will bring talks of interest to RSEs and those in RSE adjacent fields on a regular basis in a hybrid setting.



Contact Us

BEAR Software

For information about BEAR Software, see <https://www.birmingham.ac.uk/bear-software>

For any help with anything related to research software at the University of Birmingham please email bear-software@contacts.bham.ac.uk

BEAR / Advanced Research Computing

For information about BEAR, see <https://www.birmingham.ac.uk/bear>

For general help or information about any BEAR service (BlueBEAR HPC, Research Data Store, ...) contact the BEAR team by email at bearinfo@contacts.bham.ac.uk

Follow the Advanced Research Computing's twitter feed @uob_bear at https://twitter.com/uob_bear

Join our Research Computing Community slack workspace at <https://bham-rse.slack.com/signup>

Requests, Faults, Complaints

The IT Service Desk is your route to find answers. Find all Advanced Research Computing items here: <https://intranet.birmingham.ac.uk/bear/sd/bear>

For all other IT Services items/help or to log faults and any complaints visit: <https://www.itservicedesk.bham.ac.uk/>