The scientific and economic impacts of ARCHER2

Final evaluation report

November 2025

LE
London Economics

With support from Technopolis and Frazer-Nash Consultancy

About London Economics

London Economics is one of Europe's leading specialist economics and policy consultancies headquartered in London. We advise clients in both the public and private sectors on economic and financial analysis, policy development and evaluation, business strategy, and regulatory and competition policy.

Our consultants are highly-qualified economists with experience in applying a wide variety of analytical techniques to assist our work, including cost-benefit analysis, multi-criteria analysis, policy simulation, scenario building, statistical analysis and mathematical modelling. We are also experienced in using a wide range of data collection techniques including literature reviews, survey questionnaires, interviews and focus groups.

Head Office: Somerset House, New Wing, Strand, London, WC2R 1LA, United Kingdom.

w: londoneconomics.co.uk e: info@londoneconomics.co.uk 𝕏: @LondonEconomics
t: +44 (0)20 3701 7700

Acknowledgements

We would like to acknowledge the useful guidance and feedback provided by the Steering Committee throughout this research. We would like to thank EPCC and UKRI for the support and the data provided to enable this research. We would also like to thank all the stakeholders that the project team had met with over the course of the project for the valuable advice and information they have provided. We would also like to thank all of the stakeholders surveyed for their informative questionnaire responses. The research also benefited from contributions by Frazer-Nash Consultancy and Technopolis, whom we would also like to thank for their inputs and support. GPT-4o, an AI model developed by OpenAI, was used to support desk research, assist in brainstorming ideas, and aid in drafting and refining the text. All content was reviewed and edited by the authors to ensure accuracy and appropriateness. Responsibility for the contents of this report remains with London Economics.

Authors

Daniel Herr, Associate Director

Dr. Charlotte Duke, Partner

Moritz Profanter, Economic Consultant

Phoebe Worsley, Research Assistant

Grace Fradgley, Research Assistant

Wiktor Przyborowski, Intern

Table of Contents

Page

Glossary

Term	Description
AIRR	AI Research Resource
BBSRC	Biotechnology and Biological Sciences Research Council
CU	ARCHER2 Compute Unit (one node hour on ARCHER2)
DSIT	The Department for Science, Innovation and Technology
EPCC	Formerly known as the Edinburgh Parallel Computing Centre, EPCC is a supercomputing centre part of the University of Edinburgh and host of ARCHER2.
ESRC	Economic and Social Research Council
EPSRC	Engineering and Physical Sciences Research Council
HPC	High-Performance Computing
IPCC	Intergovernmental Panel on Climate Change
MRC	Medical Research Council
NERC	Natural Environment Research Council
STFC	Science and Technology Facilities Council
TRL	Technology Readiness Level
UKRI	UK Research and Innovation

OVERVIEW OF THE REPORT

This report is structured as follows:

- **Executive summary** provides a summary of the study and key findings.
- **PART I** introduces the study as well as the ARCHER2 system and explains its place within the UK HPC ecosystem; briefly sets out the study approach including caveats and limitations; and provides an overview of how ARCHER2 is being used.
- **PART II** explores the scientific and economic impacts flowing from ARCHER2. Four sections provide details on the importance of ARCHER2 in enabling high impact, high quality computational science; ARCHER2's role in supporting a highly skilled community; the economic impacts of ARCHER2; and wider benefits associated with delivery of a high-quality service as well as improved local societal impacts.
- **PART III** presents the results of the socioeconomic evaluation and provides final reflections. This includes a discussion on the impact of a loss of Tier-1 HPC capability on UK science, and the overall conclusions of the study.
- A number of **ANNEXES** provide additional detail.

Image source: Dmitry Aleynik and Sam Jones, Scottish Association for Marine Science, Physics Department: West Scotland Coastal Ocean Modelling System (WeStCOMS) - 3D currents animation

Executive summary

About this study and ARCHER2

In October 2024, the Engineering and Physical Sciences Research Council (EPSRC), part of UK Research and Innovation (UKRI), commissioned a team led by London Economics to undertake an independent study to demonstrate the impact and success of the ARCHER2 supercomputer. The London Economics study team was supported by leading experts from Technopolis and Frazer Nash Consultancy.

ARCHER2 is the UK's most powerful supercomputer, replacing its predecessor ARCHER. It was funded by EPSRC and NERC and serves as the current Tier-1 (national) High-Performance Computing (HPC) service for academic users within EPSRC and NERC research remits. A share of ARCHER2 capacity can also be accessed by other users.

Study objectives and counterfactuals

The study focuses on the socioeconomic and research benefits enabled by ARCHER2. It sought to quantify the likely economic impact from the use of ARCHER2 over its five-year intended lifetime since its launch in 2021 and provide a return on public sector investments (ROI) figure. Further, it sought to document the wider and spill over benefits arising from the use of ARCHER2 and provide robust case studies/impact narratives from the research supported.

The evaluation assessed the benefits and costs of ARCHER2 against two counterfactual scenarios:

- A **do-nothing scenario** under which the ARCHER2 capabilities would not be provided. Under this scenario there would be no EPSRC/NERC funded national supercomputing capability. This enables us to assess the total benefits associated with ARCHER2.
- A **business-as-usual scenario** under which the ARCHER2 investment would have not taken place. Instead, the previous ARCHER capability would have been retained[1]. This enables us to value the additional benefits investment in ARCHER2 has enabled.

Socioeconomic valuation of benefits flowing from ARCHER2

This independent evaluation study highlights that the scientific and **economic benefits enabled by ARCHER2** are substantial. Total benefits to the UK **are estimated at £4.3 billion**. This is equivalent to **a return on public sector investments** in ARCHER2 (and public funding for R&D (R&D) undertaken on ARCHER2) **of £8.3 per £1 of public funding invested**. These estimates are on the conservative side. The 90% confidence range resulting from the sensitivity analysis around the central estimates shows that benefits are plausibly in the region of between £4.6 and £24.1 per £1 invested.

Given the nature of ARCHER2 as a science capability that serves predominantly fundamental academic R&D, the overall benefits estimates are driven by the **spillover impacts resulting from academic R&D**. These are **estimated at £3.7 billion** using the central assumptions. However, valuing

[1] Note, in practice, it may not have been possible or viable to keep ARCHER1 operational. Due to the age of ARCHER1 as an end-of-life machine, the required maintenance may have been significantly higher and/or additional capital investments may have been needed to retain ARCHER1 operational. In addition, there may have also been other issues such as with cybersecurity, software support, or availability of replacement hardware.

benefits from academic R&D is inherently difficult. Therefore, significant uncertainty surrounds these central estimates.

Nevertheless, excluding research spillover impacts and the intrinsic value of publications and citations (and associated funding costs) suggests sizeable economic benefits from ARCHER2's other activities.[2] These are estimated to be **in the region of £517 million**. Compared to the (projected) total funding for ARCHER2 itself of around £99.6 million, this suggests **a return of £5.2 for every £1 of public money** invested into ARCHER2.

Table 1 Central valuations of economic impacts (£$_{2025}$ prices)

Metric	ARCHER2 (vs do-nothing)	Continuation of ARCHER1 (vs do-nothing)	ARCHER2 vs Continuation of ARCHER1 (business-as-usual)
Total monetised economic benefits	**£4,265.8M**	**£1,463.8M**	**£2,801.9M**
Of which spillover impacts from scientific R&D	£3,719.5M	£1,243.2M	£2,476.2M
Benefits excluding spillover impacts of research and the intrinsic value of publications and citations*	£517.1M	£203.5M	£313.6M
Total costs accruing to the public purse	**£511.5M**	**£217.2M**	**£294.3M**
Capital expenditure and operating costs for the ARCHER(2) service**	£99.6M	£56.5M	£43.1M
Grant funding accruing to the public purse	£411.9M	£160.8M	£251.2M
Total return on public sector investments	**£8.3 : £1**	**£6.7 : £1**	**£9.5 : additional £1**
Return on public sector investments excluding spillover impacts of R&D and the intrinsic value of publications and citations	£5.2 : £1	£3.6 : £1	£7.3 : additional £1

Note: Column 1 presents the results of the economic evaluation of the benefits and costs of ARCHER2 vs. the do-nothing scenario under which no EPSRC and NERC funded national Tier-1 capability exists. Column 2 presents the results of the economic evaluation of the scenario where ARCHER1 continues to operate vs. the do-nothing scenario under which no EPSRC and NERC funded national Tier-1 capability exists. Column 3 presents the additional benefits and costs of ARCHER2 vs. the scenario where ARCHER1 continues to operate (the business-as-usual-scenario). (*) Research benefits included in this figure capture only the direct benefits to researchers in terms of costs avoided compared to accessing cloud (for research that could have been undertaken on the cloud) and the value of additional citations compared to similar publications (for research that could not have been undertaken without ARCHER2). (**) Costs for continuation of ARCHER1 represent ongoing costs only, no additional capital investment for acquisition of hardware assumed.

Source: London Economics

Comparison to previous evaluation of EPSRC's investments in HPC

London Economics' previous 2019 evaluation of EPSRC's investments into HPC[3] showed benefits in the range of £3.9bn to £11.6bn in 2025 prices [original estimates of £3bn to £9bn in 2018 prices]. The return on investment estimates at the time were estimated at between £6.5 : £1 to £19.5 : £1.

[2] Research benefits included in this figure capture only the direct benefits to researchers in terms of costs avoided compared to accessing cloud (for research that could have been undertaken on the cloud) and the value of additional citations compared to similar publications (for research that could not have been undertaken without ARCHER2).
[3] London Economics (2019). The impact of EPSRC's investments in High Performance Computing infrastructure

Estimates from this study are not directly comparable to the previous study as the previous study considered all EPSRC's investments into HPC including ARCHER1, ARCHER1's predecessor, HECToR, and the Tier-2 HPC centres as well as due to methodological differences (see Section 7.6).

Nevertheless, the results of the present study fall within the range reported in 2019. The central estimate for ARCHER2 stands at £4.3 billion implying a return on public sector investments of £8.3 per £1 of public sector monies invested in ARCHER2. The 90% confidence range around the central estimate is estimated at £1.9bn to £15.0bn with return on public sector investments estimates ranging from £4.6 : £1 to £24.1 : £1.

While the central estimate sits towards the lower end of the previous range, the higher end of the 90% confidence range suggests ARCHER2 could potentially generate benefits comparable to, or even exceeding, those of the two earlier national systems and Tier-2 centres combined.

Benefits of a continuation of ARCHER1

To understand the additional benefits ARCHER2 enables, the research also considered a scenario in which ARCHER1 would continue to operate. Estimated benefits under this scenario are around two-thirds (66% or £2.8 billion) lower than estimated benefits of ARCHER2.

However, this scenario would also imply a reduction in costs. Cost reductions are estimated at around £294 million (58%) compared to investment in ARCHER2. Together, this implies a return on public sector investments, were ARCHER1 to continue operation, of £6.7 per £1 invested (a reduction of around 19% compared to ARCHER2).

Considering only non-research-spillover benefits (and associated costs) suggests a much more sizeable reduction in the return on public sector investments of 31% (from £5.2 per £1 invested for ARCHER2 to £3.6 per £1 invested under the continuation of ARCHER1 scenario).

Additional benefits of ARCHER2 over and above benefits of continuation of ARCHER1

Comparing the total estimated benefits and costs of ARCHER2 to the total estimated benefits that could have been achieved had ARCHER1 continued to operate provides an estimate of the additional benefits of ARCHER2 relative to the business-as-usual scenario. As highlighted in the previous section, this comparison indicates additional benefits of ARCHER2 of £2.8 billion and additional costs of ARCHER2 of £294 million over and above what may have been achieved had ARCHER1 continued to operate. Together, **this implies additional benefits of £9.5 per additional £1 invested in ARCHER2 over and above the costs had ARCHER1 continued to operate.**

Excluding the spillover benefits of science, and associated costs, implies additional benefits of ARCHER2 of £314 million. This compares to additional costs of £43 million, implying a return on investment of £7.3 per additional £1 invested over and above the costs incurred had ARCHER1 continued to operate. This suggests that the additional benefit produced by investing in a higher-capacity Tier-1 service far outweigh the additional costs incurred.

The importance of ARCHER2 in enabling high impact, high quality computational science

Direct benefits to scientific R&D

HPC underpins modern computational modelling and simulation, which are now fundamental tools - alongside theory, experimentation and observation - used extensively across a wide range of scientific disciplines. HPC has long been used by the scientific community to enhance research accuracy, granularity and complexity.

As a Tier-1 service, ARCHER2 enables researchers to undertake the largest and most complex computational simulations. This allows researchers to explore scientific questions that would not be feasible through traditional scientific methods alone, and to develop and improve simulations that more accurately reflect the real-world. Ultimately, this gives researchers the tools needed to undertake high quality, high impact computational science and remain competitive internationally.

ARCHER2 also improves the efficiency of R&D by enabling users to undertake complex simulations in significantly shorter timeframes than would be possible with less powerful compute capabilities. This enables greater scientific throughput, thereby enabling more users to benefit from HPC resources and ultimately enabling UK researchers to undertake more computational science.

In total, **more than 2,100 ARCHER2-related publications** that could be matched to bibliometric databases were identified to date. These were spread **across 20 different fields of research** and have been cited 27,486 times in total within 23,756 other papers, an average of **12.8 citations per publication.**

Supporting collaborations and software development efforts

In addition to directly supporting high-quality R&D, ARCHER2 also plays an important role in positioning the UK as a credible and competent partner in the international HPC landscape. This in turn contributes to academic knowledge exchange and domestic and international collaborations. In total, **more than 157 collaborations involving ARCHER2** were recorded within ResearchFish™, UKRI's research outcome database. However, this is likely an underestimate. Analysis of author-affiliations of ARCHER2-related publications (Section 3.2.1) found that contributing authors are affiliated to **1,177 different institutions**, spread **across 88 different countries**.

Further, research outputs not only include publications, but also **development of software and other technical products.** This is important as modern computational science depends not just on powerful compute hardware, but also on the software that unlocks its potential. ARCHER2 supports software development through two routes:

- First, to support the research community to develop and optimise software in a sustainable manner, EPCC run regular Embedded Computational Science and Engineering (eCSE) support calls, providing funding to researchers to develop software on ARCHER2. A total of **93 eCSE projects** have been supported on ARCHER2 to date, funding a total of **1,114 project months** across nearly **35 institutions**.

Second, users engage in independent software development efforts, which benefit from access to ARCHER2. UKRI's research outcome database, ResearchFish™, contained information on development efforts by the research community benefitting from ARCHER2 relating to **a total of 151 software and other technical products**.

ARCHER2's role in supporting a highly skilled community

Without people who possess the right skills and knowledge to exploit the HPC hardware, much of the benefits of HPC would not be realised. Specialised HPC training, including advanced computational techniques and programming languages applicable across different systems is important to addressing the UK's existing computing skills gap[4].

ARCHER2 plays a crucial role in the development of advanced computational skills within the research community. This is achieved by providing a variety of training programs designed to equip researchers with the technical skills they need, as well as through hands-on use of ARCHER2 by students and postdoctoral researchers.

In total **145 training events** were provided by EPCC, who host ARCHER2, between 2020, when the early ARCHER2 four-cabinet system was installed, and 2025. Across these training events, a total of **270 training days** were provided with **more than 2,800 attendees** (~2,000 unique individuals) recorded across all training events. In terms of hands-on training, **more than 2,650 students and postdoctoral researchers have used ARCHER2** to date, gaining first-hand experience in advanced computational R&D.

In addition, to benefits to the individuals themselves, in terms of better career outlooks and higher salaries, HPC skills are in high demand across industry, contributing valuable expertise to businesses, supporting economic growth, and generating returns for the UK exchequer.

Further, ARCHER2 plays a key role in securing the UK's international standing in high performance computing, in scientific software development, as well as the UK's international science competitiveness. While these benefits are difficult to monetise, and so are not captured in the benefits estimates produced for this study, they are no less important.

Economic impacts

ARCHER2 is primarily a capability targeted at supporting academic R&D. Nevertheless, there are a range of economic benefits flowing from ARCHER2.

Spillover impacts from academic R&D

Publicly funded R&D generate significant benefits to the private sector, which in turn lead to improvements in social welfare through increased private sector innovation, productivity gains, job creation, and contributions to GDP growth. These benefits are achieved both through direct channels such as knowledge transfer and direct commercialisation of academic R&D as well as through knowledge, market or network spillover effects. Public research investment also stimulates private sector R&D, thereby increasing total R&D spend and potentially further boosting the benefits of economic growth.

Given the nature of ARCHER2 as a capability for scientific R&D, it is not surprising that a substantial share of benefits are ascribed to the economic benefits flowing from academic R&D undertaken on ARCHER2. Spillover impacts to the UK economy from academic R&D undertaken on ARCHER2 account for 87% of estimated benefits from ARCHER2 in this analysis.

Other economic impacts

In addition to spillover impacts, ARCHER2 delivers economic benefits through a number of channels. This includes direct benefits to firms accessing ARCHER2, either directly or via scientific collaborations; firms benefitting from staff trained in computational R&D that move into industry,

[4] The Lloyds Bank Consumer Digital Index 2024, commissioned by the Department for Education, found that 18% of UK adults lacked the essential digital skills that are needed for the workplace. See: https://www.lloydsbank.com/consumer-digital-index.html?srnum=4

as well as more direct outcomes of R&D in the form of spin-outs and the development of new products and services:

- There were 151 active user accounts linked to industrial and commercial organisations between 2021 and 2025. Excluding non-commercial users such as research institutes, publicly funded or government-owned centres and bodies, **123 commercial organisations that access ARCHER2** were identified, either directly or indirectly.
- Of the around 70% of students trained on ARCHER2 that remain in the UK, around four in ten (44%) pursue careers in industry. This means **ARCHER2 has already equipped more than 800 future skilled professionals with in-demand computational skills sought by UK businesses.**
- UKRI's research outcomes databases show that there were **at least five spin-outs** that have benefitted from ARCHER2 and **at least six trademarks, patents and other intellectual property** resulting from ARCHER2 R&D. However, user survey findings suggest that these figures are substantial underestimates.

Wider benefits flowing from ARCHER2

In addition to aforementioned benefits, there are a wide range of wider benefits flowing from ARCHER2. This includes environmental benefits in terms of a reduced carbon footprint associated with HPC, increased Value-for-Money from continuous improvements, the importance of maintaining a strong HPC service skill-set, and the benefits of high-quality user support. Further, there are improved local and societal impacts arising in the local Lothian economy from hosting the ARCHER2 service in Edinburgh as well as through public outreach activities.

The need for continued investment in HPC

ARCHER2 enables substantial scientific and economic benefits to the UK. This study also highlights the need for continued investment in UK HPC infrastructure, echoing findings from the recent Future of Compute Review[5].

The UK has historically been a global leader in a number of computing domains. This includes areas such as software development, computational modelling, data analytics, cybersecurity, AI and machine learning. However, the UK's high-performance computing infrastructure has been falling behind that of other major global economies.[6]

Continued investments into HPC infrastructure are needed to ensure UK science competitiveness as well as the UK's position as a trusted partner in HPC in the future. This includes not only AI-focused GPU-based systems, such as the recent AI Research Resource (AIRR), but also continued provision of CPU-based HPC systems for large-scale complex scientific simulations.

In this regard, it is positive to note that the UK has recently announced a £750 million investment into the UK's next national supercomputer, also located at EPCC at the University of Edinburgh.[7]

[5] DSIT (2023). Independent Review of the Future of Compute: Final report and recommendations. Available at: https://www.gov.uk/government/publications/future-of-compute-review/the-future-of-compute-report-of-the-review-of-independent-panel-of-experts [accessed 21/05/2025]

[6] GO-Science (2021). Large-scale computing: the case for greater UK coordination. Available at: https://assets.publishing.service.gov.uk/media/654a4025e2e16a000d42aaef/UK_Computing_report_-_Final_20.09.21.pdf [accessed 27/05/2025]

[7] See https://www.ed.ac.uk/news/university-set-to-host-ps750m-national-supercomputer

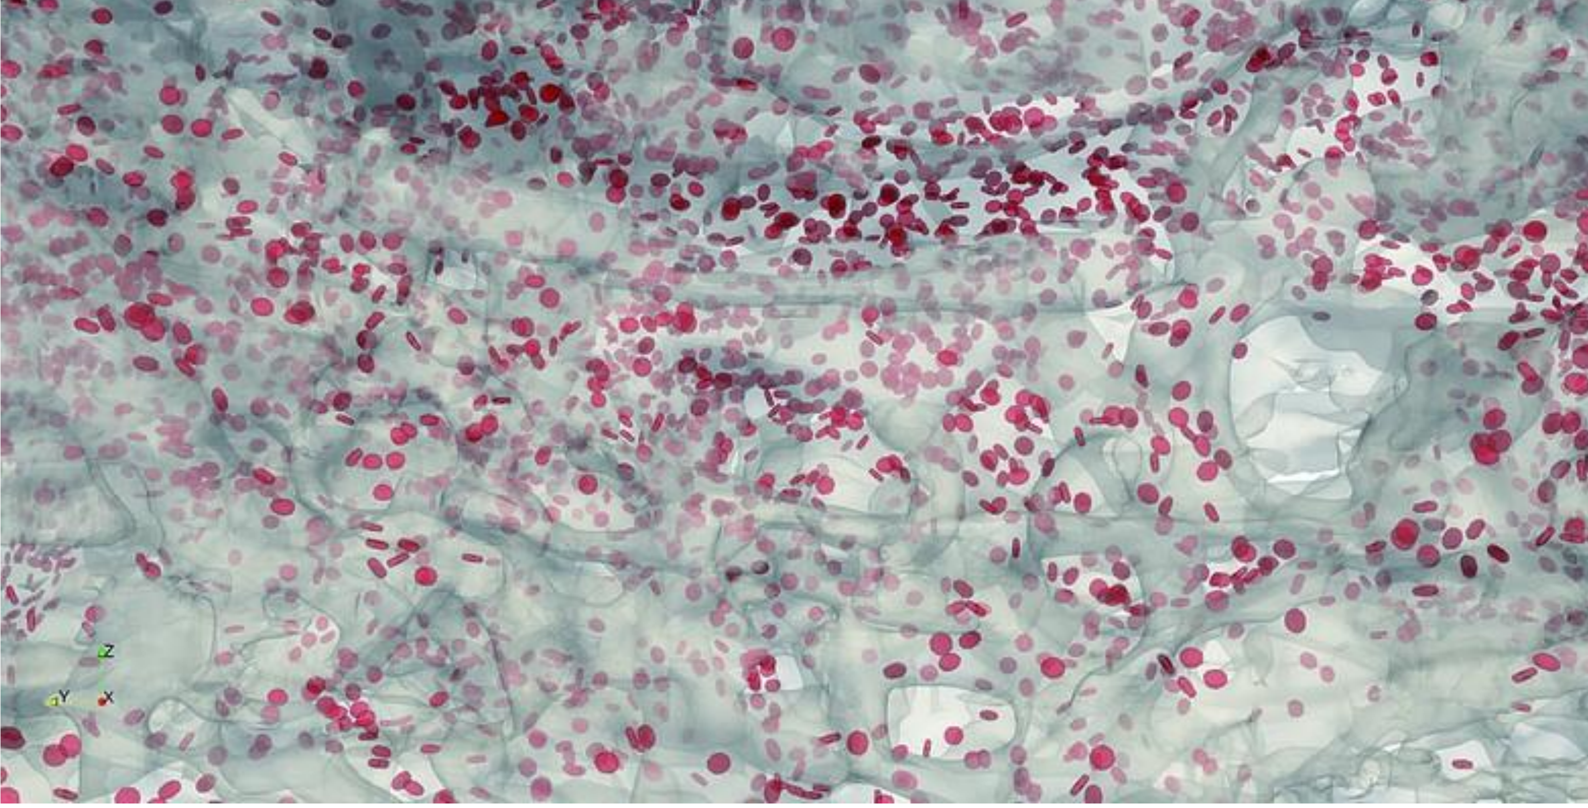

PART I: INTRODUCTION

This part provides:

- **Section 1** introduces the study and its aims and objectives, provides a brief introduction to ARCHER2 and its place in the UK HPC ecosystem, and briefly outlines the overall approach taken and caveats and limitations.
- **Section 2** provides an overview of how ARCHER2 is being used, and by whom, including the range of research disciplines accessing ARCHER and its geographical reach.

Image source: Qi Zhou, The University of Edinburgh, School of Engineering, Institute for Multiscale Thermofluids: Maternal blood flow through the intervillous space of human placenta

1 Background and context

1.1 About this study

In October 2024, the Engineering and Physical Sciences Research Council (EPSRC), part of UK Research and Innovation (UKRI), commissioned a team led by London Economics to undertake an independent study to demonstrate the impact and success of the ARCHER2 supercomputer. The London Economics study team was supported by leading experts from Technopolis and Frazer Nash Consultancy.

The study focuses on the socioeconomic and research benefits enabled by ARCHER2. It sought to quantify the likely economic impact from the use of ARCHER2 over its five-year intended lifetime since its launch in 2021 and provide a return on public sector investments (ROI) figure. Further, it sought to document the wider and spill over benefits arising from the use of ARCHER2 and provide robust case studies/impact narratives from the R&D supported.

To evidence the socioeconomic and research benefits flowing from ARCHER2, the study sought to understand the impact of ARCHER2 on various dimensions of scientific, academic, socio-economic and environmental activity in the UK. Specifically, the study:

- collated evidence on ARCHER2's impact on the UK's scientific productivity, focusing on high-quality publications, academic knowledge exchange, and software and data management;
- explores ARCHER2’s influence on the computational research community, including skill development, career paths, diversity, and the broader HPC ecosystem; and
- examines ARCHER2’s economic impacts, such as industry-academic collaborations, sectoral reach, and employment growth.

1.2 ARCHER2 and its place within the UK HPC ecosystem

At the time of writing, ARCHER2 is the UK’s most powerful supercomputer, replacing its predecessor ARCHER[8]. It was funded by EPSRC and NERC and serves as the current Tier-1 (national) High-Performance Computing (HPC) service for academic users within EPSRC and NERC research remits. A share of ARCHER2 capacity can also be accessed by other users.

The full ARCHER2 capability was launched in November 2021.[9] ARCHER2 was designed to provide a five-times increase in scientific throughput compared to ARCHER. At its launch ARCHER2 was ranked as the 22nd most powerful supercomputer in the world. In the most recent available data, November 2024, ARCHER2 is ranked 62nd.[10]

HPC underpins modern computational modelling and simulation, which are now fundamental tools - alongside theory, experimentation and observation - used extensively across a wide range of scientific disciplines. HPC has long been used by the scientific community to enhance research

[8] The full Isambard-AI system, part of the AI Research Resource, is currently being deployed. Once the full system becomes operational, it is expected to become the UK’s next most powerful supercomputer.
[9] Early access to an initial smaller system began in October 2020. The original ARCHER system was decommissioned on January 27, 2021, after which ARCHER2 became the primary national HPC service. However, the full 23-cabinet ARCHER2 system was only made available to users on November 22, 2021. To ensure continuity of service, a less powerful capability, the “ARCHER2 4-cabinet system” was installed in July 2020, with early access beginning in October 2020. For reporting purposes, ARCHER2 usage is considered to begin from January 27, 2021 unless otherwise stated.
[10] Top500.org (2024) Available at: https://www.top500.org/system/180036/

accuracy, granularity and complexity. It enables researchers to undertake complex computational simulations allowing them to explore scientific questions that would not be feasible through traditional scientific methods alone. Moreover, HPC can reduce the need for numerous costly experiments by enabling researchers and innovators to simulate a large number of scenarios. This enables researchers and innovators to maximise the impact of a more limited number of experiments that are conducted while lowering overall costs.

Fundamentally, HPC has two key impacts on research: it enables R&D that would not otherwise be possible, and it improves both efficiency and quality of existing R&D. It enables researchers to explore complex topics at a larger scale and with increased precision as well as enables them to address a wider range of scientific challenges. The research benefits enabled by ARCHER2 are discussed in greater detail in Section 3. In addition, there are a wide range of wider socioeconomic benefits flowing from ARCHER2; these are discussed in Sections 4 to 6.

As a Tier-1 service, ARCHER2 is designed to run the largest and most complex computational simulations. It sits alongside a number of other UK HPC capabilities. This includes DiRAC-3, another Tier-1 system serving users within the STFC particle physics, astrophysics, cosmology, solar system and planetary science and nuclear physics theory communities. It also includes a range of less powerful supercomputers intended for simulations with lower computational demands.

These latter systems broadly comprise local-scale (Tier-3) HPC facilities, providing HPC access to users within a specific institution or department, and mid-scale (Tier-2) facilities, providing targeted access at a regional level or for specific disciplines. Figure 1 below provides an overview of large-scale compute provision accessible by UK researchers and innovators.

Figure 1 Large-scale compute provision accessible by UK researchers and innovators

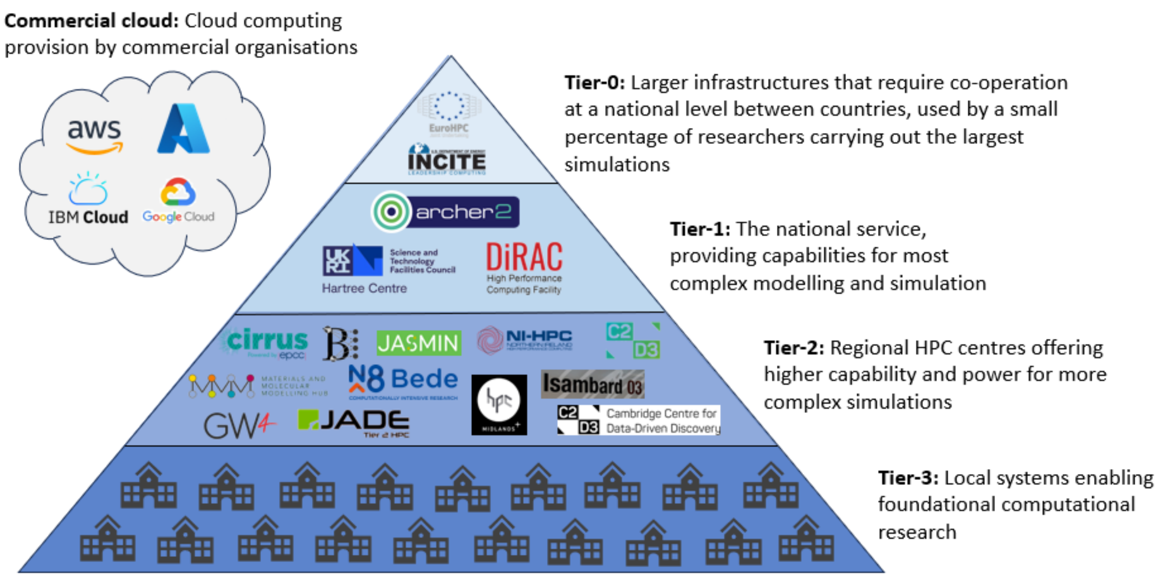

Note: DiRAC-3 was classified as Tier-1 in line with the Future of Compute Review (DSIT (2023). Independent Review of the Future of Compute: Final report and recommendations. Available at: https://www.gov.uk/government/publications/future-of-compute-review/the-future-of-compute-report-of-the-review-of-independent-panel-of-experts [accessed 21/05/2025])

Source: London Economics

In addition to national HPC capabilities, UK researchers can also access international HPC systems through partnerships. The primary route within Europe is through Euro HPC, which brings together supercomputing resources from across 35 countries. However, researchers can also access other international facilities such as the US Department for Energy's HPC capabilities through the INCITE programme. Finally, in the light of the rapid rise of AI, the UK (as well as other countries) is increasingly investing in GPU based systems, which will provide alternative supercomputer capacity in the future. Further discussion on these systems, as well as an exploration of the extent to which these systems provide viable alternatives to ARCHER2 under the (hypothetical) counterfactual scenario where ARCHER2 does not exist, is provided in Section 8.

The different tiers of HPC systems provide important and complementary functions across the full spectrum of innovation. HPC supports R&D from early-stage theory (TRL 1–3), through technology prototyping and simulation (TRL 4–6), to product development and performance optimisation (TRL 7–9). Small-scale theory development codes typically begin on Tier 3 or Tier 2 systems, where early models and algorithms are tested and refined. As these codes mature and are optimised for parallel performance, they scale up and are promoted to Tier 1 systems.

A similar process occurs on the hardware side experimental technologies, such as ARM CPUs, GPUs, low-latency interconnects, and parallel filesystems, are first introduced and tested on Tier 2/3 platforms. Once proven useful and robust within the community, they are considered for integration into Tier 1 system procurements. Similarly, researchers and innovators also grow their skills across the tiers as they progress, building the expertise needed to tackle increasingly complex computational challenges.

1.3 Approach, caveats and limitations

The research followed a multi-modal approach, beginning with a review of the existing Theory of Change and mapping of ARCHER2's activities and benefits. An initial scoping phase involved desk research, literature review, and consultations with key stakeholders to ensure a clear understanding of ARCHER2's outputs, outcomes, and impacts.

This informed the development of a Monitoring and Evaluation (M&E) framework and the further study design including the definition of counterfactuals, selection of impact indicators and evaluation methods, and further data collection and research activities to be undertaken.

Evidence gathering combined quantitative and qualitative methods, including a user survey, rolled out to all ARCHER2 users, further stakeholder interviews, and case studies, to inform the evaluation of ARCHER2's economic and non-economic benefits.

This research has been conducted by a team of independent professional economists. Estimates of economic impacts are based on best practice and best judgment to calculate the most robust and fair estimates. Annex 3 describes in detail the methodological approach including the specific evaluation approaches used for each monetised benefit stream and assumptions made, as well as the caveats and limitations of the analysis, where appropriate.

However, there are some limitations and caveats which should be considered whilst reading this report:

- Some of the analysis draws on stakeholder consultations, including an online survey of ARCHER2 users. As such, typical limitations associated with these methods, such as selection bias and lack of representativeness, should be kept in mind when using findings based on these sources in the analysis. Where survey data informs the analysis, this is clearly indicated in the relevant sections. A discussion of these methods and an assessment of the survey's representativeness are discussed in Section A2.1.

- Some elements of the analysis are based on data from UKRI's research outcome system, ResearchFish™. Since this system collects information solely from UKRI grant holders, any analysis using ResearchFish™ data is limited to UKRI-funded activities and does not capture benefits beyond these. Additionally, because ResearchFish™ depends on self-reporting by grant holders, the data may underestimate actual outcomes due to under-reporting. Evidence based on ResearchFish™ outcomes are treated as underestimations, or a 'minimum achieved benefit' in this report.
- While effort has been made to quantify benefits wherever possible, not all benefits were quantifiable. Unquantifiable benefits are highlighted through qualitative discussions and case studies.
- There is also considerable uncertainty when quantifying the benefits of scientific R&D investments. Where appropriate, benefits were estimated as a range (with a low and high estimate provided) to take these uncertainties into account. The real impact of ARCHER2 will lie somewhere between these estimates.
- HPC should be understood as an ecosystem rather than a single facility. It encompasses different levels of provision (e.g. the tiers of HPC), as well as training and networking activities (e.g. CCPs/HECs). See Section 1.2 for a more detailed discussion of the different types of tiers.

2 How is ARCHER2 being used?

2.1 Number of users, and user growth over time

Figure 2 shows the trend in requested user accounts over time. Figure 3 presents new accounts created in each year between 2021 and 2025. Note that data for 2025 is partial only.

As of September 2025, ARCHER2 had nearly 9,400 registered users. Of these, around 5,750 accounts were active users on ARCHER2, that is users that have run jobs on ARCHER2, between 2021 and 2025 (see Figure 4).

The number of registered users has grown steadily over time from around 2,500 early registered users when the full ARCHER2 service launched in 2021.

New account registrations have remained broadly stable over this time. Annual user registrations ranged from around 2,100 to 2,900 new users each year, with moderate year-on-year variation.

Figure 2 Total number of requested user accounts over time

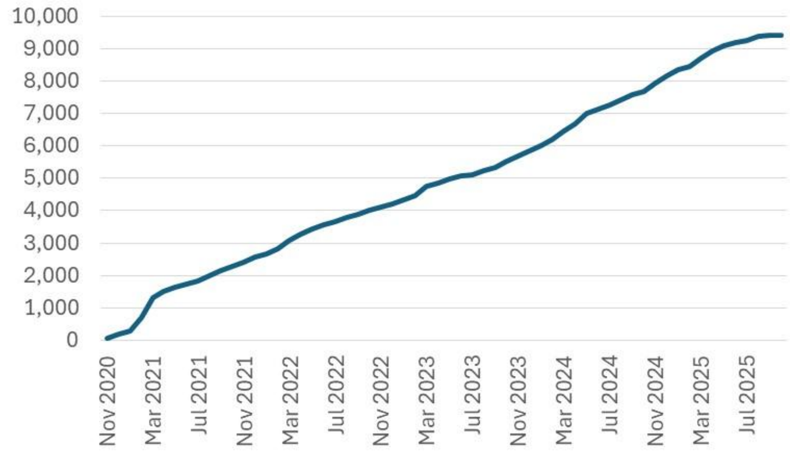

Figure 3 New accounts created by year

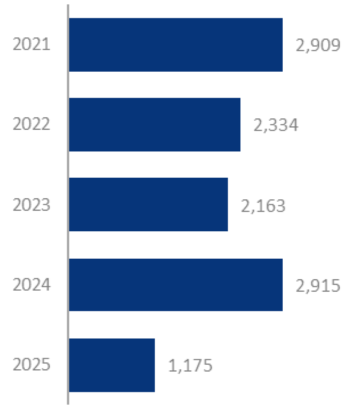

Note: 2025 data is partial only.

Source: EPCC usage data

2.2 Types of users

ARCHER2 users are predominantly academic, with over 3,800 of the approximately 5,750 active user accounts belonging to academic users. Active accounts are accounts that have run jobs on ARCHER2 between 2021 and 2025.

The majority of academic users are postgraduate researchers (1,366), followed by early-career and experienced postdoctoral researchers (747 and 223, respectively), permanent academics (695) and other R&D-focused roles such as research fellows, software engineers and support staff. Smaller numbers of undergraduate and master's students also access the service.

Non-academic users account for a smaller share of the total. These include industrial and commercial users (151 active accounts), service and administrative staff (196 and 20 respectively) and a large number of users (1,524) whose role is recorded as unknown.

In-line with funding arrangements for ARCHER2 (see Section 7.4), a large proportion of academic users undertake R&D within EPSRC's remit. EPSRC-associated users accounted for 44% of total active users between 2021 and 2025, and 80% of usage over the same time period. NERC accounted for 9% of active users and 15% of usage.

A smaller share of ARCHER2 usage is reserved for EPCC who host ARCHER2. This is known as Director's time. While Director's time accounted for a large proportion of active users (42%)[11], usage is much lower at 3% of ARCHER2 compute time.

[11] Over half (56%) of these accounts are academic users, with the vast majority being students and early-career researchers. Around 31% of active accounts accessing ARCHER2 under Director's time did not have a career stage associated with them. The remainder are commercial users and administrative and service staff.

Figure 4 Total active users by funding body, 2021-2025

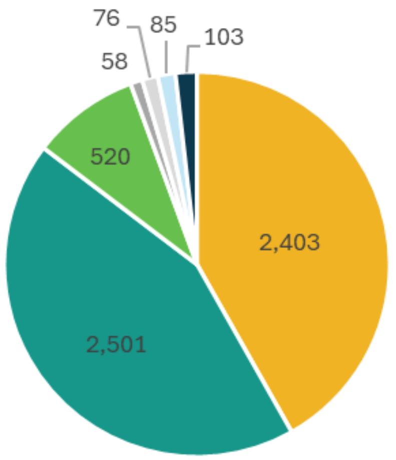

Figure 5 Usage by funding body, 2021-2025

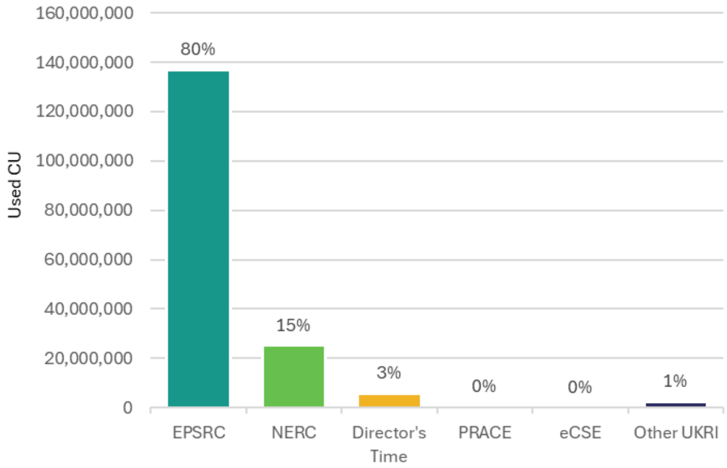

Note: Active accounts are accounts that have run jobs in the reporting period. Usage is measured by used ARCHER2 Compute Units (CU).

Source: London Economics analysis of EPCC usage data

Figure 6 Active accounts between 2021-2025 by career stage

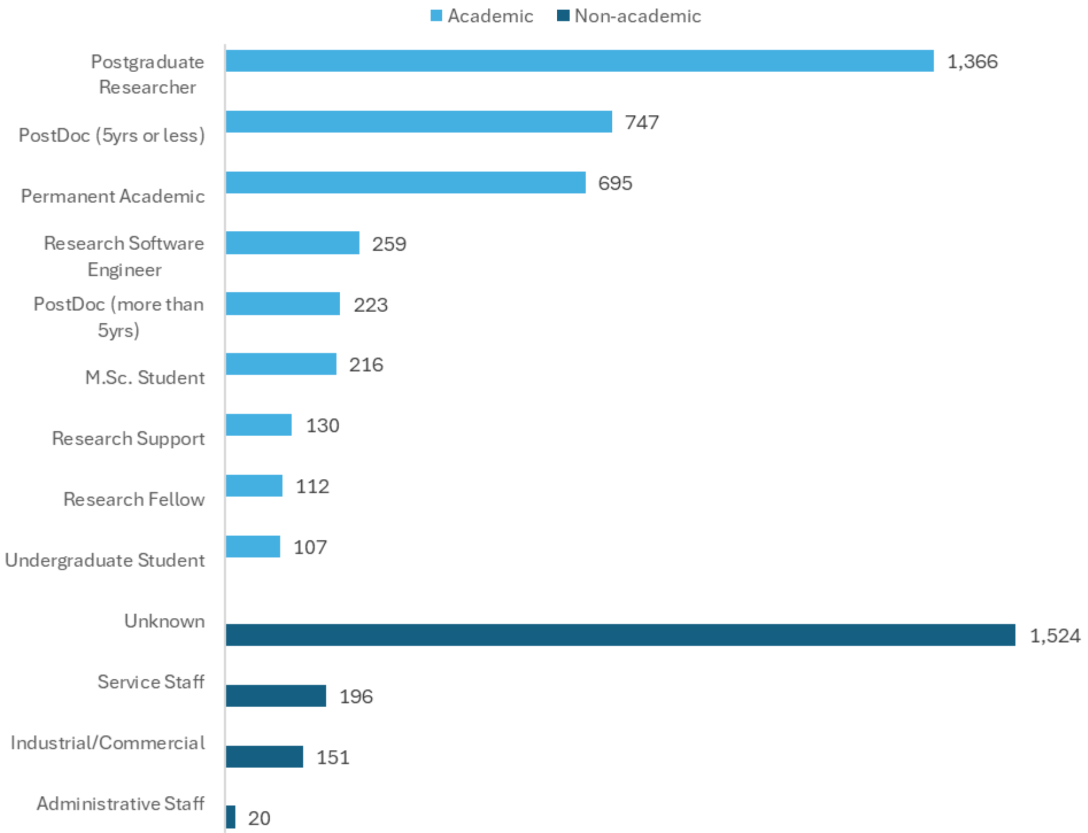

Note: Active accounts are accounts that have run jobs in the reporting period.

Source: London Economics analysis of EPCC usage data

2.3 Usage by research discipline

Materials science accounts for the largest share of usage on ARCHER2 at 42% of compute units used. This high level of usage highlights the value of computational modelling in materials science, where performing numerous simulations to identify materials with desirable properties offers a cost-effective alternative to experiments. It also removes health and safety barriers by allowing the study of hazardous or radioactive materials and enables breakthroughs in catalysis research with potentially transformative impacts, such as improving processes akin to the Haber–Bosch reaction.

Materials science is followed by computational fluid dynamics at 17% and climate science at 10%, the latter primarily supported by NERC. Other significant areas include biomolecular modelling (6%), plasma physics (4%), geophysics and seismology (3%), combustion (3%), and ocean science (2%).

Figure 7 Usage by research discipline (% of CU used between 2021 and 2025)

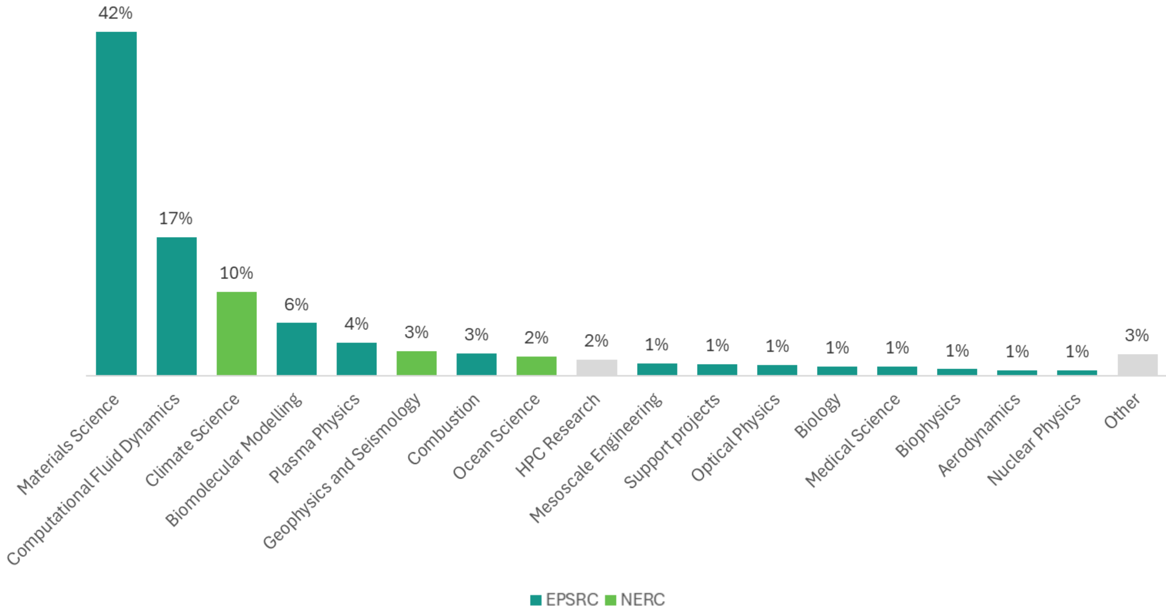

Note: Usage is measured as a % of ARCHER2 Compute Units (CU) used between 2021 and 2025. Dark green bars represent research areas with predominantly EPSRC-funded R&D. Light green bars represent research areas with predominantly NERC-funded R&D. Grey bars represent areas that cannot be attributed to either EPSRC or NERC.

Source: London Economics analysis of EPCC usage data

2.4 High-End Computing Consortia (HECs) and Collaborative Computational Communities (CCPs)

Two important research communities accessing ARCHER2 are the High-End Computing Consortia (HECs) and the Collaborative Computational Communities (CCPs).

The HECs are UKRI (EPSRC and NERC) funded research communities that bring together computational researchers from similar research domains and serve as a forum to communicate research, share experience and best practice, develop computational software and applications to support research within their respective disciplines, and provide support to their research community.

The HECs have large allocations of time on ARCHER2 that they can distribute out to their members. As such, they provide a main access channel to ARCHER2 resources for computational researchers within the EPSRC and NERC domains, quality control for the work being done on ARCHER2, and support for best practice.

In addition, ARCHER2 is also used by various Collaborative Computational Communities (CCPs). These are networks of researchers, software developers, and domain experts who work together to develop and distribute computational tools, models, and infrastructure to advance research across scientific disciplines.

Software and tools developed by the CCPs provide important enabling capabilities for individual research. This includes some of the most used codes on ARCHER2 (see Section 2.6). This ensures research on ARCHER2 and other HPC systems maximise the potential of the capabilities and thus the return on public investment in these systems.

The HECs comprise eight consortia undertaking science within EPSRC's remit and three consortia undertaken science within NERC's remit. Current HECs are listed in Table 2, while Table 3 lists current CCP communities.

Table 2 High End Computing (HEC) Consortia

EPSRC HEC Consortia	NERC HEC Consortia
UK Turbulence Consortium (UKTC)	Oceanography and Shelf Seas Consortium
Materials Chemistry Consortium (MCC)	Atmospheric and Polar Sciences Consortium
UK Car-Parrinello Consortium (UKCP)	Mineral and Geophysics Consortium
HEC Biomolecular Simulation Consortium (HECBioSim)	
Plasma HEC Consortium	
UK Consortium on Mesoscale Engineering Science (UKCOMES)	
UK High-End Computing Consortium for X-Ray Spectroscopy (HPC-CONEXS)	
High End Computing Consortium for Wave Structure Interaction (HEC WSI)	

Source: ***https://www.archer2.ac.uk/research/consortia/***

Table 3 Collaborative Computational Communities (CCPs)

Name	Research Area
CCP4	Macromolecular x-ray crystallography
CCP5	Computer simulation of condensed phases
CCP9	Computational electronic structure of Condensed matter
CCPBioSim	Biomolecular simulation
CCP-EM	Electron cryo-Microscopy
CCPi	Tomographic imaging
CCPN	Nuclear magnetic resonance
CCP-NC	Nuclear magnetic resonance for crystallography
CCP-NTH	Nuclear thermal hydraulics

Name	Research Area
CCP-QC	Quantum computing
CCPSyneRBI	Synergistic reconstruction for biomedical imaging
CCP-Turbulence	Turbulence modelling for fluid dynamics
CCP-WSI	Wave structure interaction

Source: https://www.cosec.ac.uk/communities/current-ccps-and-hecs

2.5 Usage by job size

Figure 8 and Figure 9 show the total number of jobs run on ARCHER2 and the total ARCHER2 Compute Units (CUs) used, between 2021 and 2025, by job size in terms of the number of nodes used. The figures show that ARCHER2 use is dominated by a large number of very small jobs. Jobs using only one node make up around half (50.3%) of all jobs run, while jobs using four or fewer nodes accounted for around seven out of ten (71.4%) of jobs.

However, the distribution in terms of CUs used is centred around medium-sized jobs. This suggests that many of the smaller jobs may be automated tasks such as test/debugging jobs, monitoring tasks, parameter sweeps, or usage for training purposes. In contrast, medium-sized jobs dominate in terms of the actual usage of the system, while large simulations using substantial numbers of nodes are rarer.

This suggests many ARCHER2 users are running many medium-sized simulation runs with very large-scale simulations more limited in number. An example of the latter are the Grand Challenge runs, ambitious, large-scale computations that push the boundaries of scientific understanding and that require substantial computational resources to address.

However, it is also a reflection of the popularity of ARCHER2 and resulting allocation policy, with ARCHER2 serving increasingly more users and therefore running very large jobs using up a substantial amount of available system capacity, even for shorter periods of time, is more difficult to service.

To give users access to large-scale resources despite these constraints, ARCHER2 periodically offers Capability Days. These are periods where a large fraction of the system is reserved for very large (512 nodes or more) capability jobs. Capability days provide users an opportunity to test and understand their applications at scale to help prepare software and communities for future resources; enable capability use cases not possible on other UK HPC services; and enhance science undertaken on ARCHER2 by enabling modelling and simulation at scales not otherwise possible.

In addition, ARCHER2 also facilitates Pioneer Projects calls, which enable users to apply for large amounts of computational resource to conduct computationally intensive modelling, simulation and calculations. The aim of these calls is to support ambitious and pioneering projects that significantly push the boundaries in computational research in their field and deliver outputs with significant potential for a high future impact. This report provides examples of such projects including:

- Investigating unexplained behaviours in simulations of molecules exposed to intense radiation pulses (Box 3)
- Supporting zero-emission aviation through High-Performance Computing (Box 9)
- Advancing the development of compact electron accelerators through High-Performance Computing (Box 16)

Figure 8 Number of jobs by job size in nodes (% of jobs between 2021 and 2025)

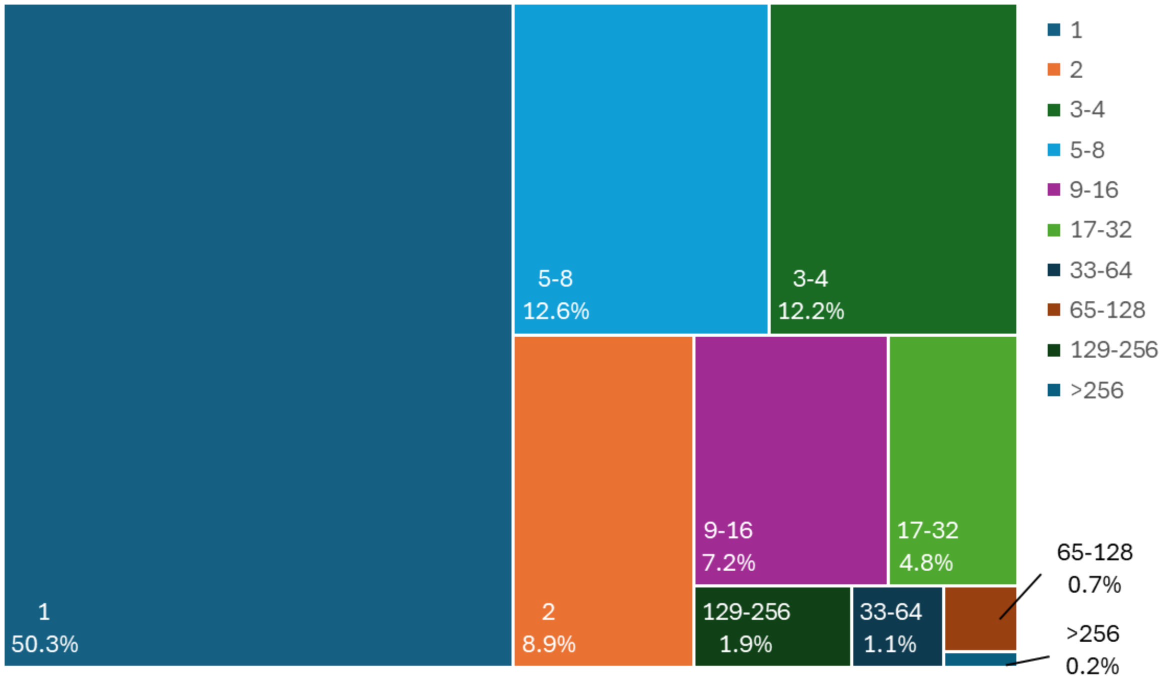

Note: Usage is measured as a % of total jobs used between 2021 and 2025.

Source: London Economics analysis of EPCC usage data

Figure 9 CUs used by job size in nodes (% of CUs between 2021 and 2025)

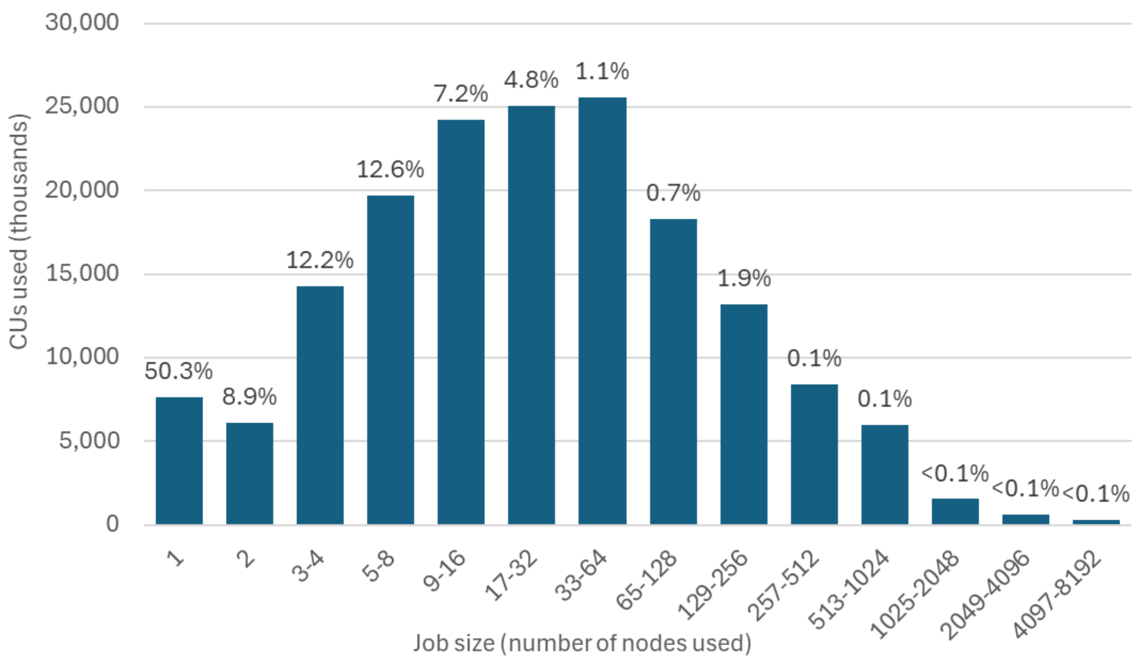

Note: Usage is measured as a % of ARCHER2 Compute Units (CUs) used between 2021 and 2025.

Source: London Economics analysis of EPCC usage data

2.6 Geographical reach

As expected from a national service, the vast majority (98.6%) of ARCHER2 users are located in the UK. The highest share of users are located in London, with 30.5% of users accessing ARCHER2 based there. Scotland accounts for 15.3% of users, with 87% of Scottish users based in Edinburgh, where ARCHER2 is located. This is closely followed by the South East (12.3% of users) and North West (10.5% of users). Table 4 provides a breakdown of users by region. In terms of cities outside of London and Edinburgh, a significant share of users are located in Cambridge (6.6% of users), Manchester (6.1% of users) and Oxford (5.7% of users) (Table 5). ARCHER2's reach also extends beyond the UK with 1.4% of users accessing ARCHER2 across Europe and beyond (Figure 10).

Table 4 Breakdown of users by region

Region	% of users
London	30.5%
Scotland	15.3%
South East	12.3%
North West	10.5%
East of England	6.7%
Yorkshire and the Humber	6.0%
West Midlands	4.1%
South West	4.0%
North East	3.8%
East Midlands	2.3%
Wales	1.6%
Northern Ireland	1.4%
Non-UK	1.4%

Source: London Economics analysis of EPCC data

Table 5 Cities with the most users

Cities	% of users
London	30.5%
Edinburgh	13.3%
Cambridge	6.6%
Manchester	6.1%
Oxford	5.7%
Southampton	2.6%
Leeds	2.5%
Newcastle	2.4%
Warwick	2.3%
Reading	2.2%
Daresbury	2.1%
Bristol	2.0%
York	1.8%

Source: London Economics analysis of EPCC data

Figure 10 Geographical location of ARCHER2 users

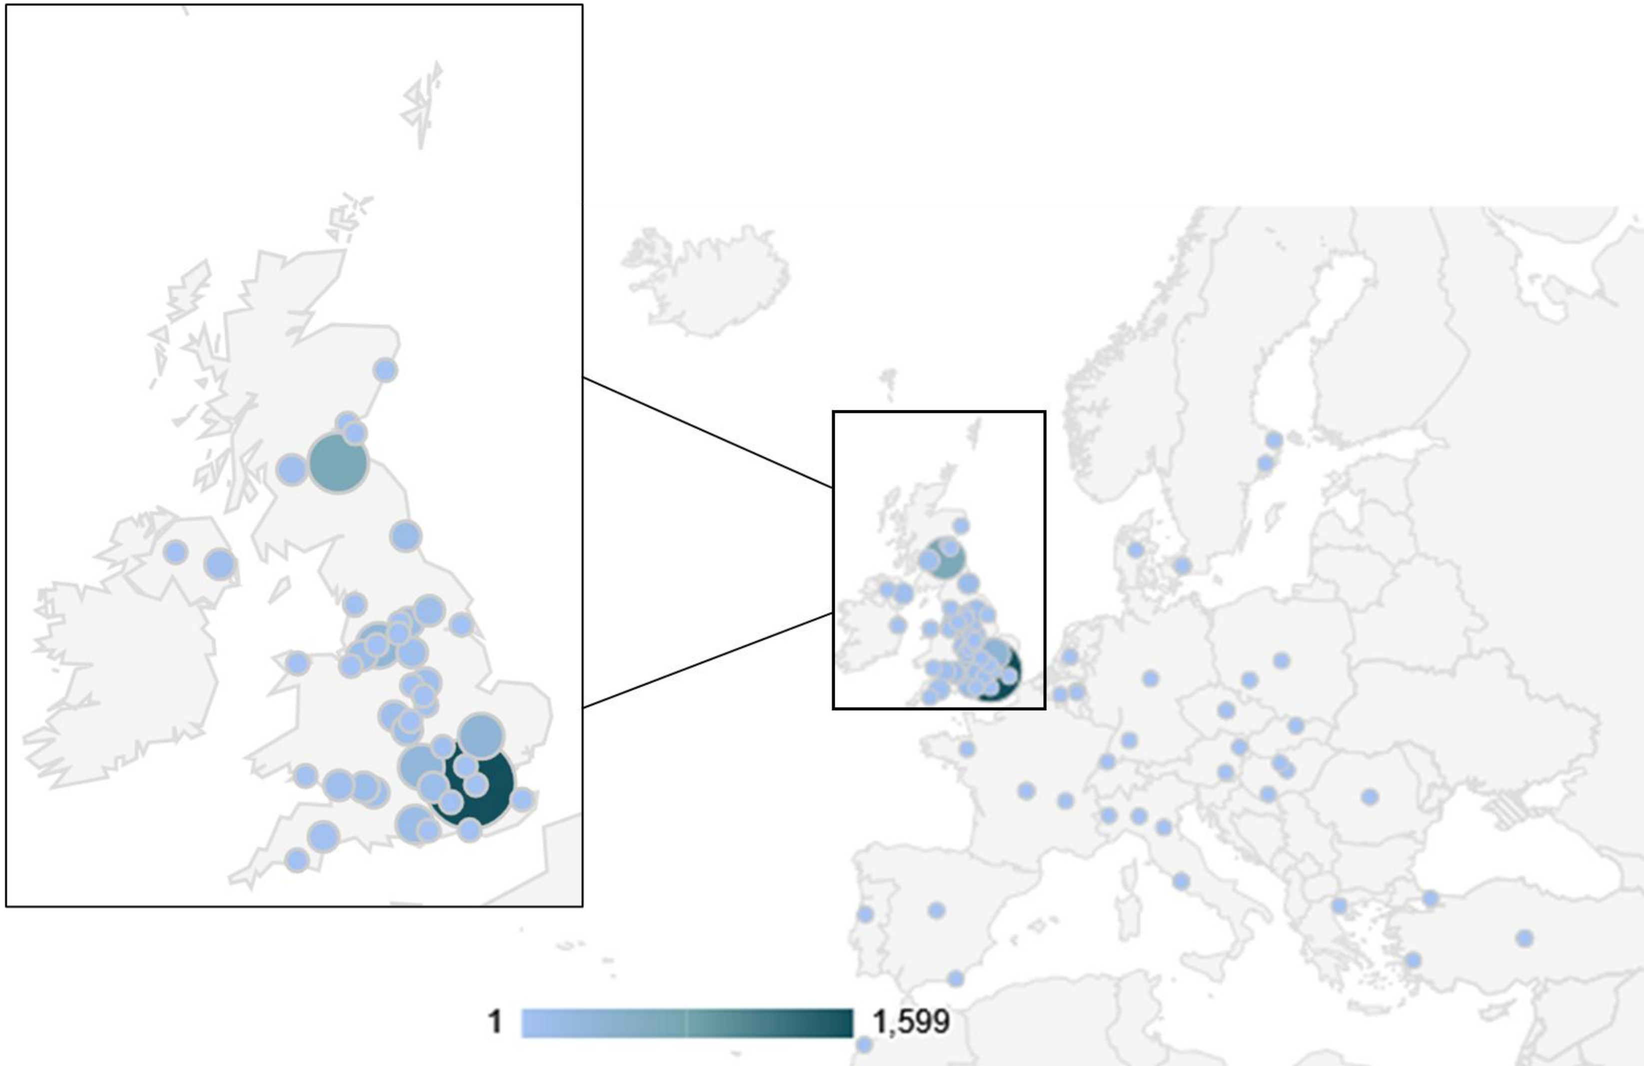

Source: EPCC

Box 1 Most used software codes run on ARCHER2

ARCHER2 provides essential capacity for both high-frequency, low-intensity tools and those supporting intensive, domain-specific modelling. These utilisation patterns highlight the diversity of scientific challenges being tackled and underscore the system's role in advancing research across multiple disciplines.

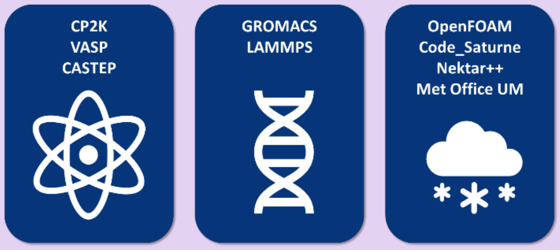

Figure 11 shows the top ten software codes run on ARCHER2. Together these top ten codes account for more than half (56%) of total ARCHER2 usage in terms of node hours.

VASP, a software package for materials modelling at the atomic scale is the most popular code on ARCHER, accounting for nearly 23% of ARCHER2 time across 140 users and 15 projects.

Other popular software packages include CP2K and CASTEP (also for atomic-scale simulations); GROMACS and LAMMPS (for molecular dynamics); and OpenFOAM, Code_Saturne, Nektar++, and the Met Office Unified Model (for large-scale system modelling).

Further details on these software codes can be found in the full case study in Annex A5.1.

Figure 11 Top Ten codes run on ARCHER2

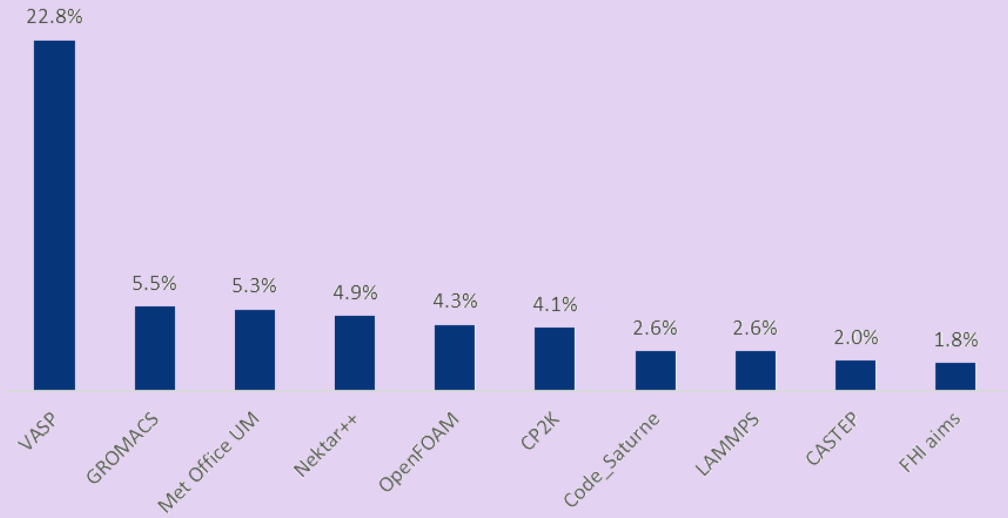

Source: EPCC data

Source: London Economics analysis based on EPCC data

PART II: SCIENTIFIC AND ECONOMIC IMPACTS

This part discusses:

- **Section 3** discusses the importance of ARCHER2 in enabling high impact, high quality science, for example the benefits of increased R&D, collaborations and knowledge exchange facilitated by ARCHER2 and the impact of ARCHER2 on software development/building software codes.
- **Section 4** discusses the role of ARCHER2 in supporting a highly skilled research community, for example by attracting talent to the UK and providing training to researchers.
- **Section 5** discusses the economic impacts of ARCHER2, such as the economic impacts of R&D enabled, for example through increased collaborations, industry impacts, resulting spin-outs, patents and publications.
- **Section 6** discusses the wider benefits of ARCHER2, such as providing a high-quality service to UK researchers, improved local societal impact and promoting outreach, diversity and inclusion.

Image source: Eric Breard, School of Geoscience, University of Edinburgh: Turbulent times

3 The importance of ARCHER2 in enabling high impact, high quality computational science

High-Performance Computing is a fundamental pillar of modern scientific R&D. As the UK's most powerful supercomputer, ARCHER2 provides researchers the capabilities needed to conduct complex large-scale simulations to tackle scientific questions that would be impossible to answer without the compute capacity of a national Tier-1 service.

High-performance compute at national scale enables researchers and innovators to develop and improve simulations that more accurately reflect the real-world in terms of accuracy, granularity and complexity, allowing for deeper and more realistic exploration of key questions. ARCHER2 also improves the efficiency of research by enabling users to undertake complex simulations in significantly shorter timeframes than would be possible with less powerful compute capabilities.

ARCHER2 also contributes to scientific progress. Access to a national Tier-1 compute capability enables UK researchers to remain competitive in the global scientific landscape, fosters academic collaboration, and supports the development and optimisation of advanced scientific software and methods. For example, a number of new mathematical models are developed and tested on ARCHER2, before being simulated at scale.[12]

Box 2 Modelling extreme wave conditions for floating wind turbines

ARCHER2 is used for new wave modelling techniques which are reducing the time and cost of designing floating wind turbines.

A team of researchers at the University of Plymouth, led by Deborah Greaves and Scott Brown, is using ARCHER2 to improve the safety and speed of designing floating offshore wind turbines. Their work focuses on testing new modelling approaches that will speed up how floating wind turbines are designed to survive harsh marine environments, addressing a growing challenge in renewable energy: how to build more floating turbines, faster, without compromising safety.

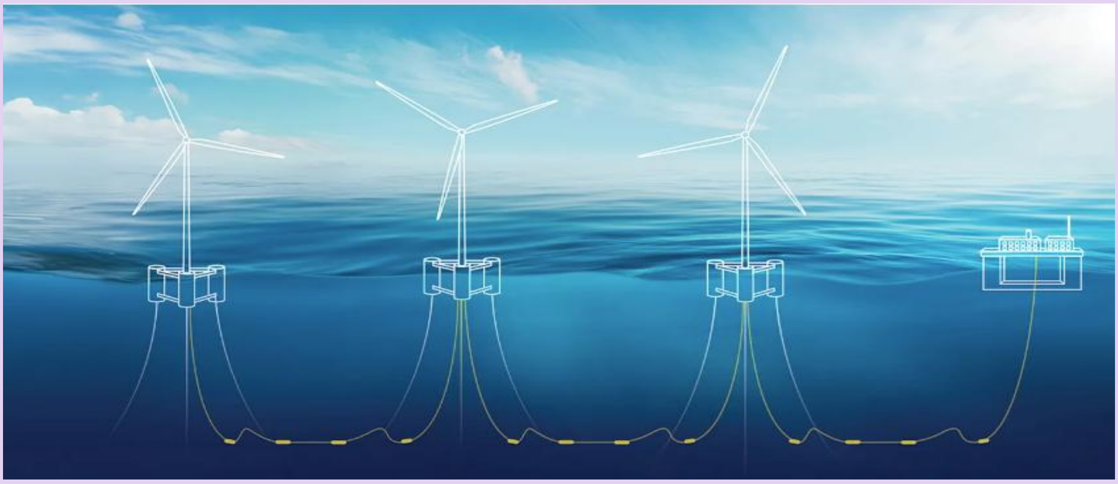

Further details can be found in the full case study in Annex A5.4.

Source: Technopolis' analysis based on an interview with Deborah Greaves and Scott Brown.

[12] See Computational Mathematics and High-Performance Computing - UCD School of Mathematics and Statistics

3.1 Greater scientific throughput and higher quality scientific output

At its core, high-performance computing drives scientific R&D in two main ways: it makes possible lines of inquiry that would otherwise be out of reach, and it enhances the speed and rigour of existing R&D. It thereby enables researchers to tackle a broader array of scientific problems and to investigate them with far greater detail and accuracy.

To understand the importance of ARCHER2 for scientific R&D, a survey of ARCHER2 users was undertaken for this study. The survey asked academic users to what extent they 'agreed' or 'disagreed' with a range of scientific benefits enabled by ARCHER. Results of the user survey are presented in Section 3.1.1, below.

Further, bibliometric analysis of research outputs (publications and citations) was undertaken by Technopolis. The bibliometric analysis explores publication patterns and trends associated with ARCHER2 use and provides additional quantitative evidence of its role in supporting research. The results of the bibliometric analysis are presented in Section 3.1.2.

"ARCHER2 has significantly benefitted my personal and group's research outputs in terms of important career metrics (publications and conferences) and, more importantly, scientific discoveries."

"ARCHER2 has been a crucial and reliable resource that has benefitted my research through access to large-scale HPC, training events and general support for computing tasks."

"ARCHER2 enables me to model large scale systems that I would not be able to tackle without this facility, which would significantly reduce my research output and potentially make some aspects impossible."

- **ARCHER2 user survey**

Box 3 Investigating unexplained behaviours in simulations of molecules exposed to intense radiation pulses

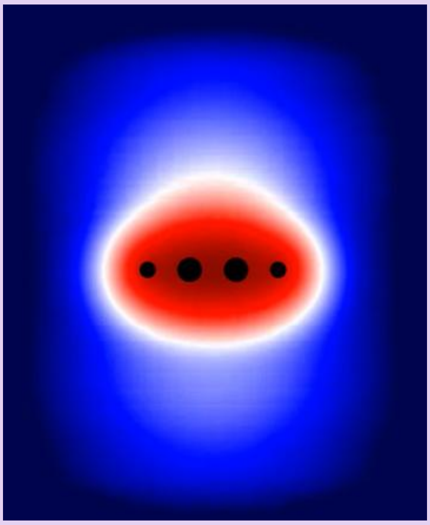

Using ARCHER2, scientists have simulated a new and unexplained oscillation of the electric fields in molecules subjected to ultra-fast laser pulses, which has important implications for our understanding of how atoms and molecules transition between states, and for novel materials science more broadly.

The team of researchers was led by Dr Daniel Dundas from Queen's University Belfast. If their simulations can be validated using alternative underpinning theories, there is strong grounds for having predicted a previously unknown physical phenomenon, which can be tested experimentally. On the other hand, if the oscillations turn out to be no more than a quirk in this particular scientific theory, the results are important for the application of this widely used theory.

Further details can be found in the full case study in Annex A5.3.

Photo Credit: ARCHER2 animation depicting the moment that an Acetylene molecule is subjected to a 50 terawatt, 30 femtosecond(fs) ultraviolet laser pulse. Source: https://www.archer2.ac.uk/about/gallery/2022-image-comp/

Source: Technopolis' analysis based on an interview with Dr Daniel Dundas

3.1.1 ARCHER2 user survey insights

"I have so far published three first-author papers using ARCHER2 resources (with two more planned), which would not be possible (or at least a lot less detailed analysis) using lower tier HPCs."

"ARCHER2 access has enabled multiple publications in leading scientific journals and contributed to novel insights in the field of materials science."

"ARCHER2 [and ARCHER] is our workhorse machine, vital to most of my 300+ publications."

- **ARCHER2 user survey**

To understand the overall importance of ARCHER2 to R&D, academic users were asked in the user survey undertaken for this study to what extent they agree or disagree that ARCHER2 is important to their research. In response, **92% of respondents agreed that ARCHER2 is important to their research** - with 89% strongly agreeing with this statement (Figure 12).

The importance of ARCHER2 was also consistently highlighted across all stakeholder consultations undertaken, reinforcing ARCHER2's central role in supporting and enabling high-quality scientific R&D.

To understand the specific R&D benefits enabled by ARCHER2, academic users were asked, in the user survey undertaken for this study, to rate their level of agreement with a series of statements about ARCHER2's specific contribution to their work (Figure 12):

- **85%** of users responding to the user survey strongly agreed that ARCHER2 **improved the efficiency of their research**,
- **84%** strongly agreed that ARCHER2 enabled them to **do more research**,
- **82% strongly agreed** that ARCHER2 **increased the quality of their research,**
- **83%** strongly agreed that ARCHER2 allowed them to **explore more questions** than if they did not have access to ARCHER2,
- **80%** strongly agreed that ARCHER2 enabled them to **do work they would have otherwise been unable to do**[13], and
- **81% agreed**, either slightly or strongly, that **ARCHER2 enabled new innovations**, with nearly three-fifths (58%) strongly agreeing.

Further, to understand satisfaction with ARCHER2's capabilities, users responding to the user survey were asked to what extent ARCHER2 met their current and future scientific needs. In response, **91% agreed**, either slightly or strongly, **that ARCHER2 met their current scientific needs**. This proportion drops to 84% over the next near-term (1-2 years), and to 60% over the medium-term (3-5 years), highlighting the need for continued investments into national Tier-1 HPC infrastructure. (Figure 13)

In this regard, it is positive to note that the UK has recently announced a £750 million investment into the UK's next national supercomputing service, also located at EPCC at the University of Edinburgh.[14]

[13] For a breakdown of jobs by job size see Section 2.4.
[14] See https://www.ed.ac.uk/news/university-set-to-host-ps750m-national-supercomputer

Figure 12 Benefits of ARCHER2 to research experienced by users

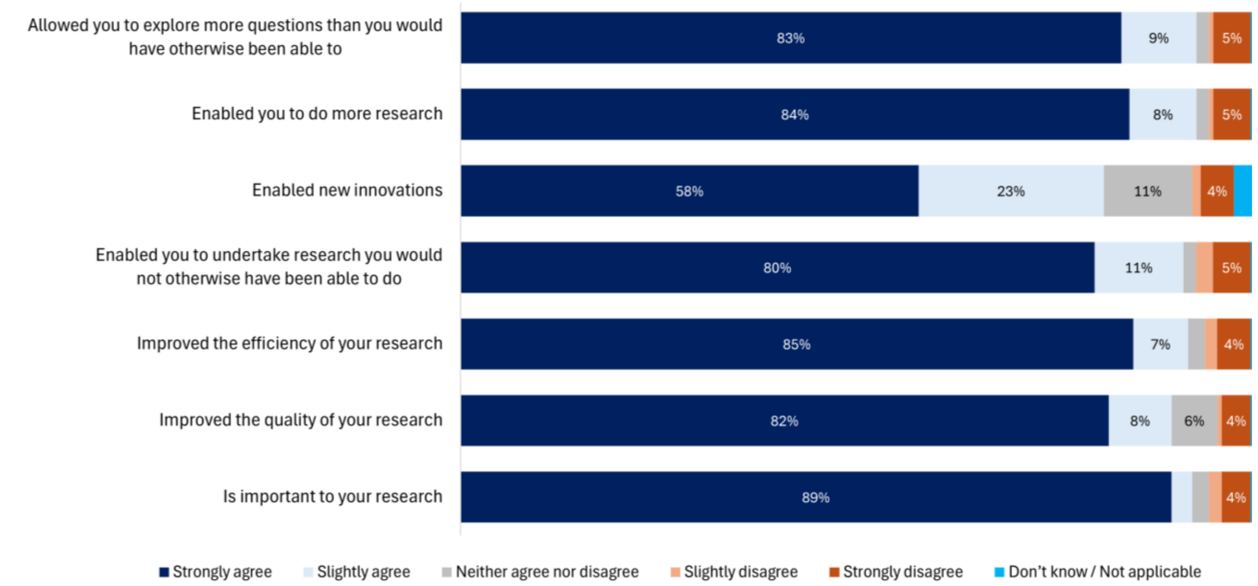

Note: Based on between 188-190 respondents (varies by question). Approximately 1% of respondents answered 'Don't know/not applicable' for each question, which are not displayed in the graph. Question asked: 'In terms of benefits to your research, to what extent do you agree or disagree that your use of ARCHER2 …?'

Source: London Economics ARCHER2 user survey

Figure 13 User perceptions of ARCHER2's ability to meet scientific and other requirements

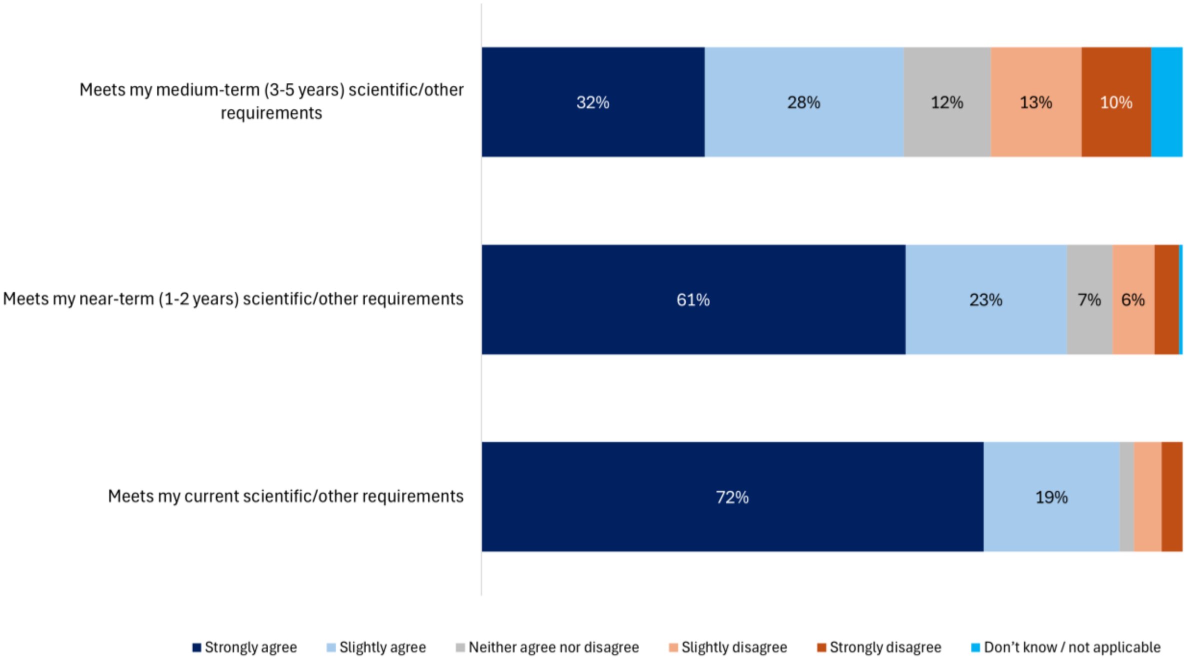

Note: Based on 200-201 respondents. Question asked: 'Do you agree that ARCHER2 has met your scientific/other requirements?'

Source: London Economics ARCHER2 user survey

Box 4 Calibrating a whole-heart electromechanics model to patient data

Access to ARCHER2 has empowered Imperial College's Cardiac Electromechanics Research Group to perform large-scale, high-resolution simulations of personalised heart 'digital twins,' enabling safer, data-driven treatment decisions and advancing cardiovascular research, clinical guidelines, and regulatory science.

The Cardiac Electromechanics Research Group (CEMRG) at Imperial College London develops personalised 'digital twin' simulations that mimic the electrical and mechanical functions of patients' hearts, using detailed medical imaging and pressure data. Running these complex models on ARCHER2 allows the team to safely test different treatment options virtually, leading to valuable clinical insights, influencing regulatory frameworks, and securing significant research funding.

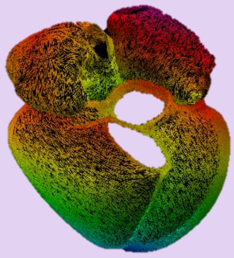

The research addresses a significant public health challenge: cardiovascular disease (CVD) remains a leading cause of death in the UK and disproportionately affects disadvantaged communities. Conditions such as atrial fibrillation require more accurate diagnostic and treatment strategies, which traditional clinical tools do not provide. CEMRG's digital twin models offer an alternative by enabling more tailored, data-driven decisions and reducing unnecessary interventions.

The team has used ARCHER2 to conduct large-scale, high-resolution simulations that would not have been feasible using standard institutional computing resources. For instance, one study required 500 simulations, each running on 512 cores for five hours. The group also improved their simulation code and modelling processes through ARCHER2 access. Their work contributed to securing an £8.8 million EPSRC programme grant, CVD-Net, focused on pulmonary arterial hypertension.

CEMRG's research has been cited in the European Society of Cardiology's guidelines and referenced by the US Food and Drug Administration in the context of regulatory science. Outputs include studies on cardiac resynchronisation therapy, atrial fibrillation and sex-based treatment responses.

Alongside research outputs, the group is active in public engagement, participating in youth and community events and featuring in media coverage. Looking ahead, CEMRG's work highlights how high-performance computing and digital twin models can improve cardiovascular care while contributing to more efficient and cost-effective healthcare delivery.

Further details can be found in the full case study in Annex A5.5.

Photo credit: Strocchi, M. (n.d.). Heart fibres. This image shows the orientation of the cells in the heart included in the model to accurately simulate tissue electromechanical behaviour. This image won the British Heart Foundation Reflection of Research image competition. Retrieved from https://www.bhf.org.uk/what-we-do/our-research/reflections-of-research/previous-winners; Baptiste, T. (n.d.). Fat heart. This image shows a whole-heart model with the fat surrounding the heart.; Rodero, C. (n.d.). Hypertrophy. This image is a cross-section of a digital model of a heart with hypertrophic cardiomyopathy (thickened heart muscle).

Source: London Economics analysis based on information gathered via desk research and consultation with Marina Strocchi.

3.1.2 Bibliometric analysis

ARCHER2-related publications

In total, **3,990 publications reported in UKRI's research outcome database ResearchFish™** by ARCHER2-related awards[15] could be identified in Open Alex[16]. This represents an average of **19 publications per award**. Figure 14 shows their distribution over time, by publication year.

These publications will have resulted from ARCHER2 as well as its predecessors. To isolate the impact of work undertaken on ARCHER2, only the **2,142 papers published in 2022 or later**[17] have been considered for the remainder of the analysis. Note that this underestimates the ultimate number of publications that will result from work undertaken on ARCHER2 over its lifetime due to the time delay from grant award to research to final publication of the work.

The 2,142 ARCHER2 publications were spread across 20 different fields of research, but with nearly three-quarters (73%) classified into just four areas: Engineering (25%), Material Science (17%), Environmental Science (16%), and Earth & Planetary Sciences (15%). (Figure 15)

Figure 14 Number of publications from ARCHER2-related awards, by year of publication

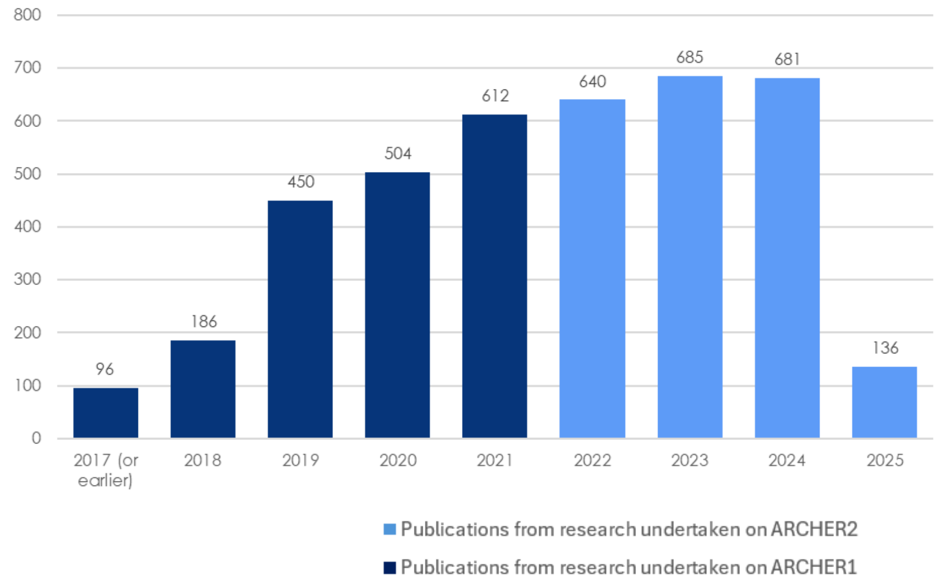

Note: Publications reported by awards that mentioned ARCHER2 in their grant application and / or in outcomes reporting of facilities used. Excludes publications that could not be identified in Open Alex based on their DOI. Light blue indicates publications from 2022 onwards (used in the analysis as the basis for identifying publications resulting from ARCHER2, rather than its predecessors).

Source: Bibliometric analysis undertaken by Technopolis based on data from UKRI's research outcomes database ResearchFish™

[15] As a starting point, 209 ARCHER2-related awards were identified. This includes EPSRC awards that mentioned ARCHER2 within their grant application (n=73), plus awards (from any council) that have mentioned ARCHER2 in the 'Use of Facilities' outcomes data reported to ResearchFish (n=152, of which 16 were already identified from applications). Of these ARCHER2 awards, 152 have reported publication information to ResearchFish (as of June 2025), with 6,302 publication records reported in total. This list includes 5,210 unique publications that have a DOI. Of these, 3,990 publications could be identified in Open Alex based on their DOI (including 2,142 papers published in 2022 or later). Figure 1 shows their distribution by publication date.

[16] A bibliographic catalogue of scientific papers, authors and institutions, accessible in open access mode. Over 250 million scholarly works from 250,000 sources are indexed in the catalogue, including books, book chapters, conference proceedings, datasets, journal articles, papers, repositories, theses and preprints.

[17] The analysis includes publications resulting from ARCHER2 as well as its predecessors. To isolate the impact of work undertaken on ARCHER2, only work published in 2022 or later was considered. It should be noted that the ARCHER2 four-cabinet system was in operation before this date, and so some earlier publications may be linked to work undertaken on ARCHER2. Similarly, it is possible that some of the publications in 2022 onwards are resulting from work undertaken on ARCHER2's predecessor. Nevertheless, this was chosen as a reasonable cut-off date to distinguish publications that are likely to still be related to work undertaken on ARCHER2's predecessor.

Figure 15 Proportion of ARCHER2-related publications by field of research

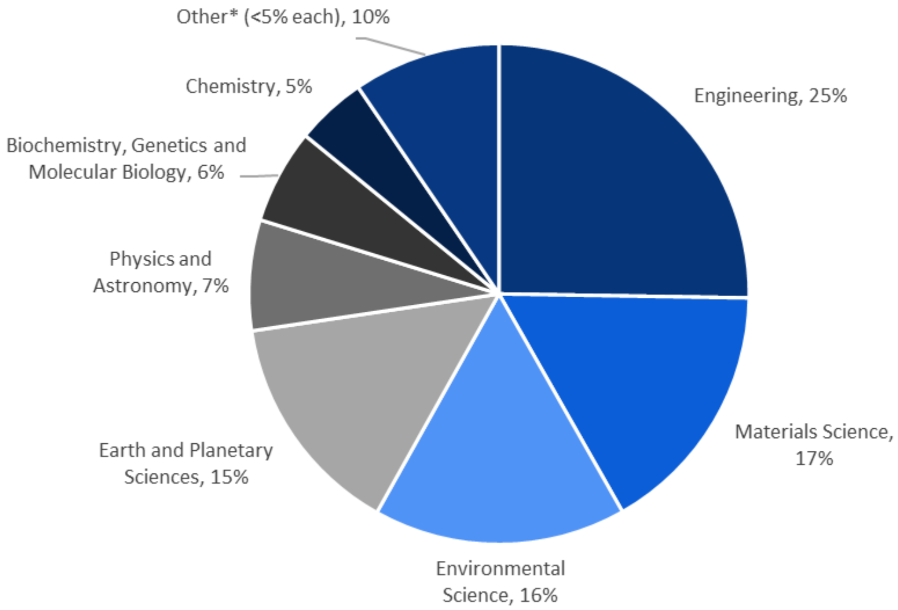

Note: *Other includes: Energy (2%), Chemical Engineering (2%), Medicine (2%), Computer Science (2%), Mathematics (0.5%), Agricultural and Biological Sciences (0.3%), Immunology and Microbiology (0.3%), Neuroscience (0.2%), Decision Sciences (0.2%), Economics, Econometrics and Finance (0.2%), Social Sciences (0.2%), Pharmacology, Toxicology and Pharmaceutics (0.05%), and No field assigned (0.2%).

Source: Bibliometric analysis undertaken by Technopolis

Uptake (Citation) by other academic publications

Most (86%) of the 2,142 ARCHER2 publications identified in Open Alex have already been cited at least once by another paper. The proportion that are cited increases to 96% if we look at just papers published in 2022 and 2023 (i.e. excluding the most recently published papers).

The average number of citations for ARCHER2 papers is 12.8 (the average is 15.0 when just considering those that have been cited).

A majority (71%) of ARCHER2 papers have been cited between 1 and 20 times each, but there are 22 (1% of the total) that have been cited more than 100 times.

Considering only papers published in 2022 or 2023, i.e., those that are more likely to have already been cited than newer papers, the average number of citations for ARCHER2 papers rises to 17.2 (the average is 16.5 for all papers, including those that have not yet been cited).

The figure overleaf shows the distribution of citation rates across the portfolio of ARCHER2 papers published in 2022 or 2023. For readability, the figure excludes five papers with 200+ citations.

Figure 16 Number of ARCHER2-related publications with x number of citations

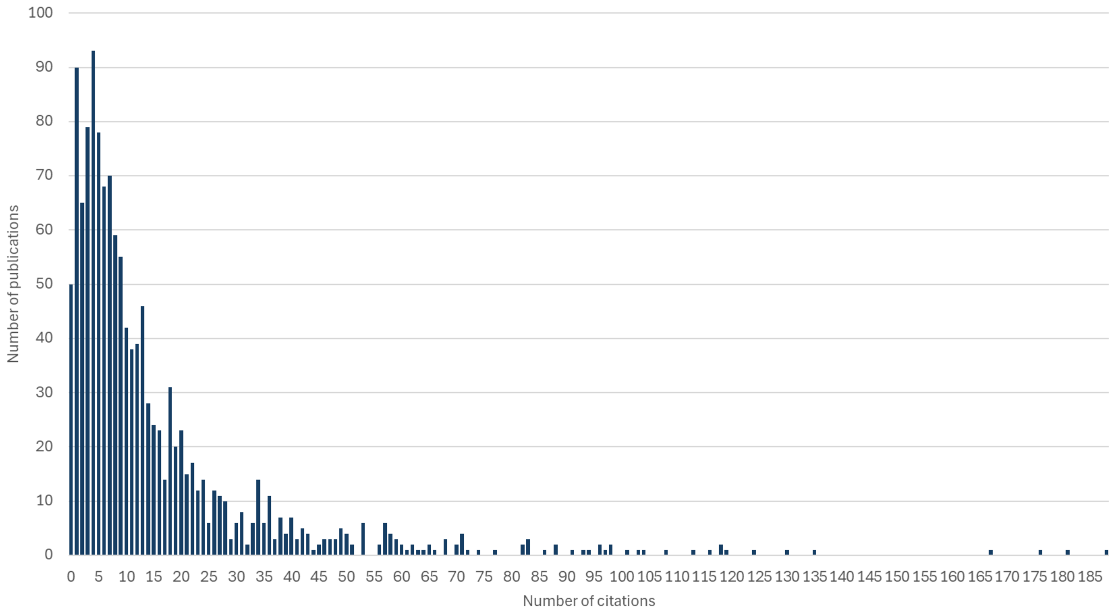

Source: Bibliometric analysis undertaken by Technopolis

The 2,142 ARCHER2-related publications have been cited 27,486 times in total within 23,756 other papers (meaning some of these papers cite multiple ARCHER2 publications). The citing papers are spread across 26 different fields of research, but with nearly three-quarters (72%) accounted for by papers that are classified into four areas: Engineering (26%), Environmental Science (22%), Material Science (13%) and Earth & Planetary Science (11%). The spread of fields is very similar to that of the original ARCHER2-related publications.

Figure 17 Fields of publications citing ARCHER2-related papers

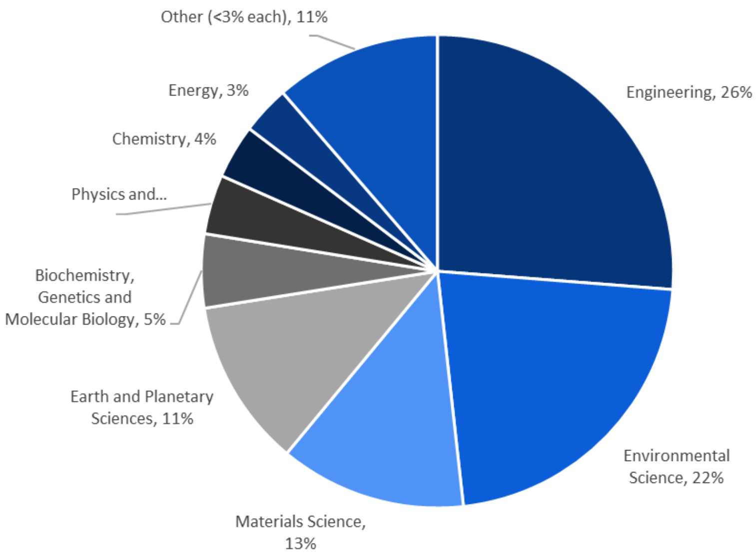

Note: *Other: Medicine (2%), Agricultural & Biological Sciences (2%), Chemical Engineering (2%), Economics, Econometrics & Finance (1%), Computer Science (1%), Social Sciences (1%), Immunology & Microbiology (0.3%), Nursing (0.3%), Health Professions (0.2%), Mathematics (0.2%), Business, Management & Accounting (0.2%), Neuroscience (0.2%), Decision Sciences (0.2%), Pharmacology, Toxicology & Pharmaceutics (0.1%), Psychology (0.1%), Arts & Humanities (0.04%), Veterinary (0.01%), No field (1%).

Source: Bibliometric analysis undertaken by Technopolis

Considering only those fields with >30 ARCHER2 papers, there is one **(Environmental Science) with particularly high rates of citation** (with 24 citations per paper, compared with the overall average of 12.8).

The average Field Weighted Citation Index (FWCI, a measure of how citations received compare to the expected number of publications for similar publications) **across the full portfolio of papers is 2.6** (i.e. on average 2.6x the expected citation rates, given the year and field of publication). Across all of the fields with >30 papers, the average FWCI is greater than 1, and is particularly high in relation to Environmental Science (4.2), Earth & Planetary Science (3.5) and Computer Science (3.2).

Figure 18 Average FWCI of ARCHER2-related papers

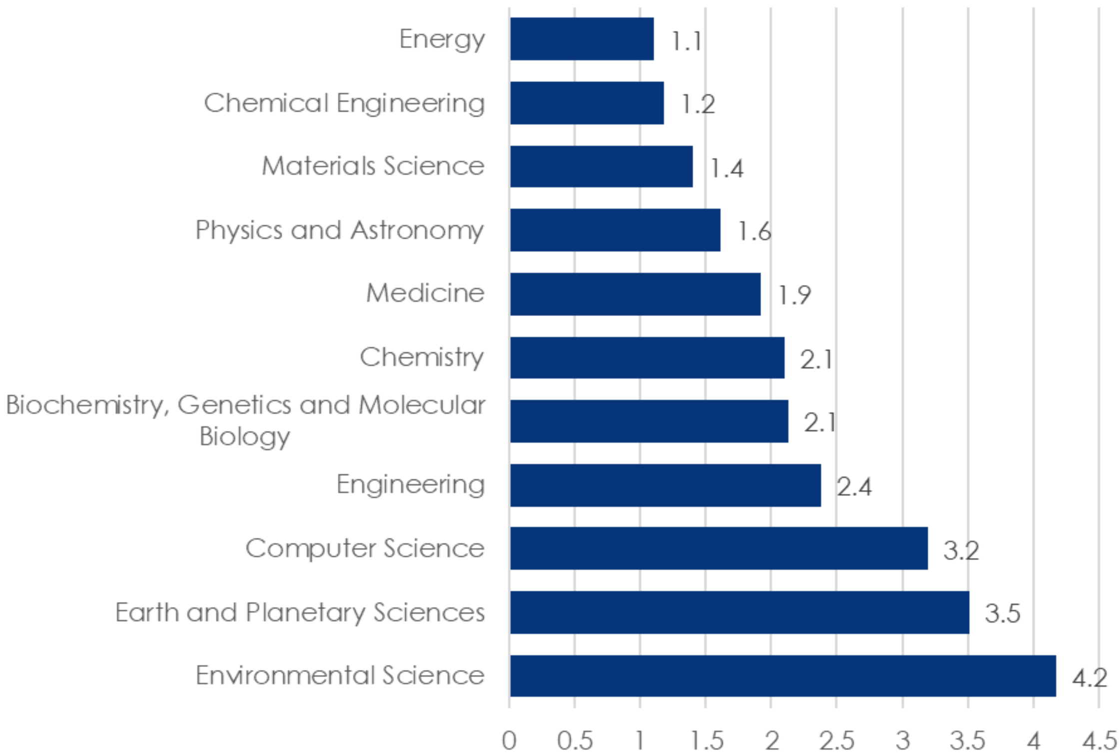

Note: Only includes fields with > 30 papers.

Source: Bibliometric analysis undertaken by Technopolis

Box 5 Using HPC to simulate bacterial membranes for antibiotic research

Antibiotic-resistant infections contributed to over 40,000 deaths in the UK in 2019[i] and cost the NHS an estimated £1.7 billion in 2023.[ii] Continued access to ARCHER2 is enabling work to reduce these costs and improve clinical outcomes.

Researchers at the University of Oxford have used ARCHER2 to produce one of the most detailed computational models of a bacterial outer membrane. This work, funded through an EPSRC Established Career Fellowship, is helping to improve understanding of how antibiotics interact with bacterial defences. It supports the early stages of drug discovery and aims to address the growing challenge of antibiotic resistance.

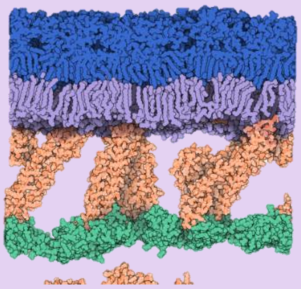

Many harmful bacteria are protected by an outer membrane that is poorly understood and difficult for antibiotics to penetrate. The Oxford team used ARCHER2 to run simulations of the membrane, generating new biological insights into how it behaves and how antibiotics might overcome its defences. High-performance computing is essential to this work, enabling researchers to simulate structures of realistic size and composition, and to run calculations for sufficient time to observe how molecules interact.

Findings published in *Science Advances*[iii] show that some lipids in the membrane bind tightly with proteins to form stable, less permeable structures, likely strengthening the bacteria's defences. The team is now using high-throughput methods and AI to test how potential new antibiotics interact with the membrane, in collaboration with partners in the US.

Further details are provided in the full case study in Annex A5.2.

Note:

(i) Department of Health and Social Care. (2024). *Confronting antimicrobial resistance: 2024 to 2029*. https://www.gov.uk/government/publications/uk-5-year-action-plan-for-antimicrobial-resistance-2024-to-2029/confronting-antimicrobial-resistance-2024-to-2029.

(ii) UK Health Security Agency. (2024, November 14). *Antibiotic resistant infections continue to rise*. GOV.UK. Retrieved June 6, 2025, from https://www.gov.uk/government/news/antibiotic-resistant-infections-continue-to-rise

(iii) Webby, M. N., Oluwole, A. O., Pedebos, C., Inns, P. G., Olerinyova, A., Prakaash, D., Housden, N. G., Benn, G., Sun, D., Hoogenboom, B. W., Kukura, P., Mohammed, S., Robinson, C. V., Khalid, S., & Kleanthous, C. (2022). Lipids mediate supramolecular outer membrane protein assembly in bacteria. Science Advances, 8(44), eadc9566. https://doi.org/10.1126/sciadv.adc9566

Source: London Economics analysis based on information gathered via desk research and consultation with Prof. Syma **Khalid.**
Picture credit: **Khalid Lab. (2025).** ***Molecular structure with different layers represented in various colours*** **[Photograph]. Khalid Lab, University of Oxford. https://khalidlab.web.ox.ac.uk/research**

3.2 Increased academic knowledge exchange and collaborations

In addition to directly supporting high-quality R&D, ARCHER2 also plays an important role in positioning the UK as a credible and competent partner in the international HPC R&D landscape. This in turn contributes to academic knowledge exchange and domestic and international collaborations.

These partnerships generate significant benefits for the UK, including the transfer of knowledge and skills across sectors. This section offers additional insights into the nature of collaborations supported by ARCHER2.

Further discussion on the channels and pathways through which collaborations and academic knowledge exchange contributes to the UK economy is provided in Section 5.1. ARCHER2's role in strengthening the UK's international scientific competitiveness is explored further in Section 4.1.

3.2.1 Collaborations and academic knowledge exchange on ARCHER2

To understand the scale of academic collaborations, analysis of data on the number of collaborations reported to UKRI's research outcome data, ResearchFish™, was undertaken. In total, **more than 157 collaborations involving ARCHER2** were recorded within ResearchFish™.

Naturally, given the academic nature of ARCHER2, most of the collaborations on ARCHER2 were academic collaborations, with 63% (184 of 292 for which a type was provided) of collaborations reported to UKRI's research outcome data being classified as academic/university collaborations. However, academics using ARCHER2 also collaborated with other private, public or non-profit organisations (see Figure 19).

Further, this figure is likely an underestimate as i) ResearchFish™ data relies on self-reporting and ii) ARCHER2 has not yet reached its end of life and so the number of collaborations taking place on ARCHER2 may still increase. Therefore, the extent of collaborations was also explored through the user survey and by examining author affiliations on ARCHER2-related publications (see overleaf).

Figure 19 Number of collaborations on ARCHER2, by type

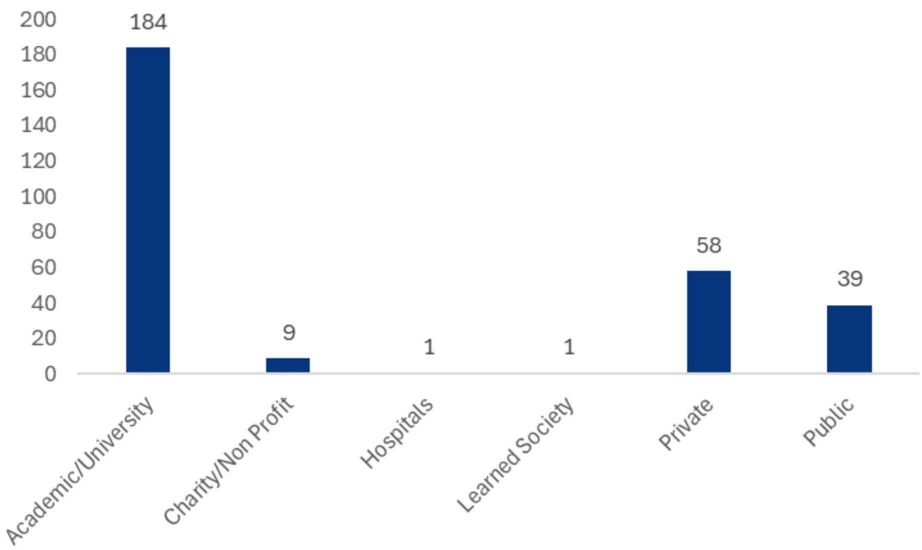

Note: Based on 157 collaborations on ARCHER2 established between January 2021 and June 2025.

Source: London Economics analysis of ResearchFish™ data

Survey findings on scale of collaborative R&D on ARCHER2

To understand the potential scale of underestimation within UKRI's research outcome data, ARCHER2 users were also asked about their involvement in collaborations involving ARCHER2 in the user survey undertaken.

The results indicate that, a substantial proportion of ARCHER2 users engage in collaborative R&D. Among respondents to the user survey, **nine in ten (91%) users reported that they were involved in collaborations as part of their work on ARCHER2**.

Of these, 84% of respondents have been involved in at least one collaboration with other UK-based researchers. 54% of respondents also reported being part of academic collaborations with researchers internationally.

Moreover, ARCHER2 users report collaborating not only with fellow academics both domestically and around the world, but also with partners across industry (36% of respondents), the public sector (15% of respondents), as well as charities and other organisations. (Figure 20)

These findings highlight that collaboration is a widespread and integral aspect of ARCHER2-supported R&D.

"Without ARCHER2, we would have to try and collaborate with an HPC centre abroad, which almost certainly would require us to compromise on the efficacy and safety of our projects." (discussing collaborations with NGO's on ARCHER2)

"Particularly engaging with other academics about the Cray systems, and being able to provide support, was valuable as a collaborator."

"Some of things we have been able to do with collaborators would simply not have been possible (without ARCHER2)."

"Using ARCHER2 to generate data has allowed for many fruitful collaborations."

- **ARCHER2 user survey**

Figure 20 Proportion of user survey respondents involved in collaborations

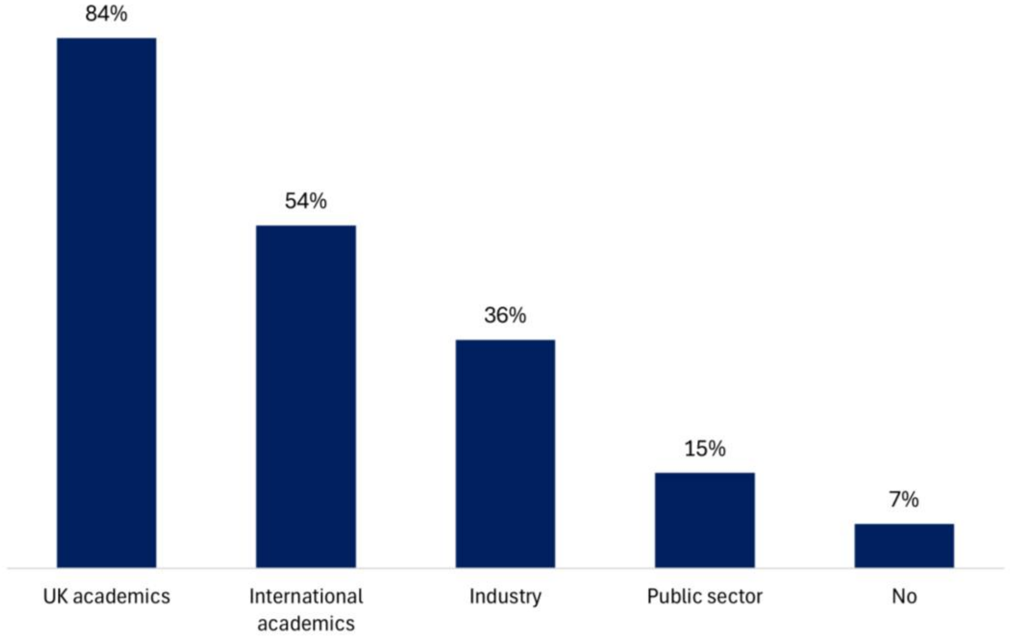

Note: Based on 200 respondents. 167 respondents said they had been involved in a collaboration with UK academics. Percentages do not sum to 100% as users could select multiple answers. 3% answered 'don't know / not applicable', which is not shown in the graph. Question asked: 'As part of your ARCHER2 work were you part of any collaborations with industry, the public sector or other academics (in the UK or abroad)?'

Source: London Economics ARCHER2 user survey

Author-affiliations of ARCHER2-related publications

In addition, to the user survey author-affiliation of ARCHER2-related publications were also explored. The **1,342 ARCHER2-related publications** published in 2022 and 2023 that could be matched to bibliometric databases include **9,495 named authors (i.e. 7 authors per paper on average).**

Most of these authors (9,119 of the 9,495) have an institutional affiliation listed (in fact some are listed against multiple institutions – in which case we have taken the first listed for the analysis). Authors on ARCHER2-related publications are affiliated to **1,177 different institutions, spread across 88 different countries.** This includes, most commonly the following institutions (each with 100+ authorships on ARCHER2-related publications):

- Imperial College London (538)
- University of Oxford (373)
- University College London (372)
- University of Manchester (301)
- University of Leeds (270)
- University of Warwick (219)
- University of Cambridge (211)
- University of Edinburgh (174)
- University of Bristol (172)
- Cardiff University (154)
- National Centre for Earth Observation (154)
- University of Reading (141)
- University of Southampton (124)
- The Faraday Institution (109)

Geographic breakdown of collaborations

Of the more than 292 collaborations captured in ResearchFish™, 53% were with partners in the United Kingdom, 12% were with the United States, 6% with France and 4% with Germany. Figure 21 provides a geographical breakdown of the number of collaborations, by country, benefitting from ARCHER2 since 2021.

Figure 21 Collaborations on ARCHER2 by country

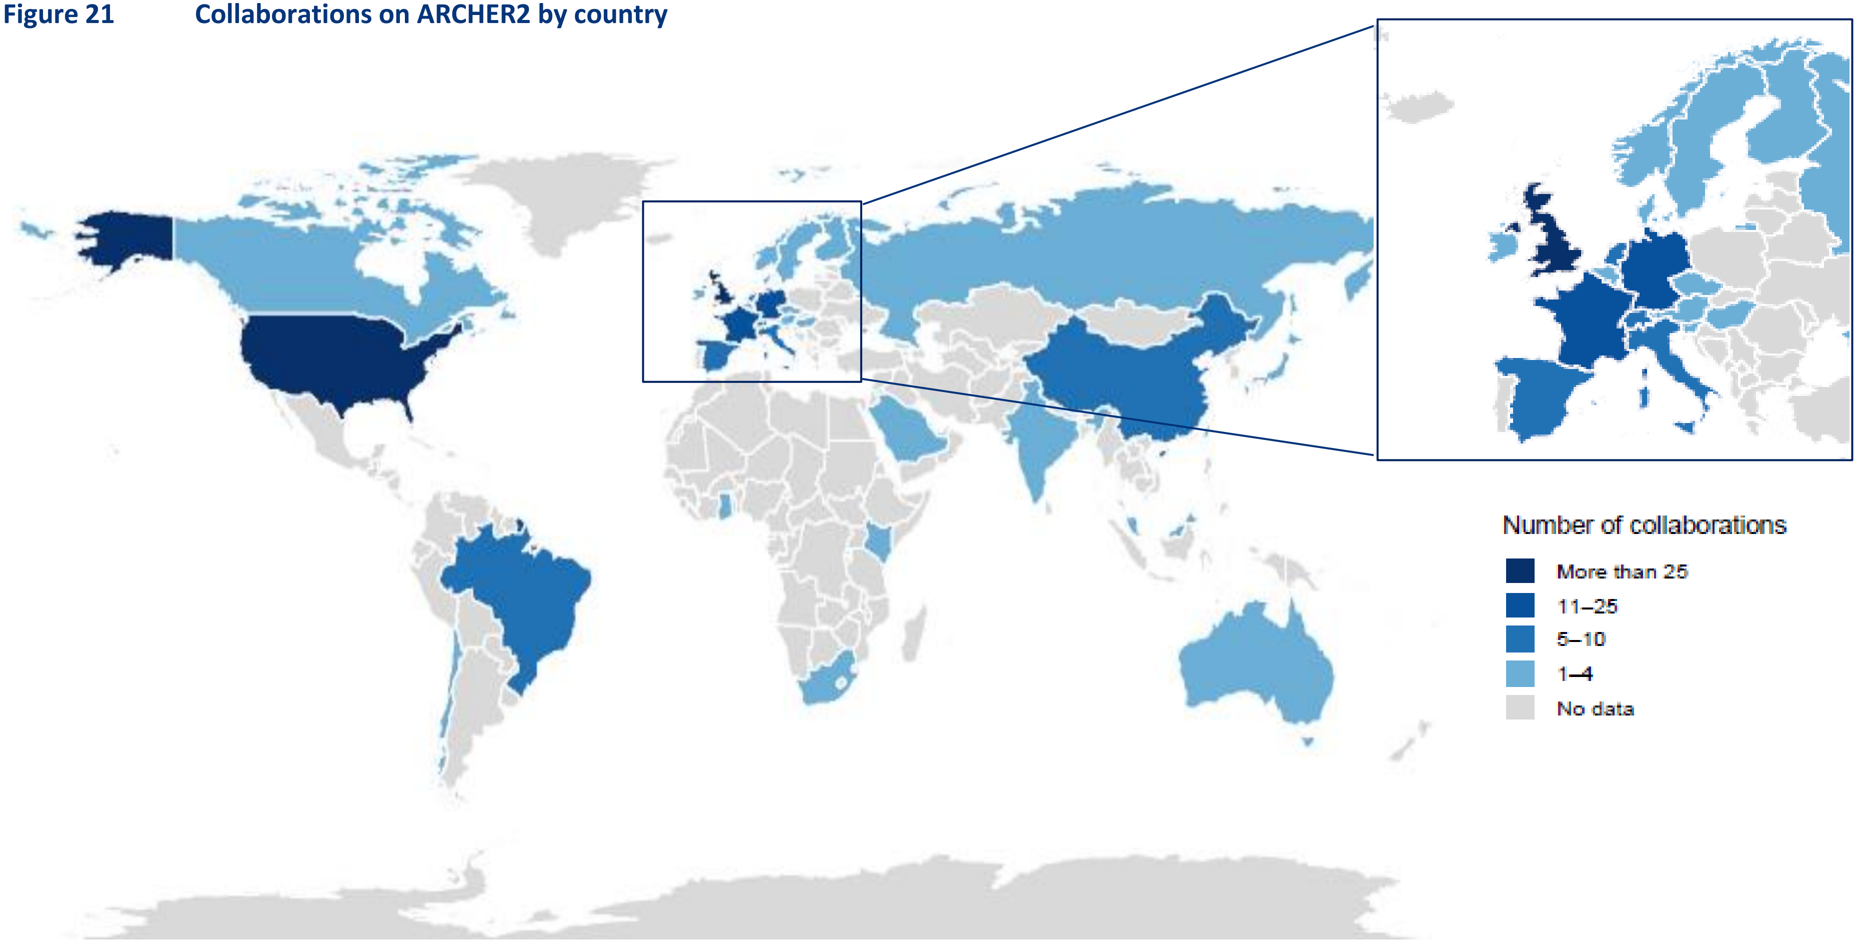

Note: This analysis is based on EPSRC ResearchFish™ collaborations data, so it doesn't include collaborations involving NERC.

Source: London Economics analysis of ResearchFish™ data

3.2.2 ARCHER2's role in supporting and enabling collaborative R&D

Academic stakeholders consulted as part of this R&D frequently highlighted ARCHER2's central role in enabling high-profile international collaboration through raising the UK's profile and providing researchers access to compute capacity of the level of a national Tier-1 facility, thereby making them more attractive partners.

The UK's participation in the Nucleus for European Modelling of the Ocean (NEMO) Consortium is a high-profile example of this. NEMO is a state-of-the-art modelling framework used for oceanic research, ocean forecasts and climate studies. NEMO comprises five EU institutions, including the UK Met Office (UKMO) and the UK's National Oceanographic Centre (NERC-NOC), who ensure the sustainable development of NEMO to support user needs. ARCHER2 provides a key underpinning capability for the UK's contribution to NEMO by enhancing the value the UK can bring to the consortium through access to Tier-1 level HPC capabilities.

Box 6 NEMO (Nucleus for European Modelling of the Ocean)

UK researchers are using ARCHER2 to run high-resolution ocean simulations that are providing insights to support coastal resilience and climate change adaptation.

Ocean modelling is a crucial tool for understanding and addressing climate change, safeguarding marine ecosystems, and anticipating extreme weather events. UK researchers are using the NEMO (Nucleus for European Modelling of the Ocean) framework to simulate the ocean's physical and biological processes, including sea level change, circulation patterns, and marine resource dynamics i.e. how ocean life and composition change over time. These complex simulations are powered by ARCHER2, the UK's national supercomputing service. This effort has directly led to upgrades in the UK Met Office's operational numerical weather forecasting system.

Further details can be found in the full case study in Annex A5.6.

Source: Technopolis analysis based on information gathered via desk research and consultation with Andrew Coward and Jeff Polton.

This is also reflected in the responses received to the user survey. **73% of researchers who participated in collaborations on ARCHER2 agreed that these collaborations would not have been possible without ARCHER2 access, with just under half strongly agreeing**. There was also broad agreement among survey respondents that ARCHER2 facilitates these collaborations (77% of respondents), and that ARCHER2 was a key facilitator of academic knowledge exchange (74%). (Figure 22).

There was slightly less agreement around ARCHER2's role in increasing the number of collaborations with other HPC researchers internationally (with 32% neither agreeing nor disagreeing) and that ARCHER2 encouraged users to collaborate (with 30% neither agreeing nor disagreeing). Though, only a small proportion of respondents actively disagreed with these statements.

Figure 22 Benefits to knowledge exchange and collaboration experienced by users

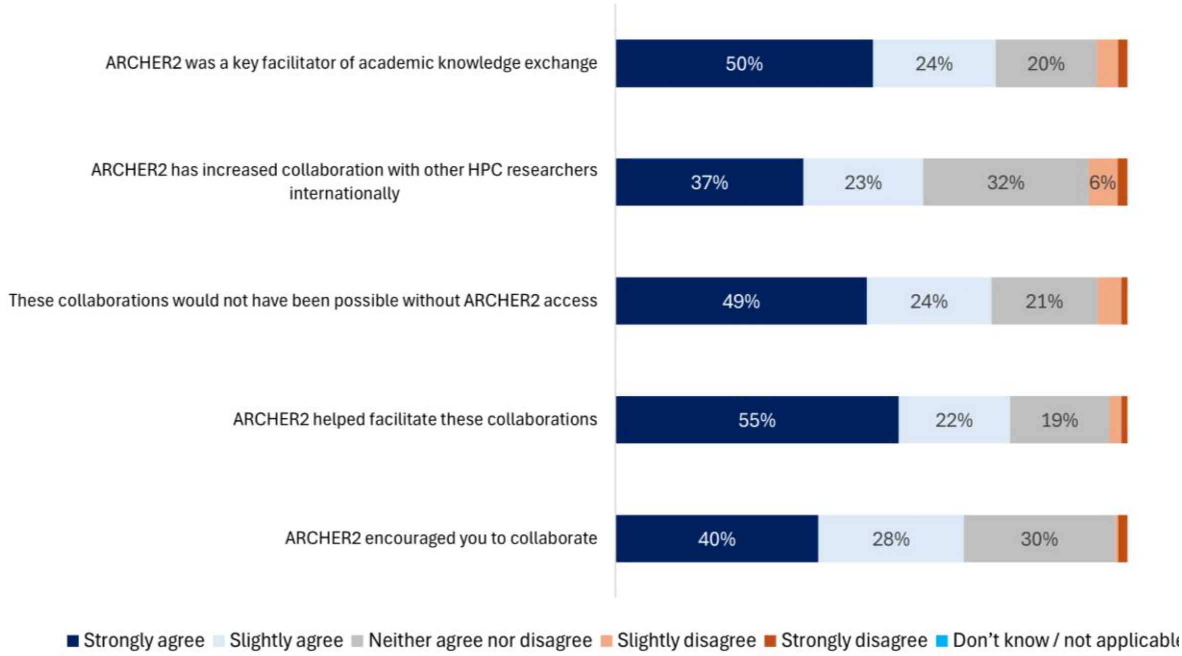

Note: Based on between 158 and 173 responses, depending on the question. Survey participants were asked this question if they answered that they had engaged in collaborations on ARCHER2. Question asked: 'To what extent do you agree or disagree that …?'

Source: London Economics ARCHER2 user survey

3.3 Maintaining world class software code, improved data management

Modern computational science depends not just on powerful compute hardware, but also on the software that unlocks its potential. High-performance computing systems like ARCHER2 can only deliver maximum value for the scientific community when they are supported by well-optimised and scalable software that fully harnesses the available compute capacity.

By maximising the potential of the available hardware, optimised software codes enable researchers to undertake their work in less time. This frees up computational resources for others, allows more detailed and granular investigations, supports more complex or ambitious simulations, and helps reduce energy consumption, thereby making scientific R&D more efficient and sustainable.

"The access to the software environment is also very useful, being able to test with the Cray compilers enables a broader range of testing which is important for research software quality. They additionally have some nice features that are not available elsewhere that can assist in ongoing development."

- **ARCHER2 user survey**

In addition to software, efficient data management also plays a key role in fully leveraging high-performance computing hardware. Well-designed data management architecture and processes reduce data latency and improve input/output efficiency, thereby ensuring that compute resources spend more time processing and less time waiting on data transfers.

The UK has a strong legacy in computer science and software engineering, with recognised strengths in research software development and parallel computing. Continued investment in developing and refining HPC software ensures that the UK research community remains competitive and that public infrastructure like ARCHER2 delivers maximum scientific and economic return.

Embedded CSE (eCSE) Programme

To support the research community to develop and optimise software in a sustainable manner, EPCC run regular Embedded CSE (eCSE) support calls. The eCSE programme provides funding to researchers to develop software on ARCHER2 that advances research aligned with UKRI's digital research infrastructure priorities.[18]

eCSE support is future-focused, helping researchers implement new features and enhance software performance not only for current computing architectures but also to ensure readiness for future Exascale supercomputing services and UK national AI platforms, as well as to ensure software works efficiently on national Tier-2 HPC facilities.

A total of **93** eCSE projects have been supported on ARCHER2 to date, funding a total of **1,114 project months** across nearly **35 institutions**. Of these, **12 projects** were specifically **aimed at optimising software for GPU-based systems**, helping to ensure these codes efficiently utilise these increasingly common compute architectures.

Efficiency improvements from eCSE funded software are not currently quantified, due to the time-effort involved. However, for historical eCSE projects leading up to 2018, the benefit from efficiency improvements, in terms of time freed up on ARCHER2's predecessor, at the time was estimated at £4 for every £1 invested in the programme. This highlights the substantial value and return that software optimisation efforts can deliver to the HPC ecosystem.[19]

Community software development efforts

Alongside funded software development projects, the ARCHER2 community also engages in a broad range of independent software development activities. The Computational Science Centre for Researc Communities (CoSeC) supports collaborative computational research communities that are funded across UKRI. These communities use multiple structures such as High-End Computing (HEC) Consortia and play a significant role, among other things, for the Development of theory, algorithms, and research software: This is a key driver for computational communities, resulting in long-term, continued expansion and updating of the UK's research software capability.[20]

UKRI's research outcome database, ResearchFish™, contained information on development efforts by the research community relating to a total of 151 software and other technical products. (Figure 23)

Moreover, in response to the user survey, 25% of respondents said they had developed optimised software for improved efficiency on HPC as part of their R&D on ARCHER2, and 25% developed other software which provides new functionality (Figure 24).

[18] See https://www.archer2.ac.uk/ecse/ for more information.
[19] See London Economics (2019). The impact of EPSRC's investments in High Performance Computing infrastructure. Available at: https://www.ukri.org/wp-content/uploads/2022/07/EPSRC-050722-ImpactEPSRCInvestmentsHighPerformanceComputingInfrastructure.pdf [Accessed 10/06/2025]
[20] See https://www.cosec.ac.uk/what-is-cosec/ for more information.

Figure 23 Software and technical products developed on ARCHER2

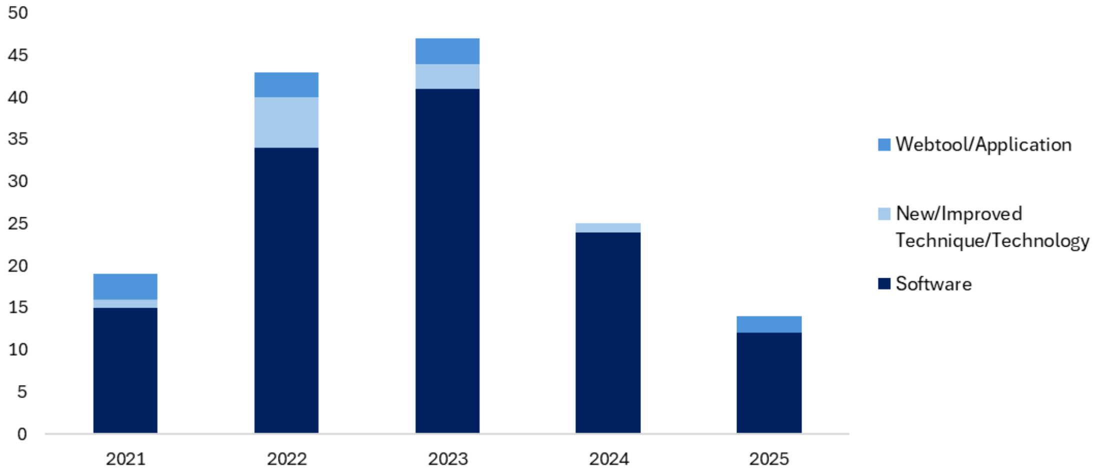

Note: There were a total of 151 technical products/software reported to EPSRC from 2021 to 2025. 126 of these were software products. This analysis excludes one grid application product reported and two physical model/kits reported.

Source: London Economics based on ResearchFish™ data provided by EPSRC

Figure 24 Share of users reporting work on ARCHER2 leading to the creation of software products, formation of spin-outs, or intellectual property

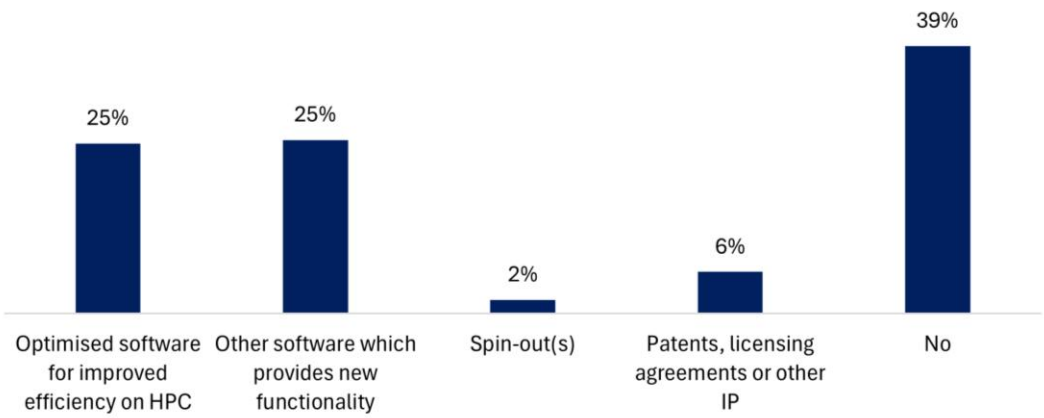

Note: Based on 197 respondents. Percentages do not sum to 100% as users could select multiple answers. 26% answered 'don't know / not applicable', which is not shown in the graph. Question asked: 'Has your work on ARCHER2 led to the creation of any software products, formation of spin-outs, or intellectual property?'

Source: London Economics ARCHER2 user survey

Knock-on benefits to the wider HPC community

Software developed on ARCHER2 not only enables more efficient usage of ARCHER2 itself but also brings benefits to the wider HPC ecosystem.

Many of the software products developed by the ARCHER2 research community are published on an open source or free-at-the point of use basis, enabling other researchers to utilise the software on other HPC systems and benefit from improvements made.

Indeed, nearly two-fifths (39%) of ARCHER2 users responding to the survey have used software developed on ARCHER2 on other HPC systems themselves. This includes other UK HPC systems such as Young, CSD3, CIRRUS, Isambard3, JASMINE as well as international HPC facilities.

Figure 25 Share of users that have used software developed on ARCHER2 on other HPC systems

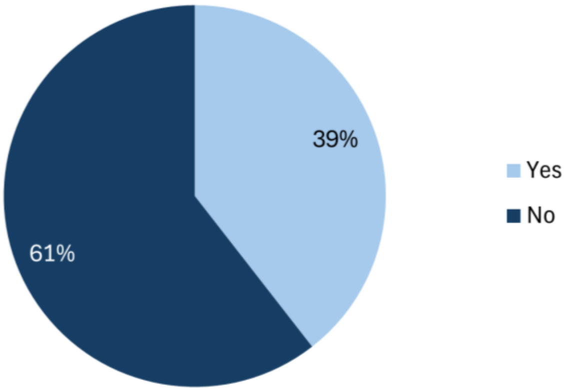

Note: Based on 152 respondents. Question asked: 'Have you used software developed on ARCHER2 on any other HPC systems? E.g. by using the newly scaled HPC simulation code after being developed on ARCHER2?'

Source: London Economics ARCHER2 User Survey

The importance of ARCHER2 for software development efforts

To understand the importance of ARCHER2 for software development efforts, users who developed software on ARCHER2 were asked to what extent being able to access ARCHER2 has contributed to the development of these software products.

The vast majority (88%) of users developing software reported that ARCHER2 had contributed to these development efforts to some extent. Three in ten (30%) users reported that ARCHER2 had made either a small or medium contribution, while 33% said ARCHER2 had made a large contribution. Furthermore, more than **one in four (28%) said these software products would have been impossible to develop without access to ARCHER2.** (Figure 26)

Figure 26 Contribution of ARCHER2 to the development of software products

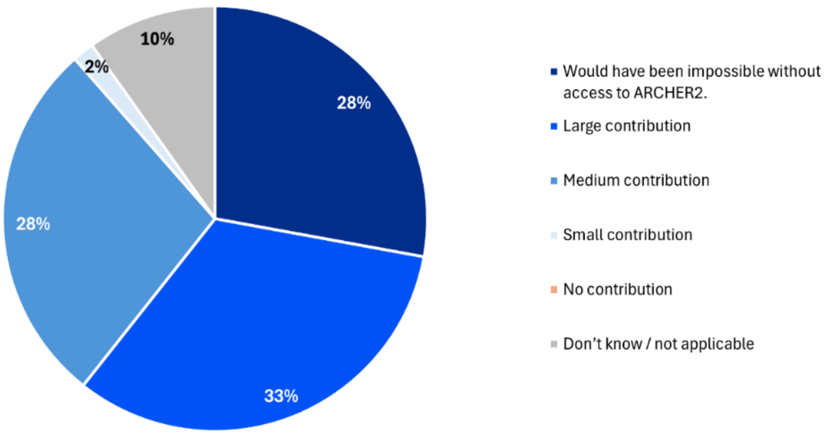

Note: Based on 61 respondents who responded that their work on ARCHER2 led to the creation of software products. Question asked: 'To what extent has being able to access ARCHER2 contributed to the development of these software products?'

London Economics ARCHER2 user survey

Box 7 VAMPIRE (Versatile Atomistic Monte Carlo, Integrator, and RElaxation)

Atomic simulations of magnetic materials developed using ARCHER2 have led to advanced understanding of materials and sparked innovation in industry.

Richard Evans, researcher at the University of York, developed VAMPIRE, an open-source software that simulates atomic behaviour in magnetic materials. Due to the high cost of materials and manufacturing, these computational simulations offer a more affordable alternative for research and development.

The simulations have provided unprecedented insights into interactions and properties of materials at the atomic scale. The software has driven innovation in industry such as the development of Heat-Assisted Magnetic Recording (HAMR), a new technology which boosts data storage capacity from 1TB to 30TB. VAMPIRE also supports other emerging applications in electronics and in healthcare.

The potential for wide industry use and impact is high, as the software continues to attract a growing user base- currently around 500-600 users across the UK. The software also strengthens the UK's position as a leader in scientific research and innovation.

Further details can be found in the full case study in Annex A5.7.

Source: Technopolis analysis based on information gathered via desk research and consultation with Richard Evans from the University of York.

4 ARCHER2's role in supporting a highly skilled community

4.1 HPC & computational skills training provision

Beyond its primary function as a computational resource, ARCHER2 also plays a crucial role in the development of advanced computational skills within the research community.

Without people who possess the right skills and knowledge to exploit the HPC hardware, much of the benefits of HPC would not be realised. Specialised HPC training events, including advanced computational techniques and programming languages applicable across different systems is important to addressing the UK's existing computing skills gap[21]. ARCHER2 plays a key role in this by providing a variety of training programs designed to equip researchers with the technical skills they need.

Direct training, alongside hands-on use of ARCHER2, by students and postdoctoral researchers ensures that the next generation of scientists is well-equipped with the necessary skills to leverage HPC and other advanced computational capabilities in their work and future career. In addition to training of early-career researchers, these training events also support established researchers and other users such as computational research software engineers in further enhancing their skills.

[21] The Lloyds Bank Consumer Digital Index 2024, commissioned by the Department for Education, found that 18% of UK adults lacked the essential digital skills that are needed for the workplace. See: https://www.lloydsbank.com/consumer-digital-index.html?srnum=4

Direct training provision

Distinct training pathways are provided for researchers, data scientists, and research software engineers / developers. These training initiatives provide hands-on experience and practical knowledge, enabling researchers to effectively utilise HPC resources, and ensuring each type of user receives tailored training to their objectives and skill level.

Each pathway offers training courses ranging from fundamental skills, such as how to connect to remote HPC systems and transfer data, and basic computing skills like program design, version control, testing, and task automation, to intermediate and advanced level courses. The latter include a range of advanced HPC and wider important computational skills including parallel computing, GPU programming, machine learning, and many more.[22]

Figure 27 provides an overview of the ARCHER2 training courses provided and the different pathways.

Figure 27 ARCHER2 training courses provided by EPCC

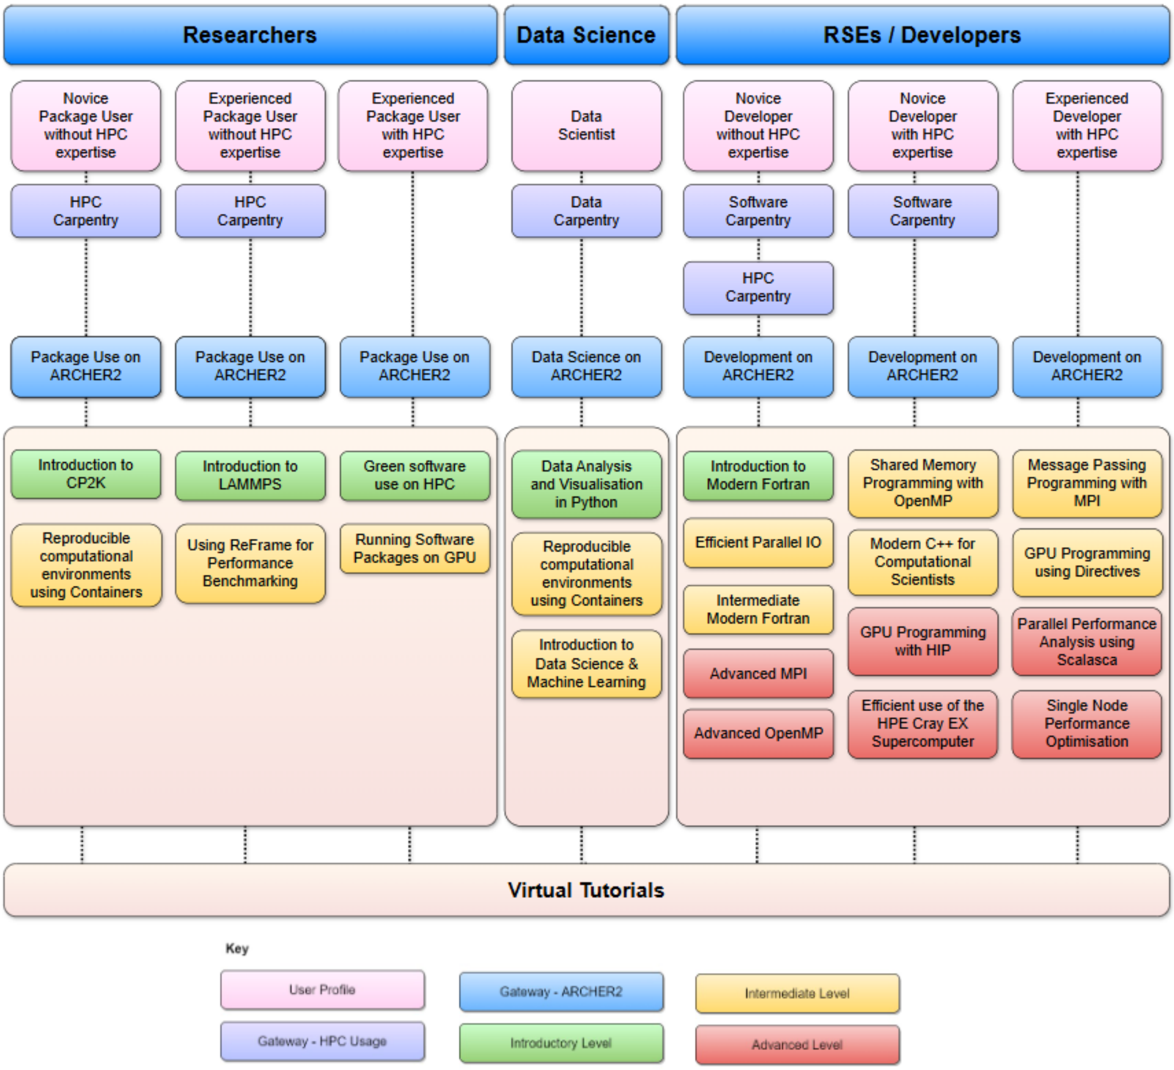

Source: EPCC: https://www.archer2.ac.uk/training/courses/

[22] An overview of training offerings on ARCHER2 can be found at: https://www.archer2.ac.uk/training/courses/

In total **145 training events** were provided by EPCC, who host ARCHER2, between 2020, when the early ARCHER2 four-cabinet system was installed, and mid-2025. Across these training events, a total of **270 training days** were provided with **around 2,800 attendees** (~2,000 unique individuals) recorded across all training events[23], and an average attendance of 19 individuals per training course.

In addition, EPCC also provides various virtual tutorials and webinars covering a wide range of topics and levels to support users learning. This includes talks on research using ARCHER2 and HPC in general, technical talks of interest to users as well as more general talks on areas such as diversity and inclusion.[24]

Figure 28 Number of ARCHER2-related training days provided by EPCC, by year

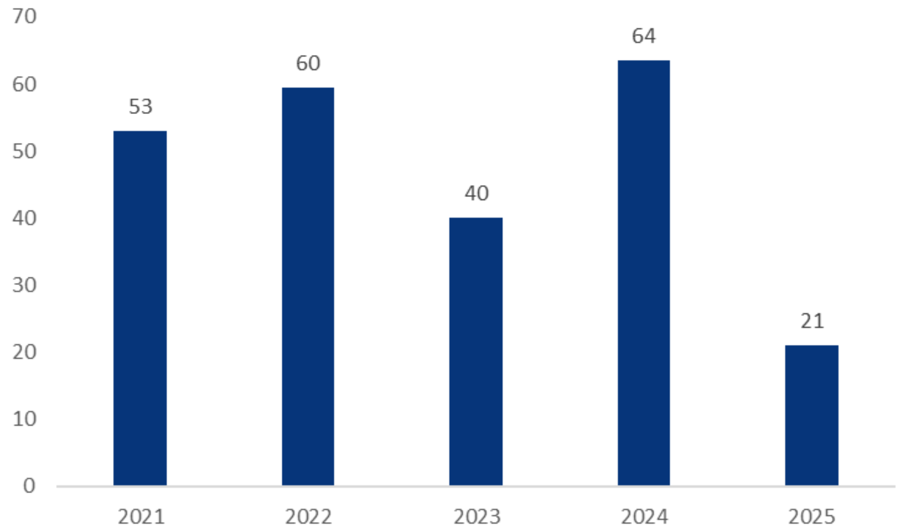

Source: London Economics analysis of EPCC data

Hands-on training of students and postdoctoral researchers

In addition to direct training, **more than 2,650 students and postdoctoral researchers have used ARCHER2 to date**, gaining hands-on experience in advanced computational R&D.

Hands-on HPC usage, in combination with the training offerings discussed in the previous section as well as guidance and support provided by their principal investigators, the HEC consortia and ARCHER2 team, helps prepare the next generation of researchers with essential computational skills for academic careers. This in turn helps to sustain the UK's position as a global leader of scientific R&D.

In addition, individuals trained in HPC are in high demand across industry, contributing valuable expertise to businesses, supporting economic growth, and generating returns for the UK exchequer.

[23] Note, researchers may attend multiple training events and so this figure does not equate to 2,750 distinct individuals.
[24] See https://www.archer2.ac.uk/training/courses/

Box 8 Supercharging Futures: How ARCHER2 is Empowering Early Career Researchers at Imperial College London

Access to ARCHER2 has equipped early-career researchers at Imperial College London with advanced computational skills leading to cutting-edge biomedical research and prestigious academic PhD opportunities. The advanced simulations of complex biological systems developed by Sarah Rouse's team have driven innovation in drug delivery and gene therapy.

ARCHER2 has provided new opportunities to simulate biological systems more efficiently and realistically. The hands-on experience of working with ARCHER2 has been especially valuable for PhD and MRes students, who are directly contributing to cutting-edge computational research. A key example is a PhD and MRes researcher team in Rouse's group who used ARCHER2 to develop a molecular simulation pipeline for AstraZeneca in the field of gene therapy—a discipline with applications in treating conditions like cystic fibrosis and cancer. The simulations investigated how adeno-associated virus (AAV) packaging mechanisms influence therapeutic efficiency, helping to address a critical manufacturing bottleneck. This work not only advanced industrial collaboration but also enabled the researchers to acquire deep technical expertise in molecular modelling and high-performance simulation. In the absence of ARCHER2, these simulations would need to be scaled down and students would miss out on realistic system complexity.

For postgraduate students, ARCHER2 has served as a platform to build confidence and capability in computational research. Former MRes student Pao Pipatpadungsin reflected on the importance of ARCHER2:

> *"My MRes at Imperial provided a key opportunity to cultivate in-demand computational skills by using the ARCHER2 supercomputer. The work produced meaningful and interesting results, leading to my first publication and demonstrating my ability to generate original research. This accomplishment was a significant factor in my successful application for a DPhil at the University of Oxford."*

By working on real-world, high-impact projects with direct industrial relevance, students have gained experience in using supercomputing infrastructure that is typically inaccessible at the early stages of a research career. This access has equipped them with highly transferable skills in simulation, data analysis, and scientific computing—capabilities that are increasingly in demand across academia, biotech, and the pharmaceutical sector.

Further details can be found in the full case study in Annex A5.8.

Note: (i) See https://www.imperial.ac.uk/news/245036/new-partnership-will-molecular-glues-tackle/.

Source: London Economics analysis based on information gathered via desk research and consultation with Dr. Sarah Rouse

Figure 29 Student and postgraduate researchers using ARCHER2, active accounts 2021-2025

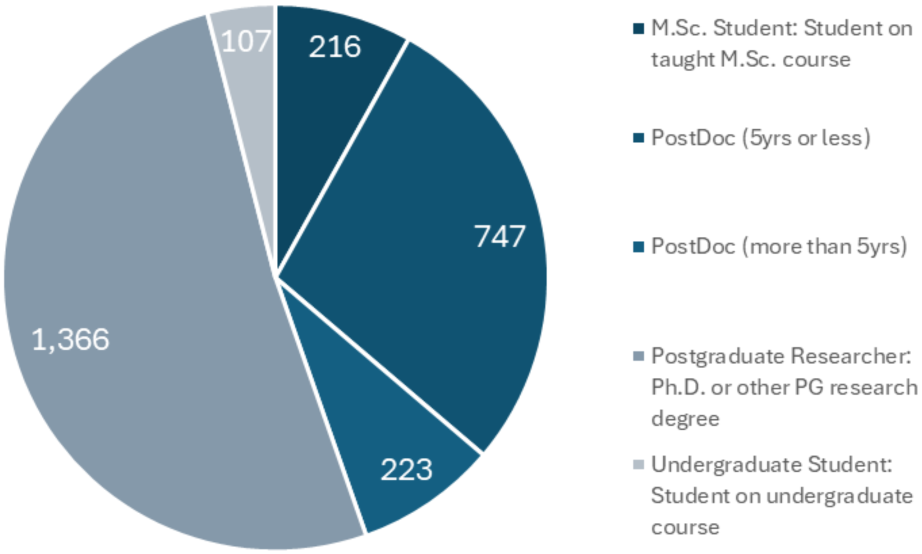

Source: London Economics analysis of EPCC usage data

4.2 The impact of training on skills development and career paths

Advanced computational skills are in high demand across various sectors. This is reflected in the high employability rates of and high wages received by employees with such skills.

Stakeholder feedback on the value of ARCHER2 training

The value of ARCHER2 training was frequently emphasised in both the user survey and stakeholder consultations. This included both the quality and value of the training courses provided by EPCC as well as the value to students and early-career researchers of using ARCHER2 in their research. Stakeholders emphasised that ARCHER2 training improved their HPC and wider computational skills and abilities, which in turn improves the quality of their research and benefits their professional development and future career-paths.

To help understand the impact of ARCHER2 usage on skills development and career paths, the user survey asked respondents to what extent they agree or disagree with a range of statements about the impact of ARCHER2 usage on their own skills development and career progression. Further, to assess the impact of ARCHER2 training on early-career researchers' career progression, the user survey also asked researchers who had taught students and postdoctoral researchers on ARCHER2 about the impact of using ARCHER2 on their student's career path and salary prospects.

Among respondents asked to rate the impact of ARCHER2 use on their own skills development and career progression (Figure 30):

- 86% agreed that their use of ARCHER2 improved their computational and technical skills,
- 88% agreed that it provided them with skills and knowledge that are of benefit in their current area of work, and
- 88% agreed that it provided them with skills and knowledge that are useful across a range of jobs and industries, making them more attractive to potential employers
- 89% agreed that usage of ARCHER2 better prepared them for a career in academia.

While there is slightly less agreement that ARCHER2 better prepares users for a career in industry, with 67% agreeing that this was the case, no respondent actively disagreed with the statement. Moreover, a large proportion (85%) of researchers training students and early-career researchers also felt that ARCHER2 usage better equips their students for a career in industry (Figure 31).

Further, 64% of researchers training students and early-career researchers agreed that ARCHER2 training enabled their students to earn higher salaries in industry, and 97% agreed that it better equipped their students for a career in academia.

Figure 30 Benefits to skills development and career progression experienced by users

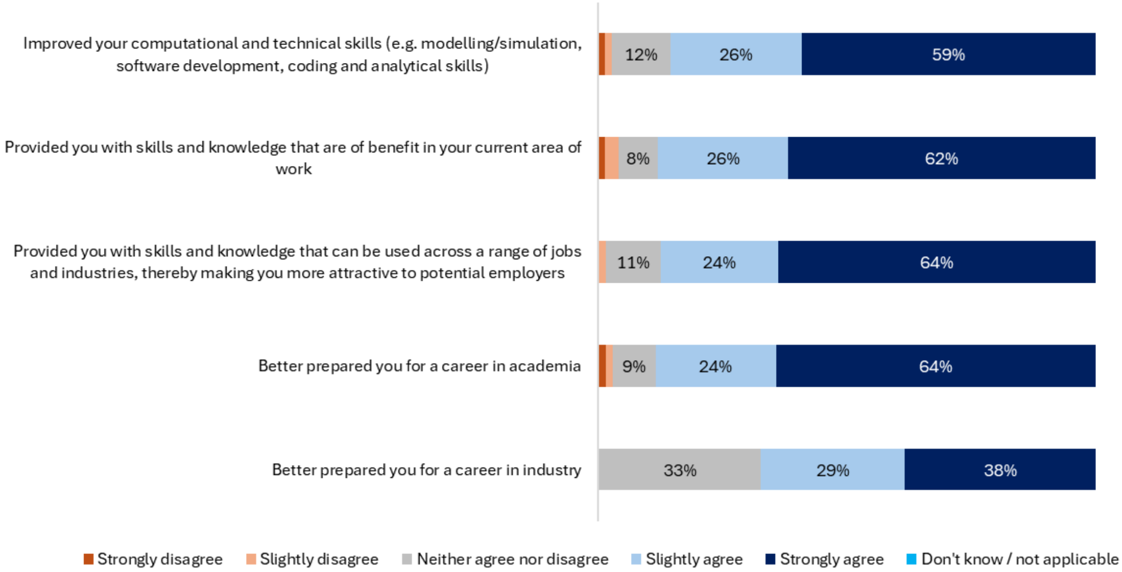

Note: Based on between 52 to 76 responses. Question asked: 'In terms of skills development and career progression, to what extent do you agree or disagree that your use of ARCHER2 …?'

Source: London Economics ARCHER2 User Survey

Figure 31 Perceived impact of ARCHER2 on student career readiness and salary prospects

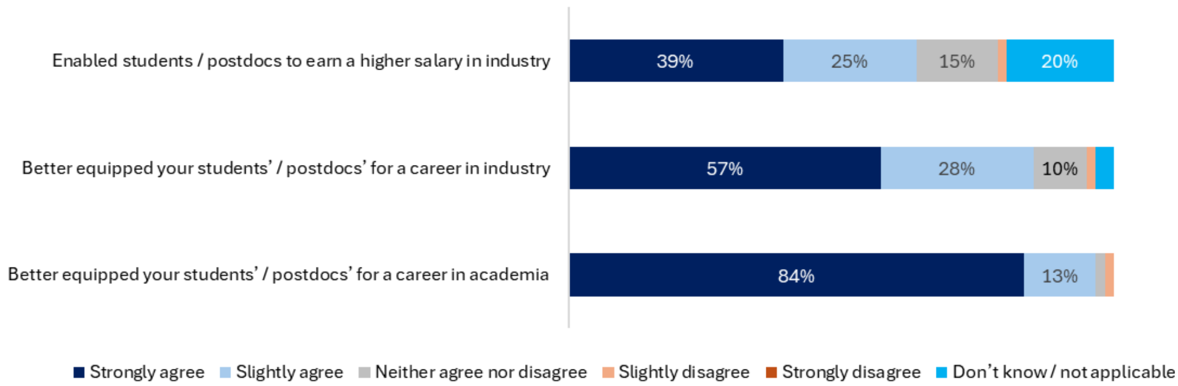

Note: Based on 61 responses. Question asked: 'To what extent do you agree that training on ARCHER2 has…? '

Source: London Economics ARCHER2 User Survey

"ARCHER2 offers regular training courses and online resources on MPI, OpenMP, performance tuning, and scientific workflows, which are crucial for improving code efficiency and scalability."

"I have benefited immensely from ARCHER2's training programs. These courses have equipped me with the skills needed to use the system efficiently and professionally."

"ARCHER2's software support and training resources help optimise performance and improve my modelling skills, directly benefiting my PhD research outcomes"

"They (EPCC) offer training courses such as for LAMMPS, which have been useful for myself as well as PhD students in our group. Their guidance on using the HPC is very well explained, and is a great source of support, as a lot of the information extends to general HPC use."

"I have found the access to software training invaluable as I have learned useful software development skills, particularly developing parallel programs. The supercomputing resources of ARCHER2 make the UK a major player in the field of materials science."

"ARCHER2 also provides an important training ground for our PhD students in terms of HPC skills."

- **ARCHER2 user survey**

Destinations of students and postdoctoral researchers trained on ARCHER2

Training students and early-career researchers, and supporting the development of highly demanded computational skills, directly benefits both UK science and industry. Insights from researchers who trained students or postdoctoral researchers on ARCHER2 indicate that the majority pursue careers in academia or the private sector, with around four in ten (44%) pursuing each of these paths, respectively. A smaller proportion, around 6% each, enter the public sector or other areas. In terms of geographic destination, a large proportion (around 70%) of students trained on ARCHER2 remain in the UK, while around 30% pursue careers overseas.

Figure 32 Destination of students by sector

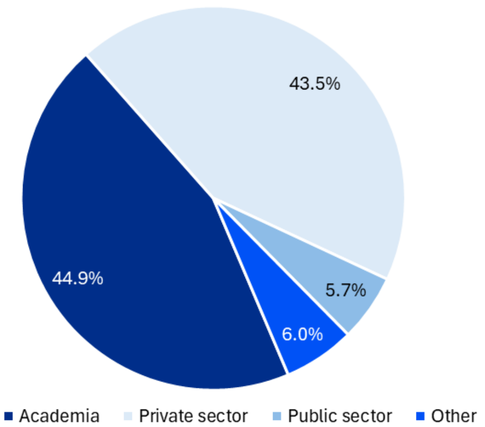

Note: Based on 283 responses. Question asked: 'How many of these stayed in academia, how many moved outside of academia?'

Source: London Economics ARCHER2 User Survey

Figure 33 Share of students remaining in the UK

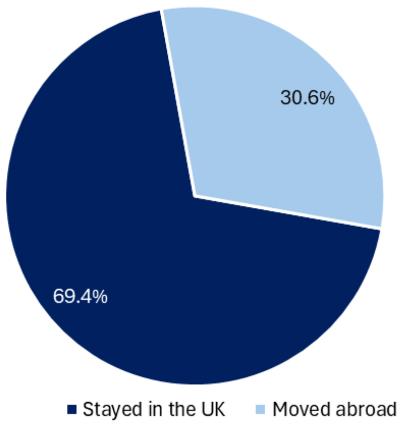

Note: Based on 232 responses. Question asked: 'How many of these stayed in the UK, how many moved abroad?'

Source: London Economics ARCHER2 User Survey

4.3 Contribution to international standing and competitiveness of UK science and attraction of inward talent

Contribution to international standing and competitiveness of UK science

A strong national HPC ecosystem, encompassing high-quality infrastructure such as ARCHER2 and the new AI Research Resource (AIRR), recognised expertise in parallel computing and HPC software development, and comprehensive training opportunities, all contribute to the UK's international standing in HPC.

Computational modelling and complex simulations, enabled by HPC, play a critical role in modern scientific R&D. As highlighted previously, they are now integral to many disciplines, with advanced computing enabling researchers to explore complex systems, accelerate discovery, and test hypotheses at scales and speeds that would be impossible through traditional methods alone.

> *"A national supercomputer is required to ensure the UK can maintain international standing in broader HPC, not just AI."*
>
> *"Without ARCHER2 or a better class HPC we will be unable to compete with researchers in other countries, such as US, China, Japan, Germany in the field of multiscale computational materials modelling and simulation."*
>
> - **ARCHER2 user survey**

For the UK, maintaining world class national HPC infrastructure is vital to staying at the forefront of global science and technology, driving breakthroughs in areas such as climate modelling, materials modelling and biomolecular simulations. Given the interdependence between theory, experiments and computational modelling in modern research, investment in High Performance Computing plays a key role in helping UK science remain competitive internationally. This is reflected in stakeholder feedback from the user survey, where 89% of respondents strongly agreed that ARCHER2 helps make UK science more internationally competitive. (Figure 34)

Figure 34 Perceived contribution to UK science competitiveness

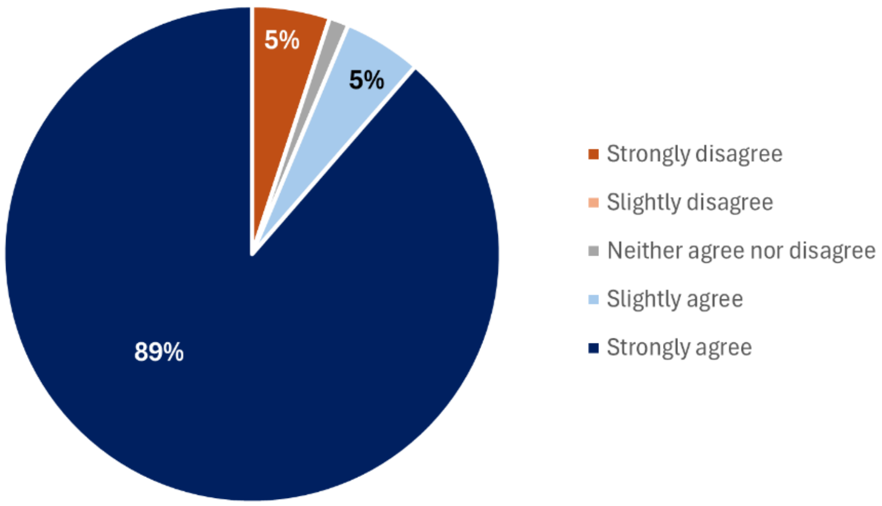

Note: Based on 79 respondents. Question asked: 'To what extent do you agree or disagree that ARCHER2 helps make UK science more competitive internationally?'

Source: London Economics survey of ARCHER2 users

However, stakeholders also emphasised that the UK has now fallen behind other countries in terms of HPC capabilities. Many academics and consortia heads emphasised that they would benefit from additional time than is available on ARCHER2, and that the UK HPC market requires significant and continued investment. This mirrors findings of the recent 2023 Future of Compute Review[25], which emphasised the need for significant and strategic investments in large-scale UK compute infrastructure. (See further discussion in Section 8)

Attraction of inward talent

An internationally competitive UK HPC ecosystem alongside strong reputation in computational modelling and excellent training opportunities, in turn contributes to attracting top scientific talent to the UK.

As highlighted previously, a significant proportion of users accessing ARCHER2 are students and early-career researchers. Access to national Tier-1 level compute resources, combined with high-quality training opportunities and a strong wider ecosystem, enables students and early-career researchers to gain hands-on experience with large-scale computing at national level. It allows them to develop skills in parallel programming and data handling, and to undertake research at a scale that would not be possible with local resources alone early in their careers.

For universities and research groups, being able to offer access to national-level compute resources enables them to strengthen their offerings to students and early-career researchers. This includes both offerings in disciplines that heavily rely on computational methods as well as programmes focused on computational training. A prominent example of the latter is the MSc in High Performance Computing with Data Science at the University of Edinburgh, which is delivered through EPCC and provides students with hands-on experience on ARCHER2 as part of the programme.

The user survey shows that, among survey respondents who relocated to the UK (61% of the 199 respondents), ARCHER2 played an important or key role in their decision to move to the UK for almost two-fifths (37%) of these respondents. A further 21% said ARCHER2 played some role in their decision. Further, among all respondents to the survey, 86% agreed that ARCHER2 helps attract researchers to undertake their work at UK Universities, and 81% agree that ARCHER2 helps attract international students to study at UK Universities. This highlights the value of ARCHER2 not only as a research tool, but as an important factor in attracting talent to the UK research and higher education system.

[25] DSIT (2023). Independent Review of the Future of Compute: Final report and recommendations. Available at: https://www.gov.uk/government/publications/future-of-compute-review/the-future-of-compute-report-of-the-review-of-independent-panel-of-experts [accessed 21/05/2025]

Figure 35 Extent to which ARCHER2 played a role in users decision to study, work or undertake research in the UK

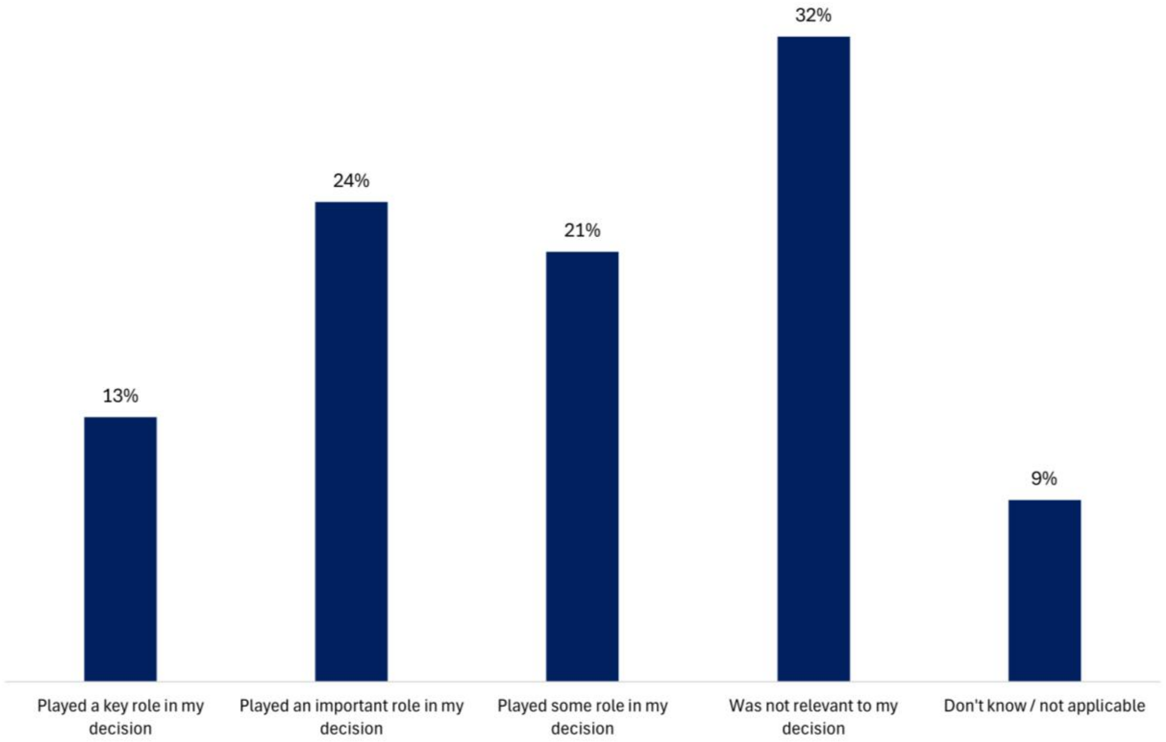

Note: Based on 121 respondents. Question asked: 'To what extent did the availability of ARCHER2 play a role in your decision to study, work, or undertake your research in the UK …?''

Source: London Economics survey of ARCHER2 users

Figure 36 Perceived role of ARCHER2 in attracting international talent

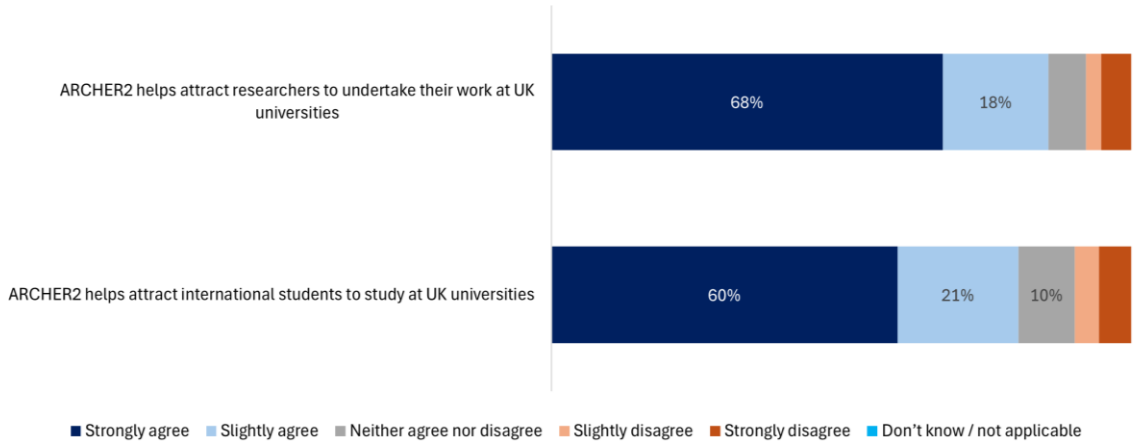

Note: Based on 72 and 77 respondents for attracting international students and researchers respectively. Question asked: 'To what extent do you agree or disagree that …?'

Source: London Economics survey of ARCHER2 users

Box 9 Supporting zero-emission aviation through High-Performance Computing (HPC)

Decarbonising aviation is important for achieving net-zero targets. Hydrogen combustion offers a promising solution and EPSRC support (through access to ARCHER2) has enabled researchers at Newcastle University and the University of Cambridge to overcome the challenge of modelling hydrogen behaviour in turbulent engine conditions. Hydrogen-powered propulsion could cut in-flight climate impact by up to 75%[i].

This research, supported through EPSRC-funded Pioneer Projects, enabled collaboration with Rolls-Royce on the development of hydrogen-fuelled aeroengines. Access to ARCHER2 facilitated early modelling work, which underpinned academic publications and industrial engagement. Dr Umair Ahmed from the research team confirmed that the modelling work would not have been possible on HPC systems other than ARCHER2. The team's findings have informed six papers[ii], advancing the understanding of flame behaviour in hydrogen-air mixtures.

New models developed during the project have been implemented in open-source software such as Code_Saturne, supporting adoption across industry and academia. These tools reduce the need for physical testing and enable more efficient combustor design. The work has also supported UK skills development. One PhD student completed their doctorate in 2023 and is now a PDRA in quantum computing at Imperial College London, with a University Fellowship at the University of Manchester starting in late 2025. Another PhD graduate is working on hydrogen combustion as a PDRA at the University of Cambridge. One of the PDRAs who used ARCHER2 is now employed as a data scientist at CPI in Durham.

Longer term, the project supports zero-emission flight and contributes to reducing global climate impacts. With global aviation CO_2 emissions projected to reach 1.5-2.0 Gt annually by 2025[i], the estimated societal cost is between £254 and £338 billion[iii, iv].

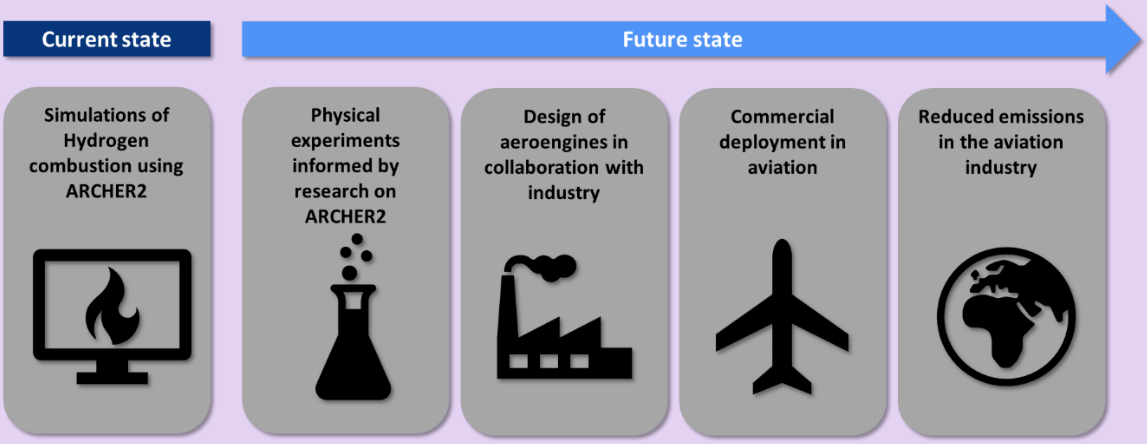

Further details are provided in the full case study in Annex A5.9.

Note:

(i) Clean Sky 2 JU & Fuel Cells and Hydrogen 2 JU. (2020). *Hydrogen-powered aviation: A fact-based study of hydrogen technology, economics, and climate impact by 2050*. Publications Office of the European Union. https://www.euractiv.com/wp-content/uploads/sites/2/2020/06/20200507_Hydrogen-Powered-Aviation-report_FINAL-web-ID-8706035.pdf

(ii) Young, F. W., Ahmed, U., & Chakraborty, N. (2025). Influence of preferential diffusion on the distribution of species in lean H2-air laminar premixed flames at different equivalence ratios. Fuel, 381(Part B), 133363. https://doi.org/10.1016/j.fuel.2024.133363; Ghai, S. K., Ahmed, U., & Chakraborty, N. (2025). Effects of wall temperature on scalar and turbulence statistics during premixed flame–wall interaction within turbulent boundary layers. Flow, Turbulence and Combustion, 114, 421-448. https://doi.org/10.1007/s10494-024-00603-w; Mohan, V., Young, F. W., Ahmed, U., & Chakraborty, N. (2025). Influence of equivalence ratio on the statistics of the invariants of velocity gradient tensor and flow topologies in turbulent premixed lean H_2-air flames. International Journal of Hydrogen Energy, 98, 35-51. https://doi.org/10.1016/j.ijhydene.2024.11.477; Mohan, V., Ahmed, U., & Chakraborty, N. (2025). Distributions of wall heat flux and wall shear stress and their interrelation during head-on quenching of premixed flames within turbulent boundary layers. Flow, Turbulence and Combustion, 114, 1361-1376. https://doi.org/10.1007/s10494-024-00633-4; Ahmed, U., Chakraborty, N., & Klein, M. (2024). Effects of laminar burning velocity to friction velocity ratio on turbulent premixed flame-wall interaction within turbulent boundary layers. Physical Review Fluids, 9(11), 113201. https://doi.org/10.1103/PhysRevFluids.9.113201; Ghai, S. K., Ahmed, U., Chakraborty, N., & Klein, M. (2024). Multiscale analysis of Reynolds stresses and its dissipation rates for premixed flame–wall interaction. Physics of Fluids, 36(10), 105199. https://doi.org/10.1063/5.0232629

(iii) USD values were converted to GBP using the Bank of England's average 2022 spot exchange rate (1 USD = 0.808930594 GBP), then uprated to 2025 prices using GDP deflators from HM Treasury (2024) (2022 = 0.88607225, 2025 = 1.00). Sources: Bank of England. (n.d.). *Data viewer*. Retrieved May 2025, from https://wwwtest.bankofengland.co.uk/boeapps/database/fromshowcolumns.asp?Travel=NIxAZxSUx&FromSeries=1&ToSeries=50&DAT=RNG&FD=1&FM=Jan&FY=2015&TD=19&TM=May&TY=2025&FNY=&CSVF=TT&html.x=176&html.y=34&C=DMY&Filter=N; HM Treasury. (2024). *GDP deflators at market prices, and money GDP: March 2024 (Quarterly National Accounts)*. Retrieved from https://www.gov.uk/government/statistics/gdp-deflators-at-market-prices-and-money-gdp-march-2024-quarterly-national-accounts

(iv) The estimated annual cost to society of $278 to $370 billion is derived by multiplying projected global aviation emissions of 1.5 to 2.0 gigatonnes of CO_2 in 2025 by the central estimate of the social cost of carbon, which is $185 per tonne (Rennert et al., 2022). This results in a range of $277.5 billion to $370 billion. Rennert, K., Errickson, F., Prest, B. C., Rennels, L., Newell, R. G., Pizer, W., Kingdon, C., Wingenroth, J., Cooke, R., Parthum, B., Smith, D., Cromar, K., Diaz, D., Moore, F. C., Müller, U. K., Plevin, R. J., Raftery, A. E., Ševčíková, H., Sheets, H., Stock, J. H., Tan, T., Watson, M., Wong, T. E., & Anthoff, D. (2022). *Comprehensive evidence implies a higher social cost of CO_2.* Nature, 610, 687–692. https://doi.org/10.1038/s41586-022-05224-9

Source: London Economics analysis based on information gathered via desk research and consultation with Dr Umair Ahmed.

5 Economic impacts of ARCHER2

In addition to the scientific and skills benefits enabled, there are also a range of economic benefits flowing from ARCHER2. This includes direct benefits to firms accessing ARCHER2, either directly or via scientific collaborations (see Section 5.2); firms benefitting from staff trained in computational skills that move into industry (see earlier discussion in Section 4); spillover impacts from scientific R&D and discovery (Section 5.1); as well as more direct outcomes of R&D in the form of spin-outs and the development of new products and services (Section 5.4).

5.1 Economic impacts of scientific R&D

Publicly-funded R&D generate significant economic and wider social welfare benefits through a number of channels. This includes **direct channels**, where publicly-funded R&D directly benefit private companies and, in turn, lead to wider economic and social impacts, for example:

- **Knowledge transfers** which can arise through direct R&D collaborations between universities and other organisations, the publication and dissemination of R&D outputs, through university graduates entering the labour market, joint-academic-industry research centres or industrial PhD / doctoral training partnerships, and through consulting activities undertaken by academics.
- Direct **commercialisation of academic research**, for example, through private companies licensing university-owned intellectual property (IP), technology transfer offices (TTOs) facilitating transfer of research outputs into markets, or the formation of academic spin-outs or start-ups based on academic R&D.

Direct channels to industry are also built into the ARCHER2 embedded Computational Software Engineering (eCSE) development activities discussed in Section 3.3. A key component of the eCSE projects was to ensure that academic use of HPC on ARCHER2 has a line of sight from both the scientific community and industry so that industry can benefit more directly from research undertaken on ARCHER2, as well as for industry needs to be able to be reflected in R&D efforts on ARCHER2.

These activities generate private benefits to companies, for example, through the introduction of new products, services, methodologies and concepts. These innovations in turn can lead to improved revenues for firms, a reduction in costs as well as improvements in private sector productivity, for example, by reducing the resources required to produce a product or service. These private sector benefits in turn lead to improvements in social welfare through new and/or improved products and services, job creation and contributions to GDP growth.

In addition, academic research can also lead to economic and societal benefits more widely through **spillover economic impacts**, which arise when economic activities in one part of a market have effects elsewhere in the market. Spillover economic impacts can arise through:

- **Knowledge spillovers** when skills and techniques arising from R&D activity are diffused more widely. This can occur when skills and techniques developed through R&D are exchanged between academic groups and industry, for example at conferences; through labour mobility, i.e., researchers or highly trained students and graduates moving institutions or moving to industry; or when research is disseminated more broadly, e.g., via publications.

- **Market spillovers** when there are wider benefits to the economy or society as a result of commercialisation of innovation in excess of the benefit that buyers of a new product or a product made with new processes receive. In these cases, markets may undervalue the innovation, because firms cannot capture all of the benefits through prices. For example, ARCHER2 has supported research that improves climate modelling. While the immediate beneficiaries are the researchers who use the service, the wider societal gains (e.g., better forecasting for policy and industry) extend far beyond what the original users pay for.
- **Network spillovers** when the value of participating in an activity depends on the number of other participants. For example, consider charging stations for electric vehicles (EVs). An individual may be hesitant to buy an EV if there are only few charging stations available, and charging station providers may be reluctant to build stations if not enough people own EVs. This is known as a co-ordination problem[26] and can lead to under-investment. However, funding from sources such as the government can incentivise unilateral investments, helping to overcome the coordination problem. A similar effect can be seen with ARCHER2: the more researchers and institutions that use the system, the more valuable it becomes. Shared expertise, common software tools, and cross-disciplinary collaborations create benefits that go well beyond individual projects, diffusing innovation into related fields and magnifying the economic and societal returns.

While they are challenging to quantify, economic spillovers from scientific R&D and discovery are well evidenced in the economic literature. Most prominently, research by Haskel and Wallis find that research funded by UK public research councils significantly filters through to the private sector, suggesting a strong role for knowledge spillovers.[27] More recently, a study commissioned by the Department for Science, Innovation and Technology (DIST) and conducted by Frontier Economics estimated substantial returns to public R&D investment, further reinforcing the case for publicly funded research as a driver of wider economic benefits.[28] An overview of the limited quantitative evidence on the size of economic spillovers from academic R&D is provided in Box 20 in Annex A3.2.3.

Public research investment also **stimulates private sector R&D**, thereby increasing total R&D spend and potentially further boosting the benefits of economic growth. For example, Haskel et al.[29] find that public sector funding of science is consistent with “crowding in” of private sector investment i.e. private sector and public sector R&D investment are frequently complements, not substitutes. Other studies (e.g. Jeaumotte and Pain[30]) confirm this finding.

Note, however, that academic R&D can also lead to negative spillovers such as obsolescence. Research could create new ideas, technologies, products and services, rendering existing solutions obsolete. As a result, firms’ investment or intellectual property could lose its value and firms could lose out from innovation. This phenomenon is known as creative destruction[31].

The above research applies to all forms of scientific R&D and discovery undertaken. Evidence of spillover impacts of HPC-specific R&D is more limited. However, one study, in the US, finds that investment into HPC is associated with a contemporaneous increase in the number of academic

[26] See, for example, Farrell, J., and Klemperer, P. (2007). Coordination and lock-in: competition with switching costs and network effects. Handbook of Industrial Organisation, Vol(3) Ch.31, edited by M. Armstrong and R. Porter

[27] Haskel J, Wallis G (2010). Public Support for Innovation, Intangible Investment and Productivity Growth in the UK Market Sector

[28] Frontier Economics (2024). Returns to Public R&D. Report for the Department for Science,Innovation and Technology (DSIT)

[29] Haskel J, Haskel J, Hughes A, Bascavusoglu-Moreau Eet al. (2014). The economic significance of the UK science base: a report for the Campaign for Science and Engineering. Publisher: Imperial College Business School

[30] Jeaumotte, F. and Pain, N. (2005). Innovation in the Business Sector. OECD Economics Department Working Paper No.459, OECD, Paris

[31] Schumpeter, J. (1942). Capitalism, socialism and democracy. New York: Harper, 1975 (original publication 1942), pp. 82 – 85.

publications[32]. Another study, by the same lead author, find that, in the US, locally available HPC resources enhance the technical efficiency of research output in subject areas related to Chemistry, Civil Engineering and Physics, amongst other subjects[33].

Box 10 Computer Modelling of Liquid Repellent Surfaces

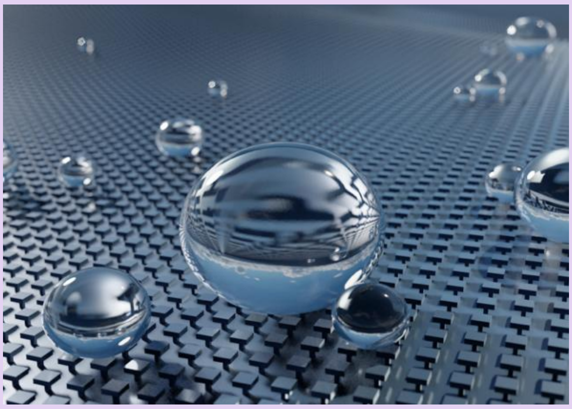

Powerful simulations developed at ARCHER have led to new innovations in the design and manufacture of liquid repellent surfaces for industry.[177]

Researchers at Durham University collaborated with Procter and Gamble to develop simulation tools that are capable of modelling how liquids interact with fibrous materials. The findings have enabled P&G to optimise their designs across their product range.[178]

The simulation has broad, transformative potential across other industries, with additional collaborations also supported by EPSRC. For example, with ExxonMobil in designing novel carbon capture systems[183] and Infineon to design more robust coatings for semiconductors.[184]

The impacts are widespread, including strengthening UK's industrial competitiveness across a range of growing sectors, boosting collaboration and investment from large multi-national companies and driving innovations that are more environmentally sustainable.[177]

Further details can be found in the full case study in Annex A5.10.

Photo Credit: Water droplets on a textured surface, showing the interaction of liquid against a hydrophobic material at high-resolution

Source: Technopolis based on information gathered via desk research and consultation with Prof. Kai Luo and Dr. Halim Kusumaatmaja.

[32] Apon et al. (2010), High Performance Computing Instrumentation and R&D Productivity in US Universities, JITI Journal of Information Technology Impact

[33] Apon et al. (2015). Assessing the Effect of High Performance Computing Capabilities on Academic R&D Output, Empirical Economics

5.2 Direct benefits of access to ARCHER2 to industry

While ARCHER2 is primarily intended to support academic R&D, it also provides valuable access for a number of industry users. There are two main routes through which industry benefits from access to ARCHER2:

- First, companies can **apply directly for compute time** through EPCC, which has an allocation of director's discretionary time, which can be utilised for such purposes. As noted in Section 2.2, there were 151 active user accounts linked to industrial and commercial organisations between 2021 and 2025.
- Second, industry also gains access through **collaborative partnerships with academic institutions** and research groups, allowing them to benefit from shared use of ARCHER2 resources. Data from UKRI's research outcome database (ResearchFish™) indicates that there were at least 38 collaborations between academic users and the private sector, while 36% of respondents to the user survey said they were involved in industry-academia collaborations (See Section 3.2).

> *"While we have the option to use cloud services like AWS, ARCHER2 proves to be significantly more cost-effective and convenient for data transfer and storage." (industry user)*
>
> *"ARCHER2 was incredibly important for industrial collaboration as a trusted resource."*
>
> - **ARCHER2 user survey**

To better understand the nature and extent of use of ARCHER2 by commercial organisations, data from EPCC, NERC and ResearchFish™ were combined. EPCC data captured commercial organisations accessing ARCHER2 directly, while the other sources reflected industry involvement through academic-industry collaborations.

Excluding non-commercial users such as research institutes, publicly funded or government-owned centres and bodies, **123 commercial organisations that access ARCHER2** were identified, either directly or indirectly.

Of these, 31 commercial organisations accessed ARCHER2 directly via Director's time. In addition, the collected data indicates that there were at least 89 organisations that accessed the ARCHER2 indirectly via collaborations with academic researchers. A further three commercial organisations accessed ARCHER2 both directly and via collaborative R&D with academia.[34]

Type of commercial organisations benefiting from ARCHER2 access in terms of firm size

Table 6 provides a breakdown of commercial organisations benefiting from ARCHER2 access by size and access route. Direct users are relatively evenly distributed across micro, small and large firms. Only two medium-sized firms accessing ARCHER2 directly were identified.

In contrast, among indirect users, i.e., those accessing ARCHER2 through R&D collaborations with academic partners, larger firms are overrepresented, likely reflecting their greater propensity to engage in R&D activities.

[34] The datasets were first cleaned to omit entries with an end date prior to 18 January 2021. Then duplicate records, universities, not-for-profit organisations, R&D institutes, The Crown Estate, publicly funded or government-owned centres and collaborative bodies such as working groups, consortia and panels were excluded.

Table 6 Size of commercial organisations benefiting from ARCHER2 access

Company size by employee count	Direct users		Indirect users	
	#	%	#	%
Large	12	35%	45	49%
Medium	2	6%	9	10%
Small	9	26%	12	13%
Micro	10	29%	16	17%
Data unavailable	1	3%	10	11%
Total	**34**	**100%**	**92**	**100%**

Note 1: Three large companies accessed ARCHER2 both directly and indirectly and are recorded in both categories.

Note 2: Company size bands are classified separately by number of employees and annual turnover. Employee bands are: 1–9 (micro), 10–49 (small), 50–249 (medium) and 250+ (large). Turnover bands are: up to £632,000 (micro), over £632,000 to £10.2 million (small), over £10.2 million to £36 million (medium) and over £36 million (large).

Note 3: Entries marked as "n.a." or left blank in Orbis and Fame have been grouped under the category *Data unavailable*. This reflects fields where information was either not provided, not applicable or missing due to data entry gaps. Data coverage varies by company type, size and geography, and some financial or organisational information may be incomplete or inconsistently reported.

Source: London Economics' (2025) analysis of Fame/Orbis data using lists of organisations provided by the EPCC, NERC and from ResearchFish™ Data

Box 11 Simulating accident scenarios in nuclear reactors to improve and accelerate the design process

Science and Technology Facilities Council (STFC) and Électricité de France (EDF) are partnering to model hypothetical accident scenarios in the new generation of nuclear reactors using ARCHER2 and are accelerating the development and deployment of these safer, more efficient reactors in the process.

ARCHER2 enabled high-resolution CFD simulations of High Temperature Gas-cooled Reactors (HTGRs) under hypothetical accident scenarios, such as Loss of Flow Accidents (LOFAs). These simulations would have been computationally unfeasible without access to a supercomputer of ARCHER2's capacity. The outputs served as benchmark data to validate SubChCFD, confirming its accuracy in replicating complex reactor behaviour. By facilitating this validation work, ARCHER2 contributed to:

- Reducing future reliance on supercomputers by supporting the development of a more efficient modelling tool (SubChCFD).
- Accelerating the design and deployment of Advanced Modular Reactors (AMRs), including HTGRs, which support the UK's goal of increasing nuclear capacity to 24 GW by 2050.
- Supporting industrial decarbonisation targets, such as at EDF's Hartlepool site.
- Strengthening academic–industry collaboration and raising STFC's profile in nuclear research.

Further details can be found in the full case study in Annex A5.11.

Source: Technopolis analysis based on information gathered via desk research and consultation with Bo Liu

5.3 Sectoral reach of ARCHER2

Sectoral breakdown of firms accessing ARCHER2

Figure 37 shows the sectoral breakdown of companies accessing ARCHER2 directly and indirectly, via academic collaborations, between 2021 and 2025. Commercial organisations benefitting from ARCHER were spread across **15 sectors**, predominantly in research-intensive, energy-related and digital sectors:

- Professional, scientific and technical activities accounted for the largest share, 34 out of 123 companies. Within this sector, the most frequent activities were other professional, scientific and technical activities not elsewhere classified (10 firms), R&D on natural sciences and engineering (8 firms), and engineering and technical consultancy (7 firms).
- The information and communication sector accounted for 29 companies. These companies were primarily engaged in computer programming activities (10 firms), other IT and computer service activities (6 firms), and IT consultancy activities (5 firms).
- From manufacturing, 20 companies accessed ARCHER2. Access was concentrated in high-tech subsectors such as manufacture of air and spacecraft (5 firms) and manufacture of electronic components (5 firms).

  The electricity, gas, steam and air conditioning supply sector accounted for 10 out of 123 companies. These firms were all involved in electric power generation.
-

Figure 37 Sectoral breakdown of industry users accessing ARCHER2 directly and indirectly

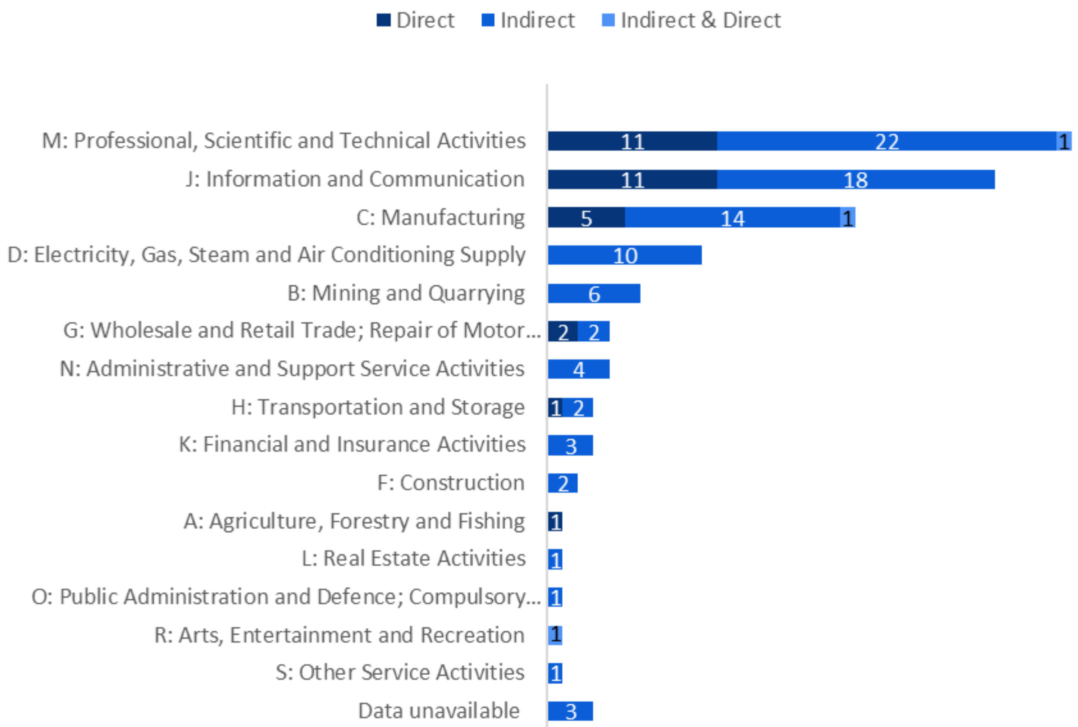

Source: London Economics' (2025) analysis of Fame/Orbis data using lists of organisations provided by the EPCC, NERC and from ResearchFish™ Data

Box 12 Improving combustion research with High-Performance Computing (HPC)

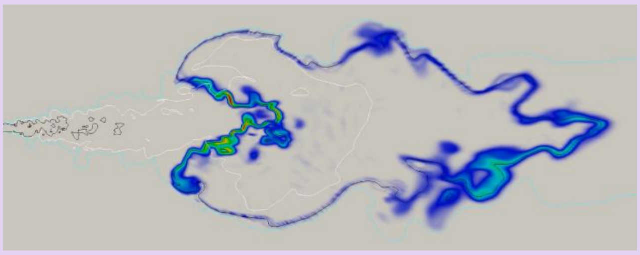

ARCHER2 supported research led by Professor Nedunchezhian Swaminathan at the University of Cambridge led to the development of high-fidelity combustion models which can save companies such as Mitsubishi Heavy Industries around £2 million annually.

Combustion modelling is computationally intensive, particularly for low-carbon fuels such as hydrogen. Simulations conducted on ARCHER2 allow researchers to test and refine models under realistic operating conditions. This would not be feasible using university systems alone, due to limitations in size, cost and processing time. Access through the EPSRC Pioneer Projects has supported PhD training, enabled collaboration with industry and underpinned **16 research papers.**

These models are now used by transport, power generation and manufacturing companies including Rolls-Royce, Siemens and Mitsubishi Heavy Industries (MHI) to improve fuel efficiency, reduce emissions and design quieter engines. MHI reported savings of **£2 million** and over one year in engine development by applying the tools generated through this research. The company's estimates suggest the use of the models can reduce annual computational costs by **£100,000** and testing costs by **£1 million**.

The work has also generated open-access datasets and applied machine learning to accelerate simulation techniques. Future development will focus on hydrogen-methane blends and combustion-induced noise, both of which are increasingly relevant to net-zero and regulatory goals.

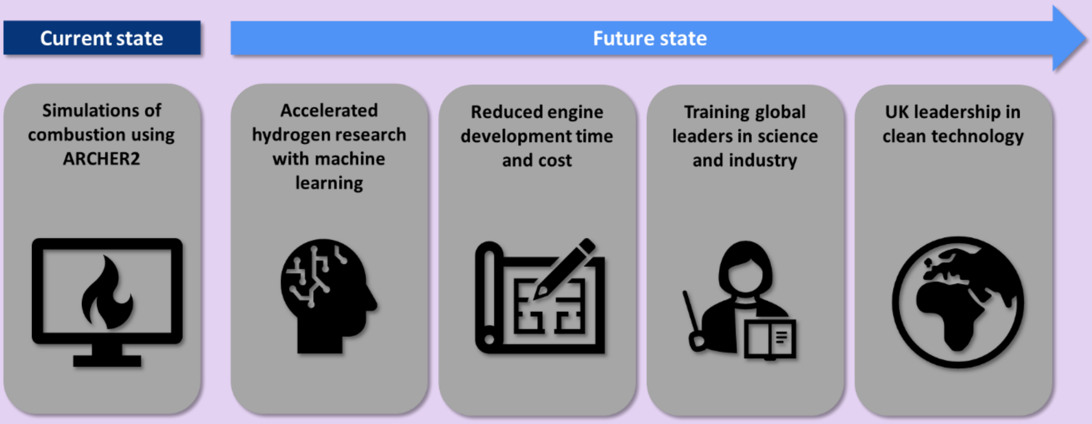

Further details are provided in the full case study in Annex A5.12.

Photo Credit: Image adapted from Schumann, C. D. K., Massey, J. C., Li, C. J., & Swaminathan, N. (2025). Large eddy simulation of transient leading-edge propagation in a turbulent lifted hydrogen jet flame. Proceedings of the Combustion Institute; the image shows the evolution of the ignition kernel toward its final position where the fully lit flame has emerged. The black, cyan, and white lines represent the rich flammability limit, lean limit, and stoichiometric line, respectively. The colour scale indicates the fuel consumption rate. The hydrogen jet flows from left to right.

Source: London Economics analysis based on information gathered via desk research and consultation with Prof. Nedunchezhian Swaminathan.

The following breakdown shows the number of companies in more granular sectoral categories within Manufacturing:

Table 7 Breakdown of industry users within manufacturing

Sector	No. of companies
Manufacture of air and spacecraft and related machinery	5
Manufacture of electronic components	5
Manufacture of refined petroleum products	2
Manufacture of engines and turbines, except aircraft, vehicle and cycle engines	1
Manufacture of other fabricated metal products n.e.c. (not elsewhere classified)	1
Manufacture of other parts and accessories for motor vehicles	1
Manufacture of paints, varnishes and similar coatings, printing ink and mastics	1
Manufacture of pharmaceutical preparations	1
Manufacture of soap and detergents, cleaning and polishing preparations	1
Manufacture of tubes, pipes, hollow profiles and related fittings, of steel	1
Precious Metal Production	1

Source: London Economics' (2025) analysis of Fame/Orbis data using lists of organisations provided by the EPCC, NERC and from ResearchFish™ Data

Wider sector reach of R&D undertaken on ARCHER2

To understand the wider sectoral reach of R&D undertaken on ARCHER2, survey respondents were asked which sectors (if any) benefit from the R&D they undertake using ARCHER2. The survey responses, outlined in Figure 38, indicate that R&D conducted using ARCHER2 has had a broad cross-sectoral impact, with benefits extending across **20 sectors**.

Figure 38 Sectoral reach of R&D undertaken on ARCHER2

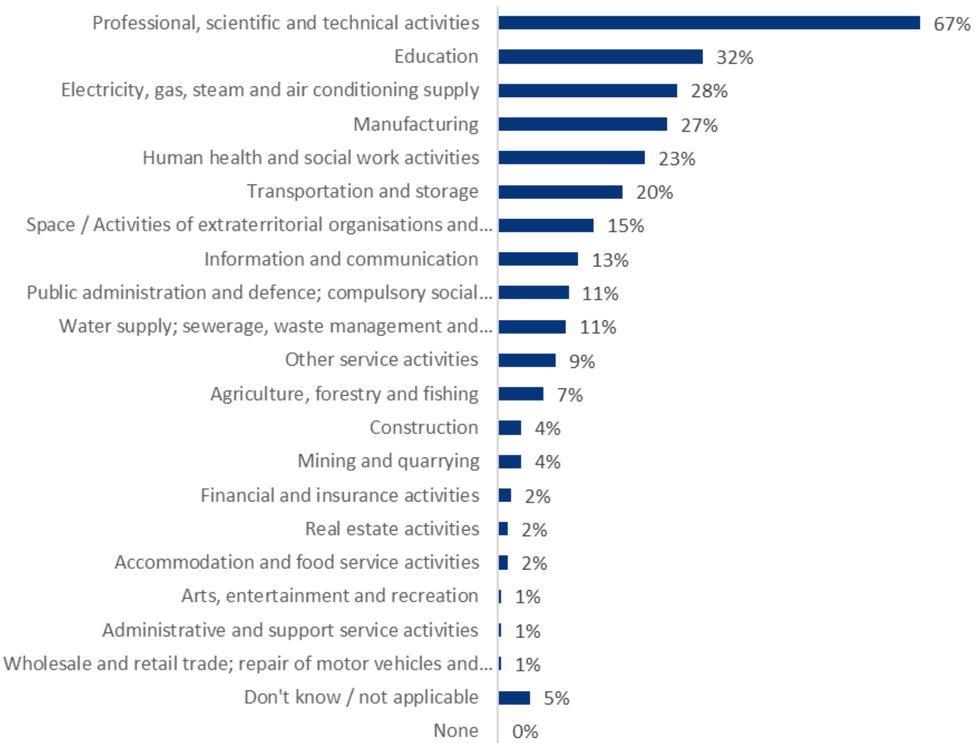

Note: Based on 197 respondents. Respondents could select more than one area. Question asked: 'Which sectors (if any) benefit from the research you undertake using ARCHER2?'

Source: London Economics ARCHER2 user survey

The most frequently selected sector was professional, scientific and technical activities, identified by 132 respondents (67%). Education (32%) was also a key beneficiary. Several industrially-linked sectors benefit from R&D conducted on ARCHER2, including electricity, gas, steam and air conditioning supply (28%), manufacturing (27%) and transportation and storage (20%). Human health and social work activities (23%) also featured prominently.

5.4 Formation of spin-outs and generation of intellectual property

High-performance computing enables cutting-edge R&D that contributes to the development of new technologies, methods, and insights. These breakthroughs can form the basis for spin-out companies and the generation of intellectual property, translating R&D into commercially viable innovations. This, in turn, contributes to economic growth by creating high-skilled jobs, attracting investment, and stimulating innovation across different parts of the economy.

Spin-outs

Spin-outs contribute to the UK economy by developing innovative products and services, that support economic growth. They contribute to UK gross value added (GVA) and create high-skilled jobs, deliver broader economic benefits such as knowledge transfer and market expansion, and enhance competitiveness across industries.

In addition, spin-outs attract private investment, helping to drive innovation and business growth. For example, Wright and Fu[35] analysed the returns to university spin-outs using FAME and Zephyr data, examining whether they secured venture capital (VC) funding, received business angel investment, or underwent a trade sale:

- 24% of university spin-outs received VC funding.
- 6.6% of university spin-outs received business angel funding.
- 6.6% underwent a trade-sale to an existing corporation.

However, the evidence on returns of university spin-outs is mixed. For example, Wright and Fu[36] (citing Bobelyn et al.[37]) suggest that academic spin-offs from the UK were significantly less likely to earn high returns than other young high-tech firms that undergo tech sales. Nevertheless, Wright's and Fu's analysis of spin-outs from UK universities showed that whilst under 1% of spin-outs achieved an IPO by July 2015, these were predominantly from scientific R&D spin-outs.

Evidence of spin-out formation from R&D undertaken on ARCHER2 is limited as this evidence relies on information that is reported back to EPSRC and NERC. Nevertheless, UKRI's research outcomes databases show that there were **at least five spin-outs** that have benefitted from ARCHER2. While four of these were spun-out from R&D undertaken on ARCHER, ResearchFish™ suggests they have continued to benefit from ARCHER2. It is likely that this figure is an underestimate. This is illustrated by insights from the user survey, where 2% of respondents said their work on ARCHER2 led to the creation of spin-outs (see Figure 24 in Section 3.3).

Box 13 provides an overview of the six known spin-outs benefitting from ARCHER2.

[35] Wright, M., & Fu, K. (2015). University spin-outs: What do we know and what are the policy implications? Evidence from the UK. Journal of Innovation Management, 3(4), 5-15. https://doi.org/10.24840/2183-0606_003.004_0002
[36] Wright, M., & Fu, K. (2015). University spin-outs: What do we know and what are the policy implications? Evidence from the UK. Journal of Innovation Management, 3(4), 5-15. https://doi.org/10.24840/2183-0606_003.004_0002
[37] Bobelyn, A., Clarysse, B., & Wright, M. (2015). Returns to technology commercialization and the market for firms: Evidence from venture capital backed young technology based firms. ERC Working Paper.

Box 13 Known spin-outs benefitting from ARCHER2

- **Devito Codes Ltd**: Founded in 2020, Devito Codes originates from Imperial College's open-source software 'Devito' – a Python package primarily used for geophysical simulations. The spin-out provides products, including the more advanced DevitoPRO software packages, and services such as performance optimisation for computer architectures, software development projects, consultancy, and training.[38] At the end of 2023, Devito Codes had a financial net worth of around £580,000 and one recorded full-time employee.[39]
- **Breath Battery Technologies Ltd**: Founded in 2019, Breathe Battery Technologies is an Imperial College London spin-out that develops battery management solutions. They offer two products: Breathe Design – a battery cell design and simulation platform – and Breathe Charge – a battery charging software.[40] By May 2024, Breath Battery Technologies Ltd had a financial net worth of around £4 million and 31 full-time employees.[41]
- **Cognition Energy Ltd**: Founded in 2018, Cognition Energy is an Imperial College London spin-out that specialises in battery testing. They offer products, such as battery holders designed to support testing processes and services, including performance and storage testing to analyse cell behaviour under varying conditions and over time.[42] By October 2024, Cognition Energy's statement of net worth stood at around £330,000 and had 13 full-time employees.[43]
- **About:Energy Ltd**: Founded in 2021, About:Energy is a spin-out from Imperial College London and the University of Birmingham that focuses on battery modelling. The company offers a range of battery models (thermal, electrical, electrochemical and degradation) which are used to simulate battery behaviour with the aim of accelerating the development of new battery technologies.[44] By the end of 2023, About:Energy had a financial net worth of around £1 million and 19 full-time employees.[45]
- **Bind Research**: Founded in 2025, Bind Research is a Focused Research Organisation (FRO) dedicated to making disordered proteins druggable. These biomolecules are linked to incurable conditions like cancer and have long been considered 'undruggable' by the pharmaceutical industry. Bind aims to tackle this challenge by screening millions of disordered protein–drug molecule pairs using advanced biology, engineering, and AI techniques to accelerate the development of new drugs and tools.[46]

Source: London Economics based on ResearchFish™ Data and desk research

Intellectual property and patents

Economic literature suggests that patents and IP are associated with increases in UK firm value and growth, as well as potentially being linked to GDP growth. For example, studies by European Union

[38] See: https://www.devitoproject.org/
[39] Based on the financial information of spinouts available in Researchfish data.
[40] Breath Battery Technologies. (2025) Battery Industry Supplier. Available at: https://battery-tech.net/breathe-battery-technologies/
[41] Based on the financial information of spinouts available in Researchfish data.
[42] See: https://www.cognition.energy/
[43] Based on the financial information of spinouts available in Researchfish data.
[44] See: https://www.aboutenergy.io/2025
[45] Based on the financial information of spinouts available in Researchfish data.
[46] See: https://bindR&D.org/

Intellectual Property Office[47], Willoughby[48], and Yun et. al.[49] all found a positive statistical relationship between patents and increased financial performance by firms. Yıldız and Görkey[50] analysed the relationship between the growth in number patent grants and GDP growth, and found a small, but positive, association.

6% of survey respondents said their work on ARCHER2 led to the creation of patents, licensing agreements or other intellectual property. In total, **6 trademarks, patents and other intellectual property** resulting from ARCHER2 R&D were captured in UKRI's research outcomes databases.

Box 14 provides a summary of some of these patents.

Box 14 Examples of patents resulting from R&D undertaken on ARCHER2

- **Edinburgh turbine-blade**: In collaboration with key tidal and wind energy companies, researchers at the University of Edinburgh have developed a morphing tidal turbine blade that could reduce the levelized cost of energy (LCOE) by over 10%. Designed to address fatigue caused by fluctuating tidal loads, the blade passively adjusts its pitch without mechanical mechanisms - improving durability, reliability and operational lifetime – while remaining compatible with active control systems.[51]
- **Passive Pitch Control Mechanism for Wind Turbine Blades**: researchers at the University of Edinburgh have developed a passive turbine blade pitch technology that reduces aerodynamic load fluctuations, improving efficiency and reducing structural fatigue. The simple mechanical system works with most blade designs and can replace or complement active pitch control. It offers greater reliability, longer operational life, and the potential to lower the levelized cost of energy.[52]
- **AirSea Flux Code**: researchers at the National Oceanography Centre developed AirSeaFluxCode – a Python package for calculating air-sea fluxes of heat and momentum. The project was motivated by the fact that these fluxes play a crucial role in Earth's climate system, but direct measurements are scarcely possible or inadequate. In response, the team created an open-source Python software package that enables robust and flexible computation of turbulent heat and momentum exchanges between the atmosphere and the ocean.[53]

Source: London Economics based on ResearchFish™ Data and desk research

[47] EUIPO. (2021). Intellectual property rights and firm performance in the European Union. Available at: https://www.euipo.europa.eu/en/publications/intellectual-property-rights-and-firm-performance-in-the-european-union

[48] Willoughby, K. (2013). What impact does intellectual property have on the business performance of technology firms? Available at: https://www.R&Dgate.net/publication/258514677_What_impact_does_intellectual_property_have_on_the_business_performance_of_technology_firms

[49] Yun, J. J., Jin, S., & Kim, D. (2021). The effects of patents on the relationship between R&D activities and business management performance: Focus on South Korean venture companies. Journal of Open Innovation: Technology, Market, and Complexity, 7(4), 210. https://doi.org/10.3390/joitmc7040210

[50] Yıldız and Görkey. (2024). Causal links between patents and economic growth: empirical evidence from OECD countries. Available at: https://www.R&Dgate.net/publication/383471188_Causal_links_between_patents_and_economic_growth_empirical_evidence_from_OECD_countries

[51] University of Edinburgh. (no date). Flexible Blade Design for unsteady Load Mitigation in Tidal Turbines. Available at: https://edinburgh-innovations.ed.ac.uk/technology/flexible-blade-design-for-unsteady-load-mitigation-in-tidal-turbines

[52] University of Edinburgh. (no date). Passive Pitch Control Mechanisms for Wind Turbine Blades. Available at: https://edinburgh-innovations.ed.ac.uk/technology/passive-pitch-control-mechanism-for-wind-turbine-blades

[53] Biri, S., Cornes, R. C., Berry, D. I., Kent, E. C., & Yelland, M. J. (2023). AirSeaFluxCode: Open-source software for calculating turbulent air-sea fluxes from meteorological parameters. Frontiers in Marine Science, 9, 1049168.

6 Wider benefits of ARCHER2

In addition to benefits discussed in the previous sections, there are a wide range of wider benefits flowing from ARCHER2. While many of these benefits are difficult to quantify and monetise, they are no less important.

This section explores two broad benefit categories:

- Section 6.1 discusses the benefits flowing from delivering and maintaining a high-quality national Tier-1 service. This includes environmental benefits in terms of a reduced carbon footprint associated with HPC, increased Value-for-Money from continuous improvements, the importance of maintaining a strong HPC service skill-set, and the benefits of high-quality user support.
- Section 6.2 then discusses improved local and societal impacts arising in the local Lothian economy from hosting the ARCHER2 service in Edinburgh as well as through public outreach activities.

6.1 High quality service

EPCC is committed to delivering a high-quality ARCHER2 service that is safe, accessible, and inclusive for all users. This commitment is supported by a range of policies and ongoing initiatives aimed at reducing the environmental impact of HPC, improving value for money through continuous enhancement, and sustaining the specialist skills required to operate and support a world-class HPC ecosystem.

6.1.1 Reduced HPC carbon footprint

EPCC's commitment to environmental sustainability is reflected in various areas of ARCHER2's services. These lead to significant reductions in its carbon footprint, establishing ARCHER2 not only as a high-performance computing (HPC) resource but also as a benchmark for sustainable supercomputing.

ARCHER2 operates under a 100 % certified renewable electricity contract via the UK National Grid, effectively bringing its operational (Scope 2) greenhouse gas emissions down to zero.[54] While baseline calculations of grid-based emissions remain available for comparison, the service's renewable arrangement ensures that real-world Scope 2 impact is nil.

ARCHER2 is more energy-efficient than its predecessor, ARCHER, delivering more computing power per unit of energy. Smart systems help manage how energy is used. For example, processors can run at slightly lower speeds during less demanding tasks, which significantly reduces power consumption without making the system noticeably slower. These kinds of adjustments have helped cut energy use across the system by over 20%.[55] Furthermore, the ARCHER2 data centre also uses 'free cooling' (outside-air or water cooling), and plans are underway to redirect waste heat to nearby buildings, further reducing the facility's energy use.

While the system runs on renewable electricity, some carbon emissions are linked to building and delivering the supercomputer hardware. These "Scope 3" emissions are carefully measured and

[54] EPCC. (no date). Environmental Sustainability. Available at: https://www.archer2.ac.uk/community/sustainability/
[55] Jackson, A. (no date), Assessing and reducing the environmental impact of HPC. Available at: https://www.archer2.ac.uk/community/events/celebration-of-science-2024/slides/NetZeroTalk.pdf

averaged across the expected lifetime of ARCHER2, with a typical figure of around 0.023 kg CO_2e per compute unit hour. These estimates are also included in the user-facing emissions tools. Lastly, ARCHER2's facility supports biodiversity with green infrastructure such as tree groves, wetlands, and pollinator meadows. These efforts recognise that environmental responsibility extends beyond energy into land stewardship.[56]

To support more sustainable computing, ARCHER2 provides tools that help users calculate the carbon footprint of their work. These tools—such as the emissions command and CU calculator[57]—allow users to estimate Scope 3 emissions. This helps researchers understand the climate impact of their jobs and adjust their workflows where possible.

Box 15 The economic value of reducing ARCHER2's carbon footprint

The carbon savings achieved by ARCHER2 can be valued using HM Treasury's Green Book guidance, by comparing its emissions to the business-as-usual counterfactual in which ARCHER1 remains in operation. This comparison focuses on two main types of emissions: Scope 2 emissions from electricity use and Scope 3 emissions from the production and maintenance of computing equipment (also known as embodied emissions).[58]

To estimate the difference in economic impact between the two systems, annual carbon emissions for ARCHER1 and ARCHER2 were calculated. ARCHER2 figures are based on published estimates available online[59]. For ARCHER1, emissions were calculated by multiplying annual electricity consumption estimates (provided by EPCC) by official carbon conversion factors from the Department for Energy Security and Net Zero (DESNZ).[60] For example, in 2022, one kilowatt-hour of electricity was associated with 0.193 kilograms of carbon dioxide equivalent.

These emissions were then assigned a monetary value using the carbon prices outlined in the Green Book. These prices reflect the estimated societal cost of emitting one tonne of carbon dioxide equivalent. In 2023, for instance, this value was £294.46 per tonne (in 2025 prices).[61] Subtracting ARCHER2's annual carbon costs from the counterfactual costs were ARCHER1 to continue operating gives the total economic benefit in terms of the value of reduced emissions.

ARCHER2 runs on 100% certified renewable electricity and therefore produces no Scope 2 emissions. In contrast, ARCHER1 was powered by electricity from the UK grid, thereby generating emissions. The value of avoided Scope 2 emissions from switching to ARCHER2 is estimated at £2.49 million over ARCHER2's lifetime.

However, ARCHER2 has higher Scope 3 emissions due to its larger number of compute nodes—5,860 compared with ARCHER1's 4,920—and a higher embodied emissions rate per node-hour (0.02 kgCO_2e versus 0.015). As a result:

- The additional embodied emissions amount to approximately 3,400 tonnes of CO_2e over ARCHER2's lifetime.

[56] EPCC (no date). Machine Room and Hardware Theme Sustainability at the Advance Computing Facility (ACF). Available at: https://eng.ox.ac.uk/media/evyozujf/netdrive_hardware_sustainability_v2.pdf?utm_source=chatgpt.com#page=20.00
[57] See: https://www.archer2.ac.uk/support-access/cu-calc.html
[58] The Scope 1 class of emissions is not relevant for ARCHER1 and ARCHER2 because it is very small and negligible. Scope 1 emissions are direct GHG emissions from sources owned or controlled by the organisation — e.g., from combustion in boilers, vehicles, or on-site fuel use. See https://docs.archer2.ac.uk/user-guide/energy/
[59] See https://docs.archer2.ac.uk/user-guide/energy/
[60] See https://www.gov.uk/government/publications/greenhouse-gas-reporting-conversion-factors-2025
[61] See Table 3 here https://www.gov.uk/government/publications/valuation-of-energy-use-and-greenhouse-gas-emissions-for-appraisal

- The corresponding increase in economic cost is estimated at £1.01 million.

Taking both Scope 2 and Scope 3 emissions into account, the total net discounted reduction in emissions from using ARCHER2 instead of ARCHER1 is valued **at approximately £1.47 million over ARCHER2's lifetime.**

It is important to note that this estimate assumes that energy use remains constant across both systems over time, and that ARCHER1 would continue to rely on grid electricity without a transition to renewable energy. If around 60% of ARCHER1's energy use was based on renewable energy sources, the value of carbon savings of ARCHER2 would be zero or negative.

Moreover, it is assumed that the utilisation rate of ARCHER2 and ARCHER1 remains constant across ARCHER2's lifetime. This implies that embodied carbon emissions remain constant throughout that period.

Source: London Economics analysis based on information provided by EPCC and desk research

6.1.2 Increased VfM through continuous improvement

ARCHER2 is underpinned by a continuous improvement philosophy that ensures the service evolves to meet the changing needs of its users. Regular improvements in response to ongoing system monitoring, user feedback, and performance evaluations, as well as the training and software development efforts discussed in earlier sections, drive incremental enhancements to ensure ARCHER2 delivers maximum value for money by enabling a better, more streamlined service for users, minimising downtimes and by enabling users to undertake more research with the same resources.

Regular feedback from users is collected through user satisfaction surveys, service desk queries, the website feedback form, as well as direct engagement via UKRI, High End Computing (HECs) consortia, community meetings and events, and virtual user groups. User feedback, satisfaction and trends in turn trigger improvement projects to enhance the service. This could include system updates, provision of additional software codes and research tools, user support improvements, and enhanced training offerings, among others.

ARCHER2's performance is continually monitored to identify bottlenecks, inefficiencies, or emerging needs. Software updates, hardware adjustments, and scheduler optimisations are implemented to maintain high utilisation and responsiveness. In addition, EPCC supports users in optimising and enhancing software through embedded Computational Software Engineering (eCSE) projects (discussed further in Section 3.3), helping to ensure continuous improvements in how software utilises the hardware. While efficiency improvements from eCSE funded software are no longer quantified, historical eCSE projects leading up to 2018, delivered efficiency improvements, in terms of time freed up on ARCHER2's predecessor, of £4 for every £1 invested in the programme.

As user needs evolve, ARCHER2 provides updated training courses and materials, workshops, and online resources. This supports a growing and diversifying user base. For example, in 2024, EPCC have collaborated with leading technology companies, including AMD and Codeplay, to provide training courses designed to exploit the advanced capabilities of GPUs.

This commitment to deliver continuous improvements is reflected in user satisfaction responses, where all service aspects of ARCHER2 are consistently rated very highly (see Table 8).

Table 8 User satisfaction with ARCHER2 service aspects

Service aspect	2022-2024 ARCHER2 mean score (out of 5)	2021-2022 ARCHER2 mean score (out of 5)
Overall	**4.4**	**4.3**
Hardware	3.9	3.9
Software	3.8	3.9
Service desk	4.6	4.8
Documentation	4.1	4.2
Website	4.1	4.1
Training	4.3	4.2
Webinars and virtual tutorials	4.2	4.0
Online training	4.2	4.1

Note: Scores ranging from 1-5 with 1 representing "Very Unsatisfied" and 5 representing "Very Satisfied"

Source: ARCHER2 user survey reports, 2021-22 and 2022-24. Available at: https://www.archer2.ac.uk/about/reports/

In addition, to continuous improvements of the ARCHER2 service itself, the ARCHER2 team has a set of **initiatives designed to prepare users and their software for the next National HPC service**, ensuring the new service can be used effectively and maximises value for money achieved from day one. These activities are further described below:

- **Training in GPU technologies:** The ARCHER2 training team has implemented a comprehensive programme of GPU training activities to equip users with the skills necessary for efficient use of modern high-performance computing (HPC) systems. These activities include targeted courses on GPU programming using both directive-based approaches (such as OpenMP and OpenACC) and kernel-based programming (such as CUDA and HIP), enabling users to develop portable and high-performance code. In collaboration with technology partners, such as AMD and HPE, ARCHER2 has also hosted a range of supplementary events including webinars, hackathons, and specialised training courses. These sessions have provided users with hands-on experience and expert guidance on optimising applications for GPU architectures, ensuring a smooth transition to future GPU-based platforms. As part of the Excalibur project (UK's efforts to prepare for the arrival of exascale HPC), dedicated training was provided to Excalibur and HEC funded RSEs tasked with porting software onto GPUs. These trained RSEs then pass on their gained knowledge to their HEC communities.
- **Funding programme to develop software for future infrastructure:** Recently the eCSE (Embedded Computational Science and Engineering)[62] programme has expanded, with a new £4 million programme of eCSE calls focusing on the development of software for GPU-based architectures. This has funded 12 successful projects, across all of UKRI's remit and embedded across the UK. These projects are developing software that facilitates research targeted at UKRI's digital research infrastructure including future National supercomputing services and the UK national AI services.

  Examples of the work involved include: transforming existing software that targets CPU architectures into software that can exploit GPU-based architectures; implementation of algorithmic improvements to allow new research to be carried out; improving the portability, performance and scalability of existing software, improving sustainability and maintainability of software; adding new functionalities to existing GPU-based software for use in new research areas/workflows.

[62] See Section 3.3 for more background on the eCSE programme.

- **Capability Days:** Capability Days are periods on the ARCHER2 system where a large fraction of the system is reserved for very large, capability jobs which run free of charge. This allows researchers to test and understand the performance of their applications at large-scale. The aim is to enhance world-leading science from ARCHER2 by enabling modelling and simulation at scales that are not otherwise possible. This will let users prepare their software to realise their scientific ambition on future national services. Four capability days have taken place over the last 2 years, with each of these allowing users to run large capability jobs across 2-3 days.
- **Environmental sustainability of future services - tools to help users:** Environmental sustainability is a core priority for ARCHER2, both in daily operations and in planning future services. Meeting Net Zero targets is essential, and the service supports users in addressing these challenges. Work to date has focused on measuring and communicating emissions within the Greenhouse Gas (GHG) protocol to improve carbon efficiency and raise user awareness, as well as promoting biodiversity at the ACF data centre. EPCC has also developed a command-line tool that lets users estimate the GHG emissions from their ARCHER2 use and compare them with other sources.

6.1.3 Maintain HPC service skill set and improved HPC ecosystem

The operation of a large-scale HPC service such as ARCHER2 is critical for maintaining the UK's specialist skill base in high-performance computing. Delivering and supporting such a capability requires the coordination and delivery of large-scale infrastructure capable of servicing a wide range of users and supporting the most complex scientific computations. This includes deep technical expertise in areas such as system architecture, software and hardware optimisation, job scheduling, and user support. Without an active large-scale service, this skill set risks being lost or diminished.

ARCHER2 enables both the retention of experienced personnel and the development of new talent through hands-on involvement in service delivery, training, and collaborative support. This continuity is vital to ensure the UK remains prepared to manage and effectively exploit future HPC infrastructure, including potential exascale deployment.

National Tier-1 services do not operate in isolation, rather, ARCHER2 is part of a wider domestic UK HPC ecosystem that includes regional HPC centres (Tier-2), local university and research clusters (Tier-3), emerging GPU-based systems, and, increasingly, cloud-based HPC platforms. This interconnected landscape enables collaboration, resource sharing and the development of complementary capabilities across different levels of infrastructure. For example, close coordination between Tier-1 and Tier-2 centres allows rapid sharing of threat intelligence, security patches, and best practices in response to emerging cyber threats. Joint monitoring initiatives and simulated attack exercises help detect and prevent hacking attempts more effectively, ensuring consistent and robust security standards across the UK HPC ecosystem.

Therefore, a strong domestic skillset in high-performance computing not only benefits the national service, but in turn also has spillover benefits to Tier-2 systems and supports the wider HPC ecosystem across research and industry. These benefits are realised through staff mobility, collaborative projects, and the active exchange of knowledge and best practices. As skilled personnel move between institutions, support different levels of HPC infrastructure, or engage with external partners, they help to elevate capability, efficiency, and innovation across the UK's entire HPC ecosystem.

Furthermore, many ARCHER2 users also access other HPC facilities. Six in ten users responding to the user survey say that they have accessed other HPC facilities in addition to ARCHER2 (Figure 39).

Almost half (45%) of respondents who had used other systems had accessed Tier-2 capabilities. 26% had used international HPC systems, 11% used cloud computing, and 1% of users had already accessed the new AI Research Resource. In addition, 18% of users said that they had accessed other HPC capabilities beyond those listed. (Figure 40)

Users being able to access the right capability, whether that is Tier-2 services, GPU-based systems, or smaller-scale local systems or commercial cloud resources where this is appropriate, is crucial to ensure that computational jobs are matched to the most appropriate level of compute and that large-scale compute resources are used effectively.

Figure 39 Proportion of ARCHER2 users who have used other HPC facilities

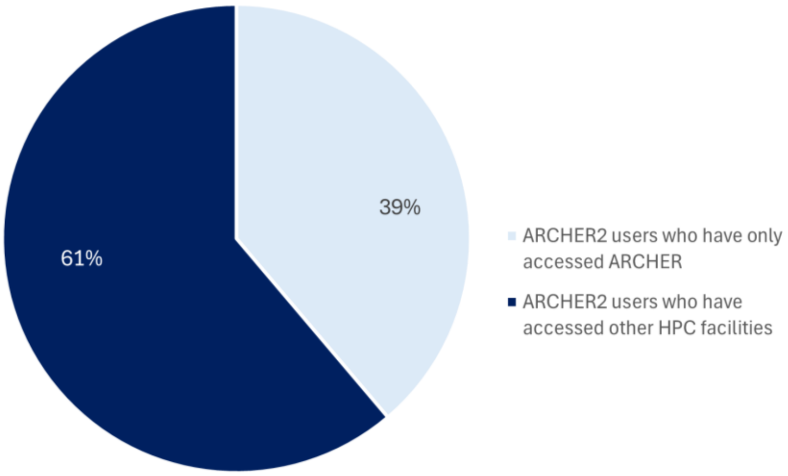

Note: Based on 201 respondents. 3 respondents who said they did not access ARCHER2 were manually removed.

Source: London Economics ARCHER2 user survey

Figure 40 Proportion of ARCHER2 users using alternative HPC facilities, by type of facility

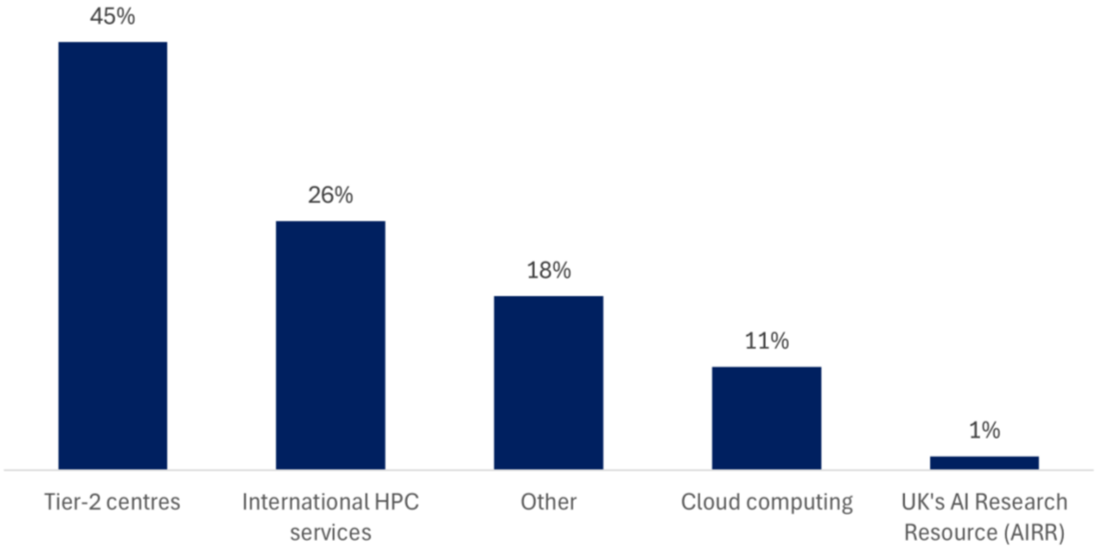

Note: Based on 201 respondents. Percentages do not sum to 100% as users could select multiple answers.

Source: London Economics survey of ARCHER2 users

Box 16 Advancing the development of compact electron accelerators through High-Performance Computing (HPC)

Researchers at the University of Oxford and Ludwig Maximilian University of Munich are developing a novel plasma-based accelerator, the Plasma-Modulated Plasma Accelerator (P-MoPA), which has the potential to be significantly smaller, cheaper and more energy-efficient than existing systems.

Particle accelerators are used in a diverse range of fields, such as cancer radiotherapy, industrial inspection and scientific imaging. Traditional systems are typically large and costly, limiting their broader use. P-MoPA seeks to address these limitations by using high-frequency electron pulses and built-in machine learning to deliver stable performance. It is projected to be around **1,000 times** more energy-efficient and capable of firing pulses **100 times** more frequently than current models.

Access to EPSRC's national HPC facilities through the Pioneer Projects enabled the research team to run complex simulations using ARCHER and ARCHER2. These simulations supported the design of a Gemini laser experiment and helped optimise experimental conditions in advance. The results, obtained within a single six-week session, informed published research and saved significant resources, with one laser session estimated at **£500,000**.

The project has since attracted further investment, including a **£2.16 million** EPSRC grant and international funding, and has led to a patent filing and early commercial exploration. A spinout company is under consideration. In the near term, the technology shows strong promise in advanced imaging and scientific applications, such as compact X-ray sources and Free Electron Lasers, offering performance gains over current accelerator systems.

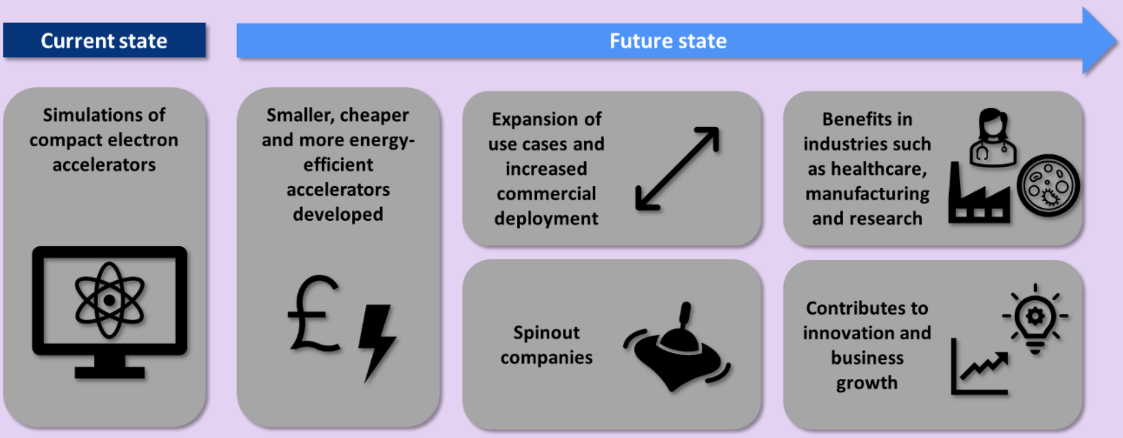

Further details are provided in the full case study in Annex A5.14.

Source: London Economics based on information gathered via desk research and consultation with Prof. Roman Walczak.

Commonly mentioned systems within the UK that users had accessed are:

- Cirrus was the most commonly mentioned alternative, appearing in over 40 responses.
- Isambard, including its variants such as Isambard 2 and Isambard-AI, was cited in approximately 20 responses.
- Thomas and Young, part of the MMM Hub, were mentioned around 12 and 26 times, respectively.
- CSD3 was noted in 15 responses.
- The JADE and JADE2 GPU clusters were cited in 9 responses.
- Sulis (Warwick Tier-2) and JASMIN, the data analysis environment maintained by CEDA, were referred to in 13 and 11 responses each
- Bede (Durham Tier-2) appeared 11 times.

Baskerville (Birmingham Tier-2), Grace (UCL), Kathleen (UCL), ARC (UCL), and Iridis (Southampton) were also mentioned, though each in fewer than 10 responses. Beyond the UK, researchers reported access to a wide range of international systems. European HPC centres featured prominently: MareNostrum (Spain) was cited in 8 responses, while JUWELS and SuperMUC (Germany), Leonardo and Marconi (Italy), LUMI (Finland) and Piz Daint (Switzerland) appeared in a handful of responses.

A number of respondents also referenced North American systems, particularly those operated by the National Center for Atmospheric Research (NCAR), such as Cheyenne, Derecho, and Casper. NERSC facilities, Summit, Frontier, and Stampede2 also featured. Japanese systems such as Fugaku and K computer appeared, as did other international centres including CSCS, CINECA, and HLRS.

6.1.4 User support

EPCC provides user support through a dedicated help desk. The helpdesk acts as the primary point of contact for users encountering issues, handling everything from registration and administrative changes to technical inquiries and complex problem resolution.

Among respondents to the user survey undertaken for this project, **70% of respondents had used the ARCHER2 helpdesk**. (Figure 41)

Out of those who had used the helpdesk, **97% said that the helpdesk saved them time** compared with if they had tried to solve the issue by themselves.

A large proportion (69%) reported that it had saved them several hours or more per query. (Figure 42)

"The ARCHER2 technical support team has been invaluable. Due to the complexity of my simulations, I have encountered various technical challenges. The support team has been consistently responsive and dedicated—organizing virtual meetings, reviewing my simulation setup in detail, and even involving domain experts when necessary to provide tailored solutions."

"Their guidance on using the HPC is very well explained, and is a great source of support, as a lot of the information extends to general HPC use. The helpdesk staff have also been very quick to respond to queries and issues I have had previously, and problems are often resolved swiftly."

"The Helpdesk always responds quickly and helpfully to queries."

- **ARCHER2 user survey**

Figure 41 Proportion of users that have used the ARCHER2 helpdesk

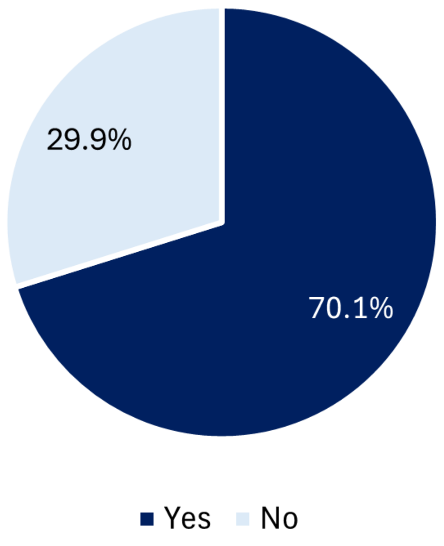

Note: Based on 201 responses.

Source: London Economics ARCHER2 user survey

Figure 42 Reported time saving from helpdesk support

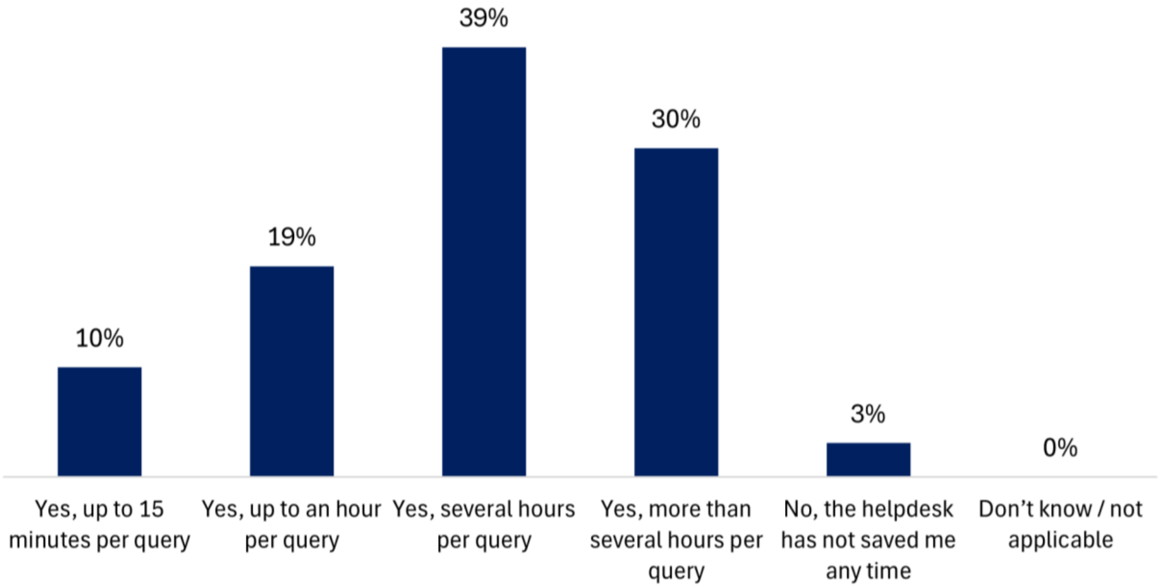

Note: based on 132 respondents. Was asked to users who said they had used the helpdesk. Question asked: 'Has the helpdesk saved you time compared to if you had to solve the issue by yourself? How much time do you think the helpdesk has saved you on average per query?'

Source: London Economics ARCHER2 user survey

Volume of support queries and response times

Between 2021 and 2024, the number of user support queries resolved by the ARCHER2 service increased each year. In 2021, just over 5,000 queries were resolved across all support levels. This figure exceeded 20,000 in 2024.

Support queries are categorised into three levels depending on their estimated complexity (see Table 9). Level 1 queries are the simplest queries to resolve, typically involving administrative or straightforward technical issues, and also the most common accounting for 93% of all queries. Level 2 queries ranged from 195 in 2021 to 1,174 in 2022, then remained relatively stable through 2024. Level 3 queries were consistently low in volume, ranging from 2 to 25 per year, making up less than 0.1% of all queries.

Table 9 Explanation of user support query types

Level	Service	Description
Level 1	SP	Simple administrative, job management, software licensing/access to software issues, account management and/or low level system administration matters (e.g. account/allocation management or administration, disk quota management, project account management etc.)
	CSE	Simple administrative, job management, registration for or enquiries about courses provided as part of the CSE Training Services
Level 2	SP	Moderately complex issues and/or matters including investigations of system performance and throughput, scientific analysis, package updates or installations, in-depth systems administration, etc
	CSE	Moderately complex issues and/or matters, including queries in relation to compilers and system software, porting of codes, data transfer, batch systems and queues
Level 3	SP	Complex or technically challenging longer-term issues and/or matters including those requiring extended input from Related Service Providers and/or changes to the Facility, filesystems, software packages, libraries etc.
	CSE	Complex or technically challenging longer-term issues and/or matters including those requiring extended input from Related Service Providers and/or relating to third party software (including End User Software), complex User programmes and software, performance and/or scaling.

Note: The EPCC SP support provides system administration, user and project management, and service desk management. The EPCC CSE support provides in-depth technical support for users, extensive training, embedded CSE support to fund software development, and outreach to the wider public.

Source: EPCC

Mean Time to Resolution (MTR) for Level 1 queries decreased significantly after 2021. In 2021, the average MTR was approximately 4 minutes and 42 seconds. From 2022 onward, MTR for Level 1 queries remained under one minute. Level 2 MTRs ranged between 33 and 40 minutes across the five years. Level 3 MTRs were much higher, ranging from approximately 37 to 67 hours, reflecting the complexity of these issues.

Table 10 ARCHER2 user support metrics by service level (2021-2025)

Service level	2021		2022		2023		2024		2025	
	MTR	#	MTR	#	MTR	#	MTR	#	MTR	#
Level 1	0:04:42	4868	0:00:16	10703	0:00:29	13289	0:00:35	19269	0:00:16	6615
Level 2	0:39:14	195	0:40:23	1174	0:36:16	1004	0:38:19	1141	0:33:43	454
Level 3	42:31:07	4	39:13:52	25	41:00:29	12	67:03:35	8	37:31:10	2
Total		**5067**		**11902**		**14305**		**20418**		**7071**

Note: Figures for 2025 represent partial-year data (January-May).

Source: EPCC data

Table 11 presents the number and percentage of queries responded to within three hours over the five years from 2021 to 2025. The data shows variation across years in both the absolute number and proportion of timely responses.

The highest proportion of queries responded to within three hours occurred in 2022, with 2,852 queries representing 24.0% of the total for the year. This contrasts with 2021, which recorded the lowest proportion, with only 325 queries (6.4%) meeting the same threshold. The number of queries responded to within three hours peaked in 2024 at 3,032, despite a lower percentage (14.8%) compared to 2022.

Table 11 Queries responded to within 3 hours by number and percentage of total (2021-2025)

2021		2022		2023		2024		2025	
#	%	#	%	#	%	#	%	#	%
325	6.4%	2852	24.0%	2572	18.0%	3032	14.8%	1244	17.6%

Note: Figures for 2025 represent partial-year data (January-May). ***Source: EPCC data***

User satisfaction with helpdesk support

Across all years, most users rated the service as "Excellent," with this category consistently receiving the highest proportion of responses, ranging from 78% in 2021 to 89% in 2025, suggesting users have a predominantly very positive experience with ARCHER2. Ratings of "Very Good" and "Good" also made up a large portion of responses, while "Poor" and "Very Poor" ratings were minimal or absent throughout the period.

The overall positive view of the helpdesk is reflected in both the 2021-22 and 2022-24 EPCC ARCHER2 user survey reports, where the service desk scored the highest in terms of user satisfaction out of categories including software, hardware and training, scoring 4.6/5 and 4.8/5 respectively.[63]

Table 12 User satisfaction rating by number and percentage of total (2021-2025)

Rating	2021		2022		2023		2024		2025	
	#	%	#	%	#	%	#	%	#	%
Very poor	1	1%	0	0%	1	0%	0	0%	0	0%
Poor	0	0%	0	0%	0	0%	0	0%	0	0%
Good	11	14%	42	10%	38	11%	22	6%	8	6%
Very Good	5	7%	45	11%	37	11%	29	8%	6	5%
Excellent	60	78%	319	79%	274	78%	302	86%	112	89%

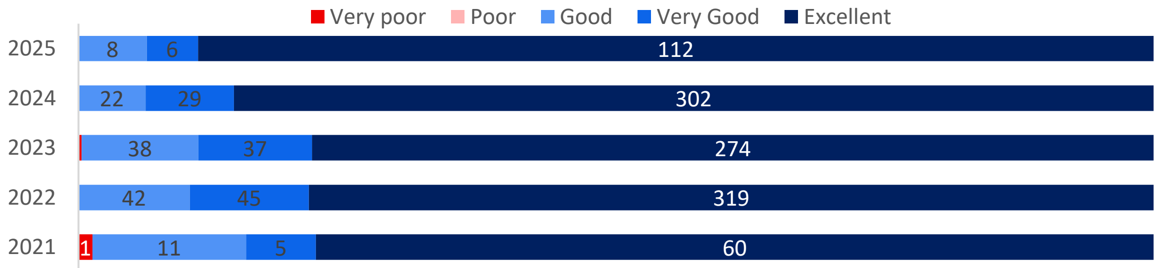

Note: Figures for 2025 represent partial-year data (January-May). ***Source: EPCC data***

[63] ARCHER2 user reports, 2021-22 and 2022-24. Available at: https://www.archer2.ac.uk/about/reports/UserSurvey_Report_2023_v1.0.pdf

6.2 Improved local and societal impact

6.2.1 Local economic impact of ARCHER2

Operating a supercomputing facility of the scale of ARCHER2 creates important economic impacts in the local Edinburgh area and neighbouring councils.

Most directly, EPCC, who host ARCHER2, employ around 150 staff who live within the local Edinburgh and Midlothian areas as well as neighbouring councils such as Fife, East and West Lothian, Scottish Borders, creating highly skilled technical, administrative, and support staff employment opportunities to the local area.

EPCC and ARCHER2 are committed to promoting and growing the skills required to develop and operate complex supercomputing and data science systems, and to help young people find their way into careers in the HPC field. For this reason, EPCC offers a range of work placement and graduate training opportunities. EPCC's aim is to provide opportunities for everyone, including those traditionally under-represented in higher education.

EPCC are involved in Employ.ed, a University of Edinburgh Careers Service programme that provides undergraduate students with **paid internships** during the summer months. The primary goal of the programme is to help promote and grow the skills required to develop and operate complex supercomputing and data science systems, while helping young people find their way into careers in the HPC field. For example, EPCC recently welcomed three undergraduate Engineering students to offer them the opportunity to engage with the large-scale, complex mechanical and electrical systems at EPCC's Advanced Computing Facility.

EPCC have also offered **graduate apprenticeships** at their Advanced Computing Facility for many years. In 2025, EPCC offered two graduate apprenticeships.[64] These apprenticeships allow students to work with EPCC in parallel while completing the final two years of their degree, with one day a week reserved for study leave in addition to gaining credit towards their degree for project work that is completed while working for EPCC.

For students in the senior phase of their school education, EPCC offer short **work placement opportunities**, for example EPCC have hosted students from the Nuffield Research Placement (now STEM learning research placements), the Career ready programme and the In2STEM programmes during their summer break. This year EPCC have hosted around 15 summer placement students, on placements ranging from 1 to 3 weeks.

Beyond direct job creation and workplace training opportunities, operating ARCHER2 brings wider benefits to the Edinburgh area and neighbouring councils, for example:

- **Supply chain spending:** Local companies benefit from providing goods and services to EPCC to support operation of the computing centre, bringing increased revenue to local firms and contributing to local economic growth.
- **Enhanced postgraduate training:** Access to ARCHER2, alongside CIRRUS, a Tier-2 HPC system hosted at EPCC, benefits students enrolled in the University of Edinburgh's High Performance Computing postgraduate programmes, equipping them with hands-on experience and practical advanced computational skills.

[64] See: https://www.epcc.ed.ac.uk/whats-happening/articles/introducing-our-new-system-administration-interns-and-graduate-apprentices?utm_source=chatgpt.com

- **Visitor spending:** Events hosted at or around EPCC such as workshops and training events bring visitors who spend on hotels, restaurants, and transport, benefitting the local economy.
- **Reputation and investment attraction:** Being home to a leading supercomputing facility boosts the region's profile, and can help attract further research funding, technology companies, and skilled professionals to the area.

6.2.2 Contribution to policy making

R&D undertaken on ARCHER2 includes many areas of national importance and policy relevance in diverse areas such as climate science; public health and epidemiology (e.g. large-scale simulations of disease spread); energy system modelling, battery research, and fusion energy research; and materials science and drug discovery, among others.

Insights gained from this research can, in turn, inform innovation policy and regulation, support government priorities such as the UK's net zero targets, and contribute to international policy efforts. For example, large-scale climate simulations undertaken on ARCHER2 help inform the IPCC's climate scenario analyses, which assess the projected future impacts of different global emissions pathways and guide climate mitigation and adaptation strategies.

In response to the user survey, 10% of respondents said their research or work on ARCHER2 has contributed to policy impacts or fed into policy making. Where these impacts are reported back to UKRI, they are collated in their research outcome databases. Analyses of these databases identified at least 70 policy contributions. A breakdown of these by type of contribution is provided in Table 13.

Table 13 Breakdown of policy impacts of ARCHER2 work

Type of policy contribution	#
Citation in other policy documents	3
Contribution to a national consultation/review	11
Contribution to new or improved professional practice	9
Implementation circular/rapid advice/letter to e.g. Ministry of Health	6
Influenced training of practitioners or researchers	22
Membership of a guideline committee	6
Participation in a guidance/advisory committee	13
Grand Total	**70**

Source: London Economics analysis of UKRI ResearchFish™ data

Uptake / citation of ARCHER2 publications in policy-related literature

Publications from scientific research undertaken on ARCHER2 are also frequently cited in policy-related literature (PRL). Of the 1,342 ARCHER2-related publications that were identified in Open Alex, 87 (6.5%) were identified through Overton[65] as having been cited in PRL documents. Most (87%) of these PRL-cited papers were in the fields of Environmental Science (66% of the papers cited) and Earth & Planetary Science (22%).

[65] The largest policy and grey literature database, containing more than 18 million documents from across the world.

The publications were cited 322 times in total across 245 different PRL documents[66]. These documents were published by a range of different organisations and groups, including Inter-Governmental Organisations (IGOs) (108 PRL documents citing ARCHER2-related publications), Government departments agencies and committees (75), Think Tanks (37), Legislative Bodies (13) and NGOs (12).

The citing PRL documents that were not published by IGOs (incl. EU institutions) were associated with 20 different countries, including most commonly the UK (28), the USA (28), and France (14). There were 28 UK-based organisations in total, including most commonly:

- UK Parliament Select Committees (8 citations)
- The UK Committee on Climate Change (3)
- The Met Office (3)

Box 17 HPC modelling for next-generation battery materials in electric vehicles

New modelling using ARCHER 2 will enable identification of sustainable materials capable of producing low cost, reliable and long-range batteries for electric vehicles.

Development of cheaper, more reliable batteries is essential to support the UK's transition to electric vehicles (EVs). Lithium-rich materials are promising candidates for the positive electrode (cathode), potentially offering longer ranges at lower costs. However, they are prone to structural degradation, resulting in performance loss over time. The CATMAT consortium, led by the University of Oxford, used the ARCHER 2 supercomputer to simulate battery processes of different lithium-rich material compositions at the nanoscale. Their work led to greater insights into these materials that are now helping to guide academic and industry partners in material design and manufacturing of batteries.

Further detail is provided in Annex A.6.13.

Source: Technopolis analysis based on information gathered via desk research and consultation with Prof. Saiful Islam.

6.2.3 Public outreach

EPCC runs a range of public outreach and engagement activities to demonstrate to people across the UK the relevance of high-performance computing and how everyone benefits from it. In particular, outreach activities seek to reach young people and show them that careers in science, technology, engineering and mathematics (STEM) are accessible by them. To achieve this, EPCC undertakes a number of activities including:

- Generate impact at large science festivals around the UK, such as the Big Bang Fair and New Scientist Live

[66] There may be a degree of double counting in relation to the analysis of citations within policy-related literature. Policy documents usually lack identifiers such as ISBNs or DoIs that can be used to uniquely identify them. Therefore, Overton does regularly run a manual disambiguation process to check for and remove duplicates (e.g. where a website lists multiple language versions, or where a document is hosted in multiple parts of the same organisation). However, by design, Overton does not disambiguate across different policy sources. As such, the same policy document can be included twice, as long as two different organisations are hosting it. To understand the degree of overestimation, a manual check of the titles of the 374 PRL documents that are included within the analysis was undertaken, looking for instances where similarly titled documents have been identified in different sources. This indicated that the number of unique PRL documents is probably closer to 320.

- Attend multiple smaller (local) events including school visits and careers events
- Hold workshops, insight days, talks and drop-in sessions
- Provide online materials accessible on demand
- Offer work experience for high school students

Between 2022[67] and 2025 EPCC ran **24 outreach events** across 68 days. On average, five staff members were involved in each event, with a combined effort of 374 days between staff members across all events. Staff members carrying out outreach activities are all experienced staff who work in HPC and data science technical, managerial, and training roles from a range of different scientific disciplines.

In total, events attended or held by EPCC attracted **approximately 245,000 visitors** across all events. Some of these events were aimed at school children, some were aimed at students, and others targeted the general public. A breakdown of the target audience of events is provided in Figure 44.[68]

Figure 43 Public outreach efforts from EPCC

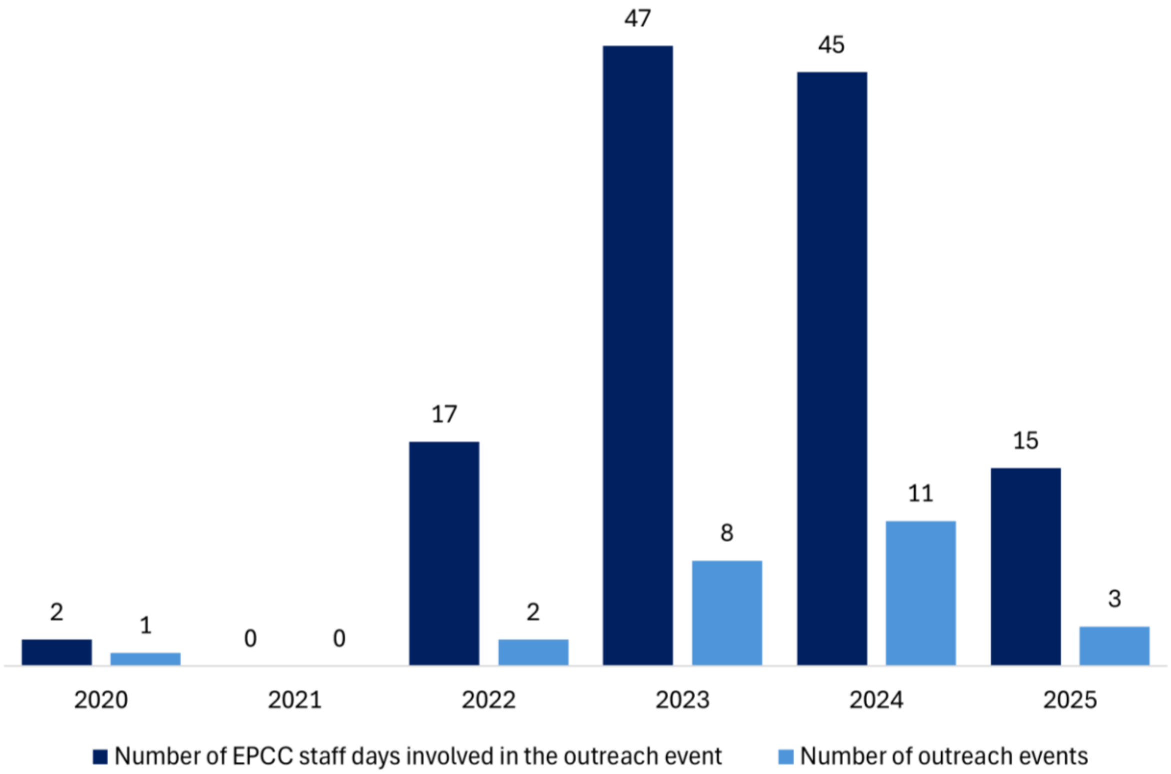

Note: There were 29 outreach events in total across 69 days, with a combined attendance of 245,347 between 2020 and 2025.

Source: London Economics analysis of EPCC data

[67] No outreach events were recorded in 2021.

[68] It should be noted that work experience provided to students, sometimes running for weeks at a time, is counted under public outreach and explains why despite there only being 363 attendees to student-only outreach, it makes up 35 (51%) of the total outreach days. In contrast, the events open to the public are often large science fairs, with the 28 outreach days attracting approximately 175,000 visitors in the events attended by EPCC between 2020 and 2025.

Figure 44 EPCC outreach events by target demographic

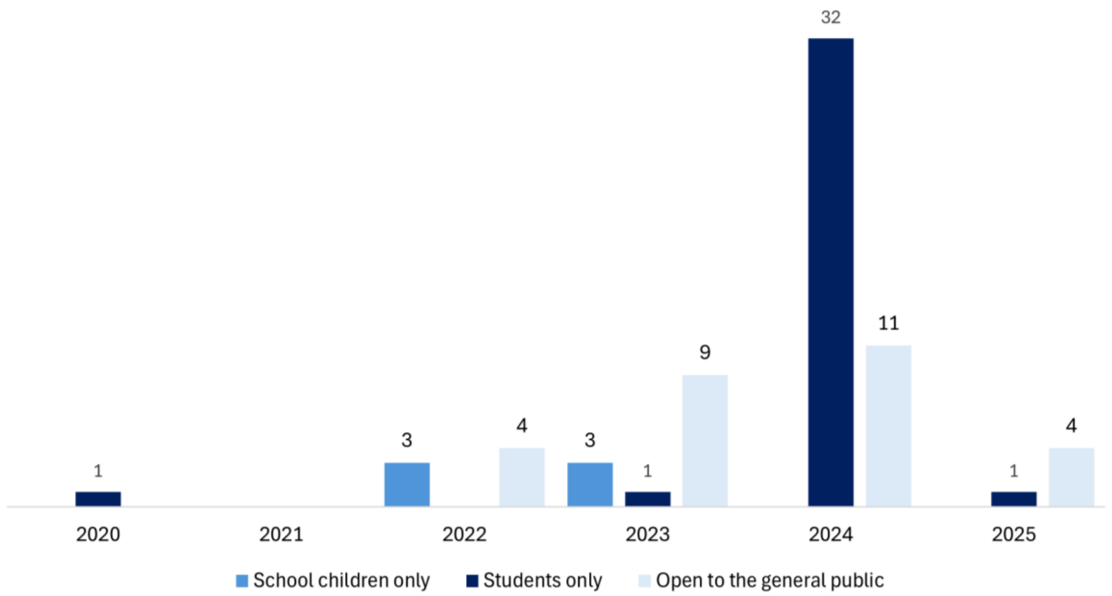

Note: Based on the 69 outreach days provided by EPCC.

Source: London Economics analysis of EPCC data

Box 18 Outreach Highlight: Wee Archie and Public Engagement

Some time ago, EPCC built a Raspberry Pi cluster called Wee Archie to demonstrate the principles behind a full-sized supercomputer.

EPCC designed the suitcase sized cluster to be portable as they cannot carry around a real supercomputer and just looking at a laptop "connected" to a remote system is not that engaging.

The Wee Archie unit is small, light and far less powerful than real supercomputers but it does allow EPCC to show how the large real supercomputers work, what they are made of, and the basics of how they are assembled.

Wee Archie has been a key part of EPCC's outreach activities over the lifetime of the ARCHER2 service. The original Wee Archie design was so popular that EPCC had to build multiple units to meet demand. To ensure a consistent and usable experience across the entire Wee Archie family, the software stack has been redesigned to allow attendees at events to be able to try running a simulation from their phone or tablet any time they are near a Wee Archie.

Wee Archie has been used at numerous events and science festivals across the UK (e.g. London, Birmingham, Liverpool, Edinburgh, etc.) to demonstrate the value of supercomputing to society and to show how supercomputers work. For example, the Big Bang Fair in Birmingham has 20,000+ school children engaging in science activities.

About Wee Archie

Wee Archie is made-up of 18 Raspberry Pi computers, which behave like supercomputer nodes. Each one is a tiny, single-board computer, size 85mm x 56mm x 17mm. Each has a central processing unit (CPU) to execute computer code. 16 of the nodes in Wee Archie are compute nodes, the other two are controller nodes.

The Raspberry Pis are able to communicate with each other with a series of Ethernet cables and 3 switches. In order to share data with other nodes, they must send messages to one another across the Ethernet network.

To ensure a consistent and usable experience across the entire Wee Archie family, the new software stack means a custom application is no longer required to use Wee Archie, all that is needed is a web browser and a WiFi connection.

Prior versions of Wee Archie software did not showcase a good example of a scheduling system or multi-user capabilities. The new software can handle multiple users at a time, and users may have to wait their turn for their job to run, just like with a real supercomputer.

Games

There are multiple games available for Wee Archie for users to play. They are inspired by real life simulation experiments that were run on the previous ARCHER system, and designed to show users the kind of real-world problems scientists use supercomputers to solve. They are also good examples of how computer simulations make design and testing quicker, safer, and cheaper than doing the same experiments in real life.

For example:

- **Design a Wind Turbine:** The user uses three sliders to experiment with different angles, curvatures, and thicknesses of a turbine blade before simulating how much power their turbine blade, turbine and wind farm can generate. Attendees then see how many houses their wind farm can heat and light.
- **Coastal Defences:** The user is given a budget, and a coastal town prone to tidal waves, and is asked to place defences in the sea to reduce the size much as possible. The smaller the waves, the smaller the cost of repairs to damaged buildings after the emergency. The further from the coast and the deeper the ocean, the more expensive each defence becomes, so it is important to weigh up the cost versus effectiveness of each placement. Once the user has placed their defences, Wee Archie simulates a wave approaching, and how the height is affected by the user's design.

Source: EPCC

PART III: SOCIOECONOMIC VALUATION AND FINAL REFLECTIONS

This part discusses:

- **Section 7** outlines the socioeconomic valuation of the benefits from ARCHER2, including both individually valued benefit streams, aggregated economic impacts, and cost-benefit/ROI analysis.
- **Section 8** provides a qualitative exploration of what a loss of a national Tier-1 HPC capability would entail for UK science in practice.
- **Section 9** provides concluding remarks.

Image source: Dr Sébastien Lemaire, EPCC, University of Edinburgh: Effects of "wake steering" on wind turbine flows. The upwind turbine on the right is rotated to steer its wake away from the downstream turbine increasing the overall efficiency.

7 Socioeconomic valuation of benefits flowing from ARCHER2 – cost benefit analysis

This section presents the economic valuation of the benefits to the UKand costs to the public purse associated with ARCHER2 over its lifetime. All costs and benefits are presented in 2025 real discounted prices, using the HM Treasury Green Book discount rate of 3.5%.

7.1 Counterfactuals underlying the evaluation

The evaluation assessed the benefits and costs of ARCHER2 against two counterfactual scenarios:

- A **do-nothing scenario** under which the ARCHER2 capabilities would not be provided. Under this scenario there would be no EPSRC/NERC funded national supercomputing capability. This enables us to assess the total benefits associated with ARCHER2.
- A **business-as-usual scenario** under which the ARCHER2 investment would not have taken place. Instead, the previous ARCHER capability would have been retained. This enables us to value the additional benefits that investment in ARCHER2 has enabled.

The do-nothing scenario, while not realistic in practice, is the most intuitive counterfactual. It is the typical counterfactual used in many evaluations as it **allows for assessment of the total benefits flowing from ARCHER2**. This means it aligns best with the key evaluation question ("What are the scientific and economic benefits of ARCHER2?"). It was also the counterfactual used in the previous evaluation and so allows direct comparison to that study.

The **business-as-usual** scenario is useful to showcase the additional benefits of ARCHER2 over the previous generation of ARCHER. However, it means that many of the benefits flowing from ARCHER2 will not be assigned a monetary value, as only additional benefits over and above ARCHER1 are counted.

A qualitative exploration of what a counterfactual scenario where ARCHER2 does not exist would entail in practice is provided in Section 8.

7.2 Assessment of overall economic benefits and costs

7.2.1 Central valuations of benefits

Table 14 provides the central valuations of the benefits flowing from ARCHER2 and associated costs accruing to the public purse. The table should be read as follows:

- Column 1 presents the results of the economic evaluation of the benefits and costs of ARCHER2 vs. the do-nothing scenario under which no EPSRC and NERC funded national Tier-1 capability exists.
- Column 2 presents the results of the economic evaluation of the scenario where ARCHER1 continues to operate vs. the do-nothing scenario under which no EPSRC and NERC funded national Tier-1 capability exists.
- Column 3 presents the additional benefits and costs of ARCHER2 vs. the scenario where ARCHER1 continues to operate (the business-as-usual-scenario).

A breakdown of benefits by benefit stream is provided in Section 7.3. Section 7.4 provides a more granular breakdown of the costs. To assess the total benefits, and costs, over ARCHER2's lifetime (until 21st November 2026), the underlying data was projected forward in-line with historical trends, with an adjustment to account for 2026 being a partial year.

Table 14 Central valuations of economic impacts (£$_{2025}$ prices)

Metric	ARCHER2 (vs do-nothing)	Continuation of ARCHER1 (vs do-nothing)	ARCHER2 vs Continuation of ARCHER1 (business-as-usual)
Total monetised economic benefits	**£4,265.8M**	**£1,463.8M**	**£2,801.9M**
÷ Benefits excluding spillover impacts of research and the intrinsic value of publications and citations*	£517.1M	£203.5M	£313.6M ÷
Total costs accruing to the public purse	**£511.5M**	**£217.2M**	**£294.3M**
Capital expenditure and operating costs for the ARCHER(2) service**	£99.6M	£56.5M	£43.1M
= Grant funding accruing to the public purse	£411.9M	£160.8M	£251.2M =
Total return on public sector investments	**£8.3 : £1**	**£6.7 : £1**	**£9.5 : additional £1**
Return on public sector investments excluding spillover impacts of R&D and the intrinsic value of publications and citations*	£5.2 : £1	£3.6 : £1	£7.3 : additional £1

Note: Column 1 presents the results of the economic evaluation of the benefits and costs of ARCHER2 vs. the do-nothing scenario under which no EPSRC and NERC funded national Tier-1 capability exists. Column 2 presents the results of the economic evaluation of the scenario where ARCHER1 continues to operate vs. the do-nothing scenario under which no EPSRC and NERC funded national Tier-1 capability exists. Column 3 presents the additional benefits and costs of ARCHER2 vs. the scenario where ARCHER1 continues to operate (the business-as-usual-scenario). (*) Research benefits included in this table capture only the direct benefits to researchers in terms of costs avoided compared to accessing cloud (for research that could have been undertaken on the cloud) and the value of additional citations compared to similar publications (for research that could not have been undertaken without ARCHER2). (**) Costs for continuation of ARCHER1 represent ongoing costs only, no additional capital investment for acquisition of hardware assumed.

Source: London Economics

Central valuations of ARCHER2 over its lifetime

The **total present benefits** of ARCHER2 over its lifetime are estimated to be **in the region of £4.3 billion**. This compares to **total present costs** accruing to the public purse **of around £0.5 billion**. The central estimates therefore indicate a **return of £8.3 per £1 of public money invested** into ARCHER2 itself and the R&D undertaken on ARCHER2.

Given the nature of ARCHER2 as a science capability that serves predominantly fundamental academic R&D, the overall benefits estimates are driven by the **spillover impacts resulting from academic R&D**. These are **estimated at £3.7 billion** using the central assumptions (see Section 7.3).

However, valuing benefits from academic R&D is inherently difficult. Therefore, significant uncertainty surrounds these central estimates. To characterise this uncertainty, as well as uncertainty surrounding central estimates of the other benefit streams evaluated, sensitivity analysis, in the form of Monte-Carlo simulations, around the central estimates was undertaken. The results of the sensitivity analysis are presented in Section 7.5.

Excluding R&D spillover impacts and the intrinsic value of publications and citations (and considering only public funding for ARCHER2 itself), **suggests sizeable economic benefits** from ARCHER2's other activities. These are estimated to be **in the region of £517 million**[69]. Compared to the total funding for ARCHER2 itself of around £99.6 million, this suggests **a return of £5.2 for every £1 of public money** invested into ARCHER2.

The sensitivity analysis further shows that the **central estimates provided here are on the conservative side**. The 90% confidence range around the central estimates resulting from the sensitivity analysis shows **that benefits are plausibly in the region of between £4.6 and £24.1 per £1 invested**.

While not directly comparable, estimates are also in line with findings from London Economics' previous 2019 evaluation of EPSRC's investments into HPC[70], which considered all EPSRC's investments into HPC including ARCHER1, ARCHER1's predecessor, HECToR, and the Tier-2 HPC centres. This showed benefits in the range of £3.9bn to £11.6bn in 2025 prices [original estimates of £3bn to £9bn in 2018 prices]. The return on investment estimates at the time were estimated at between £6.5 : £1 to £19.5 : £1.

While the central estimate of this study sits towards the lower end of the previous range, the higher end of the 90% confidence range suggests **ARCHER2 could potentially generate benefits comparable to, or even exceeding, those of the two earlier national systems and Tier-2 centres combined** (see Section 7.6).

Central valuations of continuation of ARCHER1 scenario

Estimated benefits under the scenario in which ARCHER1 would continue to operate are around two-thirds (66% or £2.8 billion) lower than the estimated benefits of ARCHER2.

However, this scenario would also imply a reduction in costs. Cost reductions are estimated at around £294 million (58%) compared to investment in ARCHER2. This estimate captures both reduced costs of hardware investments as well as reduced funding for the lower volume of R&D that could be undertaken on a machine of the size of ARCHER1.

Together, this implies a return on public sector investments, were ARCHER1 to continue operation, of £6.7 per £1 invested (a reduction of around 19% compared to ARCHER2).

Given the large share of benefits ascribed to spillover impacts of scientific R&D, the magnitude of this reduction is predominantly driven by the proportional reduction in science spillovers and associated costs.

Excluding spillover benefits of science, and associated costs, suggests a much more sizeable reduction in the return on public sector investments of 31% (from £5.2 per £1 invested for ARCHER2 to £3.6 per £1 invested under the continuation of ARCHER1 scenario).

[69] Note: Research benefits included in this figure capture only the direct benefits to researchers in terms of costs avoided compared to accessing cloud (for research that could have been undertaken on the cloud) and the value of additional citations compared to similar publications (for research that could not have been undertaken without ARCHER2). The reason for this is that, in the absence of ARCHER2, the grant funding for research that could not have been undertaken without ARCHER2 would likely have been spent anyway on, for example, other (non-HPC) research, which would have also led to publications. Therefore, in this calculation we exclude the intrinsic value of these publications and only consider the additional value (in terms of additional citations compared to similar publications) generated from ARCHER2 simulations.

[70] London Economics (2019). The impact of EPSRC's investments in High Performance Computing infrastructure

The assumptions underlying the benefit estimations for this scenario are summarised in Table 16. Assumptions underlying the reduction in costs are discussed in Section 7.4. Further details are provided in the methodological annex (Annex 3). In practice, keeping ARCHER operational may not have been feasible due to hardware failures and the difficulty and high costs of finding replacement parts.

Additional benefits of ARCHER2 over and above benefits of continuation of ARCHER1

Comparing the total estimated benefits and costs of ARCHER2 to the total estimated benefits that could have been achieved had ARCHER1 continued to operate provides an estimate of the additional benefits of ARCHER2 relative to the business-as-usual scenario.

As highlighted in the previous section, this comparison indicates additional benefits of ARCHER2 of £2.8 billion and additional costs of ARCHER2 of £294 million over and above what may have been achieved had ARCHER1 continued to operate. Together, **this implies additional benefits of £9.5 per additional £1 invested in ARCHER2 over and above the costs had ARCHER1 continued to operate.**

Excluding the spillover benefits of science, and associated costs, implies additional benefits of ARCHER2 of £314 million. This compares to additional costs of £43 million, implying a return on investment of £7.3 per additional £1 invested over and above the costs incurred had ARCHER1 continued to operate. This suggests that the additional benefit produced by investing in a higher-capacity Tier-1 service far outweigh the additional costs incurred.

7.3 Central valuations of economic impacts by benefit stream

Table 15 provides a breakdown of the central valuations of economic benefits for each benefit stream that could be monetised. In addition to the benefits presented in this table, there are significant unmonetiseable benefits. For example, this includes benefits to the UK's international standing in high-performance computing and international science competitiveness.

Further, some of the monetised benefits are likely to be underestimates. For example, the value of software was estimated using a production cost approach quantifying the time-cost of developing the software. Actual benefits to the scientific community and private organisations using the software are likely to be substantially higher, though very difficult to monetise.

The central valuations therefore represent a conservative estimate of the overall benefits flowing from ARCHER2. The approaches used to value each of these benefit streams, as well as important considerations underlying the economic assessments, are briefly summarised in Table 16. Further details are provided in the methodological annex (Annex 3).

Table 15 Central valuations by benefit stream (£$_{2025}$ prices)

Benefit stream	ARCHER2 (vs do-nothing)	Continuation of ARCHER1 (vs do-nothing)	ARCHER2 vs Continuation of ARCHER1 (business-as-usual)
High impact & quality science	**£224.0M**	**£57.4M**	**£166.5M**
Direct benefit of access to researchers	£188.9M	£37.8M	£151.1M
Value of publications and citations *(Value of additional citations compared to similar publications)*	£29.5M *(£0.4M)*	£17.3M *(£0.2M)*	£12.2M *(£0.2M)*
Value of software	£5.2M	£2.2M	£3.0M
Economic Impact	**£3,866.4M**	**£1,290.8M**	**£2,575.6M**
Spillover impacts of scientific R&D	£3,719.5M	£1,243.2M	£2,476.2M
Contribution of direct access by industry to UK GVA	£75.5M	£0.0M	£75.5M
Value of spinouts and IP	£71.5M	£47.6M	£23.8M
Highly skilled community	**£161.2M**	**£106.7M**	**£54.5M**
Direct provision of training	£7.4M	£2.4M	£5.0M
HPC graduate premium and associated exchequer benefits	£153.8M	£104.3M	£49.6M
Improved local societal impact & high-quality service	**£14.2M**	**£8.9M**	**£5.3M**
Public outreach	£7.6M	£7.6M	£0.0M
Specialised helpdesk support	£6.6M	£1.3M	£5.3M
Total monetised economic benefits	**£4,265.8M**	**£1,463.8M**	**£2,801.9M**

Note: Column 1 presents the results of the economic evaluation of the benefits and costs of ARCHER2 vs. the do-nothing scenario under which no EPSRC and NERC funded national Tier-1 capability exists. Column 2 presents the results of the economic evaluation of the scenario where ARCHER1 continues to operate vs. the do-nothing scenario under which no EPSRC and NERC funded national Tier-1 capability exists. Column 3 presents the additional benefits and costs of ARCHER2 vs. the scenario where ARCHER1 continues to operate (the business-as-usual-scenario). ***Source: London Economics***

Table 16 Summary of evaluation approaches used

Grouping	Benefit stream	Evaluation approach	Important considerations	Adjustments for Continuation of ARCHER1 scenario
High impact & quality science	Direct benefits of access to ARCHER2 for academic users	To monetise the direct benefits to researchers using ARCHER2, an **avoided cost approach** was used. This approach compares the notional cost of access to ARCHER2 to the higher cost of undertaking R&D using commercial cloud alternatives. The approach only quantifies avoided costs for the share of R&D that could have been undertaken in the absence of ARCHER2, as indicated by the user survey, and for which researchers indicated that they would have considered commercial cloud as an alternative.	The survey results indicate that for the share of R&D that could be undertaken without ARCHER2, some of this research could be undertaken using alternative UK publicly funded or international HPC capabilities, while for some research cloud could be used as an alternative. Only the share for which cloud could be used as an alternative was monetised in the evaluation of costs avoided. For the share that could use alternatively publicly funded HPC, additionality of ARCHER2 was assumed to be zero and no benefits were monetised. In practice, other HPC may not provide sufficient capacity in the absence of ARCHER2 for researchers to move their work to these machines. Therefore, the central results are likely an underestimate.	Assumed reduction in capacity of five times in line with the extra capacity provided by ARCHER2 over ARCHER1. No degradation over time was assumed as core-hours stay constant.
	Value of publications and citations	Existing **literature estimates** for the value of academic publications were used to monetise publications. These are predominantly based on **a production cost approach**, using typical time effort of researchers to value the cost of producing academic publications. A small number of studies used other methods including the prices of publications and Willingness-to-Pay estimates. To value citations, the bibliometric results from Technopolis were combined with estimates of researchers' time to read a publication were used.	Collaborations directly impact other outcome metrics. For example, production cost estimates of publications typically already account for the time-value of collaborators. Further, collaborations lead to higher quality publications, which are, as a result, cited more often. The economic value of collaborations is therefore already implicitly captured in other monetised benefit streams. To avoid double counting collaborations are therefore not valued separately.	Assumed that ARCHER1 has reached capacity. Therefore, 2020 levels of publications are assumed throughout. In addition, a small degradation was assumed over time (5% in the central case). To represent a reduction in quality of science, an effective reduction of 17% in the number of citations per publication was assumed.
	Value of software	A **production cost approach** was used to value software. This uses insights from the user survey on the average time spent on developing software products combined with researchers' salary data.	While the production cost of software underestimates the actual benefits of the software (i.e., improving speed of simulations and allowing more granular modelling), part of these wider benefits are already implicitly captured in the estimates of other benefit streams where the improved software is used for further R&D on ARCHER2. Nevertheless, the production cost should be seen as a conservative estimate as efficiency and quality improvements from software on other HPC systems are not monetised.	Similarly to publications, it was assumed that ARCHER1 had reached capacity, so 2020 values for software products produced were used throughout and a small 5% degradation over time was applied.

Grouping	Benefit stream	Evaluation approach	Important considerations	Adjustments for Continuation of ARCHER1 scenario
Economic impacts	Spillover impacts of scientific R&D	Spillover impacts of scientific R&D are notoriously difficult to monetise. Best estimates were produced using the **spillover estimates from the economic literature** presented in Box 20 in Annex A3.2.3 Given the significant uncertainty surrounding spillover estimates from R&D, a relatively wide uncertainty range was placed around the central estimates for the sensitivity analysis and a conservative central estimate on the lower range of this interval chosen. The analysis also accounted for a time-lag between R&D taking place and spillover impacts materialising based on best literature estimates.	As a scientific resource, economic spillovers from R&D are the most important benefit flowing from ARCHER2. However, given the significant uncertainty surrounding estimates of R&D spillovers, return on public sector investment figures were also calculated excluding R&D impacts.	Similarly to publications, it was assumed that ARCHER1 had reached capacity, so 2020 values for science funding were used throughout and a small 5% degradation over time was applied. A proportional reduction in cost was also assumed. In addition, a reduction in line with that assumed for citations was applied to spillover multipliers to account for a reduction in quality of science.
	Contribution of direct access by industry to UK GVA	Industry receives benefits in terms of reduced costs of access compared to commercial services, as well as potential long-term benefits in terms of cost reductions as a result of HPC usage, efficiency gains, improvements to existing products or services, or by contributing to the introduction of new products or services, or enhanced business growth. Industry benefit estimates were based on best available evidence of average annual **profit increases to firms accessing HPC** and subsequently converted to GVA contributions. As with spillover impacts of scientific R&D, the analysis also accounted for a time-lag between R&D taking place and impacts materialising.	Given the early nature of R&D undertaken on ARCHER2, these benefits are difficult to monetise accurately. A reasonable range was constructed based on the limited estimates that were available. The low end of the range is based on avoided costs relative to investments into commercial HPC, while the top end of the range is based on obtained estimates of a limited sample of profit increases due to HPC use obtained previously in the 2019 study of ARCHER1. Given the significant uncertainty, a central value towards the lower end of the range was chosen. Moreover, given the uncertain returns of early-stage R&D, it was assumed that profit increases only materialise for a share of firms.	It was assumed that due to the reduced capacity of ARCHER1, the full capability would be used for academic R&D purposes. Therefore, benefits of industry access were assumed to be nil.
	Value of spinouts and IP	The value of spinouts and intellectual property was estimated based on best available estimates from the **economic literature** on the value of a patent and the economic contribution of spinouts.	Only limited data on spinouts and intellectual property was available in UKRI's research outcome database, ResearchFish^{TM}. Estimates obtained from the user survey suggest that data in ResearchFish^{TM} are significant underestimates, though the precise number of spinouts benefitting from and intellectual property resulting from R&D undertaken on ARCHER2 is difficult to estimate precisely. To capture this uncertainty in the analysis, a	Given the low and sporadic incidence of spinouts and IP in ResearchFish^{TM} reductions in the number of spinouts and IP were assumed to be proportional to the assumed reduction in software products.

Grouping	Benefit stream	Evaluation approach	Important considerations	Adjustments for Continuation of ARCHER1 scenario
			right skewed distribution (truncated normal) was used to characterise this uncertainty.	
Highly skilled community	Direct provision of training	Benefits of training are estimated using **an avoided cost approach**. That is, the benefit per attendee, per training day is calculated as the cost of attending a similar training from a commercial provider.	Given the specific nature of training provided by EPCC it was difficult to obtain direct one-to-one comparator courses provided by commercial organisations. Therefore, a range of the most similar offerings were selected. These included cloud computing, data science, and advanced programming courses.	A reduction in training days in line with an assumed reduction in user numbers was assumed. The reduction was estimated by assuming user numbers of ARCHER1 had peaked at its end of life.
	HPC graduate premium and associated exchequer benefits	Benefits of PhD and postdoctoral students with HPC training are calculated as the **lifetime earnings premium HPC graduates obtain** compared to a counterfactual group. The counterfactual used in the analysis are students with any postgraduate degree aged between 25 and 34 (i.e. recent graduates) in the same professions. Benefits are split into two parts: benefits accruing to graduates themselves (the HPC earnings premium) **and benefits accruing to the UK exchequer**, calculated as the additional income tax, national insurance contributions, and VAT accruing to the exchequer.	The analysis only estimated a graduate premium for graduates moving into UK industry. No graduate premium was assumed for graduates staying in academia or students obtaining jobs outside of the UK. To do this the analysis relied on estimates of career destinations and the HPC skills premium obtained by graduates obtained from the user survey. Given the self-reported nature of the graduate premium, a 50% uncertainty band was placed around the central estimate obtained from the user survey.	A reduction in the number of PhDs and postdoctoral students trained in line with the assumed reduction of users was used. No reduction to the graduate premium was assumed with the rationale that training with a lower capacity machine would not impact the skills learned.
Improved local societal impact & high-quality service	Public outreach	To monetise the value of public outreach events, this study estimated the non-market benefits associated with outreach events by estimating the **opportunity cost** for visitors to attend these events.	This approach may underestimate the true value of public outreach events, as it only captures visitors' time cost and not the broader educational, social, or inspirational benefits. It also assumes that all visitors value their time similarly, which may not be the case, though adjustments were made to differentiate between adults, students and children (for which no opportunity cost was assumed).	Outreach activities were assumed to be unaffected. That is, no change in the number of outreach events or visitors reached were assumed.
	Specialised helpdesk support	To quantify the value of helpdesk support, the study estimated the **value of time savings to researchers** from helpdesk support.	Average time savings were based on self-reported figures by users obtained from the user survey. A weighted average time saving was obtained as a central estimate and combined with data on the number of helpdesk queries obtained from EPCC.	A reduction in line with the assumed reduction in science capacity was assumed. No degradation was assumed over time.

Source: London Economics

7.4 Breakdown of costs accruing to the public purse

To assess the net socio-economic benefits of ARCHER2, it is important to understand the costs to the public purse associated with ARCHER2. Costs captured in this study include the direct costs associated with funding of ARCHER2. This includes both capital expenditure for the hardware and associated infrastructure and ongoing operational expenditure. In addition, this study also considers the costs to the public purse of grant funding for R&D undertaken on ARCHER2.

Funding for ARCHER2

Excluding VAT, ARCHER2 cost a total of **£79.1 million** between 2020/21 and 2024/25. This is broken down into **£40 million** of capital funding and **£39.1 million** of resource funding.

Capital funding is broken down into **£36 million** of hardware costs and **£4 million** of capital maintenance costs. Capital funding was provided by BEIS Capital.

Resource funding is broken down into **£7.2 million** of service provision funding, **£10.6 million** of CSE support funding and **£21.3 million** of funding to host the ARCHER2 service at the advanced computing facility in Edinburgh. Resource funding is jointly provided by EPSRC and NERC, where EPSRC OpEx provides 80.8% (**£31.6 million**), and NERC provides 19.2% (**£7.5 million**).[71] A breakdown of these costs per year is provided in Table 40 in Annex 4.

Figure 45 Costs of ARCHER2 (2020/21 – 2024/25)

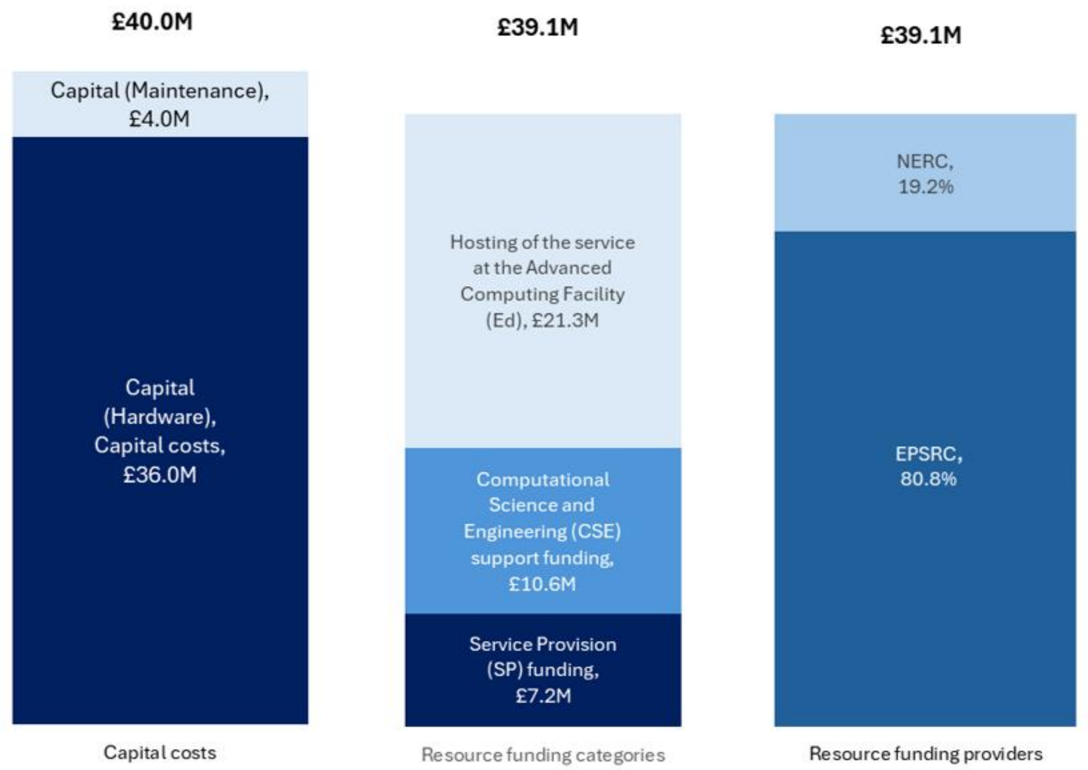

Note: These estimates exclude Capital (Maintenance) costs in 2024/25 as the Full Business Case for ARCHER2 reported these as *tba.*

Source: Full Business Case for ARCHER2

[71] Full Business Case for ARCHER2, UKRI

Grant funding for R&D undertaken on ARCHER2

In addition to the funding for hardware and operation of ARCHER2, researchers also receive funding from the research councils to support the R&D they undertake on ARCHER2. To ensure the study fully captures the costs accruing to the public purse, these funds are also considered as costs.

Table 17 summarises the total grant funding by research council for R&D on ARCHER2 between 2021 and 2025. All these costs are assumed to be borne by the public purse.

Table 17 Initial research grant funding for R&D undertaken on ARCHER2, 2021-2025

Research council	Total funding provided (£$_{2025}$-value)
BBSRC	£1.0M
EPSRC	£58.2M
ESRC	£0.4M
MRC	£9.4M
NERC	£29.8M
STFC	£0.7M
Total	**£99.5M**

Source: London Economics based on ResearchFish™ data

In addition to initial grant funding, researchers also receive further funding from a number of sources. This includes funding from universities and other academic sources, from charities and non-profit organisations, from other private and public sector entities, as well as further grant funding from UKRI.

Table 18 summarises the total further funding by source for R&D on ARCHER2 between 2021 and 2025. Column three indicates whether costs are assumed to accrue to the public purse or not.

Table 18 Further funding for academic R&D undertaken on ARCHER2, 2021-2025

Funding source (UK & international)	Total funding provided (£$_{2025}$-value)	Accruing to the public purse
Academic/University	£22.4M	Yes
Charity/Non Profit	£42.7M	No
Private	£8.6M	No
Public (UKRI)	£349.6M	Yes
Public (non-UKRI)	£27M	Yes
Total	**£450.3M**	-

Source: London Economics analysis of ResearchFish™ data

Total costs captured in the evaluation

To ensure the analysis captures the total costs over ARCHER2's lifetime, these costs were projected forward to 21st November 2026, when ARCHER2 is expected to cease operation. Historical funding values were adjusted for inflation, while future funding was discounted using the HMT Green Book discount rate of 3.5%.

Further, R&D funding was adjusted to capture the share of R&D that relies on ARCHER2 and for which spillover impacts were quantified. This is because funding for R&D that would still be undertaken under the counterfactual, i.e., were ARCHER2 not to exist, would still be incurred by the public purse under the counterfactual.

Table 19 Total costs accruing to the public purse over ARCHER2's lifetime

Type of funding	£2025-value
ARCHER2 funding	£99.6M
Research funding accruing to public purse	£411.9M
Total costs accruing to public purse	**£511.5M**

Note: Captures costs accruing to the public purse projected forward until the end of ARCHER'2 lifetime (2021 to 21st November 2026).

Source: London Economics analysis of ResearchFish™ data

Adjustments for continuation of ARCHER1 scenario

If ARCHER1 were to continue, the initial capital investment for ARCHER2 hardware would be avoided. However, capital costs for maintenance as well as other costs for service provision, CSE support and hosting of the advanced computing facility would still need to be covered. For the purpose of the analysis, these were hypothetically assumed to be in line with costs for ARCHER2.

However, in practice, it may not have been possible or viable to keep ARCHER1 operational. Due to the age of ARCHER1 as an end-of-life machine, the required maintenance may have been significantly higher and/or additional capital investments may have been needed to retain ARCHER1 operational. As such, the assumed cost should be seen as an underestimate of the likely actual cost. In addition, there may have also been other issues such as with cybersecurity, software support, or availability of replacement hardware.

Although ARCHER1 had already reached capacity before the end of its lifetime, 2020 funding values for ARCHER1 were assumed throughout the evaluation period. Further, a small adjustment to capture degradation over time was assumed (5% per year in the central case). Further details are provided in the methodological annex (Annex 3).

Table 20 Total costs accruing to the public purse under continuation of ARCHER1 scenario

Type of funding	£2025-value
ARCHER2 funding	£56.5M
Research funding accruing to public purse	£160.8M
Total costs accruing to public purse	**£217.2M**

Note: Captures costs accruing to the public purse projected forward until the end of ARCHER'2 lifetime (2021 to 21st November 2026).

Source: London Economics analysis of ResearchFish™ data

7.5 Sensitivity analysis surrounding central benefit estimations

To characterise the uncertainty surrounding central benefit estimates, Monte-Carlo simulations around the central assumptions used were undertaken. Table 21 provides a summary of the results of the sensitivity analysis. Figure 46 and Figure 47 provide the results of the sensitivity analysis for the total estimate of benefits for ARCHER2 and for the continuation of ARCHER1 scenarios, respectively. Figure 48 provides the sensitivity analysis results for economic benefits excluding spillover impacts of R&D. The central assumptions used and uncertainty assumptions surrounding the central estimates are set out in Annex 3.

The sensitivity results show that central valuations used are conservative. The average benefit estimate resulting from the sensitivity analysis was around £3.0 billion higher than the central estimate at around £7.3 billion. The average return on public sector investments was £13.4 per £1 of public money invested, compared to the central estimate of £8.3 per £1. The median estimate was £6.6 billion and the median return on public sector investments was £12.9 per £1 invested.

Benefits could be substantially higher. For example, when considering the higher end of the estimates, four-fifths of scenarios (the 80th percentile) point to benefits of around £10.6 billion or less, and ninety-five out of one hundred scenarios (the 95th percentile) point to benefits of roughly £15.0 billion or less. In these more optimistic cases, the return on public sector investments was £18.6 to £24.1, respectively, for every £1 invested.

More conservative scenarios suggest benefits of between £1.9 billion (the 5^{th} percentile) and £3.5 billion (the 20^{th} percentile) and return on public sector investments figures of between £4.6 to £7.6 per £1 invested.

The results for the continuation of ARCHER1 scenario as well as for benefits excluding spillover impacts of R&D are briefly summarised below:

- Sensitivity analysis around the estimates of the economic benefits excluding spillover impacts of R&D (presented in Table 21) indicate benefits in the region of between £0.20 billion [5^{th} percentile] and £1.8 billion [95^{th} percentile]. Return on investment figures ranged between £2.2 and £18.4 per £1 of public money invested into ARCHER2. The mean and median benefit estimates resulting from the sensitivity analysis were £0.8 billion and £0.9 billion, respectively. The mean and median return on investment figures were £7.9 and £8.7, respectively, per £1 invested in ARCHER2. This is again substantially higher than the central estimates which indicate a return on public sector investments of £5.2 per £1 invested.
- In contrast, the sensitivity results for the continuation of ARCHER1 scenario suggest a much narrower range of benefits of between £0.8 billion [5^{th} percentile] and £3.6 billion [95^{th} percentile] with return on investment estimates of between £4.2 and £14.9 per £1 invested. Benefits excluding spillover impacts of R&D under this scenario are estimated at between £0.1 billion and £0.6 billion [5^{th} and 95^{th} percentiles] implying a return on investment of between £1.9 and £9.9 per £1 invested.

Table 21 Sensitivity analysis: Summary of results

Metric	ARCHER2		Continuation of ARCHER1	
	Estimated benefits	Return on public sector investments	Estimated benefits	Return on public sector investments
Total benefits				
Central estimates	***£4,266M***	***£8.3 : £1***	***£1,464M***	***£6.7 : £1***
Mean	£7,247M	£13.4 : £1	£2,155M	£9.7 : £1
95th percentile	£14,971M	£24.1 : £1	£3,574M	£14.9 : £1
80th percentile	£10,610M	£18.6 : £1	£2,950M	£12.8 : £1
Median	£6,623M	£12.9 : £1	£2,149M	£9.9 : £1
20th percentile	£3,479M	£7.6 : £1	£1,310M	£6.5 : £1
5th percentile	£1,877M	£4.6 : £1	£792M	£4.2 : £1
Benefits excluding spillover impacts of R&D				
Central estimates	***£517M***	***£5.2 : £1***	***£203M***	***£3.6 : £1***
Mean	£864M	£8.7 : £1	£291M	£5.2 : £1
95th percentile	£1,832M	£18.4 : £1	£557M	£9.9 : £1
80th percentile	£1,286M	£12.9 : £1	£416M	£7.4 : £1
Median	£783M	£7.9 : £1	£265M	£4.7 : £1
20th percentile	£400M	£4.0 : £1	£167M	£3.0 : £1
5th percentile	£223M	£2.2 : £1	£108M	£1.9 : £1

Note: Based on 10,000 Monte Carlo simulations. Percentiles show the value below which a given percentage of simulation results fall. For example, the 50th percentile is the median – 50% of simulations resulted in benefit estimates below the median benefit estimate.

Source: London Economics

Figure 46 Sensitivity analysis results: ARCHER2

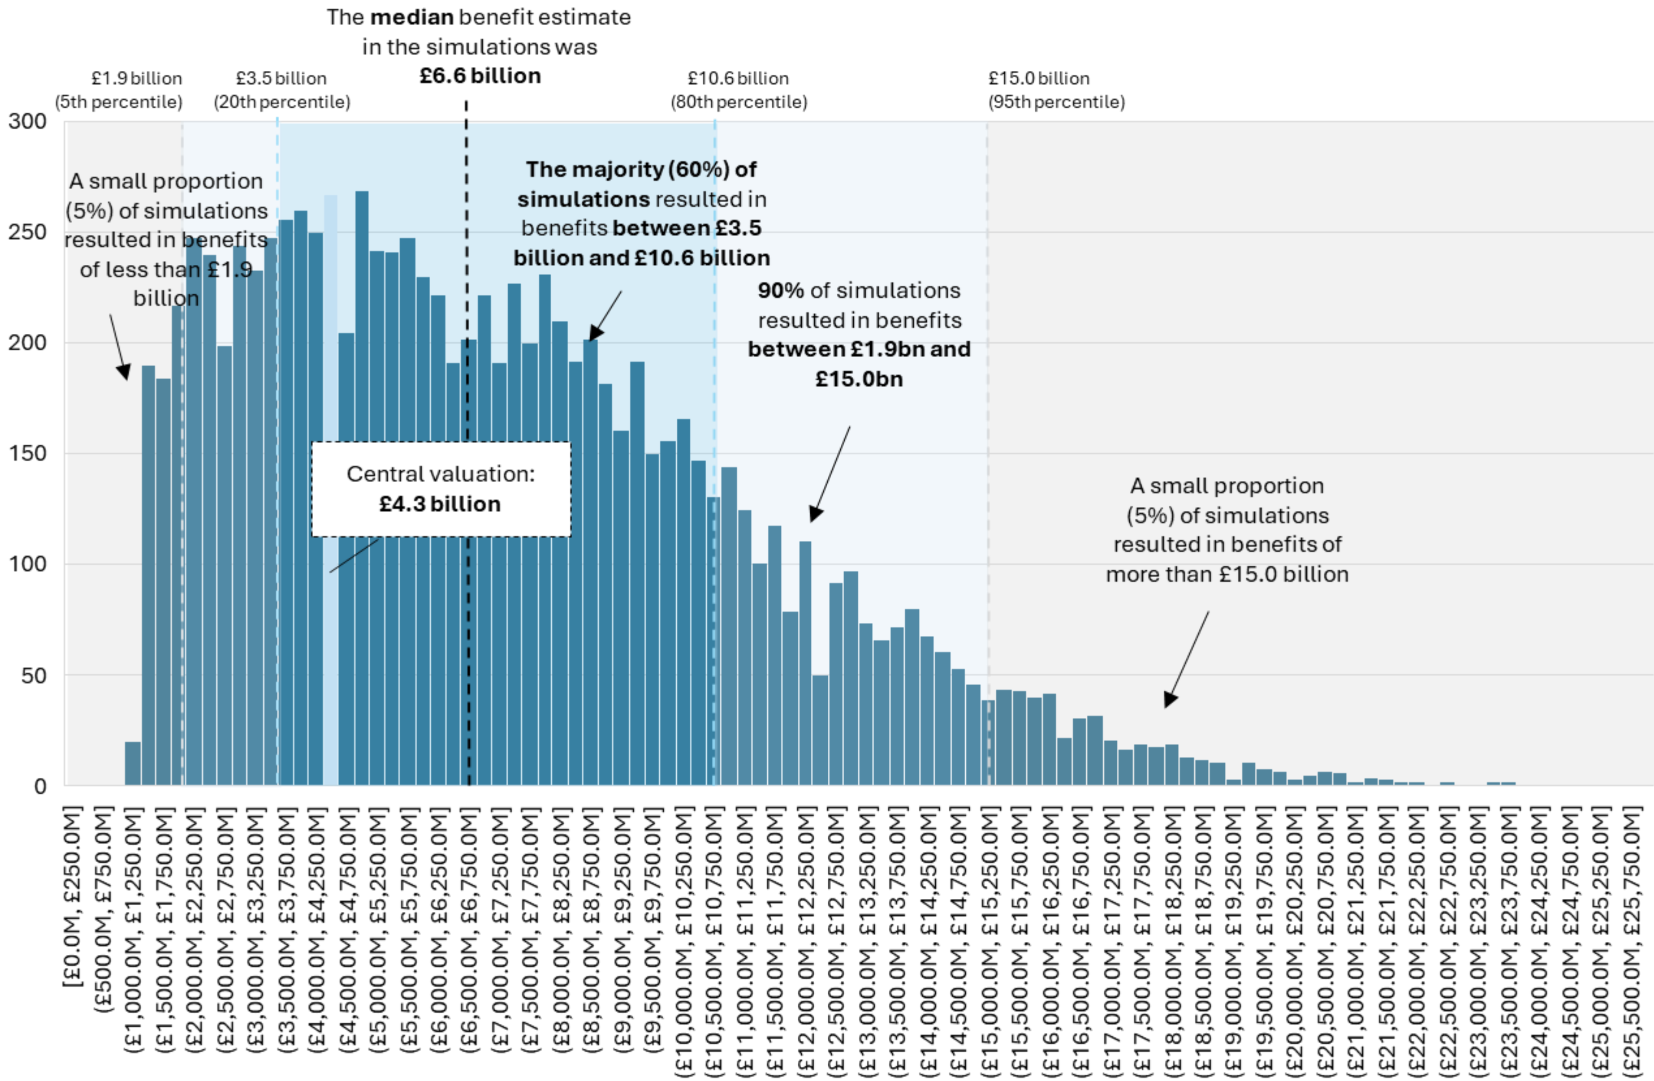

Note: Based on 10,000 Monte Carlo simulations. Percentiles show the value below which a given percentage of simulation results fall. For example, the 50th percentile is the median – 50% of simulations resulted in benefit estimates below the median benefit estimate.

Source: London Economics

Figure 47 Sensitivity analysis results: Continuation of ARCHER1

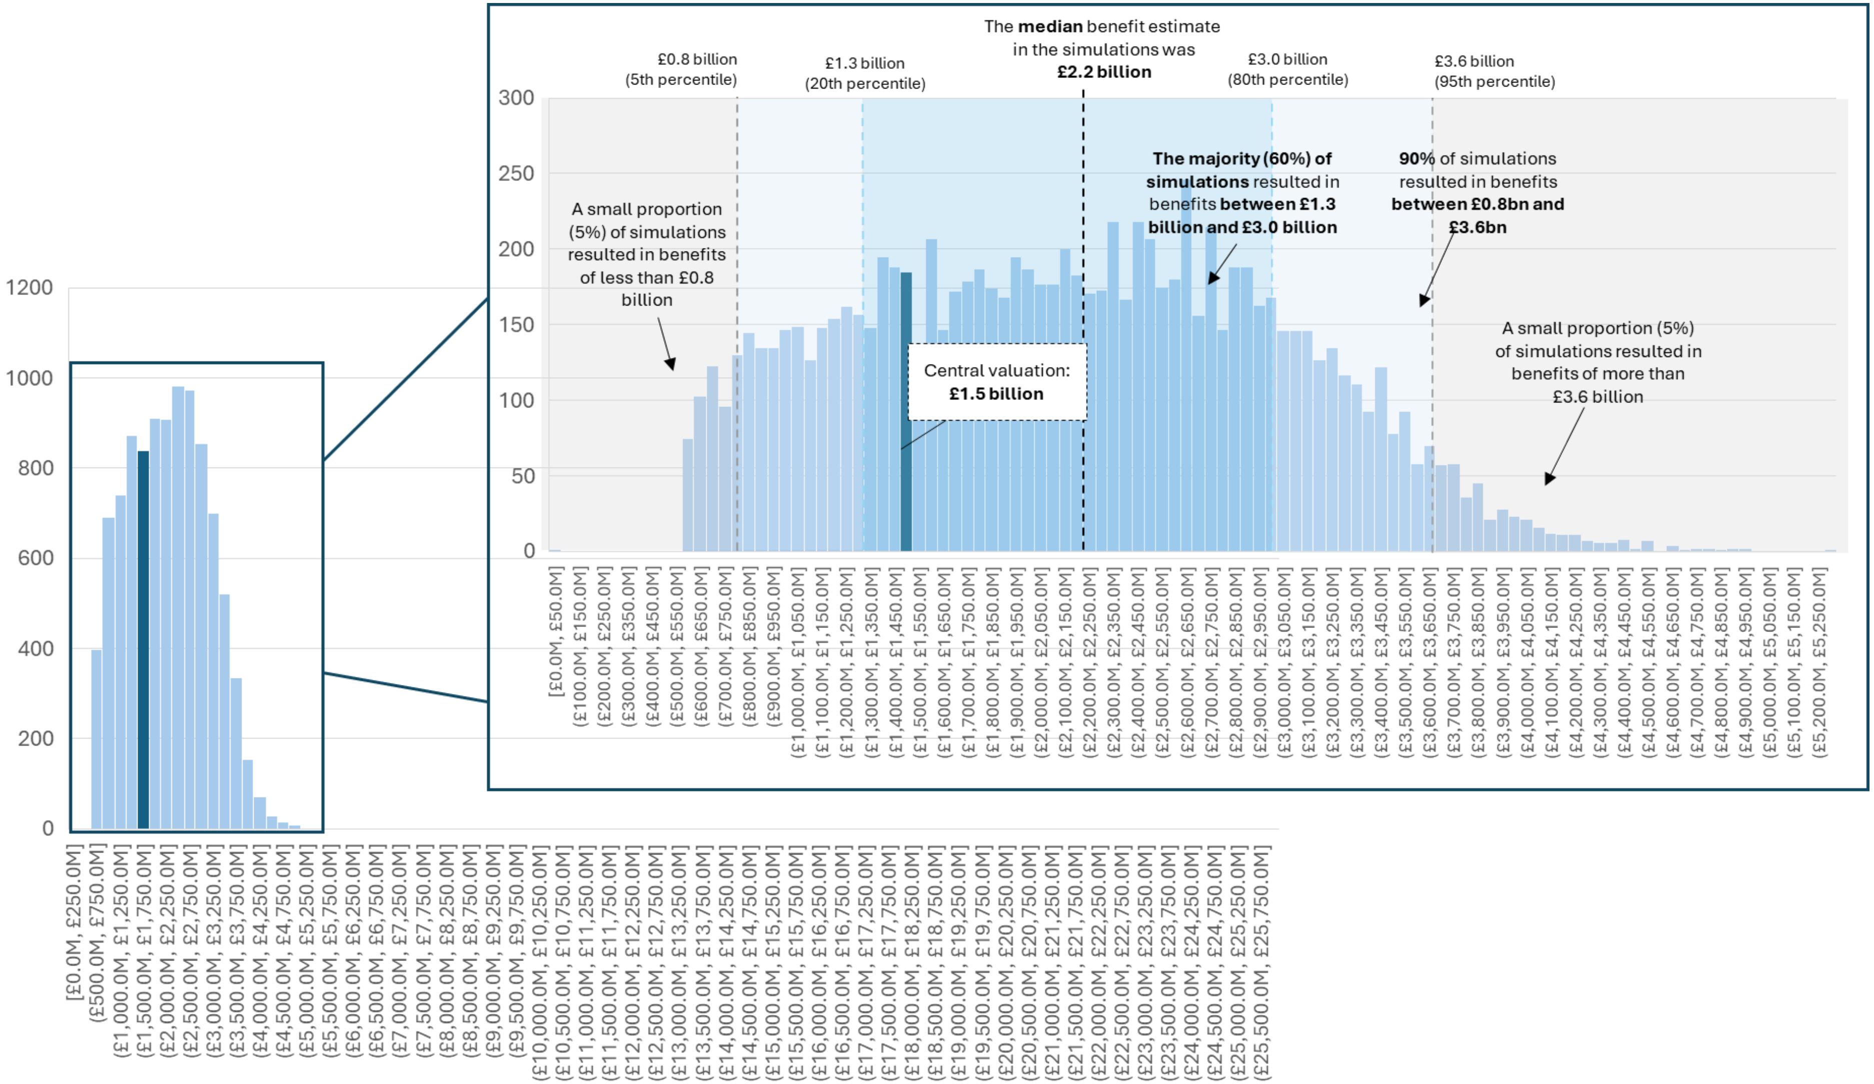

Note: Based on 10,000 Monte Carlo simulations. Percentiles show the value below which a given percentage of simulation results fall. For example, the 50th percentile is the median – 50% of simulations resulted in benefit estimates below the median benefit estimate.

Source: London Economics

Figure 48 Sensitivity analysis results for economic benefits excluding spillover impacts of R&D

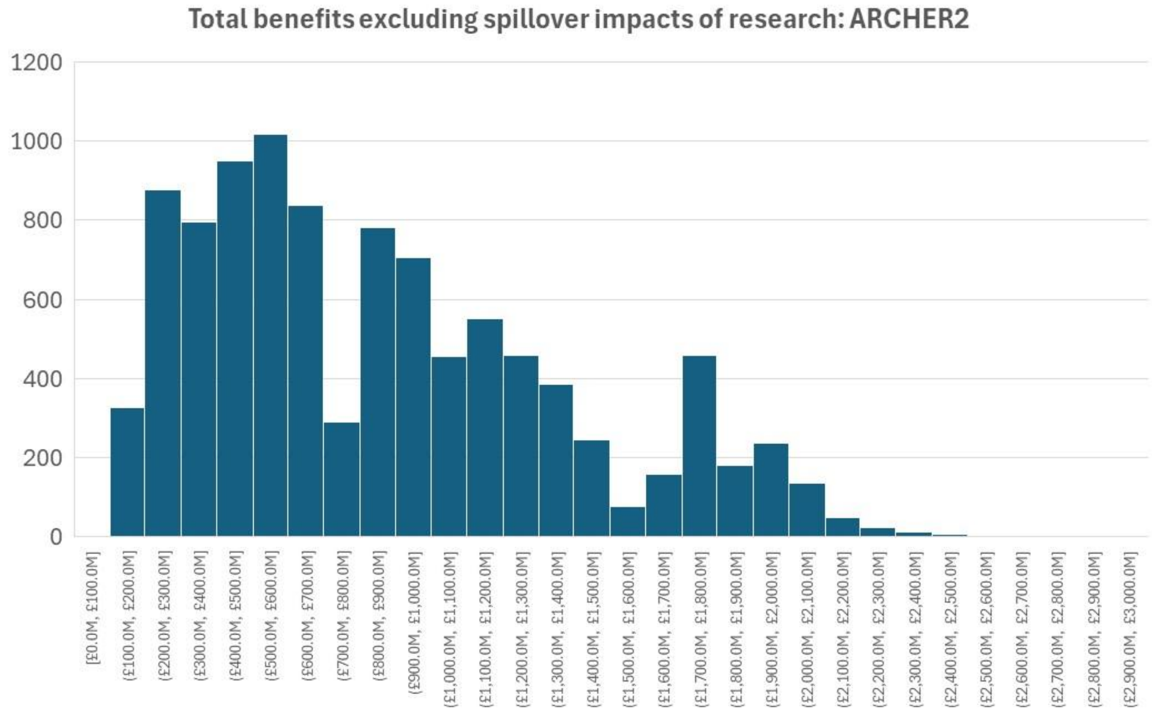

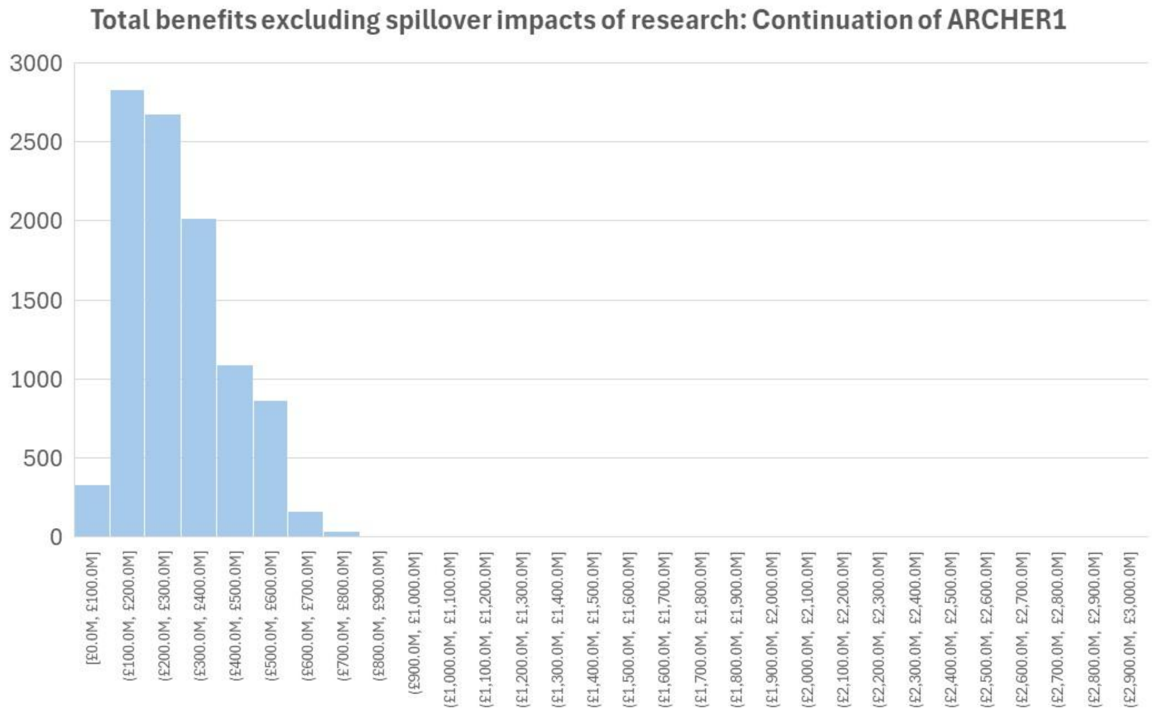

Note: Based on 10,000 Monte Carlo simulations

Source: London Economics

7.6 Comparison to estimated benefits of other evaluations of HPC and science infrastructure

Comparison to previous evaluation of EPSRC's investments in HPC

London Economics' previous 2019 evaluation of EPSRC's investments into HPC[72] showed benefits in the range of £3.9bn to £11.6bn in 2025 prices [original estimates of £3bn to £9bn in 2018 prices]. The return on investment estimates at the time were estimated at between £6.5 : £1 to £19.5 : £1.

However, there are a number of important differences between the present study and the previous evaluation. Therefore, the results are not directly comparable. Importantly:

- The previous study considered all EPSRC's investments into HPC including ARCHER1, ARCHER1's predecessor, HECToR, and the Tier-2 HPC centres. In contrast, the present study only considers benefits flowing from ARCHER2.
- Further, the previous study provided estimates for the minimum likely and maximum likely benefits. The present study adopts a different methodological approach, estimating a central benefit figure accompanied by a 90% confidence interval.
- Finally, there are methodological differences with updated methodologies used where new approaches were seen as more robust. Most importantly, the previous study used estimates of the spillover multiplier by Haskel and Wallis (2010)[73] of £12.7 per £1 invested to estimate economic spillover impacts from research council funding. In contrast, the present study derived estimates based on Frontier Economics (2024)[74]. Frontier Economics found an annual 40% (0.4) rate of return to public R&D investment in perpetuity. Much more conservative spillover multipliers were derived (see Annex A3.2.3) from this figure for the central (£4.6 : £1) and lower (£1.4 : £1) assumptions used in this study. The upper assumption in this study is the same as the maximum used in the previous evaluation (£30 : £1), based on IDC (2014)[75]. However, because of the design of the sensitivity analysis, this high value is very unlikely to be selected in the simulations[76]. This means that the results are more heavily weighted towards the conservative end of the range.

Nevertheless, the results of the present study fall within the range reported in 2019. The central estimate for ARCHER2 stands at £4.3 billion implying a return on public sector investments of £8.3 per £1 of public sector monies invested in ARCHER2. The 90% confidence range around the central estimate is estimated at £1.9bn to £15.0bn with return on public sector investments estimates ranging from £4.6 : £1 to £24.1 : £1.

While the central estimate sits towards the lower end of the previous range, the higher end of the 90% confidence range suggests ARCHER2 could potentially generate benefits comparable to, or even exceeding, those of the two earlier national systems and Tier-2 centres combined.

[72] London Economics (2019). The impact of EPSRC's investments in High Performance Computing infrastructure
[73] Haskel, J., & Wallis, G. (2010). Public support for innovation, intangible investment and productivity growth in the UK market sector (IZA Discussion Paper No. 4772). Institute for the Study of Labor (IZA). Retrieved from https://docs.iza.org/dp4772.pdf
[74] Frontier Economics. (2024). Returns to public R&D: Report for the Department for Science, Innovation and Technology (DSIT).
[75] IDC (2014). EESI-2 Special Study To Measure And Model How Investments In HPC Can Create Financial ROI And Scientific Innovation In Europe
[76] The sensitivity analysis used a truncated normal distribution and assumed the £30 : £1 value is four standard deviations away from the mean. The mean was assumed to be £4.6 : £1, and the distribution was truncated at the lower multiplier estimate of £1.4 : £1.

Comparison to other evaluation studies

Table 22 presents a number of studies which have established benefit-to-cost ratios for HPC and other science infrastructure, providing context for the results of this study. Some of the information is drawn from media releases, where the full studies were not available.

Evaluations of other HPC infrastructure are very limited and estimates vary widely. Among the limited number of studies identified ROI estimates range from £0.9 : £1 to £37.0 : £1. Broader science infrastructure studies identified report more modest returns between £1.0 : £1 and £4.6 : £1. This compares to return on public investment estimates derived in this study of between £4.6 : £1 to £24.1 : £1 with a central estimate of £8.3 : £1 for ARCHER2.

The recent evaluation of CSC's high-performance computing services is particularly noteworthy, reporting high return figures of €25–€37 return per euro invested.[77] However, at the time of reporting the full study had not been published and the media release contained insufficient detail to allow investigation of what drives these high valuations.

Table 22 Estimated benefits from other evaluation studies of HPC and science research infrastructure

Study	Region	Benefit Streams	Benefit measures	BCR / ROI
London Economics (2019). The impact of EPSRC's investments in High Performance Computing infrastructure	United Kingdom	Research & Discovery, Industry, Training and skills development	Return on EPSRC's investment in HPC.	ROI from 6.5:1 to 19.5:1
The Innovation Partnership (2015). FINAL EVALUATION OF HIGH PERFORMANCE COMPUTING WALES	Wales	Research, Industry, Training, Innovation	Return on investment (ROI) by WEFO, UK BIS, HEFCW, and other stakeholders in HPC Wales.	ROI from 2.1:1 to 8.4:1
CSC (2024). SROI of CSC's high-performance computing services studied: €25–37 return per euro invested. Media Release	Finland / Europe	Scientific and industrial innovation, economic and data economy growth, environmental sustainability, strategic infrastructure and national competitiveness	Social return on investment in CSC – IT	SROI from 25:1 to 37:1
IDC (2013). Creating Economic Models Showing the Relationship Between Investments in HPC and the Resulting Financial ROI and Innovation — and How It Can Impact a Nation's Competitiveness and Innovation. Final Technical Report for Contract: DE-SC0008540	United States	Competitiveness, productivity, and innovativeness of U.S. scientists, engineers, researchers, and companies	Sales revenue and profits and cost savings from investments in HPC (based on 208 HPC use cases)	Average Revenue of $30 per HPC $1 for academic sites and $974 per HPC $1 for industrial concerns.

[77] See: https://csc.fi/en/media-release/sroi-of-cscs-high-performance-computing-services-studied-e25-37-return-per-euro-invested/

Technopolis (2018). Hartree Centre Phase 1&2 Baseline Evaluation. Final Report	United Kingdom	Innovation, competitiveness and productivity, direct work with industry, skills and capability development, wider social and economic benefits	Commercial benefits among its phase 1 and 2 users and economic impact from operational expenditure only	Early ROI of £0.92 per £1 invested* (monetised benefits limited due to early timing of the baseline evaluation)
G. Snapp-Childs et al (2024). Evaluating Return on Investment for Cyberinfrastructure Using the International Integrated Reporting Framework	United States	Research, industry, finance, human capital, social capital, natural resources capital	Return on federal investment in the XSEDE.	ROI at least 1.
Bastianin (2018). Social Cost Benefit Analysis of HL-LHC	Switzerland/ Europe	1.The value of training (or human capital formation) for students and early-stage researchers 2.Technological or industrial spillovers for collaborating firms and other economic agents 3.Cultural effects for the public 4.Academic publications and pre-prints for scientists 5.Existence or public good value of the RI for non-users.	Benefits to cost ratio of investment in HL-LHC.	BCR of 1.7:1
RTI International (2024). Assessment of the Retrospective and Prospective Economic Impacts of Investments in U.S. Neutron Research Sources and Facilities from 1960 to 2030	United States	Focus on technological developments (with case studies on GMR Hard Drive Development, Safer Aviation Fuel, Improved Aircraft Components, Weight Loss Medications and Improved EV Performance)	BCR for U.S. neutron scattering facilities from 1960 through 2030	BCR from 2.36:1 to 4.61:1
London Economics (2020). The impact of the UK's public investments in UKAEA fusion research	United Kingdom	Creation of direct employment (e.g., research jobs), creation of direct contract and materials spending in the UK, attraction of foreign investment, generation of spin-offs in the form of new technologies and firms	The return on the UK Government's investments in UKAEA	ROI 2.9:1 (counting only impacts arising from UKAEA activities) ROI from 3.7:1 to 4.1:1 (counting additional impacts)
Diamond (2022). Diamond's socio-economic impact of over £2.6 billion on UK science and economy revealed. Media release	United Kingdom	Research, patents, software and applications, training	Cumulative, monetised economic impact	ROI of around 1.85:1**

Note: (*) Calculated by London Economics based on £27.5 million in commercial benefits among phase 1 and 2 users and £7.1 million in net economic impact as a result of operational expenditure during the first two phases compared with £37.5 million initial capital investment in the centre. (**) Calculated by London Economics based on £2.6 billion of cumulative economic impact compared with £1.4 billion investment made.

Source: London Economics

8 The impact of a loss of Tier-1 HPC capability on UK science

The critical importance of ARCHER2 as a national Tier-1 capability for the UK scientific community was universally emphasised across the range of stakeholder consultations held for this study. As has been demonstrated in the previous sections of this study ARCHER2 is a key enabler for the UK scientific community. Computational R&D, made possible by a Tier-1 capability such as ARCHER2, is firmly established as a fundamental pillar of modern scientific R&D.

8.1 Users' views on the impact of a loss of Tier-1 HPC capability on UK science

> *"ARCHER2 is the only machine in the UK suitable for our capability calculations."*
>
> - **ARCHER2 user survey**

To better understand the impact of a loss of ARCHER2 as a UK Tier-1 capability, the survey undertaken for this study asked users to consider two scenarios: First, a scenario where ARCHER2 was not funded, but its predecessor ARCHER still exists with its previous capacity; and, second, a scenario where neither ARCHER2 nor ARCHER exist.

The impact of no EPSRC/UKRI funded Tier-1 capability

Analysis of responses received to the survey mirror the sentiment raised in consultations. In the case where users were asked to consider a scenario where neither ARCHER2 nor ARCHER exist, almost three-quarters of users (74%) said the impact on their work would be catastrophic. All remaining (26%) of users said that there would be a major impact on their work.

Figure 49 Impact of the absence of a national Tier-1 HPC service

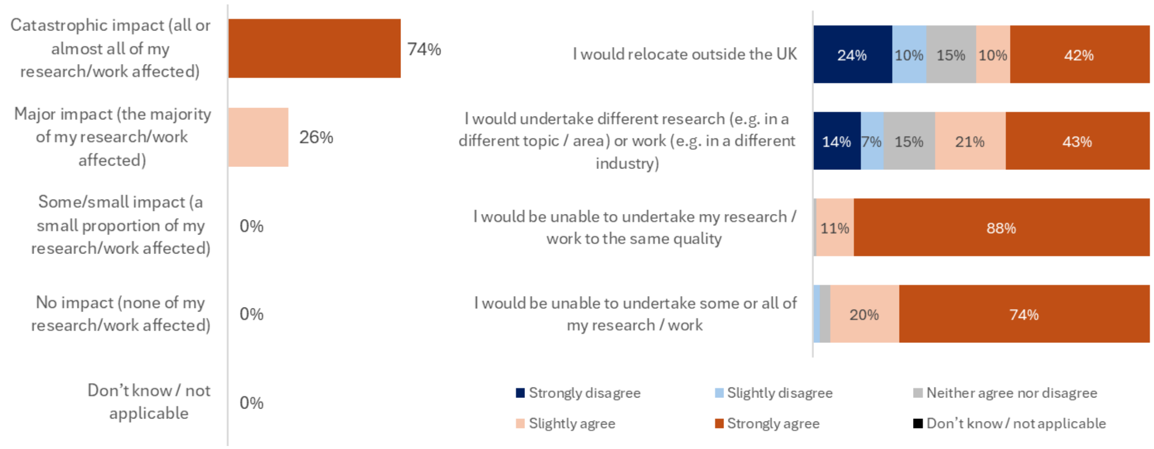

Note: Based on 201 respondents.

Source: London Economics survey of ARCHER2 users

Note: Based on between 166 and 189 respondents.

Source: London Economics survey of ARCHER2 users

In practical terms, this would mean that there would be a substantial impact on the quality of research or work UK researchers relying on computational modelling would be able to achieve:

- Almost nine in ten users strongly agreed that not having an EPSRC/NERC funded Tier-1 HPC system available would mean that they were unable to undertake their research or work to the same quality.

- Even more crucially, three in four users strongly agreed that they would not be able to undertake their work at all;
- around two in five strongly agreed that they would relocate to a different country,
- and the same number strongly agreed that they would undertake different work under this scenario.

The impact of retaining ARCHER2's predecessor ARCHER

Of course, the complete absence of a UK Tier-1 capability is a deliberately extreme hypothetical scenario, intended to test the level of dependence on national-scale HPC infrastructure. Nevertheless, in practice, the impact of not having ARCHER2, while somewhat less severe than a total absence, remains substantial.

To illustrate this, users were also asked about a scenario where ARCHER2's predecessor ARCHER was to be retained in operation. In this case, 60% of users responding to the survey said that only having access to a capability of the scale of ARCHER2's predecessor ARCHER would have a major impact that would affect the majority of their research or work. A further 29% said that the impact of this scenario would be catastrophic, affecting all or almost all of their work.

The number of users strongly agreeing that they would not be able to undertake their work at all under this scenario fell to one in four, though a further 50% of users slightly agreed with the statement. Nevertheless, three in four users continued to strongly agree that they would be unable to undertake their research or work to the same quality under this scenario.

Figure 50 Impact of a reduced national Tier-1 HPC service (continuation of ARCHER1)

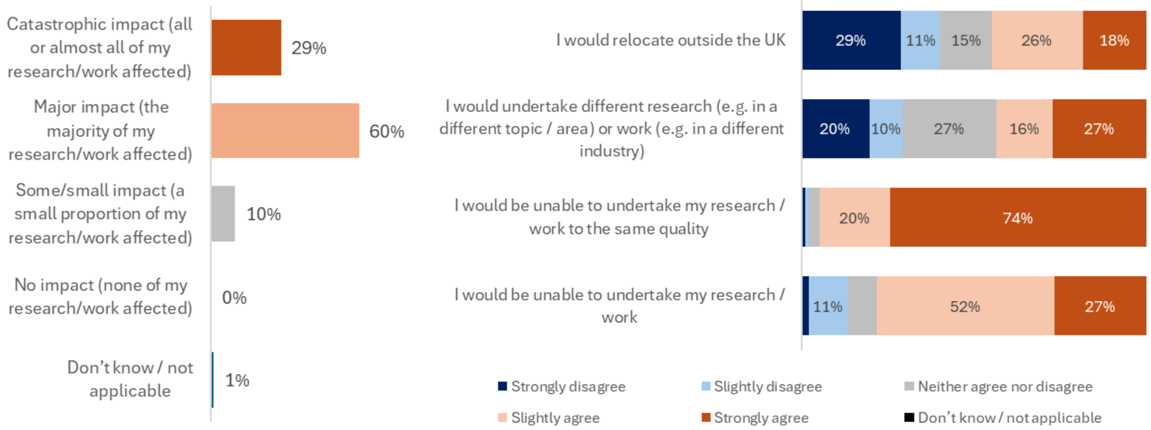

Note: Based on 189 respondents.

Source: London Economics survey of ARCHER2 users

Note: Based on between 53 to 160 respondents.

Source: London Economics survey of ARCHER2 users

8.2 The impact on the UK's international standing in HPC

The UK has a strong international reputation in compute and has historically been a global leader in a number of computing domains. This includes areas such as software development, computational modelling, data analytics, cybersecurity, AI and machine learning.[78]

However, the UK's high-performance computing infrastructure has been falling behind that of other major global economies[79]:

- In November 2007, the UK accounted for 9.4% of the top 500 HPC systems in the world and for 7.3% of the global performance share of the top 500 HPC systems. This was the highest systems share and second highest (behind Germany with 7.7%) performance share of any country outside the US, with the US at the time accounting for 56.8% of system share and 59.9% of performance share.
- By June 2021, the last available data prior to the launch of ARCHER2, the UK's system share had declined to 2.2%, while the UK's performance share had fallen to 1.3% and ARCHER, ARCHER2's predecessor, had fallen to 471st (down from 20th when it was first ranked) in the Top 500 ranking of the world's most powerful supercomputers.
- At its launch, in November 2021, ARCHER2 was ranked as the 22nd most powerful supercomputer in the world, marking a substantial improvement in the position of the UK's most powerful HPC capability and raising the UK's performance share from 1.3% in June 2021 to 1.8% in November 2021.
- However, since then the UK has continued to fall behind further. In the most recent Top 500 data, published in November 2024, the UK's performance share has declined further to 0.7% while ARCHER2's position has fallen to 62nd.
- For comparison, Germany's performance share, which was only marginally ahead of the UK in November 2007, stood at 3.4% in November 2024, and France's performance share, which was substantially lower than the UK's in 2007 but has tracked the UK's performance share closely between around 2011 and 2018, stood at 2.5% as of November 2024.

As this data highlights, the launch of ARCHER2 in 2021 was a much-needed boost to the UK's high-performance computing capability. However, it was not sufficient on its own to counter the broader decline in the UK's position in global high-performance computing. Nevertheless, without ARCHER2, the UK's international position would have deteriorated even further, with UK researchers not having access to a Tier-2 system within the top 50 globally at all.

However, it is also worth noting that the use of ARCHER2 is almost entirely UK-based research whereas a fraction of the compute time of the German and French HPC systems will be spent on supporting computational work from non-host countries via EU-HPC resource sharing.[80]

Further, it should be noted that with the investment in the AI Research Resource (AIRR), the UK's performance share has risen to 2.9% in the most recently published Top500 data from June 2025, while the UK's system's share has fallen slightly to 2.6%. However, HECs do not have a budget on AIRR and only AI workflows can be run on AIRR whereas current HECs typically use other workflows.

[78] GO-Science (2021). Large-scale computing: the case for greater UK coordination. Available at: https://assets.publishing.service.gov.uk/media/654a4025e2e16a000d42aaef/UK_Computing_report_-_Final_20.09.21.pdf [accessed 27/05/2025]

[79] Data from Top500.org

[80] The EuroHPC Joint Undertaking manage access time to EuroHPC supercomputer systems ranging from 35% of total capacity for petascale systems up to 50% of pre-exascale systems. See https://www.eurohpc-ju.europa.eu/supercomputers/supercomputers-access-policy-and-faq_en.

At the same time, investments by other countries meant that Germany's performance share has risen to 8.6%, while France's share has remained broadly stable at 2.4%.

Figure 51 Global performance and system share by country, June 2025

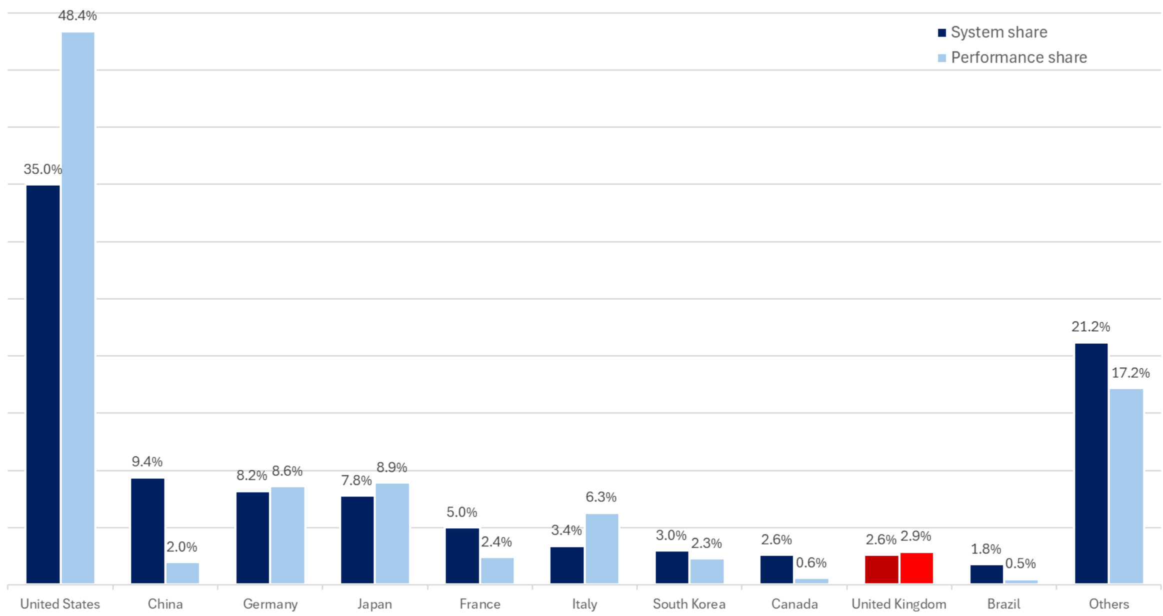

Note: Shows the top ten countries as of June 2025. Performance measure = maximal LINPACK performance achieved. A country's **system share** is the proportion of the Top500 computers which belong to that country. A country's **performance share** is the proportion of the total rmax GFlops of the combined Top500 computers held by that countries' computers.

Source: London Economics analysis based on Top500.org data

Figure 52 The UK's global performance share over time

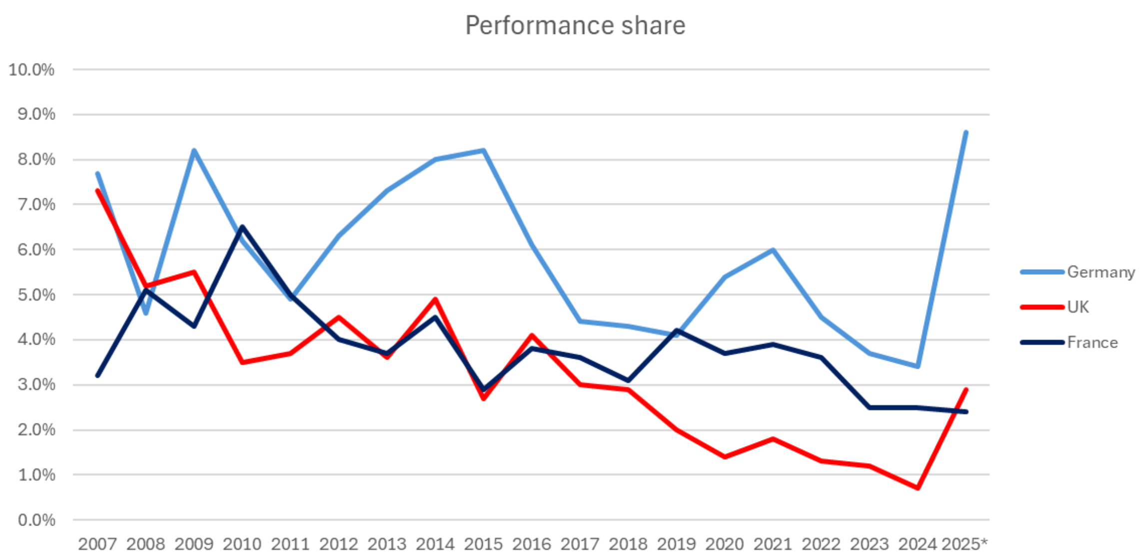

Note: Historical datapoints were taken from the November Top500 reports for each year. (*) Data for 2025 was taken from the June 2025 Top 500 report.

Source: London Economics analysis based on Top500.org data

Figure 53 The UK's global system share over time

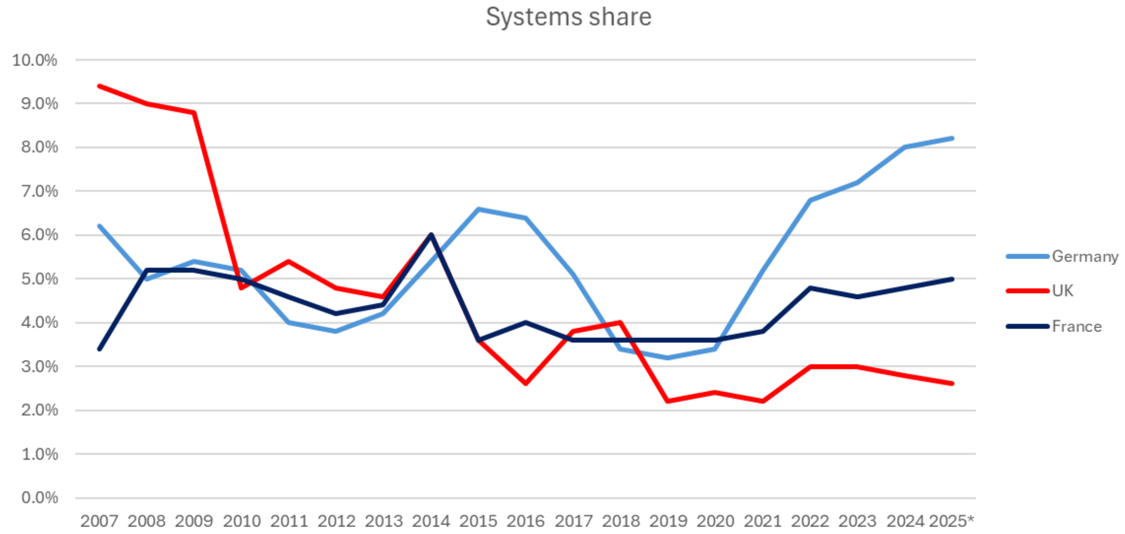

Note: Historical datapoints were taken from the November Top500 reports for each year. (*) Data for 2025 was taken from the June 2025 Top 500 report.

Source: London Economics analysis based on Top500.org data

8.3 Access to and viability of alternative computational capabilities

Another key factor in determining the impact a loss of ARCHER2 would have on the UK science community is the degree to which users could access and utilise alternative computational capabilities for their research or work. In evaluation terminology, there are two key concepts that are particularly relevant for this discussion:

- **Displacement:** The extent to which provision of ARCHER2 reduced investment in or use of other HPC systems. For example, the degree to which individual institutions or private organisations might have invested in their own HPC capabilities had ARCHER2 not been funded.
- **Substitution:** The extent to which users opt to undertake their research or work on ARCHER2 over other HPC systems not because other systems are not suitable, but because ARCHER2 is more convenient, capable or cost-effective. In this context, this involves assessing the degree to which other systems, where they exist or could have existed, provide a viable alternative to ARCHER2.

In simplified terms, the analysis needs to consider two key questions: i) To what extent would alternative systems have been provided? and ii) To what extent are these alternatives suitable?

Other publicly funded HPC capabilities in the UK

As discussed in the introduction to this report, ARCHER2 is not the only supercomputer in the UK. Rather, ARCHER2 is part of an ecosystem of publicly funded HPC services. Therefore, the most obvious alternative for users, in the absence of ARCHER2, would be to access alternative existing HPC capabilities.

In principle, other HPC systems are substitutes for ARCHER2. Of course, as ARCHER2 is the UK's most powerful supercomputer, existing alternatives may not be powerful enough to run some of the largest simulations undertaken on ARCHER2; for example, the "Grand Challenges", highly complex challenges that require a significant amount of computational power to solve. Moreover, Tier-2 HPC systems do not have the capacity to run many mid-size jobs. There are also noteworthy differences in the hardware configuration (e.g., different types of chips) and the available software stack, meaning that tweaks to the user code may have to be made. Nevertheless, these systems could in principle be used as alternatives for R&D undertaken on ARCHER2.

Indeed, when asked about the proportion of their work they could undertake without ARCHER2, on average, survey respondents indicated that 40% [median 35%] of their work could be carried out using alternatives. Further, around 70% of respondents said [that were they to use alternatives] they were "very likely" or "likely" to access alternative publicly funded HPC for the proportion of R&D that they would be able to undertake without ARCHER2. (Figure 54)

However, in practice, it is unlikely that these systems, without additional capacity provision, would be able to cope with the additional demand from users currently serviced by ARCHER2. Indeed, the 2023 Future of Compute Review[81] found that even at the time (and with the existence of ARCHER2) the then current provision of compute was not sufficient for many academic users and researchers and was already limiting the UK's scientific capability and inhibiting scientific breakthroughs. Moreover, since then, some of the Tier-2 HPC systems have neared end of life, with others scheduled for retirement in the near future (see Table 42 in Annex A4.4).

Another consideration is the extent to which ARCHER2 displaced other publicly funded HPC provision. That is, whether, in the absence of ARCHER2, EPSRC, NERC, other UK research councils, or other public bodies would have invested in additional alternative HPC capacity. While this counterfactual is plausible, such alternatives would only have constituted an improvement if they could have delivered equivalent capability at lower cost or with greater efficiency than ARCHER2. Otherwise, the overall cost to the taxpayer to deliver similar levels of scientific throughput would have remained the same or even increased and, therefore, these alternatives would not have delivered additional net benefit.

Figure 54 Alternatives considered for research that could be conducted without ARCHER2

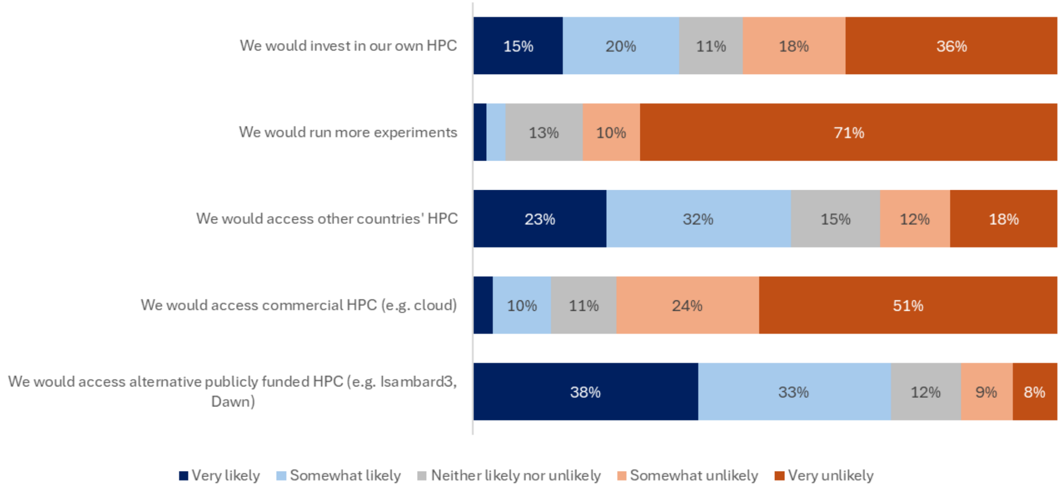

Note: Based on between 16 and 119 respondents. Question asked: For the proportion of research that you could undertake without ARCHER2, how do you think you would do this? Select how likely you think it is that you would use various alternatives.

Source: London Economics survey of ARCHER2 users

[81] DSIT (2023). Independent Review of the Future of Compute: Final report and recommendations. Available at: https://www.gov.uk/government/publications/future-of-compute-review/the-future-of-compute-report-of-the-review-of-independent-panel-of-experts [accessed 21/05/2025]

Current and future GPU based systems

A recent noteworthy development is the AI Research Resource ("AIRR") programme. AIRR is a recently established infrastructure policy programme that aims to increase the compute capacity for AI research through creating a national AI research resource. The AIRR programme was established with an initial committed funding of £250 million to establish two supercomputers dedicated to AI research at Cambridge (Dawn) and Bristol (Isambard-AI). A further, much larger increase, in public AI compute capacity of at least 20x AIRR's original capacity is envisaged by 2030.

The AIRR programme was established in light of findings and recommendations from the 2023 Future of Compute review[82]. In particular, the review found that there was a significant shortage of access to public hardware accelerators (GPUs), which cutting-edge AI models are reliant on, and recommended an immediate and significant increase compute capacity for AI research.

Given that AI methods are increasingly utilised across many disciplines, there is potential for some future substitution of ARCHER workloads by AIRR systems. Moreover, while AIRR supercomputers are primarily targeted at supporting AI research, they could also be utilised for other computationally intensive tasks, such as large-scale simulations like those undertaken on ARCHER2, whether or not AI methods are involved.

The potential for ARCHER2 workloads to be undertaken on GPU-based machines such as AIRR supercomputers (as opposed to more traditional CPU-based machines such as ARCHER2) was acknowledged in consultations undertaken for this study. Indeed, some codes used by users are already able to utilise both GPU and CPU based architectures, thereby being able to be ported with minimal efforts. However, for many of the software codes used on ARCHER2 efficient porting would require computational software engineering work to optimise these codes for GPU-based architectures. This is in principle possible, though it is not a straightforward task and would require potentially substantial computational software engineering work.

Further, while possible for many codes, it is not feasible or indeed possible to port all workloads currently undertaken on ARCHER2. In particular, some of the largest simulations are unable to be undertaken on GPUs due to limitations in GPU memory capacity, intra-node bandwidth (the communication speed between different nodes, i.e. different clusters of GPUs, within the same machine), or the complexity (and associated high cost) of adapting highly parallel, tightly coupled code to GPU architectures.

See Annex A4.1 for further discussion about the differences between GPU-based and CPU-based HPC architectures and the implications of using the one over the other.

International publicly funded HPC capabilities

UK researchers are also able to access European HPC systems. Indeed, 23% of survey respondents said [that were they to use alternatives] they were "very likely" to access international for the proportion of research that they would be able to undertake without ARCHER2. A further 32% said they were "likely" to do so.

One key route through which researchers can access international HPC is the European HPC Joint Undertaking (EuroHPC), which the UK joined in May 2024. EuroHPC brings together supercomputing

[82] DSIT (2023). Independent Review of the Future of Compute: Final report and recommendations. Available at: https://www.gov.uk/government/publications/future-of-compute-review/the-future-of-compute-report-of-the-review-of-independent-panel-of-experts [accessed 21/05/2025]

resources from across 35 countries. Membership of EuroHPC allows UK researchers to access EuroHPC supercomputers and participate in Horizon Europe funded research and innovation calls related to supercomputing.

However, as with other publicly funded UK HPC capabilities, EuroHPC calls are competitive and in practice would be unlikely to accommodate the unmet demand resulting from the absence of ARCHER2. Further, while not directly a requirement to join EuroHPC, the UK's long-standing expertise and established HPC infrastructure, exemplified by systems such as ARCHER2 and institutions such as EPCC, may have strengthened the UK's position as a credible candidate and valuable partner in the European supercomputing landscape.

Some users are also able to access international systems through other means such as through existing collaborations or partnerships with international research institutions. However, this is not a viable alternative for all users and reliance on such routes as a normalised way of accessing HPC resources would create disparities within the wider UK research community.

Box 19 Driving transition to electrification of flight

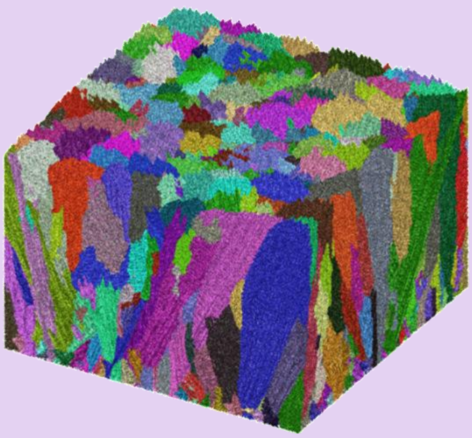

Researchers are using ARCHER2 to develop advanced simulations with Rolls-Royce that reveal how defects form during turbine blade manufacturing. These insights will enable the production of lighter, more efficient turbines—helping to reduce emissions, support sustainable aviation, and meet growing public demand for greener air travel.

The ARCANE project, a £9 million collaboration between the Universities of Greenwich, Birmingham, Oxford, and Rolls-Royce, is using the UK's ARCHER2 supercomputer to simulate how turbine blades solidify—from the atomic to the structural scale. These simulations aim to reduce defects like bicrystals, which weaken turbine performance, and support the development of stronger, lighter components crucial for future electric aircraft. This work will help Rolls-Royce improve manufacturing, reduce emissions, and meet the EU's FlightPath 2050 targets for sustainable aviation. Beyond aerospace, the model has potential applications in battery design and 3D printing in space, showcasing its broader industrial impact.

Further details can be found in the full case study in Annex A5.15.

Note: Figure represents large scale 40 mm cube, 8 billion cell simulation of polycrystalline dendritic directional solidification using HPC, capturing component scales at a microscopic resolution. The different colours represent different orientations of dendrites.

Source: Technopolis analysis based on information gathered via desk research and consultation with Prof. Andrew Kao.

Cloud and commercial HPC as a potential alternative?

Commercial cloud has established itself as a staple for many organisations seeking flexible computational compute resources. It therefore stands to reason to ask whether commercial cloud could provide a viable alternative to publicly funded HPC. Commercial cloud can have many benefits; for example, access is comparatively straightforward and quick, computational capacity is flexible and scalable, no upfront capital investment is required, etc.

As a result, commercial cloud can be a good alternative particularly for users with relatively more limited computational needs. However, given the nature of ARCHER2 as a national Tier-1 capability, and therefore used to run the most complex simulations, many of the workloads run on ARCHER require substantial computational resources which may be beyond the capabilities typically provided by cloud. Moreover, even if commercial cloud was able to provide sufficient computational resources, sustained use will incur substantial fees.

The limitations of current commercial cloud services as a replacement for traditional HPC systems is also reflected in feedback from the user survey, where only a small proportion (<15%) of respondents said [that were they to use alternatives] they were likely to access commercial cloud for the proportion of research that they would be able to undertake without ARCHER2.

Further, there are also a range of other disadvantages and risks including:

- First, reliance on commercial providers risks "lock-in". That is as codes become optimised to the infrastructure of a specific provider, the UK science base risks becoming reliant on the ecosystem of that provider and therefore being locked-into continued usage of those services.
- Second, as this report has highlighted, the benefits of publicly funded HPC are not limited to simple provision of compute capacity. Reliance on commercial cloud would risk loss of the wider benefits of publicly funded HPC such as the wider HPC community, HPC skillset and public training provision and support.
- Third, there are risks around security and compliance especially for the most sensitive workloads. Commercial ownership can also limit the ability to rapidly reallocate computing resources during national emergencies or for other strategic (and sensitive) national priorities. Moreover, the UK having its own HPC hardware also **acts as a form of insurance for rare significant events**. For example, ARCHER enabled researchers to run simulations that helped design vaccines and medicines during the COVID-19 pandemic.[83]
- Fourth, there are additional data transfer (moving large datasets to/from the cloud is costly and time-consuming), latency (incl. the time it takes for data to move between compute nodes or between storage and compute as well as the time it takes for nodes to communicate with each other) and other considerations.

Despite these limitations, cloud is certainly an important part of the computational ecosystem and can be a strong option for certain users and use-cases.

[83] See here: https://www.discover.ukri.org/the-supercomputers-that-helped-in-the-fight-against-covid-19/

9 Conclusions

This independent evaluation study highlights the substantial scientific and economic benefits enabled by ARCHER2. **Total benefits of ARCHER2 to the UK are estimated at £4.3 billion**. This is **equivalent to a return on public sector investments in ARCHER2** (and public funding for R&D undertaken on ARCHER2) **of £8.3 per £1 of public funding invested**.

Further, comparing the total estimated benefits and costs of ARCHER2 to the total estimated benefits that could have been achieved had ARCHER1 continued to operate indicates additional benefits of ARCHER2 of £2.8 billion. This compares to additional costs of ARCHER2 estimated at £294 million, implying **additional benefits of £9.5 per additional £1 invested in ARCHER2 over and above the costs had ARCHER1 continued to operate**. This suggests that the additional benefit produced by investing in a higher-capacity Tier-1 service far outweigh the additional costs incurred.

Sensitivity analysis accounting for uncertainty surrounding the magnitude of benefits suggests the central benefit estimate for this study is on the conservative side. The mean and median sensitivity results indicate a higher return on investment of £13.4 per £1 and £12.9 per £1, respectively, of public funds invested. The 90% confidence range around the central estimates shows that benefits are plausibly in the region of between £4.6 and £24.1 per £1 invested. In contrast, the sensitivity results for the continuation of ARCHER1 scenario suggest a much narrower range of benefits of between £4.2 and £14.9 per £1 invested.

Given the nature of ARCHER2 as a capability for scientific R&D, it is not surprising that a substantial share of benefits are ascribed to the economic benefits flowing from academic R&D undertaken on ARCHER2. Spillover impacts to the UK economy from academic R&D undertaken on ARCHER2 are estimated at around £3.7 billion, accounting for 87% of estimated benefits from ARCHER2. However, while best estimates were used to derive these benefits substantial uncertainty surrounds the estimation of R&D spillover impacts not least due to the fundamental nature of R&D undertaken on ARCHER2 and the associated significant time-lag between R&D and materialised impact.

Beyond the economic impacts of R&D, ARCHER2 delivers a broad range of additional benefits. This includes scientific benefits such as publications, cost savings to researchers of being able to access compute capacity free at the point of use. It also includes wider benefits such as the development of scientific software, the creation of spin-outs and intellectual property, and benefits from training and skills development, among others.

Excluding spillover impacts of R&D, shows that these wider benefits are still sizeable with central estimates standing at around £0.5 billion. This represents a return on public investment into ARCHER2 itself of £5.2 per £1 of public money invested into ARCHER2. Compared to the scenario where ARCHER1 continues to operate, additional benefits of ARCHER2 are estimated at £0.3 billion, implying a return on investment of £7.3 per additional £1 invested in ARCHER2 over and above the costs incurred had ARCHER1 continued to operate.

In addition to the monetised benefits highlighted above, ARCHER2 plays a key role in securing the UK's international standing in high performance computing, in scientific software development, as well as the UK's international science competitiveness. While these benefits are difficult to monetise and so are not captured in the benefits estimates produced for this study, they are no less important.

However, this study also highlights the need for continued investment in UK HPC infrastructure, echoing findings from the recent Future of Compute Review[84]. The UK has historically been a global leader in a number of computing domains. This includes areas such as software development, computational modelling, data analytics, cybersecurity, AI and machine learning. However, the UK's high-performance computing infrastructure has been falling behind that of other major global economies.[85]

In response to the findings from the Future of Compute Review, the UK Government has recently made a first round of investments into the AI Research Resource (AIRR), which aims to provide substantially increased UK AI compute resources for AI research. While this is a welcome step in the right direction, GPU-based machines are not a direct substitute for traditional HPC systems as has consistently been highlighted by stakeholders consulted for this study. Rather, investments in both types of systems are needed to ensure users have access to the right capabilities for their needs.

[84] DSIT (2023). Independent Review of the Future of Compute: Final report and recommendations. Available at: https://www.gov.uk/government/publications/future-of-compute-review/the-future-of-compute-report-of-the-review-of-independent-panel-of-experts [accessed 21/05/2025]

[85] GO-Science (2021). Large-scale computing: the case for greater UK coordination. Available at: https://assets.publishing.service.gov.uk/media/654a4025e2e16a000d42aaef/UK_Computing_report_-_Final_20.09.21.pdf [accessed 27/05/2025]

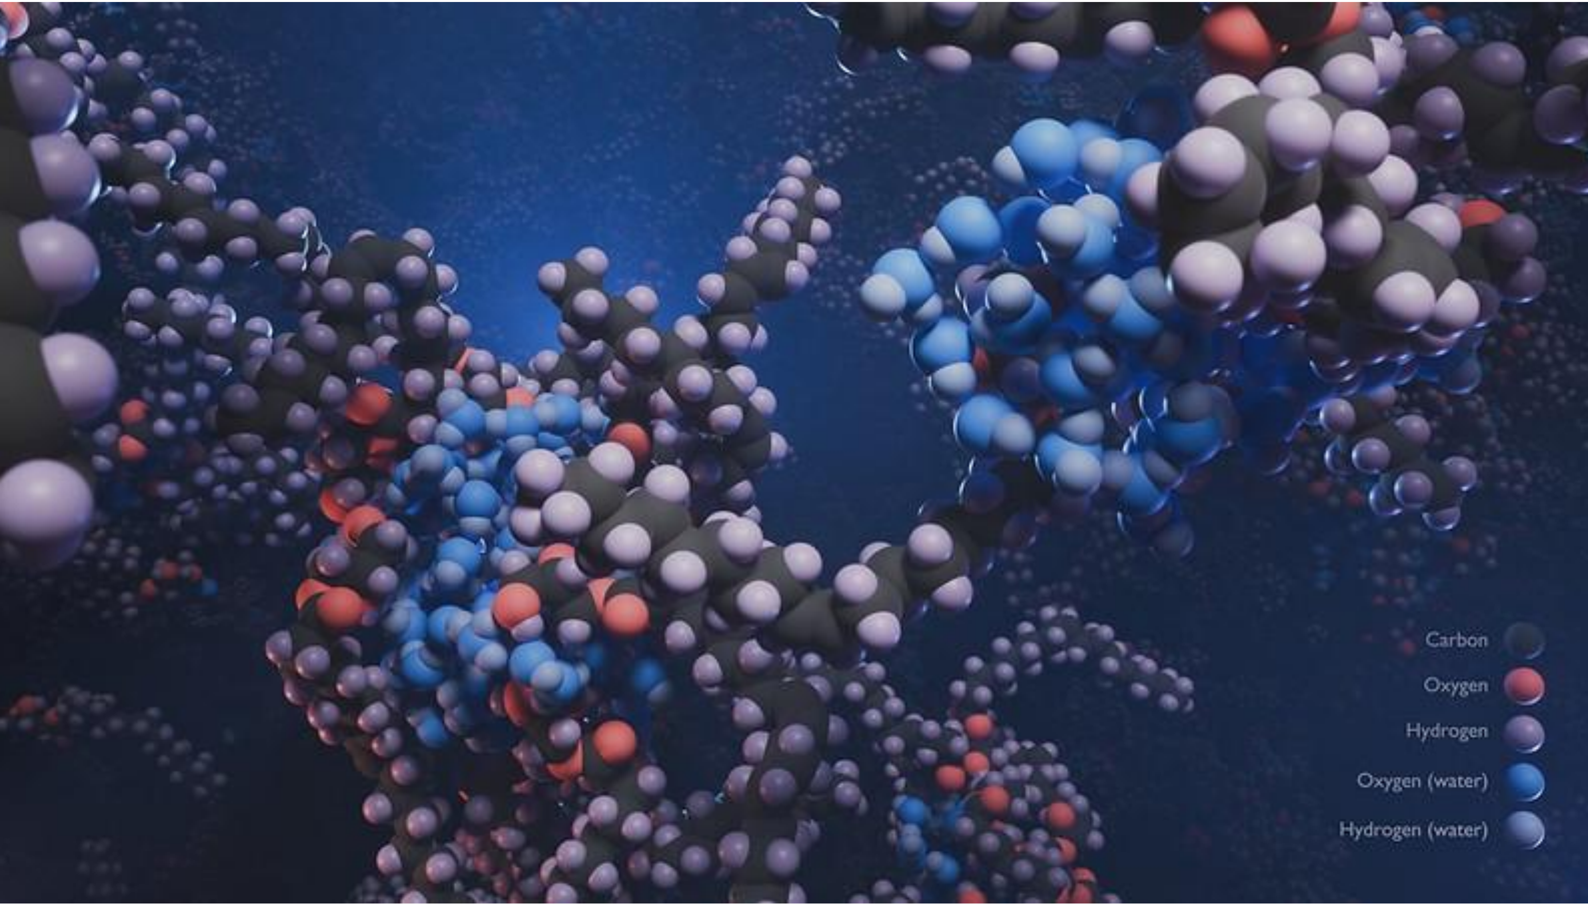

ANNEXES

The following annexes provide further detail:

- **Index of Tables, Figures and Boxes**
- **Annex 1** provides details on the Theory of Change, which provides an overview of the inputs into ARCHER2 and the resulting benefit streams.
- **Annex 2** provides further detail on stakeholder inputs and survey responses.
- **Annex 3** provides details on the evaluation approaches/methodology used to estimate economic benefits of ARCHER2. It also discusses assumptions used in the cost-benefit model. Further information on limitations of the methodology is also provided were appropriate.
- **Annex 4** provides additional material such as more in-depth breakdowns of costs mentioned in the report.
- **Annex 5** provides case studies providing evidence of the impact of ARCHER2.

Image source: Dr Sébastien Lemaire, EPCC, University of Edinburgh: Molecular dynamics computation of Glycerol monooleate (GMO) molecules gathering around water droplets to form reverse-micelles helping lubrication efficiency

Index of Tables, Figures and Boxes

Tables

Figures

Boxes

Annex 1 Theory of Change

Figure 55 provides the benefits logic map of the ARCHER2 service. The benefits logic map was shared with London Economics by UKRI and reviewed by London Economics. Benefits identified as part of the desk research and stakeholder consultation exercise undertaken for this study showed that the existing logic map already captures the breadth of benefits flowing from ARCHER2 well. Therefore, no changes to the existing logic map were needed.

Figure 55 ARCHER2 benefits logic map

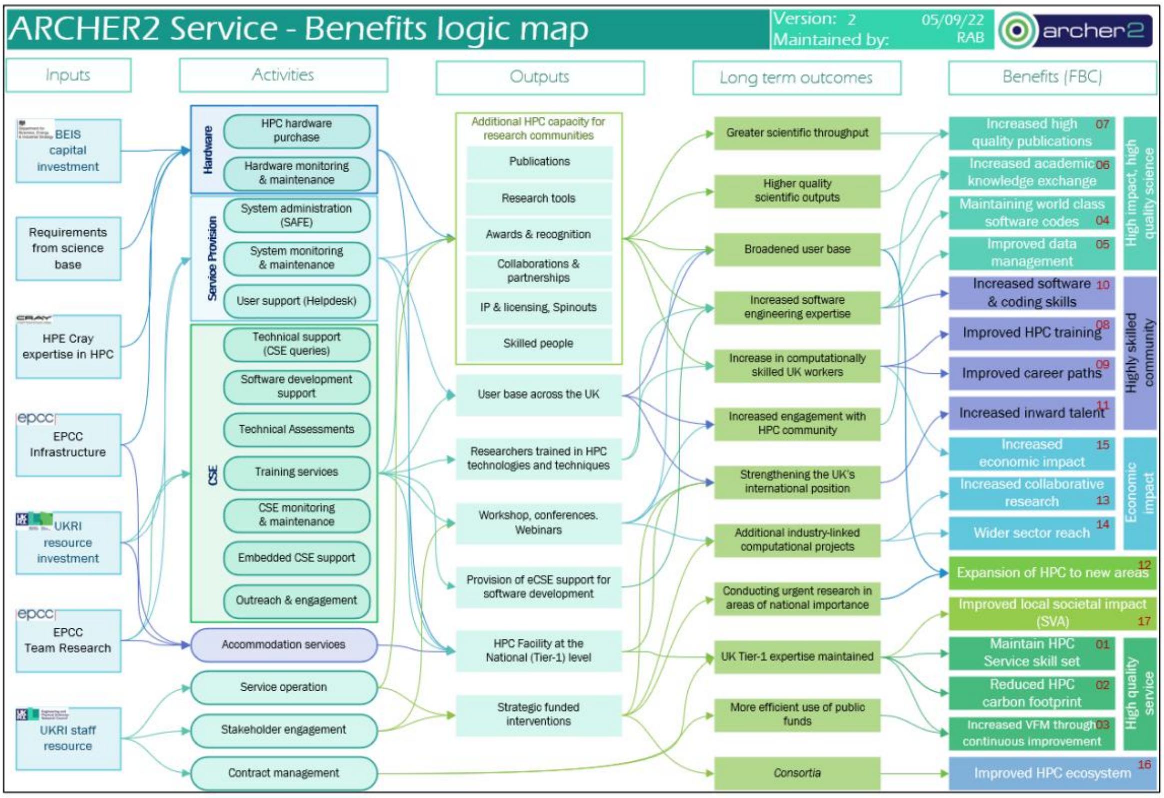

Source: UKRI

Annex 2 Stakeholder inputs, survey responses

A2.1 Representativeness of data sources used

A2.1.1 The user survey and stakeholder consultations

To understand the importance of ARCHER2 for academic R&D, a wide range of stakeholders were consulted as part of this study. Numerous interviews with leading academics in the HPC field, including High End Computing (HEC) consortia leads were conducted to gain insights, as well as an online survey of numerous groups using ARCHER2 as part of their research or work.

As with any survey or stakeholder feedback, the results may not fully reflect the experiences of all users, with some perspectives potentially missed or over/under-represented.

Sample selection bias may be present if certain users are more likely to complete the survey or engage in consultation compared with other users. This can happen because the benefits of engaging in stakeholder consultations/surveys is costly, in foregone time, and the benefits, of potentially contributing to the case for future funding, are both uncertain and not immediate. Moreover, some postgraduate users of ARCHER2 (particularly graduating PhDs) may have not been captured because they may have already joined the labour market and would not have received an invite to the survey.

To try to encourage as wide participation among ARCHER2 users, we distributed the survey via HEC consortia heads, with multiple reminders sent directly from EPCC as well as the HEC consortia heads. In addition, where possible, primary evidence from the survey is combined with secondary evidence.

To gain a better understanding of the potential biases, Section 1.5.2 provides a comparison between respondents to the online survey and the population of registered users on ARCHER2.

A2.1.2 User survey representativeness

Figure 56 shows the broad research areas in which survey respondents undertake R&D. For comparison, Figure 57 shows the actual usage of ARCHER2 by academic field.

It should be noted the share of users that undertake R&D in a specific field does not necessarily directly correlate to the volume of usage of ARCHER2. Further, the survey allowed users to select multiple research areas. Therefore, the data only provides a rough indication of the representativeness of the survey sample.

To enable understanding of the representativeness of the user account by types of users, Figure 58 and Figure 59 further provide a breakdown of the share of survey respondents by user type and share of active user accounts, respectively, by user type.

Figure 56 R&D undertaken by survey respondents by field

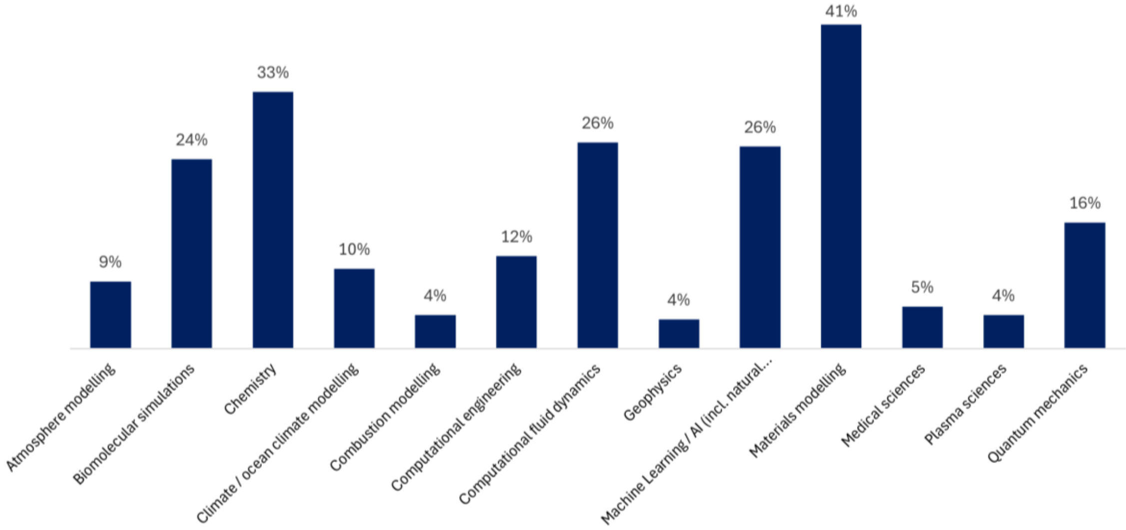

Note: Based on 186 respondents. Percentages do not sum to 100% as users could select multiple areas of research. Question asked: 'In which broad area are you undertaking research?'

Source: London Economics survey of ARCHER2 users

Figure 57 Usage of ARCHER2 by research area

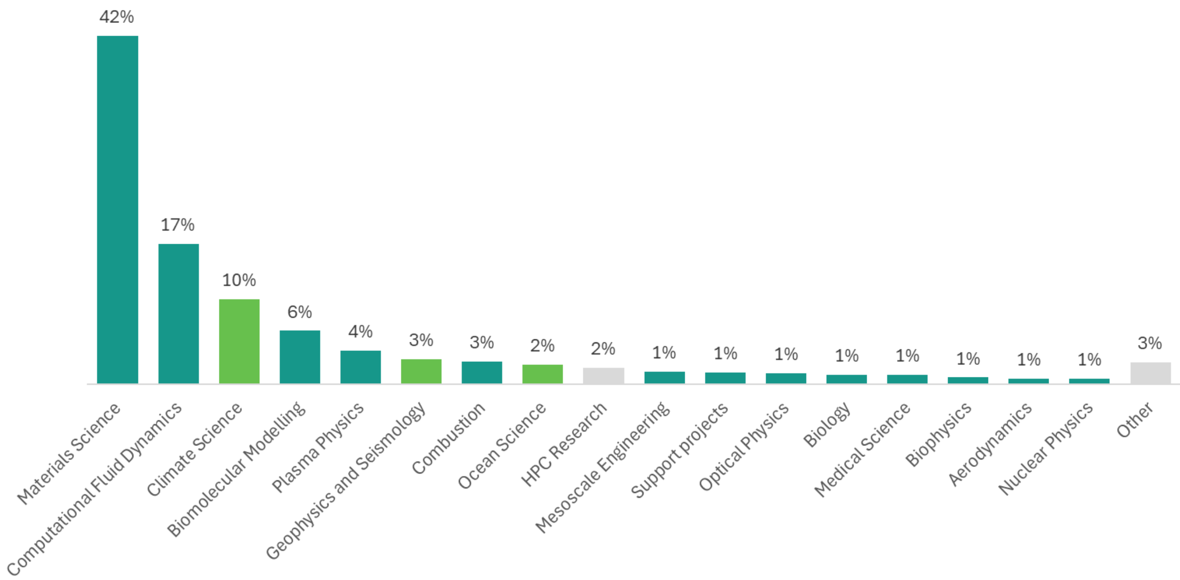

Note: The graph shows the usage by research discipline (% of CU used between 2021 and 2025)

Source: EPCC usage data

Figure 58 Share of survey respondents by type of user

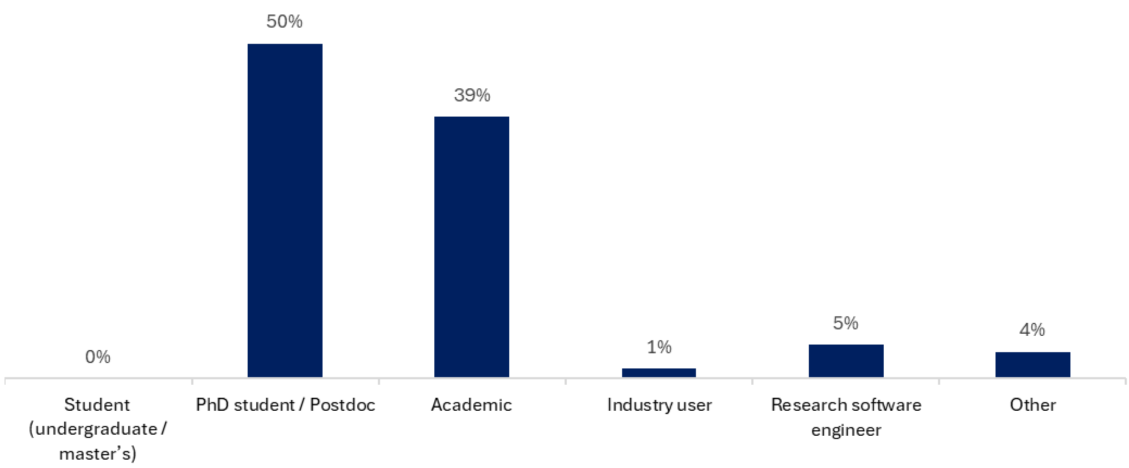

Note: User survey results based on 201 respondents.

Source: London Economics survey of ARCHER2 users; EPCC usage data

Figure 59 Active accounts between 2021-2025 by type of user

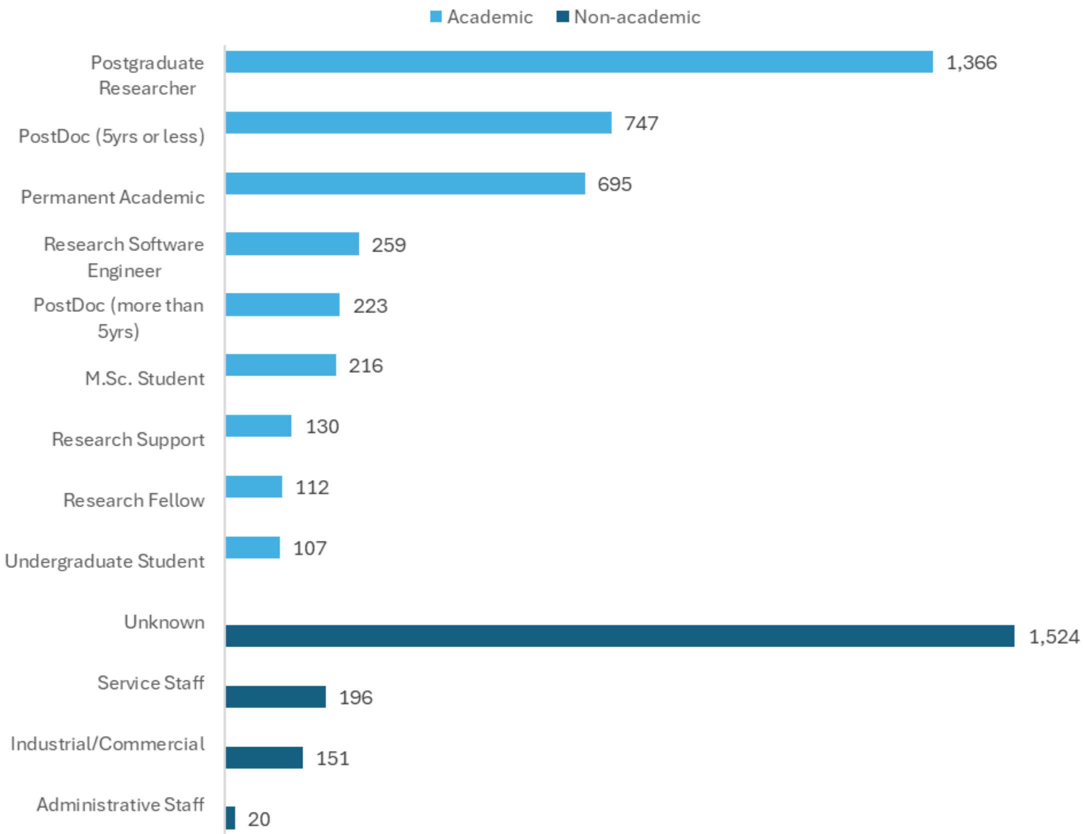

Note: Active accounts are accounts that have run jobs in the reporting period.

Source: London Economics analysis of EPCC usage data

Annex 3 Methodological annex

A3.1 Overarching approach

The analysis of economic benefits was undertaken in real-terms adjusting monetary values to reflect real 2025 £-values. Future values were discounted using the HMT discount rate of 3.5% in line with Green Book guidance. The following sections detail the evaluation approaches used for each of the benefit streams. In addition, the following methodological choices were made across benefit streams:

- **Start date of benefits estimated:** For training outreach events held in 2020, we assumed that these were already provided with the funding received for ARCHER2 (though note that no outreach events were held in 2020 and so this only applies to training in practice). For outputs from R&D activities, as there is a delay between the research and research outputs materialising, we only monetised benefits accruing from 2021 onwards. Similarly, as ARCHER1 was still live in 2020 and the reduced-service four-cabinet system only opened for early-access in October 2020 we assumed the 2020 grant funding was primarily used for research undertaken on ARCHER1.
- **Estimated benefits over ARCHER2's lifetime:** The analysis sought to estimate the total economic benefit flowing from ARCHER2 over ARCHER2's lifetime. As ARCHER2 will continue operating, under current plans, until 22nd November 2026, time-series input data was projected forward. This was done using a linear trend based on the historical data observed to fill-in values for 2025 (where only partial data was available) and 2026. 2026 figures were adjusted to account for the fact that ARCHER2 is planned to cease operation on 22nd November 2026 and 2026 will only be a partial year.

  Note, for academic publications it was assumed that a share of papers published in 2021 may have still been as a result of research undertaken on ARCHER1 due to the typical delays in the publication process of academic research. The share of publications ascribed to ARCHER1 in 2021 vs. 2022 was estimated based on the relative size of the four-cabinet ARCHER2 system, which operated alongside ARCHER1, to the full system size of ARCHER1. This was estimated at 47% of publications in 2021 ascribed to ARCHER2. Based on these assumptions and the annual publication growth trends, an assumption on the share of publications only published in the following year was estimated. This was estimated at 69%. Given the estimated delay in publications, publications data was projected forward to 2027 in-line with these assumptions. A 25% uncertainty interval was placed around both assumptions in the sensitivity analysis.
- **Accounting for underestimation of data obtained from ResearchFish™:** As discussed in the introduction to this study, data from UKRI's research outcome database, ResearchFish™, is expected to suffer from underestimation due to the self-reported nature of outcomes data. To account for this in the analysis, adjustments were made to estimates of software products, intellectual property and spinouts. Although publication data will still miss publications from authors who forget to acknowledge the use of ARCHER2, it is assumed to be captured well in ResearchFish™. No adjustment was made in the central case, though the upper-bound allowed for a small degree of underestimation. The lower-bound for publications also allowed for a degree of overestimation, for example, to capture different versions of the same publication not captured in the deduplication as well as meta-analysis papers that did not actually undertake any research on ARCHER2. The adjustments made are described in the relevant sub-sections for these benefit streams.

A3.2 Benefits of research undertaken on ARCHER2

To monetise the benefits of research undertaken on ARCHER2, three strands of analysis were undertaken: i) the direct benefits of access to ARCHER2 for researchers in terms of the avoided cost compared to commercial alternatives; ii) the value of publications and citations resulting from R&D undertaken on ARCHER2; and iii) the wider spillover impacts of R&D.

Note, in addition to these benefit streams, ARCHER2 also enables scientific collaborations which deliver scientific and economic benefits. However, collaborations directly impact other outcome metrics. For example, production cost estimates of publications typically already account for the time-value of collaborators. Further, collaborations lead to higher quality publications, which are, as a result, cited more often. The economic value of collaborations is therefore already implicitly captured in other monetised benefit streams. To avoid double counting collaborations are therefore not valued separately.

A3.2.1 Direct benefits of access to researchers

To monetise the direct benefit to researchers[86] of ARCHER2 access **an avoided cost approach** was used. This approach compares the notional cost of access to ARCHER2 to the higher cost of undertaking R&D using commercial cloud alternatives. This was calculated using the following formula:

Avoided cost of ARCHER2 access for researchers	=	*Number of core hours on ARCHER2 accessed by researchers*	x	(	*Commercial price per core hour*	-	*Notional cost of accessing ARCHER2 per core hour*	)

Note, ARCHER2 uses different notional costs for research funded by partner institutions (EPSRC and NERC) and that funded by non-partner institutions (other research councils). Therefore, the calculation distinguished the share of academic usage by partner and non-partner funded research. To compare this cost to the cost of accessing commercial cloud alternatives, prices from a range of commercial cloud service providers were obtained. Table 23 provides an overview of these costs.

Table 23 ARCHER2 vs commercial HPC cost comparison

Provider	Cost per hour	Cores	Cost per core hour
Archer (partner)	***£0.20***	***128***	***£0.0016***
Archer (non-partner)	***£0.39***	***128***	***£0.003***
Amazon (AWS) (C5a.24xlarge)	£ 3.42	48	£0.071
Microsoft Azure	£3.36	120	£0.028
Google Cloud	£7.28	96	£0.076
Oracle Cloud Infrastructure	N/A	N/A	£0.051
Sabalcore	N/A	N/A	£0.094
R Systems	N/A	N/A	£0.070

[86] Note that access to commercial computing infrastructure for researchers would likely still be publicly funded, so the cost avoided is approximately equal to the cost avoided by the public purse.

Note: Figures reported represent baseline prices for on-demand hourly use. Additional charges for support, storage, software or other services may apply. Costs are calculated on a per core-hour basis and reported to three decimal places to avoid rounding to zero. Where applicable, prices are converted from US dollars to Pound Sterling using the Bank of England's average exchange rate for 2024.

Source: LE analysis of industry pricing from Amazon Web Services; Bank of England; Google Cloud; Microsoft Azure; Oracle; R Systems; Sabalcore Computing obtained in May 2025

The average (mean) value of the commercial cloud prices **(£0.065**) was used as a **central estimate**. The **lower and upper bounds** were set to **the lowest (£0.028) and highest (£0.094) values** obtained. A triangular distribution was used as the basis for estimating uncertainty within this range in the sensitivity analysis.

Note, the prices above only take account of the avoided cost of compute time. However, there are likely other costs of using commercial providers such as data transfer and storage costs. These additional costs were not considered in the analysis and therefore, the results of the analysis should be seen as a conservative estimate of the avoided costs.

A3.2.2 The value of publications and citations

To estimate the value of publications, data on the number of unique publications were combined with literature estimates on the value of publications. Only a limited number of published estimates were found to be available. Therefore, the analysis considered both estimates for general research as well as from related scientific fields.

The literature provides estimates using three different approaches:

- Florio et al. (2016) used a production cost approach assuming that the estimated value of a publication is, on average, at least equal to its production costs. A similar approach using the marginal cost of production (MPC) was used by Gifoni et al (2025)[87] and Morretta et al (2022)[88].
- Rousseau et al (2021)[89] used the cost of making publications open access as a proxy for the value of publishing a study. This is because the price reveals people's preferences for paying for publications.
- Finally, Rousseau et al (2021) also provides willingness-to-pay (WTP) for scientific publications aiming to assess how much scientists would be willing to pay to benefit from an additional publication in their field.

Table 24 summarises the estimates obtained from these studies.

Table 24 Literature estimates of the value of publications

Source	Method	Original value	Reference year	Real value in £-2025 prices
Rousseau et al (2021)	Value of making a publication open access (mean)	€2,250	2018	£2,541
Rousseau et al (2021)	MPC value of a publication (mean)	€5,383	2017	£6,137

[87] Gifoni et al (2025). The link between large scientific collaboration and productivity. Rethinking how to estimate the monetary value of publications. Scientometrics. Available at: https://link.springer.com/article/10.1007/s11192-025-05348-5

[88] Morretta, V., Vurchio, D., & Carrazza, S. (2022). The socio-economic value of scientific publications: The case of Earth Observation satellites. Technological Forecasting and Social Change, 180, Article 121730. https://doi.org/10.1016/j.techfore.2022.121730

[89] Rousseau et al (2021). Can we estimate a monetary value of scientific publications? Available at: https://www.sciencedirect.com/science/article/abs/pii/S0048733320301918

Source	Method	Original value	Reference year	Real value in £-2025 prices
Gifoni et al (2025)	MPC value of a publication applied to High-Energy-Physics (mean)	€11,011	2013	£12,886
Gifoni et al (2025)	MPC value of a publication applied to medical research (mean)	€8,125	2013	£9,508
Morretta et al (2022)	Full MPC value of a publication applied to space research (mean)	€ 4,854	2018	£5,481
Morretta et al (2022)	Partial MPC value of a publication applied to space research (mean)	€ 17,561	2018	£19,830
Rousseau et al (2021)	WTP value of a publication (mean)	€ 2,510	2019	£2,752

Source: London Economics research

The analysis for this study used the value of making a publication open access from Rousseau et al (2021) as the **lower-bound (£2,541**) as it was the lower of the two low estimates. As the highest estimate among the identified literature, the high estimate from Morretta et al (2022) was used as an **upper bound (£19,830**). As a central estimate, the average of the three publications providing estimates in the middle of the range were used. This provided a resulting **central assumption** of **£8,168**. For reference, the average using all literature estimates identified would suggest a slightly higher central assumption of £8,448.

To estimate, the total number of publications, the unique number of publications from the bibliometric analysis undertaken by Technopolis was used as a central estimate. To allow for a degree of potential over- or underestimation, a small 5% underestimation was assumed as an upper-bound as publications are reasonably well captured within ResearchFish^{TM}, while the lower unique number of publications that could be matched to OpenAlix (instead of all unique publications) was used as a lower-bound (25% lower than the central estimate).

Monetising the value of citations

To estimate the value of citations, a production cost approach was used, assuming that the value of citations is at least equal to the value of the time that is spent reading the publication. This was done by combining the average number of citations per paper with an assumed time spent to read a paper, researcher weekly salary and the weekly hours worked by academic researchers:

- The average number of citations per paper was taken from Technopolis' bibliometric analysis and was **12.8 per publication** for all publications published after 2022. Considering only publications from 2022 and 2023, which are more likely to have already been cited than newer publications, the average number of citations for these papers was **17.2** for papers that had already been cited (the average was **16.5** for all papers, including those that have not yet been cited). For the sensitivity analysis, a triangular distribution was used using 16.5 as the central estimate, while 17.2 was used as an upper-bound and 12.8 as a lower-bound.
- ONS salary data[90] was used to estimate gross annual salary of researchers. Adjusting for inflation using ONS deflators, the average weekly salary from May 2020 to February 2025 (in 2025 terms) for those working in Education is £578 (£30,073 annually). However, this data captures the education sector as a whole, rather than salaries for academic researchers. Therefore, desk research on current academic job postings and salary bands

[90] ONS. EARN03: Average weekly earnings by industry. (2025). Available at: https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/earningsandworkinghours/datasets/averageweeklyearningsbyindustryearn03

was undertaken to verify the applicability of this estimate. This exercise suggested a significant underestimation. Therefore, a 30% uplift was applied in the central case, bringing the central salary estimates in line with early-career researchers (e.g., research assistant) salaries. A higher 60% uplift was applied as the upper-bound, bringing salary assumptions in line with research associates.

- A standard work week of 38 (37) hours was assumed as the central (lower) case for the average weekly hours worked by academic researchers. In practice, the nature of academic work means that researchers often work longer hours. Therefore, a longer week of 50 hours was assumed as the upper bound.
- An average of thirty minutes to read and incorporate an academic paper into a publication was used as the central case. A shorter time of ten minutes to skim-read and cite a paper was assumed as the lower-bound. Finally, a higher time of sixty minutes to deep-read and critique a paper was assumed as the upper bound.
- Together, these estimates indicate an average **value per citation** of **£9.98** over the study period in the **central case** with an average **lower and upper bound of £2.63 and £18.66**, respectively.

Total economic benefit excluding impacts of research

Note, in the aggregate statistics of the total economic benefit excluding impacts of research, an adjustment to the citations and publications was made to only count the additional value of citations over and above what could be expected from similar publications not benefitting from ARCHER2 use. This was done using the Field-Weighted Citation Index (FWCI) to understand the number of citations that could be expected from similar publications within the same academic fields.

A3.2.3 Spillover impacts of scientific R&D

In addition to the value of knowledge creation itself (i.e., the value of publications produced), R&D also leads to substantial spillovers to the economy and society. However, while the presence of these spillovers is well established in the economic literature estimates of the magnitude of spillovers are limited. Box 20 provides a summary of the published literature.

To estimate the spillover impacts of R&D undertaken on ARCHER2 these published spillover multipliers were combined with funding data from ResearchFish™. ResearchFish™ captures R&D funding from UKRI research councils as well as non-research council further funding from academic, charity/non-profit, private and public sources. However, the latter relies on self-reporting by grant recipients and is therefore likely to be an underestimate.

As some R&D funding started on ARCHER1 and carried over to ARCHER2, only the estimated share after January 2021 was included in the calculation. To enable allocation of funding for the share of research undertaken on ARCHER2, funding amounts for multi-year projects were allocated equally across the years within the funding period. For example, if EPSRC provided a research grant of £1 million, with a start year of 2022 and an end year of 2023, £500,000 was allocated to 2022 and £500,000 to 2023.

Due to the significant uncertainty surrounding spillover estimates of publicly funded R&D a wide uncertainty interval was placed around spillover estimates:

- The **central value** assumed a **40% productivity uplift** for publicly funded R&D in line with Frontier Economics (2024)[91]. A lower productivity estimate of 20% was assumed for private sector further funding in line with Frontier Economics (2023)[92]. These estimates represent productivity estimates that raise annual productivity in perpetuity, rather than multiplier estimates of the total value of the spillover. To derive multipliers, we assumed a **time-lag of six years** between funding and productivity impacts materialising in line with Frontier Economics' analysis. Further, rather than impacts in perpetuity, we assumed a shorter **twenty-year time horizon over which productivity stays higher**, with values in future years appropriately discounted and discounting accounting for the time-lag prior to impacts. This resulted in a spillover **multiplier of 4.6** (around one-third of the spillover estimates derived by Haskel and Wallis) **for publicly funded research** and **2.3 for private further funding**.
- To derive the **lower bound estimate**, we assumed a **lower 20% productivity uplift** in line with the literature prior to Frontier Economics' 2024 updated analysis. Further, **productivity was assumed to stay higher for only ten years** rather than twenty years. This resulted in a lower spillover multiplier of **1.4.**
- For the **upper bound estimate**, research on the HPC specific returns from IDC (2014)[93] were used. Specifically, the study suggests a much higher return from investment in European academic HPC projects of **$30 : $1**. However, given the substantial difference in magnitude of this estimate compared to the general literature on R&D spillovers presented in Box 20, this was considered to be unlikely to materialise for all R&D undertaken on ARCHER2 across the board. Therefore, the sensitivity analysis assumed a truncated normal distribution assuming the 30 : 1 return was four standard deviations away from the central value (i.e., there is a less than 1% chance of this return being drawn in the sensitivity analysis, with higher return values between the central estimate and upper bound becoming increasingly less likely to be drawn).

Box 20 Published estimates of economic spillovers from publicly funded R&D

A small body of R&D literature provides estimates of productivity spillovers from academic R&D. A study by Haskel and Wallis[94] investigates evidence of **spillovers from public funding of Research & Development**. The authors analyse productivity spillovers to the private sector from public spending on R&D by the UK Research councils[95].

Haskel and Wallis find strong evidence of the existence of market sector productivity spillovers from public R&D expenditure originating from UK Research councils[a, b]: the marginal spillover effect of public spending on R&D through the **Research councils** stands at **12.7 (i.e. for every £1 spent on university research through the Research councils results in an additional output of £12.70 in UK companies)**. The analysis also suggests that the spillover benefits of public spending on research in higher education are greater than those from other R&D areas supported by government.

[91] Frontier Economics. (2024). Returns to public R&D: Report for the Department for Science, Innovation and Technology (DSIT).
[92] Frontier Economics. (2023). Rate of return to investment in R&D: Report for the Department for Science, Innovation and Technology (DSIT).
[93] Joseph, E. C., Conway, S., & Dekate, C. (2014). EESI-2 Special Study To Measure And Model How Investments In HPC Can Create Financial ROI And Scientific Innovation In Europe. International Data Corporation. Retrieved from https://risc2-project.eu/wp-content/uploads/2015/05/EESI2_D7.4_Final-report-on-HPC-Return-on-Investment.pdf
[94] Haskel, J., & Wallis, G. (2010). Public support for innovation, intangible investment and productivity growth in the UK market sector.
[95] The authors use data on government expenditure published by the Department for Business, Innovation and Skills for the financial years between 1986-87 and 2005-06.

A more recent study by Haskel et al.[96] provides additional insight into the size of potential productivity spillovers from university research. Rather than estimating effects on the UK economy as a whole, the authors analyse the size of spillover effects from public research across different UK industries[c]. The authors investigate the correlation between the combined research conducted by the Research councils, the higher education sector, and central government (e.g. through public research laboratories)[d], interacted with measures of industry research activity, and total factor productivity within the different market sectors[e]. Their findings imply a total rate of return on publicly-funded research of **0.2 (i.e. every £1 spent on public R&D results in an additional output of £0.20 within the UK private sector).** Subsequent work by Frontier Economics (2024)[97] used the same methodology as in Haskel et al. with more granular data than was possible at the time of these previous studies and found a **40% (0.4) annual rate of return to public R&D investment** in perpetuity

Further quantitative evidence on the return to public investment in scientific research is limited. Examples include:

- A report for the (former) Department for Business, Innovation and Skills[98], which evaluates the Haskel and Wallis approach, using a different (publicly available) dataset and a slightly different methodology. They find comparable productivity spillover impacts of Research council R&D investments estimated of **10.71**.
- Comparable research by Elnasri and Fox[99] applies the Haskel and Wallis approach to assess the productivity spillovers associated with publicly funded research in Australia. The authors find a similar research spillover to Haskel and Wallis (2010), albeit with a slightly lower research multiplier of **9.7** (which may be expected given the different country studied).
- A US-based study by Jones and Summers[100] undertakes an economy-wide calculation of the average social benefits of investments in innovation, including spillovers. They find a baseline **benefit-to-cost ratio of 13.3:1**, although their estimates range from 5 to more than 20 depending on the assumptions made in relation to inflation bias, health benefits, and the discount rate (among other factors).

It is important to note that both studies examine productivity spillovers of research across a wide range of research areas and are not specific to HPC. HPC R&D is likely to generate different productivity spillovers than those of general research. However, research specifically on productivity spillovers of HPC research is even rarer, though a recent study by the IDC[101] suggests that the return from investments into academic HPC projects may be significantly higher: among the European HPC projects analysed, **every $ invested in academic HPC projects generates a return on public sector investments of $30**. Though it should be noted that the study is based on a small sample of only seven academic HPC projects. Moreover, the study only includes success stories, not projects that didn't generate economic or scientific results. This likely biases the results, resulting in inflated HPC estimates.

Notes:

(a) Based on regression of total factor productivity growth in the UK on various measures of public sector R&D spending.

(b) The authors' regressions only test for correlation, so that their results could be subject to the problem of reverse causation (i.e. it might be the case that increased market sector productivity induced the government to raise public sector spending on R&D). To address this issue, the authors not only test for 1-year lags, but for lags of 2 and 3 years respectively, and obtain similar estimates. The time lags imply that if there were a reverse causation issue, it would have to be the government's anticipation of increased total factor productivity growth in 2 or 3 years which would induce the government to raise its spending on research ; as this seems an unlikely relationship, Haskel and Walls argue that their results appear robust in relation to reverse causation.

(c) Haskel et al. (2014) use data on 7 industries in the United Kingdom for the years 1995 to 2007.

(d) A key difference to the multiplier estimate for Research Council spending in Haskel and Wallis (2010) lies in the distinction between performed and funded research, as outlined by Haskel et al. (2014). In particular, whereas Haskel and Wallis estimated the

[96] Haskel, J., Hughes, A., & Bascavusoglu-Moreau, E. (2014). The economic significance of the UK science base: a report for the Campaign for Science and Engineering.

[97] Frontier Economics. (2024). Returns to public R&D: Report for the Department for Science, Innovation and Technology (DSIT).

[98] Department for Business, Innovation and Skills (2014a). Rates of Return to Investment in Science and Innovation

[99] Elnasri, A., & Fox, K. J. (2017). The contribution of R&D and innovation to productivity' Journal of productivity analysis, 47, 291-308

[100] Jones, B., & Summers, L. (2020). A Calculation of the Social Returns to Innovation

[101] IDC (2014). EESI-2 Special Study To Measure And Model How Investments In HPC Can Create Financial ROI And Scientific Innovation In Europe

impact of research funding by the Research Councils on private sector productivity, Haskel et al. instead focus on the performance of R&D. Hence, they use measures of the research undertaken by the Research Councils and the government, rather than the research funding which they provide for external research, e.g. by higher education institutions. The distinction is less relevant in the higher education sector: to measure the research performed in higher education, the authors use Higher Education Funding Council funding (where research is both funded by and performed in higher education).
(e) The authors regress the three-year natural log difference of total factor productivity on the three-year and six-year lagged ratio of total research performed by the Research Councils, government and the Higher Education Funding Councils over real gross output per industry. To arrive at the relevant multiplier, this ratio is then interacted with a measure of co-operation of private sector firms with universities and public research institutes, capturing the fraction of firms in each industry co-operating with government or universities. The lagged independent variables are adjusted to ensure that the resulting coefficients can be interpreted as annual elasticities and rates of return.
(f) Frontier economics used a 6-year lag between public R&D investment and the impact on productivity.

A3.2.4 Accounting for additionality

To account for additionality, adjustments were made to the benefits derived using the approaches described in the previous sections; specifically:

- The value of **publications** and **spillover impacts** of R&D were only monetised for the share of research that could not have taken place in the absence of ARCHER2. The rationale for this adjustment is that the share of research that could use other facilities or cloud infrastructure instead of ARCHER2 would still have taken place in the counterfactual, that is, in the absence of ARCHER2 and therefore these impacts are not additional. However, as the discussion in Section 8 has highlighted, the degree to which other facilities or commercial cloud are in fact alternatives in practice is debatable.
- For **direct benefits of access**, it was elected to only monetise avoided costs for the share of R&D that could have been undertaken in the absence of ARCHER2, as indicated by the user survey, and for which researchers indicated that they would have considered commercial cloud as an alternative. The rationale for this adjustment is that, in the counterfactual scenario where ARCHER2 did not exist, researchers (or their funders) would only pay for commercial cloud access for R&D, which could be undertaken via cloud alternatives to ARCHER2. However, some research would instead be using other publicly funded HPC facilities, while some research could not be substituted for another HPC facility or cloud alternatives and so would not take place in the absence of ARCHER2. Therefore, no avoided costs were assumed for this share of research.
- Note, that for the share of research that could in principle be undertaken without ARCHER2 using other publicly funded HPC facilities such as the Tier-2 centres or local systems, no benefits were monetised. This is because under the counterfactual there would be no additional costs to this research (as it would simply substitute ARCHER2 for another publicly funded HPC capability) and the benefits of the research would still materialise (and so there are no additional benefits from using ARCHER2). In practice, it is possible that there may be a reduction in quality by using a less powerful capability, though these impacts were not monetised.

To inform these assumptions, the user survey asked ARCHER2 users about i) the proportion of their research they would be unable to undertake without ARCHER2 and ii) for the proportion of research that they could undertake without ARCHER2, how they think they would do this? The resulting assumptions made are summarised in Table 25 and briefly described below.

In response to the first question, the median response was 65% (i.e., 65% of research undertaken on ARCHER2 is dependent on ARCHER2), which was used as the central assumption. Note, that the counterfactual survey questions exploring the "No ARCHER2 or ARCHER1 scenario" suggested a higher additionality of 90% of work that could not be undertaken. To be conservative, we retained the lower value as a central estimate and used the 90% as the higher estimate for the share of research dependent on ARCHER2. For the lower bounds, a lower value of 40% was used to account for optimism bias in survey responses. This is equivalent to the 25th percentile of responses. These assumptions were applied to the monetised benefits of the value of publications and citations, and the spillover impacts of research.

For avoided costs of direct access, the reverse of the above assumptions was used. For example, a central assumption of 35% was assumed for research that could not be undertaken without ARCHER2. In addition, adjustments were made for the share of this research that may use commercial cloud as an alternative. 28% of survey respondents said they were either "very likely", "somewhat likely" or "neither likely nor unlikely" to use cloud as an alternative, suggesting cloud may be a feasible alternative for these cases. Therefore, this was used as a central assumption. The upper-bound assumption also included the share of respondents that said they were "somewhat unlikely" to use cloud (i.e., in total 49% of respondents), while the lower-bound assumption included only those that said they were either "very likely", "somewhat likely" (14%).

In the central case, this implies that a value is associated with around three-quarters (75%) of research undertaken on ARCHER2 – 65% for which benefits of publications/citations and spillover impacts were monetised and 10% for which avoided costs of access were monetised and 25% for which no benefits were monetised. Note, where a combined value >100% was drawn in the sensitivity analysis the share of research for which avoided costs were calculated was adjusted downward so that no more than 100% of research is valued in any simulation run.

Table 25 Additionality adjustments for research benefit streams

Assumption	Central	Min	Max
Proportion of research that would not have been undertaken without ARCHER2 *(publications/citations and spillover impacts) (a)*	***65%***	***40%***	***90%***
Proportion of research that would have been undertaken without ARCHER2 (b = 100% - a)	35%	10%	60%
Proportion that may plausibly use commercial cloud (c)	28%	14%	49%
Proportion of research that could have been undertaken on cloud without ARCHER2 *(avoided costs of access) (d = b * c)*	***10%***	***1%***	***29%****
Implied share of research using other publicly funded HPC under the counterfactual (no benefits monetised) (e = 100% - a – d)	*25%*	*59%*	*0%*

Note: (*) Where a combined value >100% was drawn in the sensitivity analysis this value was adjusted downward so that no more than 100% of research is valued in any simulation run.

Source: London Economics analysis of ARCHER2 user survey

A3.3 The value of software products, IP and spinouts

This section details the approaches used to value software products, intellectual property (IP) and spinouts. The section first describes the general approaches used for each of these streams, then discusses the estimation of the number of software products, IP and spinouts used in the evaluation, and finally discusses the approach to attribution.

A3.3.1 Evaluation approaches:

The following evaluation approaches were used to monetise benefits of software products, IP and spinouts:

- **The value of software produced on ARCHER2:** To value software product on ARCHER2, a production cost approach was used. This approach values software based on the cost of its production. This is consistent with existing accounting practices, which value software as the cost of its production, considering any accumulated amortisation and impairment losses (International Accounting Standards Board, 2024)[102].

  The production cost approach used combined data on the number of software products from UKRI's research outcome database ResearchFish™, with insights from the user survey on the average time spent on developing software products with researchers' salary data.

  Assumptions on the average time taken to develop software was obtained from the user survey, where the average time taken to produce software on ARCHER2 was **265 days.** A 25% uncertainty interval was assumed surrounding this central estimate for the sensitivity analysis.

  The average researchers' salary[103] for each year is obtained using ONS salary data for those who work in education, adjusted for inflation using ONS deflators and the same underestimation adjustments discussed in Section A3.2.2 made.

  Note, that the production cost approach may undervalue software where it is sold on the open market and future income streams are derived. Further, the approach does not value any improvements generated from the use of software, for example, by speeding up simulations on existing HPC. However, these would implicitly already be captured in the other benefit streams and so monetising them would be double counting.

- **The value of patents/IP produced on ARCHER2:** To estimate the value of patents and other IP existing literature estimates on the value of intellectual property were used and combined with data on the number of IP products resulting from research undertaken on ARCHER2 obtained from UKRI's research outcome database ResearchFish™.

  While many studies highlighted the economic value of intellectual property to companies themselves, only limited literature evidence placing a £-value on patents were identified.

  The **low value** of a patent was extracted from Bessen (2006)[104], who updated calculations of Schankerman and Pakes (1986)[105]. Schankerman and Pakes (1986) used an income-based method to estimate the value of patents in 1970, based on the assumption that patents would only be renewed if the expected future income from the patent exceeded the cost of the renewal fee. If the expected income was lower than the renewal fee, the patent would not be renewed. The value of a patent estimated by Bessen was $11,128 in 2006 dollars **(£9,666** in £-2025 prices).

  For the **central and high estimates**, the median and mean values as reported in Serrano (2005) were used, which considered the market prices patents are traded at and assumed patent value equals the discounted present value of its stream of patent returns at a given age. These were $27,895 and $86,720 (**£29,631 and £92,150** in £-2025 prices), respectively.

[102] IAS 38. Intangible assets. Available at: https://www.ifrs.org/content/dam/ifrs/publications/pdf-standards/english/2021/issued/part-a/ias-38-intangible-assets.pdf

[103] Note, that it's not just researchers developing software on ARCHER2 but also computational research software engineers. As the salary data covers the education sector as a whole, no distinction was made between the salaries of researchers and computational research software engineers, i.e., they were both assumed to be equal.

[104] Bessen (2006). The Value of U.S. Patents by Owner and Patent Characteristics. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=949778

[105] Schankerman and Pakes (1986). Estimates of the Value of Patent Rights in European Countries During the Post-1950 Period. Available at: https://www.jstor.org/stable/2233173

- **The value of spinouts benefitting from ARCHER2:** Similarly to IP, to estimate the value of spinouts benefitting from ARCHER2 spinout data from ResearchFish™ was combined with existing literature estimates on the economic contribution of a spin-out to the UK.

  Similarly to intellectual property, only limited literature evidence monetising the economic contribution of spinouts was identified.

  In this case, the central and high estimate are based on a study by Cambridge Policy Consultants (2024)[106] on the value of BBSRC attributable spin-outs. Specifically, the lower real net GVA of a spin-out throughout the company lifetime (£12.9 million as reported by the authors equivalent to £15.3 million in £-2025 prices) was used the central value. The higher real net GVA of a spin-out throughout 20 years of operation (£17.5 million as reported by the authors equivalent to £20.8 million in £-2025 prices) was used as the upper-bound estimate. These estimates were judged to be the most closely aligned to the present study as it specifically considered spinouts from academic research funded by a UK research council.

  The low estimate was based on a study by London Economics (2024)[107] on the economic value of Russel Group spin-outs, which implied a slightly lower economic contribution of £9.2 million per spin-out (£10.4 million in £-2025 prices).

A3.3.2 Estimating the number of software products developed on and the spinouts and IP resulting from research on ARCHER2

Baseline data on the number of software products developed on ARCHER2 as well as the number of IP generated and spinouts benefitting from work on ARCHER2 was obtained from ResearchFish™. A total of 151 software products, 4 patents/other IP, and 6 spinouts benefitting from ARCHER2 were identified from ResearchFish™ data (though one was in administration). However, due to the self-reported nature of outcomes data in ResearchFish™, there was expected to be significant underestimation present in these figures. Furthermore, as these figures only capture data up until early 2025, the numbers are expected to grow over the remainder of the lifetime of ARCHER.

To understand the degree of underestimation the ARCHER2 user survey undertaken for this study were asked about whether their work on ARCHER2 led to the creation of software products, formation of spin-outs, or intellectual property. In response, 32% of users reported their work leading to the creation of software products, 6% reported generating of IP based on their work, and 2% reported that their work on ARCHER2 led to the creation of spin-outs (see Figure 24 in Section 3.3).

Combining these insights with data from EPCC on the total number of active academic user accounts of ARCHER2 (3,855), indicates significantly higher numbers of software products (1,234), IP (231) and spinouts (77) resulting from work undertaken on ARCHER2. However, these figures were assumed to be overestimates as i) multiple users could be involved in, for example, the same spin-out, and ii) to account for optimism bias in the survey responses.

Therefore, similarly to the treatment of uncertainty around research spillover impacts, a truncated normal distribution was used to characterise uncertainty around estimates of the number of software products, IP and spinouts used in the analysis. Specifically, it was assumed that:

[106] See https://www.ukri.org/wp-content/uploads/2024/02/BBSRC-07022024-Economic-impact-assessment-of-BBSRC-attributable-spin-outs-February-2024.pdf.
[107] See https://www.universitiesuk.ac.uk/sites/default/files/field/downloads/2024-09/LE-UUK-Impact-of-university-TL-and-RI-Final-Report.pdf

- As the survey figures are judged to be an overestimate it was assumed that these figures are highly unlikely and therefore the analysis assumed this figure is four standard deviations away from the mean of the distribution (i.e., there is a less than 1% chance for these figures to be drawn in the simulations).
- As, the ResearchFish™ data is most likely an underestimate, it was assumed that the actual figures should lie somewhere between the ResearchFish™ estimates and those implied by the user survey insights. However, as it was not clear by how much these figures are underestimated, the analysis aimed to remain on the conservative side for central estimates. Therefore, we assumed the central assumption to only be slightly lower than the mean value implied by the uncertainty distribution – equivalent to the 40th percentile.

Based on these uncertainty assumptions the truncated normal distribution could be characterised, and a plausible mean estimate to use as a central assumption for the number of products flowing from ARCHER2 *over ARCHER's lifetime* were calculated, as well as an implied standard deviation. The resulting mean (and standard deviation) estimates used for the central case were:

- **215** software products (std. dev. ~255)
- **17** patents or other intellectual property (std. dev. ~53)
- **9** spinouts (std. dev. ~17)

The mean and standard deviation estimates were used as inputs to a truncated normal distribution from which values were drawn in the simulation results. The actual data gathered from ResearchFish™ was used as the lower bound, with the assumed -distribution left-truncated at the figures implied by the ResearchFish™ data. The estimates implied by the user-survey were used to inform the calculation of the standard deviation and therefore constrained to be very unlikely (less than 1% chance) to be drawn in the simulations, making them the effective upper bound of the uncertainty distribution.

A3.3.3 Accounting for attribution

To understand the contribution of ARCHER2 to the development of software products and the formation of IP and spinouts, ARCHER2 users stating, in response to the user survey, that their work led to the creation of these products, to what extent being able to access ARCHER2 contributed to the formation of these products.

Users were asked to rate ARCHER2's contribution on a five-point scale ranging from "No contribution", to "Would have been impossible without access to ARCHER2". To derive attribution assumptions from these responses, assumed quantitative values for the contribution of ARCHER2 were assigned to each of these responses. This was done both using a more optimistic ("higher") scale and a conservative ("lower") scale.

A triangular distribution was again used to characterise the uncertainty around these attribution assumptions for the sensitivity analysis. To mitigate optimism bias, the central attribution assumption used the midpoint between these two scales, while the higher and lower estimates were used as the upper and lower bound respectively.

Table 26 provides an overview of the results of the survey, the assumed contributions, and the derived attribution assumptions used as an input to the modelling.

Table 26 User survey responses to attribution questions for spin-outs, software and IP

Contribution of ARCHER2	Spin-outs	Software products	Intellectual property	*Assumed contribution of ARCHER2*
Would have been impossible without access to ARCHER2	25%	31%	15%	*Higher: 100%* *Lower: 75%*
Large contribution	25%	36%	46%	*Higher: 75%* *Lower: 50%*
Medium contribution	38%	31%	23%	*Higher: 50%* *Lower: 25%*
Small contribution	0%	2%	0%	*Higher: 25%* *Lower: 5%*
No contribution	13%	0%	15%	*Higher: 0%* *Lower: 0%*
Assumed contribution of ARCHER2:				
Weighted average – upper bound:	63%	74%	62%	-
Weighted average – lower bound:	41%	49%	40%	
Midpoint – central:	**52%**	**62%**	**51%**	-

Note: Based on user survey responses received to the question of "To what extent has being able to access ARCHER2 contributed to these?" from respondents that said yes to the question of "Has your work on ARCHER2 led to the creation of any software products, formation of spin-outs, or intellectual property?"

Source: London Economics analysis of ARCHER2 user survey

A3.4 The value of skills development and training

Benefits of skills development and training were monetised in two ways. First, the direct value of training offers provided by EPCC was monetised using market prices of commercial training events. Second, the benefits to of skills development to early-career researchers themselves in terms of higher lifetime wages, the graduate premium, was estimated. Associated increased taxation receipts accruing to the exchequer were also monetised.

A3.4.1 The direct value of ARCHER2 training

The direct value of the training offerings provided by the ARCHER2 team was estimated using an avoided cost approach using market prices for similar training offerings from commercial training providers. Training offered to ARCHER2 users by EPCC is free at the point of use.

Desk research was undertaken to identify comparable commercial training offerings. Private HPC training courses could not be identified. Therefore, training events in cloud computing, cloud architecture and programming skills were used as commercially available alternatives. Table 27 lists the courses identified.

The average cost per day of training cost was **£749**. This was used as the central estimate. The 25th and 75th percentiles were used as the lower and upper bounds. These were £673 and £895, respectively.

Table 27 Commercial HPC training prices

Title of course	Provider	Number of days	Cost per day (£)
Advanced Python: Best Practices and Design Patterns	Learning Tree International	4	£549
Microsoft Azure Fundamentals Training (AZ-900T00)	Learning Tree International	1	£835
Architecting on AWS (AWSA)	Learning Tree International	3	£853
Developing on AWS (AWSD)	Learning Tree International	3	£832
Introduction to C++ Programming Training	Learning Tree International	4	£383
Advanced Python	CITYLIT	2	£125
C# Intermediate to Advanced	London Academy of IT	2	£225
Architecting on AWS	QA	3	£998
Advanced AWS Well-Architected Best Practices	QA	1	£845
Networking Essentials for Cloud Applications on AWS	QA	1	£845
VMware Cloud Foundation: Deploy, Configure, Manage [V5.2]	QA	5	£843
Amazon Web Services (AWS) Certified Cloud Practitioner (CLF-C02)	Firebrand	2	£938
AWS Certified Advanced Networking: Specialty (ANS-C00)	Firebrand	1	£850
Microsoft Certified: Azure Data Scientist Associate (DP-100)	Firebrand	3	£1,000
Microsoft Certified: Azure Fundamentals (AZ-900)	Firebrand	1	£1,000
Amazon Web Services (AWS) Certified Cloud Practitioner (CLF-C02)	Firebrand	2	£938
Microsoft: Develop data-driven applications by using Azure SQL Database	Firebrand	1	£825
U8583S: Linux Fundamentals (GL120)	Hewlett Packard Enterprise training	5	£598
HPE Performance Cluster Management Foundations	Hewlett Packard Enterprise training	2	£748

Source: London Economics desk research; Learning Tree International[108], CITYLIT[109], London Academy of IT[110], Firebrand[111], QA[112], Harlett Packard Enterprise training[113]

Data on training days provided and the number of attendees was obtained from EPCC. Table 28 shows the breakdown of this data by year between the start of ARCHER2 funding and mid-2025 when the data was obtained. Note, the analysis assumed that attendees attended all days of a training. For example, if there are 40 attendees for a two-day course, the analysis assumes this equals 80 training days attended in total

108 See https://www.learningtree.co.uk/courses/
109 See https://www.citylit.ac.uk/courses/advanced-python
110 See https://www.londonacademyofit.co.uk/training-course-advanced-c-sharp
111 See https://app-uks-webprod.azurewebsites.net/uk/courses/amazon-aws
112 See https://www.qa.com/browse/courses/cloud/
113 See https://education.hpe.com/us/en/training/portfolio/hpc.html)

Table 28 ARCHER2 training days

Year	Number of training days provided	Total number of attendee days
2020	33	981
2021	53	1,077
2022	60	1,007
2023	40	676
2024	64	1,235
2025*	21	334
Total (2020-2025)	**270**	**5,310**

Note: (*) Data for 2025 is partial only.

Source: London Economics analysis of EPCC data

A3.4.2 The economic benefit of skills development for early-career researchers

To estimate the benefits of skills development, this study assumes that the economic value of the skills acquired is equivalent to their value in the labour market (Atkinson, 2005)[114]. As computational skills are in short supply, graduates with strong computational skills moving into industry typically receive a wage premium compared to other graduates not trained in these skills.

The analysis combined data on estimated HPC earnings premium, counterfactual average salary data in technical fields, taxation information, and estimates of the number of PhD and postdoctoral students trained on ARCHER2, who move into UK industry to derive the total present value of the lifetime skills premium accruing to the graduates themselves, as well as of additional tax receipts accruing to the UK Exchequer.

While firms do benefit from a more skilled workforce, this benefit is considered neutral in economic terms, as wages tend to adjust in competitive labour markets. In equilibrium, wages rise to reflect the increased productivity of workers, effectively offsetting the gains to firms and resulting in a net economic benefit of zero for employers in equilibrium.

Estimation of wage premia received by graduates

To capture the wage premia received by graduates, earnings of graduates with HPC skills are compared to the earnings of those without these skills (the counterfactual group). Granular wage data for occupations requiring HPC skills was not available for this study. Therefore, the ARCHER2 user survey was used to derive a baseline estimate of the potential skills premia. Specifically, survey respondents were asked the following question: “*Based on your knowledge of the industry, approximately how much more on average do graduates with HPC skills earn in industry compared to those without HPC training?*”

Out of 70 respondents, the average was a 26.8% increase in salary. As with any self-reported or expert-elicited estimates, it is possible that this is an optimistic estimate. In the case of this study, the survey was completed by principal investigators who may have an optimistic view of the value of the skills their training provides as well as by early-career researchers themselves who may be optimistic about the wage premia they may be able to achieve in industry.

Therefore, the **26.8%** figure was used as an **upper-bound** and **adjustments for optimism bias** were made **for the central and lower-bound estimates**. These adjustments were 30% for the central

[114] Atkinson, B. (2005). Atkinson Review: Final Report. Measurement of Government output and productivity for national accounts.

estimate and 50% for the lower-bound estimate resulting in wage premia assumptions of **18.8% (central)** and **13.4% (lower-bound)**, respectively.

These estimates are in line with skills premia for other advanced computational skills reported. For example, a recent study by the Oxford Internet Institute[115]found that AI skills attract a 23% wage premium. Similarly, a recent study PwC[116] found that wage premia from artificial intelligence-related skills such as machine learning capabilities attract a 14% wage premium on average. However, this premia rises substantially for technical roles with AI skills in applications programmer roles found to attract a 24% wage premium, systems analysts a 34% premium and database designers and administrators a 58% wage premium.

Estimating the share of graduates moving into industry and adjusting for leakage

To capture the wage premia and its corresponding benefit to the UK Exchequer, the analysis only considered **graduates who stay in the UK after graduating and enter industry**. This is because we assume that those entering academia/the public sector on average do not earn a significant wage premium, and any benefits from those who leave the UK do not contribute to the UK economy.

Specifically, the following calculation was undertaken to estimate the proportion of total PhD/postdoc students trained on ARCHER2 who move on to work for UK industry:

Estimate of PhD and postdoctoral graduates trained on ARCHER2 who go into UK industry	=	*Total no. of PhD and postdoctoral students trained on ARCHER2*	x	*Average proportion of graduates who stay in the UK*	x	*Average proportion of graduates who go into industry*

Estimates for the destination of graduates were again obtained from the user survey. Specifically, survey respondents who had trained PhD and postdoctoral students were asked about the share of students they trained that i) stayed in the UK vs moving abroad and ii) stayed in academia, moved into the private sector, entered the public sector, or chose other career roles.

Responses to the user survey suggested that **44% of students trained went into private sector** and **69% stayed in the UK.** These estimates were used as the central estimates. A 25% uncertainty range was applied to these estimates in the sensitivity analysis. This indicates that of the **2,336**[117] total PhD students and postdoctoral researchers trained on ARCHER2 from 2020-2025 approximately 30% (705) move into UK industry in the central case.

Composition of baseline (counterfactual) wage data

Granular data on the specific industry fields graduates move to was not obtained. Therefore, baseline wage data was obtained for technical sectors judged most likely to benefit from HPC skills.

[115] University of Oxford (2025). Skills-based hiring driving salary premiums in AI sector as employers face talent shortage, Oxford study finds. Available at: https://www.ox.ac.uk/news/2025-03-04-skills-based-hiring-driving-salary-premiums-ai-sector-employers-face-talent-shortage

[116] Personnel Today (2024). Employers pay 14% wage premium for AI skills. Available at: https://www.personneltoday.com/hr/wage-premium-for-workers-with-ai-skills-pwc-ai-jobs-barometer/

[117] Based on active user accounts data by career stage obtained from EPCC. To allow for overestimation and to capture additional graduates trained over the remainder of the lifetime of ARCHER2 a 25% uncertainty interval was applied to this central estimate in the sensitivity analysis.

Specifically, our **counterfactual** is composed of average salary data, obtained from the ONS, between 2020 and 2025, and adjusted for inflation, from the following technical fields:

- professional, scientific and technical activities
- manufacturing
- financial services
- information and communication

It should be noted that this counterfactual wage captures the salaries of all professionals in the industries selected irrespective of whether they have a postgraduate degree or not. Granular data for postgraduate salaries at industry level was not available for this study. This means that the counterfactual wage is therefore likely lower than that for non-HPC postgraduates in those industries. However, as the HPC graduate premium is estimated as a % uplift from the baseline (rather than as a £-figure that is compared to the baseline wage figure), a lower counterfactual wage would entail an underestimate rather than an overestimate.

Wage data for postgraduate roles economy-wide (i.e., in all industries not just the technical fields selected above) is available from the Graduate Labour Market Statistics from the ONS. This was found to be in-line with the counterfactual wage derived.

Assumptions on years in labour force

The following assumptions were made to enable calculation of the total NPV of the lifetime HPC graduate premium:

- The average age on completion of PhD is 32[118]. This was used as the central assumption. [lower and upper bounds of 30 and 34 were assumed in the sensitivity analysis]
- The average retirement age is 66.5[119]. This was used as the central assumption. [lower and upper bounds of 65 and 68 were assumed in the sensitivity analysis]

These assumptions imply that, on average, a PhD graduate spends 34 years in the labour force. Applying a 3.5 % discount rate in line with HMT Green Book guidance, enables us to derive a net present value (NVP) of the lifetime HPC skills premia accruing to PhD and postdoctoral students trained on ARCHER2 and who move into the UK industry.

Additional taxation accruing to the UK exchequer

In addition to estimating the lifetime skills premia accruing to graduates themselves, the analysis also considered the associated additional tax receipts accruing to the UK exchequer from the increased earnings associated with HPC graduates possessing a set of skills which are valued in the workplace. The analysis considered additional tax receipts in the form of income tax, national insurance and VAT associated with additional earnings of graduates with HPC skills moving into UK industry. Note, the analysis of additional taxation did not make any conjectures on potential changes to future tax rates as a result of future policy choices. Rather, the analysis assumed current tax bands remained fixed over the lifetime of the graduates.

118 See https://www.uwslondon.ac.uk/blog/how-long-is-a-phd-uk
119 See https://www.fool.co.uk/personal-finance/research/average-retirement-age-in-the-uk

Summary of estimations

Table 29 summarises the calculations made to derive the HPC earnings premium and associated tax receipts in the central case. Note, that the graduate skills premium was assumed static over the working life of graduates. That is, while a wage premium for HPC skills was assumed in each year over the working life of graduates, faster wage growth due to higher pay-rises or earlier/additional promotions were not assumed.

Table 29 ARCHER2 earnings premium

Benefit	Annual per graduate	Present value
Baseline salary in technical fields (counterfactual)	£49,230	-
Assumed salary of PhD and postdoctoral graduates with HPC skills moving into UK industry	£58,465	-
Gross HPC earnings premium (per graduate)	£9,235	£189,752
Additional income tax collected	£3,486	£71,626
Additional National Insurance contributions collected (employee contributions)	£247	£5,075
Additional VAT collected	£474	£9,746
Net HPC earnings premium	**£5,502**	**£113,051**
Additional National Insurance contributions collected (employer contributions)	£1,385	£28,463
Additional tax receipts and NI contributions accruing to Exchequer	**£5,593**	**£114,910**

Note: No data was available on the difference in earnings growth between graduates with and without HPC skills. Therefore, the earnings premium is assumed constant over the working life. Average age of completion assumed to be 65. Average age of retirement assumed to be 32. HM Treasury Green Book discount rate of 3.5% used to discount future earnings.

Source: London Economics analysis based on ARCHER2 user survey and ONS/HESA data

A3.5 Contribution of industry impacts to UK output

ARCHER2 is primarily a capability for UK scientific R&D. Industry benefits primarily from spillover impacts of research undertaken on ARCHER2 as well as through industry-academia partnerships. Nevertheless, industry users can apply for time on ARCHER2 to EPCC, which can allocate discretionary time allowances from its allocated director's time. As a result, 151 industry users from 34 commercial organisations also access ARCHER2 compute resources directly.

Establishing relevant UK firm population and associated firm characteristics

To understand the firm size distribution and account for leakage. Commercial organisations were matched to Moody's Orbis database to identify their geographic location and firm size.

This suggested that of the 34 firms, 16 firms (**47%**) are **headquartered in the UK**. However, many of the larger firms were organisations operating internationally. As such, Orbis may count them as overseas firms even if they are accessing ARCHER2 as part of their UK operations. Therefore, the 47% was seen as a **lower-bound**. The **upper bound was assumed to be 100%,** i.e., all firms accessing ARCHER2 are doing so via their UK arm. A **central assumption towards the lower end of the scale** (**60%** equivalent to 20 UK firms out of 34 accessing ARCHER2 directly) was assumed as shown in the table below.

Table 30 Industry benefits - leakage assumptions

Scenario	Share of firms that are UK based
Central	60% (40% leakage)
Lower-bound	47% (53% leakage)
Upper-bound	100% (no leakage)

Source: London Economics assumptions

Regarding the firm size distribution, the majority of firms identified were either large firms (38%) or micro firms (44%). A further 6% were medium-sized firms, while 13% were small firms. This was based on standard employment and turnover definitions commonly used to classify firms.

As turnover data for micro firms was not available and the applicable turnover for international firms was unclear from Orbis, turnover of firms in each size band was approximated using average turnover by firm size from the ONS business population estimates. Further, to enable conversion of turnover increases to GVA contributions, aGVA : turnover ratios were obtained for each firm size band from the ONS' Annual Business Survey estimates.

These assumptions are summarised in Table 31 below.

Table 31 Industry benefits – assumptions by firm size

Firm size	Proportions of UK firms accessing ARCHER2 by firm size	Average annual turnover by firm size	aGVA : turnover ratios
Large	38%	£ 312,953,474	0.31
Medium	6%	£ 25,139,188	0.30
Small	13%	£ 3,636,200	0.35
Micro	44%	£ 599,081	0.42

Source: EPCC data on firms directly accessing ARCHER2 matched to Orbis to identify firm size bands. Turnover by firm size band obtained from ONS business population estimates 2024. aGVA : turnover ratios obtained from the Annual Business Survey estimates.

Estimating benefits of ARCHER2 access

To estimate benefits by industry users, input from industry users was sought via stakeholder consultations and the user survey. Unfortunately, this target group proved very difficult to reach and no assumptions were obtained.

Therefore, the lower- and upper-bound assumptions profit increases seen by firms were assumed to be in-line with the findings from the previous 2019 evaluation of ARCHER1. Specifically, the previous study found that:

- One respondent provided a range for benefits in the previous online survey that suggested an increase in profit of around 25% over the last five years (or a CAGR of around 4.6% per year). In addition, one SME working with the EPCC as part of the Fortissimo project reported an annual increase in profits of around 3%. Therefore, the **4.6% increase in profit** was used **as the upper-bound assumption**.
- For the lower-bound, the previous study assumed no profit increases accruing to firms. Rather, any profit impacts would accrue only from cost savings by firms avoiding

investments into their own HPC capabilities. While this figure can vary significantly depending on the configuration, the previous study estimated an average cost-saving of around 0.55% of average profit assuming the purchase of a lower-end supercomputer with a five-year replacement cycle and allowing for operational costs of up to 20% of the purchase value. In line with this estimate, a **0.55%** of profit assumption deriving from cost savings was used as the **lower-bound**.

- The previous study did not assume any central assumptions. Therefore, without additional information, a **central value** towards the lower end of the range (**1% annual profit increase**) was assumed.
- In addition, 38% of respondents to the previous industry survey said they had received increases in sales / turnover or profit. Therefore, it was assumed that the average annual profit increases only accrue to 38% of firms in the central case. However, significant uncertainty surrounds this estimate. Therefore, a 50% uncertainty interval was applied around this central value in the sensitivity analysis.

Further, it was assumed that:

- There is a delay until profit increases from research undertaken on ARCHER2 materialise. This was assumed to be six years in line with estimates of the delay of productivity impacts materialising from scientific R&D identified by Frontier Economics (2024)[120].
- The economic literature suggests that firm-level profit increases from R&D persist for a number of years, though more precise estimates could not be obtained. To remain conservative, a central value of three years was chosen. The lower and upper bounds were set to one year and five years, respectively.
- Finally, desk research on the average net profit margins of firms was obtained in order to convert turnover assumptions to profit. This suggested profit margins in the region of between 5% to 15%, with an average net profit margin of UK firms of around 9%[121].

These assumptions are summarised in Figure 30, below. They imply annual turnover increases for the average firm of between 2.1% and 17.5% - with a central assumption of 4.2% - accruing after six years and persisting for between one to five years. For comparison, reported profitability increases from cloud computing were found to be up to 11.2% year-on-year[122] while a study from Google found that SMEs using the cloud grow 26% faster and are 21% more profitable than their peers[123].

Table 32 Industry benefits – assumptions to estimate profit increases

Assumption	Central	Lower-bound	Upper-bound
Average annual profit increase from R&D on ARCHER2	1%	0.55%	4.60%
No. of firms that see profit increases	38%	19%	57%
Average net profit margin	9%	5%	15%
Effective annual turnover increase for the average firm	***4.2%***	***2.1%***	***17.5%***
Delay until profit increases are felt	6	6	6
No. of years higher profit persists	3	1	5

Source: London Economics research

[120] Frontier Economics. (2024). Returns to public R&D: Report for the Department for Science, Innovation and Technology (DSIT).
[121] See, for example, https://www.venasolutions.com/blog/average-profit-margin-by-industry and https://www.annetteandco.co.uk/good-percentage-for-profit
[122] See https://www.cloudzero.com/blog/cloud-computing-statistics/
[123] See https://www.techuk.org/resource/how-the-cloud-accelerate-smes-computing-capabilities.html.

To calculate economic benefits, assumptions on average firm level profit increases were converted to effective turnover increases using the assumptions on average firm-level net profit margins and the share of firms assumed to see profit increases. The assumptions on the time-lag and persistence of profit increases were then used to calculate the effective net present turnover increase for the average firm, using appropriate discounting in line with the HMT Green Book recommended rate of 3.5%. Finally, the firm level turnover estimates were converted to contributions to UK Gross Value Added using the aGVA : turnover ratios.

A3.6 The economic value of public outreach

To monetise the value of public outreach events, this study estimated the non-market benefits associated with outreach events by estimating the opportunity cost for visitors to attend these events.

Between 2020 and 2025, a total of 29 outreach events for ARCHER2 took place, including science fairs, career fairs, workshops, insight days, talks and drop-in sessions, and work experience. The number of events, the total effort in days and the number of visitors is shown in Table 33, below.

Table 33 ARCHER2 outreach events

Year	Total number of outreach events	Total effort (days)	Number of visitors
2020	1	2	200
2021	0	0	0
2022	2	60	40,000
2023	8	149	43,149
2024	11	143	160,503
2025	3	22	1,495
2026	4	63	46,069
Total	**29**	**439**	**291,416**

Note: Total effort (days) calculates the total number of days worked on outreach events by staff combined. E.g. a 2-day outreach event with 5 staff numbers is 10 outreach effort days.

Source: London Economics analysis of EPCC outreach data

Desk research was undertaken to inform assumptions on the target audience of each event, the EPCC outreach events data was broken down into the demographic groups of attendees. The data was then sorted into weekday and weekend events based on their dates. For the public, there were 16 public outreach days during the week and 12 outreach days on the weekend. There were 35 student-only events, which all fell on weekdays.

Figure 60 Number of EPCC outreach events by target demographic

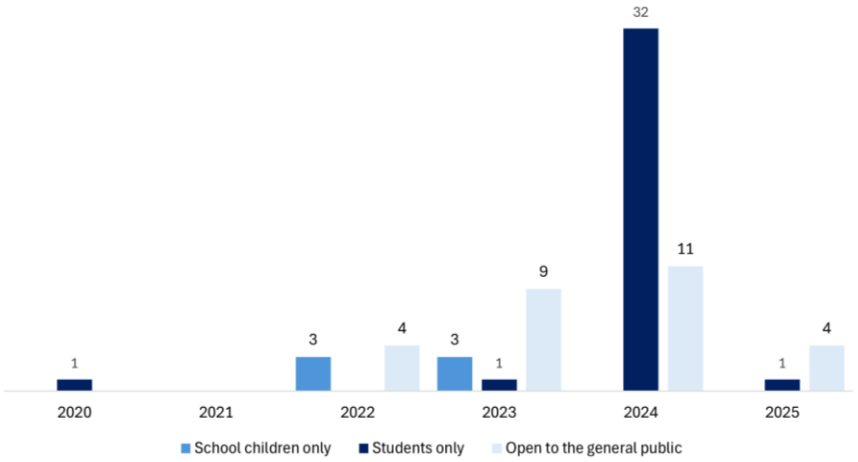

Note: This contains the number of outreach days (individual dates), rather than the number of staff effort days.
London Economics analysis of EPCC outreach data

For events that were open to all, the following attendance assumptions were made:

Table 34 Attendance assumptions for open-to-all outreach events

Target demographic	Central	Lower-bound	Upper-bound
Adults	50%	30%	60%
Students	45%	50%	40%
Children	5%	20%	0%

Source: London Economics' assumptions

The **opportunity cost** per attendee was calculated by considering the probability that if they had not attended the outreach event, they would have gone to work for a day.

Opportunity cost assumptions were assumed to vary by demographic group:

- For children no opportunity cost was assumed. For example, the ARCHER2 outreach team attended the Big Bang Fair124 which was solely for school aged children. Hence, there will not be an opportunity cost associated with these children attending the fair.
- For students, desk research was conducted to estimate the average proportion of students who work during university, and how many hours they typically work[125]. This fed into the opportunity cost parameter, which captured the proportion of students which may have chosen to work a shift instead of attending the outreach event.
- For adults, the **proportion of days attended which were weekdays was considered**, assuming that there is no opportunity cost for attending on the weekend for the majority of adults. For students, who typically do not work Monday-Friday, this is not the case, so the opportunity cost was assumed to be constant across the week, irrespective of weekday. It was also considered that some adult attendees will be part time.

[124] https://www.thebigbang.org.uk/the-big-bang-fair/
[125] Assuming 56% of students work an average of 14.5 hours a week, the proportion of a given student working on a given day (for 7hrs) is 16.6%. Source: https://www.independent.co.uk/news/uk/university-nick-hillman-government-higher-education-policy-institute-london-b2561666.html

To enable estimation of opportunity cost, it was assumed that students working would make minimum wage, while salary for adults was assumed in line with the median salary of a UK citizen.

The opportunity cost of attendance was the approximated using the following formula:

Opportunity cost of a demographic of attendees	=	*Estimate of total number of days that demographic attended ARCHER2 outreach events*	x	*Probability that someone from this demographic would work if they didn't attend the outreach event*	x	*Average wage for a day of labour from this demographic*
				Average opportunity cost of attending one day of public outreach		

A3.7 The economic value of the helpdesk

The value of the helpdesk was calculated by considering the value of researchers' time saved due to the support received. Specifically, the following calculation was used to approximate the value of the time saved by researchers by the helpdesk:

Value of the ARCHER2 helpdesk	=	*Average hourly researcher salary*	x	*Average time saved per query (hours)*	x	*Number of queries solved by the helpdesk*
				Time saved by the helpdesk		

Estimates of time savings from helpdesk support were obtained from the user survey. To approximate the average time saved per query, survey response options were coded into assumptions of time saved per query. A weighted average of the time savings (in hours) per query based on the distribution of responses obtained from the user survey was then calculated. Different codings were used to obtain low, central and high estimates (in hours). The assumptions made are shown in Table 35.

Table 35 Assumptions on average time saved per query, based on ARCHER2 user survey

Survey option	% of responses	High (hours)	Low (hours)	Central (hours)
Yes, up to 15 minutes per query	10%	0.25	0.15	0.2
Yes, up to an hour per query	19%	1	0.75	0.5
Yes, several hours per query	39%	4	2	3
Yes, more than several hours per query	30%	8	6	7
No, the helpdesk has not saved me any time	3%	0	0	0
Don't know / not applicable	0%	0	0	0
Weighted average time saving per query	**100%**	**4.1**	**2.7**	**3.3**

Source: Analysis based on London Economics ARCHER2 user survey

The assumed average time saved per query was then multiplied by the total number of queries solved by the helpdesk, obtained from data provided by EPCC.

A3.8 Adjustments for continuation of ARCHER1 counterfactual

This section outlines the adjustments made to the estimation of the benefit streams detailed in the earlier sections under the counterfactual scenario where ARCHER2 is not acquired and instead ARCHER1 continues to operate.

Table 36 provides an overview of the adjustments applied to each of the benefit streams modelled. Table 37 sets out the underlying assumptions underlying these adjustments and the rationale behind these assumptions.

Table 38 and Table 39 provide further details on responses from the user survey and the coding used to convert these responses to assumptions on the % of research affected and the % of this research where quality is affected.

Table 36 Adjustments for continuation of ARCHER1 scenario by benefit stream (central case)

Benefit stream	Adjustments for continuation of ARCHER1 scenario
Direct benefits of access to ARCHER2 for academic users	Assumed reduction in capacity of five times in line with the extra capacity provided by ARCHER2 over ARCHER1. No degradation over time is assumed as core-hours stay constant.
Value of publications and citations	Assumed that ARCHER1 has reached capacity. Therefore, 2020 levels of publications are assumed throughout. In addition, a small (5% in the central case) degradation is assumed over time. To represent a reduction in quality of science, an effective reduction of 17% in the number of citations per publication is assumed.
Value of software	Similarly to publications, it was assumed that ARCHER1 had reached capacity, so 2020 values for software products produced were used throughout and a small 5% degradation over time was applied.
Spillover impacts of scientific R&D	Similarly to publications, it was assumed that ARCHER1 had reached capacity, so 2020 values for science funding were used throughout and a small 5% degradation over time was applied. A proportional reduction in cost was also assumed. In addition, a reduction in line with that assumed for citations was applied to spillover multipliers to account for a reduction in quality of science.
Contribution of direct access by industry to UK GVA	It was assumed that due to the reduced capacity of ARCHER1, the full capability would be used for academic R&D purposes. Therefore, benefits of industry access were assumed to be nil.
Value of spinouts and IP	Given the low and sporadic incidence of spinouts and IP in ResearchFish^{TM} reductions in the number of spinouts and IP were assumed to be proportional to the assumed reduction in software products.
Direct provision of training	A reduction in training days line with an assumed reduction in user numbers was assumed. The reduction was estimated by assuming user numbers of ARCHER1 had peaked at its end of life.
HPC graduate premium and associated exchequer benefits	A reduction in the number of PhDs and postdoctoral students trained in line with the assumed reduction of users was assumed. No reduction to the graduate premium was assumed with the rationale that training with a lower capacity machine would not impact the skills learned.
Public outreach	Outreach activities were assumed to be unaffected. That is, no change in the number of outreach events or visitors reached were assumed.
Specialised helpdesk support	A reduction in line with the assumed reduction in science capacity was assumed. No degradation was assumed over time.
Costs	Initial capital investment for ARCHER2 was assumed to be saved, but capital costs for maintenance and operational expenses for service provision were assumed to continue*. In addition, a reduction in research funding was assumed in line with the adjustments made to the spillover impacts of research.

Note: (*) In practice, it may not have been possible or viable to keep ARCHER1 operational. Due to the age of ARCHER1 as an end-of-life machine, the required maintenance may have been significantly higher and/or additional capital investments may have been needed to retain ARCHER1 operational. As such, the assumed cost should be seen as an underestimate of the likely actual cost.

Table 37 Continuation of ARCHER1 modelling assumptions

Assumption	Central	Lower-bound	Upper-bound	Rationale
Reduction in science capacity	**5**	**5**	**5**	In line with additional capacity provided by ARCHER2 as outlined in the ARCHER2 business case
Reduction in quality of science:				
% of research on ARCHER2 affected	76%	58%		Estimated from user survey additionality question of continuation of ARCHER1 (Table 38)
Of research affected, % where quality is affected	91%	78%		Estimated from user survey additionality question regarding impacts on quality of continuation of ARCHER1 (Table 39)
Effective proportion of research where quality is affected	**69%**	**45%**	**91%**	% of research on ARCHER2 affected * % where quality is affected. For the upper bound it was assumed that the central quality reduction would apply to all research undertaken.
Assumed magnitude of reduction in quality	25%	0%	50%	ARCHER2 provides a 5x improvement in capacity, but several stakeholders interviewed reported that in the real-world this translated to only around a 2x improvement enabled for their simulations as the additional capacity is shared by a larger number of users (or conversely a 50% real-world reduction of ARCHER1 relative to ARCHER2). The 50% reduction in quality, chosen as an upper bound, aligns with these insights, though is unlikely to apply to all users equally. However, no other quantitative insights on the magnitude of quality-reduction that would be observed was available, though responses to the survey suggest sizeable quality impacts. Given the uncertainty a large uncertainty interval was chosen with the lower-bound assuming no quality reductions at all and the upper-bound assuming a 50% reduction as outlined above. A central assumption in the middle of the range was chosen. Given the major impacts reported by users, this was judged to be a conservative assumption and therefore justified as a central case.
Effective reduction in quality	**17%**	**0%**	**46%**	Effective proportion of research where quality is affected * assumed reduction in quality
Reduction in users	**32%**	**24%**	**40%**	Assumed in line with ARCHER2 users vs. users in final operation of ARCHER1, assuming ARCHER1 had reached capacity. 25% uncertainty around central estimate assumed
Degradation over time	**5%**	**10%**	**0%**	Stakeholders emphasised a degradation in service quality/suitability over time were ARCHER1 to continue operating. However, no concrete assumptions were available. Therefore, only a small degradation over time was assumed to remain conservative.

Table 38 % of research affected under continuation of ARCHER1

Q45. What impact would this scenario have on your research / work?	Proportion	Weighting (central)	Weighting (low)
Catastrophic impact (all or almost all of my research/work affected)	25%	100%	100%
Major impact (the majority of my research/work affected)	62%	75%	50%
Some/small impact (a small proportion of my research/work affected)	9%	25%	5%
No impact (none of my research/work affected)	0%	0%	0%
Don't know / not applicable	2%	0%	0%
% of research affected under continuation of ARCHER1 (weighted average):		**76%**	**58%**

Source: LE analysis of user survey

Table 39 Of research affected, % where quality is affected

Q45. In this scenario, to what extent do you agree or disagree that you would be unable to undertake your research / work to the same quality?	Proportion	Weighting (central)	Weighting (low)
Strongly agree	72.0%	100%	100%
Slightly agree	21.7%	75%	50%
Neither agree nor disagree	4.2%	25%	5%
Slightly disagree	1.6%	0%	0%
Strongly disagree	0.5%	0%	0%
% of research affected under continuation of ARCHER1 (weighted average):		**91%**	**78%**

Source: LE analysis of user survey

Annex 4 Additional information

A4.1 GPU vs. CPU-based HPC architectures

The **Central Processing Unit (CPU)**, often referred to as the "brain of the computer," is a general-purpose processor designed to manage a wide variety of tasks, both basic and complex. It excels at **sequential serial processing**, meaning it can handle only a limited number of computational threads simultaneously but does so at very high speeds with great precision.

In contrast, the **Graphics Processing Unit (GPU)** was originally developed to accelerate 3D graphics rendering and visual effects. However, it has since evolved into a powerful processor for **parallel computing**, making it highly effective at handling tasks that can be broken down into smaller, concurrent operations. These include **video rendering, scientific simulations, and machine learning algorithms**.[126]

Thus, the key distinction between CPUs and GPUs lies in their processing style:

- **CPUs** are more efficient at managing **a few complex and logic-intensive tasks** at high speed.
- **GPUs** are optimised for handling **many simpler tasks simultaneously**, making them ideal for parallel computing tasks.

ARCHER2 uses a CPU-based architecture. Although some codes can be ported to GPU-based systems with minimal effort, many of the software codes used require significant computational software engineering work to be efficiently adapted for GPUs. This process is possible, but not straightforward. Furthermore, some of the largest simulations cannot be run on GPUs. This is due to limitations in GPU memory capacity, intra-node bandwidth, and the complexity and high cost of adapting code to GPU architectures

[126] See https://tecadmin.net/cpu-vs-gpu-key-differences/

A4.2 ARCHER2 funding

Table 40 ARCHER2 funding breakdown 2020-2025

Funding type	2020/21	2021/22	2022/23	2023/24	2024/25	Total	Total (+VAT)
Capital (Hardware)	£36.0M					£36.0M	£43.2M
Capital (maintenance)	£1.0M	£1.0M	£1.0M	£1.0M	tba	£4.0M	£4.8M
Total capital funding	**£37.0M**	**£1.0M**	**£1.0M**	**£1.0M**		**£40.0M**	**£48.0M**
Service Provision funding						£7.2M	£8.6M
CSE support funding						£10.6M	£12.7M
Hosting of the service at the Advanced Computing Facility						£21.3M	£25.6M
Total resource funding	**£7.0M**	**£7.7M**	**7.8M**	**£8.3M**	**£8.3M**	**£39.1M**	**£46.9M**
TOTAL	**£44.0M**	**£8.7M**	**£8.8M**	**£9.3M**	**£8.3M**	**£79.1M**	**£94.9M**

Source: Full Business Case for ARCHER2

A4.3 Top Ten software codes run on ARCHER2

Table 41 Top Ten software codes run on ARCHER2

Software name	Number of jobs	Node hours	Percent of overall (%)	Number of users	Number of projects run on ARCHER2
VASP	71,386	897,194.4	22.8	140	15
Unknown*	459,092	770,578.1	19.6	439	99
GROMACS	14,191	217,271.4	5.5	46	8
Met Office UM	39,987	210,203.9	5.3	37	4
Nektar++	1,156	190,837.2	4.9	15	4
No srun**	78,066	184,589	4.7	646	101
OpenFOAM	8,817	171,084.2	4.3	52	19
CP2K	40,921	161,750.7	4.1	47	13
Python	27,137	155,682.8	4	56	28
Code_Saturne	192	103,688.4	2.6	4	2
LAMMPS	26,288	101,428.6	2.6	48	14
CASTEP	126,382	76,739.7	2	40	7
FHI aims	107134	71712.9	1.8	23	5
Overall	**1,108,954**	**3,933,431.6**	**100**	**901**	**129**

Note: * Jobs labelled as "Unknown" refer to interactive commands or applications run under testing.

** Jobs labelled as "No srun" jobs running interactively outside of the SLURM batch system because they are incompatible with SLURM or part of lots of scripted runs.

Source: **ARCHER2. (n.d.*). Service status: Software usage data*. ARCHER2.** https://www.archer2.ac.uk/support-access/status.html

A4.4 UK public HPC facilities for academic usage

Table 42 Overview of public UK HPC facilities for academic usage (Tier-1 and Tier-2)

HPC service	Location(s)	Primary remit	Start year	Expected end year
ARCHER2 (Tier-1)	**EPCC, University of Edinburgh**	**EPSRC, NERC**	**2021**	**2026**
DiRAC-3 (Tier 1*)	Universities of Cambridge, Leicester, Durham and Edinburgh	STFC	2022/23 (Edinburgh & Durham), 2023/24 (Cambridge & Leicester))	2026
Hartree Centre	Daresbury Laboratory	STFC	2012	Not announced
Baskerville (Tier-2)	University of Birmingham	EPSRC	2021	March 2026
Bede (Tier-2)	Durham University (N8 CIR)	EPSRC	2020	March 2026
Cirrus (Tier-2)	EPCC, University of Edinburgh	EPSRC	2017	August 2025
CSD3 (Tier-2)	University of Cambridge and partners	EPSRC	2017	2023/2024
Isambard-2 (Tier-2)	GW4 Alliance (Bristol, Bath, Cardiff, Exeter)	EPSRC, STFC	2020	March 2024
Isambard-3 (Tier-2)	GW4 Alliance (Bristol, Bath, Cardiff, Exeter)	UKRI	2024	Not announced
JADE-2 (Tier-2)	University of Oxford (JADE Consortium)	EPSRC	2021/22	January 2025
JASMIN	RAL, Harwell Campus, Oxfordshire	NERC (operated by STFC)	2012	Not announced
MMM Hub (Tier-2)	UCL (TYC and SES consortium)	EPSRC	2017	Not announced
Sulis (Tier-2)	University of Warwick	EPSRC	2021	November 2025
Kelvin-2 (Tier-2)	Queen's University Belfast, Ulster University	EPSRC	2022	Not announced

Note: DiRAC-3 was classified as Tier-1 in line with the Future of Compute Review (DSIT (2023). Independent Review of the Future of Compute: Final report and recommendations. Available at: https://www.gov.uk/government/publications/future-of-compute-review/the-future-of-compute-report-of-the-review-of-independent-panel-of-experts [accessed 21/05/2025])

Sources: **HPC-UK. (n.d.). *Facilities*. Retrieved May 27, 2025, from https://www.hpc-uk.ac.uk/facilities/; DiRAC. (2023). *DiRAC High Performance Computing Facility*. Retrieved from https://dirac.ac.uk/; Gross, A. (2024, November 19). *UK will fall behind supercomputer rivals, head of axed £800mn project warns*. Financial Times. https://www.ft.com/content/a8be5585-59a7-4d2a-a7c9-8106353e4594; UK Research and Innovation. (2024, February 29). *Funding opportunity: Access to High Performance Computing facilities 2024.* UKRI. Retrieved from https://www.ukri.org/opportunity/access-to-high-performance-computing-facilities-2024/; EPCC. (2017, June 19). *Cirrus transformed into Tier-2 system*. The University of Edinburgh. Retrieved from https://www.epcc.ed.ac.uk/whats-happening/articles/cirrus-transformed-tier-2-system; EPCC. (n.d.). *Cirrus powered by EPCC*. Retrieved from https://www.cirrus.ac.uk/; Thompson, S. (2021, September 22). *The Launch of Baskerville – a Tier 2 HPC system*. University of Birmingham. Retrieved from https://blog.bham.ac.uk/bear/2021/09/22/the-launch-of-baskerville-a-tier-2-hpc-system/; Zhang, Y. (2025, May 20). *May 2025 BEAR Newsletter*. University of Birmingham. Retrieved from https://blog.bham.ac.uk/bear/2025/05/20/may-2025-bear-newsletter/; JADE. (2025). *JADE HPC UK*. University of Oxford. Retrieved May 29, 2025, from https://www.jade.ac.uk/; Sulis HPC. (2025). University of Warwick. (2021). *New Tier 2 HPC platform launched by HPC Midlands+.* Retrieved May 29, 2025, from https://warwick.ac.uk/news/pressreleases/new_tier_2/; *Sulis: A Tier 2 HPC Platform for Ensemble Computing.* University of Warwick. Retrieved May 29, 2025, from https://warwick.ac.uk/research/rtp/sc/sulis/access/; Thomas Young Centre. (2025). *Materials & Molecular Modelling Hub (MMM Hub).* Retrieved May 29, 2025, from https://thomasyoungcentre.org/working-with-industry/materials-molecular-modelling-hub-mmm-hub/; GW. (2025). *BEDE service to be extended until March 2026*. N8 CIR. Retrieved May 29, 2025, from https://n8cir.org.uk/news/bede-service-to-be-extended-until-march-2026/; N8 Research Partnership. (2020). *Bede, the N8's new high performance computing platform, is now open for users.* Retrieved May 29, 2025, from https://www.n8research.org.uk/bede-the-n8s-new-high-performance-computing-platform-is-now-open-for-users/; Dunlop, J., Del Debbio, L., & Portelli, A. (2025). *DiRAC-3 Operations 2023-26 - Edinburgh*. University of Edinburgh. Retrieved May 29, 2025, from https://www.research.ed.ac.uk/en/projects/dirac-3-operations-2023-26-edinburgh-2; UKRI. (2025). *GW4 Tier-2 HPC Centre for Advanced Architectures.* Retrieved May 29, 2025, from https://gtr.ukri.org/projects?ref=EP/T022078/1; University of Cambridge. (2018). *Research Computing Services HPC Documentation*. Retrieved May 29, 2025, from https://docs.hpc.cam.ac.uk/hpc/#:~:text=The%20Cambridge%20Service%20for%20Data%20Driven%20Discovery%20(CSD3),chapter%20in%20HPC%20for%20the%20University%20of%20Cambridge.; *Woods, R. (2022). Kelvin 2 supercomputer launched at Queen's University in Belfast. Queen's University Belfast.* Retrieved May 29, 2025, from https://pure.qub.ac.uk/en/clippings/kelvin-2-supercomputer-launched-at-queens-university-in-belfast; UK Research and Innovation. (2023). *High Performance Computing: EPSRC Service Specification.* Retrieved from https://www.ukri.org/wp-content/uploads/2023/10/EPSRC-16102023-A2HPC-2023-2-HPC-Service-Specification.pdf**

Annex 5 Case studies

A5.1 Most prominent software codes run on ARCHER2

A variety of software packages were used on ARCHER2, the nine packages outlined in Table 43 used the highest share of node hours[127], making up 54% of all ARCHER2 use.

These software tools can be grouped into three broad areas based on what they simulate: atoms and electrons, molecules and large-scale systems like weather or fluids.

CP2K, VASP and CASTEP are used to study how atoms and electrons behave. This is useful in fields like chemistry, physics and materials science. They work by applying quantum mechanics, a branch of physics that describes the behaviour of very small particles. CP2K can simulate many different systems, including solids, liquids and biological molecules[128]. VASP focuses on modelling materials from first principles, using accurate equations to describe electronic structure.[129] CASTEP also works at this level and is useful for predicting the results of experiments, such as how materials absorb light or vibrate.[130]

GROMACS and LAMMPS focus on molecules and how they move. These tools track the motion of atoms over time, which is known as molecular dynamics. GROMACS is widely used in biology to study things like proteins and DNA. It is very fast, so it is also used to model plastics and other materials.[131] LAMMPS has a broader focus and is used in materials science, chemistry and engineering. It can simulate metals, polymers and large-scale systems with many particles.[132]

OpenFOAM, Code_Saturne, Nektar++ and the Met Office Unified Model simulate large-scale physical systems. These are used in engineering, weather forecasting and climate science. OpenFOAM models the flow of fluids, such as air and water, and can also be used to simulate heat and solid objects.[133] Code_Saturne is designed for more complex tasks, such as modelling combustion or tracking particles in a flow.[134] Nektar++ allows researchers to build new fluid models using advanced maths.[135] The Met Office Unified Model is used to predict the weather and simulate how the Earth's climate changes over time.[136]

---

[127] The categories "No srun", "Unknown", and "Python" have been excluded from this analysis. Jobs listed under "No srun" do not follow the standard job submission process and cannot be reliably linked to specific software usage. The "Unknown" category lacks sufficient metadata to attribute usage to any identifiable application. "Python", while widely used, serves as a general-purpose programming language and does not represent a single scientific application. Its inclusion would not allow for meaningful comparison with domain-specific software packages.

[128] ARCHER2-HPC. (n.d.). CP2K. In ARCHER2 User Documentation. Retrieved from https://docs.archer2.ac.uk/research-software/cp2k/

[129] ARCHER2-HPC. (n.d.). VASP. In ARCHER2 User Documentation. Retrieved from https://docs.archer2.ac.uk/research-software/vasp/

[130] ARCHER2-HPC. (n.d.). CASTEP. In ARCHER2 User Documentation. Retrieved from https://docs.archer2.ac.uk/research-software/castep/

[131] ARCHER2-HPC. (n.d.). GROMACS. In ARCHER2 User Documentation. Retrieved from https://docs.archer2.ac.uk/research-software/gromacs/

[132] ARCHER2-HPC. (n.d.). LAMMPS. In ARCHER2 User Documentation. Retrieved from https://docs.archer2.ac.uk/research-software/lammps/

[133] ARCHER2-HPC. (n.d.). OpenFOAM. In ARCHER2 User Documentation. Retrieved from https://docs.archer2.ac.uk/research-software/openfoam/

[134] ARCHER2-HPC. (n.d.). Code_Saturne. In ARCHER2 User Documentation. Retrieved from https://docs.archer2.ac.uk/research-software/code-saturne/

[135] ARCHER2-HPC. (n.d.). Nektar++. In ARCHER2 User Documentation. Retrieved from https://docs.archer2.ac.uk/research-software/nektarplusplus/

[136] ARCHER2-HPC. (n.d.). Met Office Unified Model. In ARCHER2 User Documentation. Retrieved from https://docs.archer2.ac.uk/research-software/mo-unified-model/

Table 43 ARCHER2 software utilisation metrics by job count, node hours, users and projects

Software name	Number of jobs	Node hours	% of overall ARCHER2 use	Number of users	Number of projects run on ARCHER2
VASP	71,386	897,194.4	22.8	140	15
GROMACS	14,191	217,271.4	5.5	46	8
Met Office UM	39,987	210,203.9	5.3	37	4
Nektar++	1,156	190,837.2	4.9	15	4
OpenFOAM	8,817	171,084.2	4.3	52	19
CP2K	40,921	161,750.7	4.1	47	13
Code_Saturne	192	103,688.4	2.6	4	2
LAMMPS	26,288	101,428.6	2.6	48	14
CASTEP	126,382	76,739.7	2.0	40	7
Total use of ARCHER2	**1,108,954**	**3,933,431.6**	**-**	**901**	**129**

Source: EPCC data

VASP is the most resource-intensive application on ARCHER2, accounting for 22.8% of total node hours. Compared to all software used on ARCHER2, it recorded the third-largest number of jobs (71,386) and was used by the highest number of users (140) across the second-highest number of projects (15).[135] This indicates a high level of computational demand and broad adoption within its user base. In contrast, CASTEP recorded the highest number of jobs (126,382) but used only 2% of total node hours, indicating that its simulations are typically smaller in scale or shorter in duration.

Nektar++ and Code_Saturne consumed a relatively high number of node hours (4.9% and 2.6%, respectively) despite having very few users (15 and 4) and projects (4 and 2). These figures imply that both are used for highly demanding simulations, likely in specialised fields such as fluid dynamics and combustion modelling; for example, Code_ Saturne was used to model hydrogen combustion in Umair Ahmed's pioneer project (see the case study in Box 9 in Section 4.2).

Mid-range tools like GROMACS, CP2K, OpenFOAM and the Met Office Unified Model each account for 4 to 5.5% of node hours. They support a moderate number of users and projects, reflecting steady demand across disciplines such as molecular biology, quantum chemistry and climate science.

CP2K is hosted on GitHub and has 953 stars, which indicate users who have marked it as useful or of interest, and 409 forks, showing that others have copied the code to work on their own versions. At least eight developers have contributed directly, and over 18,000 code updates (commits) have been made, reflecting active and ongoing maintenance. This pattern of activity suggests CP2K is a well-established tool with a focused but expert user base, in line with its application in quantum chemistry and materials modelling, where users often require specific technical capabilities. ARCHER2 supports this community by offering the scale and performance needed to run CP2K effectively on large and complex simulations. LAMMPS, with 2.6% of node hours and 26,288 jobs, shows balanced use, bridging the gap between high-frequency, low-intensity tools and those used for fewer, more intensive simulations.

A5.2 Using High Performance Computing (HPC) to simulate bacterial membranes for antibiotic research

Summary

Researchers at the University of Oxford have used ARCHER2 to create one of the most detailed computer models of a bacterial outer membrane. This work, supported by EPSRC funding, is helping to improve understanding of how antibiotics interact with the protective barrier found in dangerous bacteria, to support the discovery of more effective treatments.

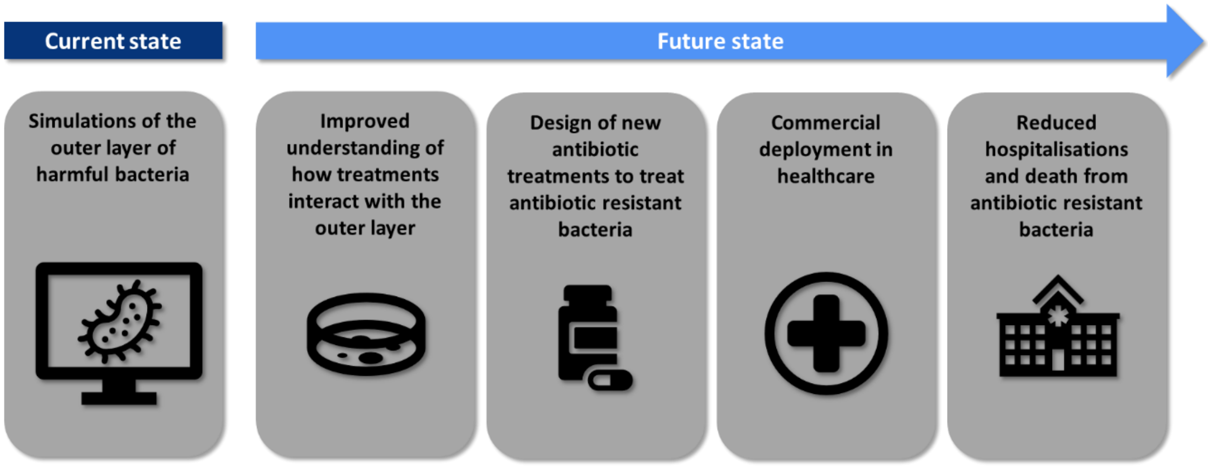

Research motivation

Some bacteria have an outer membrane that provides protection against many antibiotics and contains a compound that can trigger toxic reactions. Scientists lack important knowledge about how antibiotics can cross this membrane, limiting the development of new treatments.

The team at Oxford is using large-scale computer simulations to study the structure and behaviour of this membrane in detail. In the long term, this research will contribute to the development of new treatments for antibiotic-resistant bacteria, an increasingly urgent global health challenge. In recent years, there has been a rise in resistant strains, such as *E. coli*, due to importations from abroad, local spread and poor prescribing practices.[137]

EPSRC and ARCHER2 contribution

The research was made possible through an EPSRC Established Career Fellowship awarded to Professor Syma Khalid. This funding covered the full cost of the researchers responsible for carrying out the simulations.

The team used ARCHER2 to run large simulations of a bacterial outer membrane model. The size and power of HPC resources are required to run long-duration simulations (calculations that take a long time) that can observe the movement of proteins and lipids, to provide biological insights. ARCHER2 ensures results are achieved in a pragmatic timeframe.

[137] Webby, M. N., Oluwole, A. O., Pedebos, C., Inns, P. G., Olerinyova, A., Prakaash, D., Housden, N. G., Benn, G., Sun, D., Hoogenboom, B. W., Kukura, P., Mohammed, S., Robinson, C. V., Khalid, S., & Kleanthous, C. (2022). Lipids mediate supramolecular outer membrane protein assembly in bacteria. Science Advances, 8(44), eadc9566. https://doi.org/10.1126/sciadv.adc9566

Access to ARCHER2 ensures the continuity of funding and expertise necessary for the project by enabling the generation of preliminary data for new research directions, which supports the development of competitive grant proposals through the HEC consortia.

In the absence of ARCHER2, the researchers could have used resources in the US or Europe. However, this would have taken significantly more time given lack of familiarity with those schemes, and there is no guarantee of success. The consortium model combined with ARCHER2 provided the researchers with timely access to the only resource in the UK that is capable of handling their calculations. In the absence of ARCHER2, the simulations would have had to be scaled back significantly, which would have opposed the fundamental idea of the research: to mimic real biology. The unique characteristics of the team's research is that they can run large simulations (based on data from UK experimental scientists) and have specialist computational knowledge. Using UK resources has significant efficiency gains for the research team as there is no dependence on external institutional factors.

Key outcomes and impact

This research is outlined in the paper by Webby et al.'s (2022)[138] published in *Science Advances* that finds that some lipids, a type of fat, in the bacterial outer membrane, bind tightly with proteins to form stable structures. These structures strengthen the membrane and reduce its permeability, likely helping bacteria better resist external threats such as antibiotics. The EPSRC supported financially through the EPSRC Leadership Fellowship grant (EP/T03419X/1), the EPSRC HECBioSim award (EP/R029407/1) and the EPSRC grant awards (EP/N509577/1 and EP/K031953/1).

Building on this research, the team is using a high-throughput method to explore how potential antibiotics interact with the outer membrane of harmful bacteria. Using computer simulations and artificial intelligence in collaboration with partners in the United States, the researchers aim to identify promising molecules more efficiently as a starting point for new drug development. This approach seeks to make early-stage drug discovery faster and more focused.

"We now have a molecular model of the E. coli outer membrane that is realistic in composition and in size is large enough to visualise the impact of antibiotics – we need the HPC resource to run the simulations for long enough to utilise this model to maximum benefit." – Professor Syma Khalid

The team's research into harmful bacteria is increasingly relevant given the ongoing challenges posed by antibiotic-resistant infections, which contribute to significant health impacts and economic costs in the UK.

In 2019, 87,500 deaths involved bacterial infections in the UK. Among these, 7,600 deaths were directly caused by infections resistant to antibiotics, and 35,200 were an indirect result of infections resistant to antibiotics, i.e., contributed to death, but were not the primary cause.[139]

[138] Webby, M. N., Oluwole, A. O., Pedebos, C., Inns, P. G., Olerinyova, A., Prakaash, D., Housden, N. G., Benn, G., Sun, D., Hoogenboom, B. W., Kukura, P., Mohammed, S., Robinson, C. V., Khalid, S., & Kleanthous, C. (2022). Lipids mediate supramolecular outer membrane protein assembly in bacteria. Science Advances, 8(44), eadc9566. https://doi.org/10.1126/sciadv.adc9566

[139] Department of Health and Social Care. (2024). Confronting antimicrobial resistance: 2024 to 2029. https://www.gov.uk/government/publications/uk-5-year-action-plan-for-antimicrobial-resistance-2024-to-2029/confronting-antimicrobial-resistance-2024-to-2029

Poudel et al. (2023) conducted a systematic literature review and meta-analysis of the economic burden associated with antibiotic resistance, finding that the costs per patient episode for treating resistant infections can be up to £24,809[140] ($29,289).[141]

In England, 66,730 serious antibiotic-resistant infection cases were recorded in 2023.[142] Applying the per-episode cost to these cases suggests a total annual cost of up to £1.7 billion on the health service, particularly the NHS.[143] The significant financial burden posed by resistant infections highlights the importance of the team's research into harmful bacteria's structure and potential treatments, which will reduce the incidence of serious resistant infections.

On average, patients with resistant infections remain in hospital 7.4 days longer[140] than those with non-resistant infections. This extended hospitalisation imposes higher direct costs, exacerbates pressure on the healthcare system's capacity. In addition, antibiotic resistance is associated with increased clinical risk, with the mortality rate nearly double for patients with resistant infections compared to those that are responsive to antibiotic treatment.[140]

Looking ahead

While the research is at an early stage, it is helping improve understanding of how antibiotics interact with resistant bacteria. In the long term, the approach could contribute to the development of new treatments at a time when antibiotic resistance is becoming a growing threat to public health and, in turn, economic growth.

A5.3 Investigating unexplained behaviours in simulations of molecules exposed to intense radiation pulses

Sciences like chemistry and physics have made great progress in explaining why certain atoms and molecules bond, or how some materials generate electricity when exposed to sunlight. Much of this comes down to the behaviour of electrons—the tiny, fast-moving particles that orbit atoms. We know, for example, that solar cells work because sunlight gives electrons enough energy to 'jump', making them easier to capture and convert into electricity. In chemical reactions, we understand what goes in and what comes out – but not exactly what happens in between. That's because the steps in a reaction happen incredibly fast, in just attoseconds (there are as many attoseconds in one second as there are seconds in the age of the universe).[144] This ultra-fast scale is the focus of attoscience, a field that could help us observe and eventually control these fleeting events, paving

[140] USD was converted to GBP using the Bank of England's average 2023 spot exchange rate (1 USD = 0.8042 GBP). The resulting 2023 GBP figure was uprated to 2025 prices using GDP deflators from HM Treasury (2024) (2023 = 0.9495, 2025 = 1.00). Sources: Bank of England. (n.d.). Data viewer. Retrieved May 2025, from https://wwwtest.bankofengland.co.uk/boeapps/database/fromshowcolumns.asp?Travel=NIxAZxSUx&FromSeries=1&ToSeries=50&DAT=RNG&FD=1&FM=Jan&FY=2015&TD=19&TM=May&TY=2025&FNY=&CSVF=TT&html.x=176&html.y=34&C=DMY&Filter=N; HM Treasury. (2024). GDP deflators at market prices, and money GDP: March 2024 (Quarterly National Accounts). Retrieved from https://www.gov.uk/government/statistics/gdp-deflators-at-market-prices-and-money-gdp-march-2024-quarterly-national-accounts

[141] Poudel, A. N., Zhu, S., Cooper, N., Little, P., Tarrant, C., et al. (2023). The economic burden of antibiotic resistance: A systematic review and meta-analysis. PLOS ONE, 18(5), e0285170. https://doi.org/10.1371/journal.pone.0285170

[142] UK Health Security Agency. (2024, November 14). Antibiotic resistant infections continue to rise. GOV.UK. Retrieved June 6, 2025, from https://www.gov.uk/government/news/antibiotic-resistant-infections-continue-to-rise

[143] The estimated annual cost to the health service of £1.7 billion is derived by applying the upper-bound treatment cost per patient episode (£24,809) from Poudel et al. (2023) to the number of serious antibiotic-resistant infection cases in England in 2023 (66,730). This estimate reflects a higher bound estimate for direct treatment costs but excludes broader economic impacts such as patient's productivity losses.

[144] See https://theconversation.com/what-is-an-attosecond-a-physical-chemist-explains-the-tiny-time-scale-behind-nobel-prize-winning-research-214907

the way for breakthroughs in materials and chemistry that could greatly enhance future computing and communications technologies.[145]

To drive innovation in this field, scientific theory and experimental discovery or validation must work hand in hand. Advances in scientific theory make it possible to investigate and predict the way atoms and molecules behave as they transition between states, while new techniques and technologies are enabling the observation and manipulation of these states. ARCHER2 has been used to shed light on the predictive power of some key theories in attoscience.

Dr Daniel Dundas and his team at Queen's University Belfast work on the scientific theory used to describe how electrons move in molecules. They built code, called EDAMAME, for simulating electrons in molecules under different conditions. Meanwhile, researchers at the University of Toulouse created a similar tool, known as QDD, using different mathematical techniques to solve the underlying equations. When the Toulouse team used QDD to simulate how certain molecules respond to intense laser pulses, they observed an unexpected effect: a peculiar oscillation in the molecules' electric field caused by shifts in the surrounding electron cloud that began after the radiation pulse had ended. Intrigued by these results, Dr Dundas' team set out to investigate this further.

Solving the complex equations that underpin the EDAMAME code requires significant computational power. In 2019, the team was awarded access to ARCHER2 through a Pioneer Project, which enabled them to simulate molecular systems exposed to intense laser pulses. ARCHER2's high-performance capabilities were essential for modelling these systems over extended timeframes of over 600 femtoseconds – long enough to capture not only the immediate effects of the laser pulse (around 30 femtoseconds, above the limit of capability for non-specialised computers), but also the subsequent electric field oscillations.

The aim of the project was to determine if the oscillations observed in the Toulouse group's QDD code would also be observed in EDAMAME. If so, the project further sought to understand the oscillations and the kinds of molecular systems they occur in.

Over two years and four separate allocation periods, Dundas and his team ran simulations of different systems and translated some of these into animations. With the help of ARCHER2, they found that EDAMAME predicted the same peculiarity as QDD over a range of molecules and pulse lengths.

These findings have important implications for attoscience. Firstly, they suggest that these oscillations, known as dipole instability, are not simply a quirk in QDD's code since they also appear in EDAMAME, which uses different numerical techniques to solve the equations underlying the theory.[146] Secondly, they agree with recent findings by Dundas' colleagues based on code using an entirely different theory.[147] The latter indicates that dipole instability may be more than a theoretical artefact – it could be a physical phenomenon, with implications for materials science and chemistry, particularly as these relate to high-speed computing and communications.

The project's results have already attracted interest from researchers in France. The former are working with the Toulouse group to apply a different scientific theory, to the simulation of the pulse/molecule system. If they succeed in finding the dipole instability using this alternative theory,

[145] See https://www.imperial.ac.uk/a-z-research/quantum-optics-and-laser-science/research/laser-consortium/current-research/attoscience-for-non-specialists/
[146] P. G. Reinhard et al, Phys Rev A, 107 L020801 (2023). DOI:https://doi.org/10.1103/PhysRevA.107.L020801
[147] P. G. Reinhard et al, Euro Phys J: Spec Topics, 232 2095 (2022). DOI: https://doi.org/10.1140/epjs/s11734-022-00676-6

it would provide strong justification for investigating the phenomenon experimentally. This, in turn, could lead to novel science and innovations in computation and communication, which rely on ultra-high-speed interactions between matter and light. For example, this could enable the design of light-driven transistors (current computing technology uses electronic transistors). It could also lead to innovations that improve the stability of quantum computers and help usher in the era of quantum advantage. In this regard, the ARCHER2 simulations provide experimental scientists with insights into the conditions and mechanisms that lead to dipole instability.

However, if the dipole instability proves to be a remnant of the theoretical approach used by Dr Dundas' team and the Tolouse group, the findings remain significant. The theory is widely used across fields such as condensed matter physics and materials science, and understanding this theoretical artefact will help improve the accuracy of models based on it.

"Understanding the source of these instabilities is crucially important. They could represent a new physical mechanism or else may be due to the level of theory being used to model these systems" — Dr Daniel Dundas, Queen's University of Belfast School of Mathematics and Physics

Figure 61 ARCHER2 animation depicting the moment that an Acetylene molecule is subjected to a 50 terawatt, 30 femtosecond(fs) ultraviolet laser pulse.

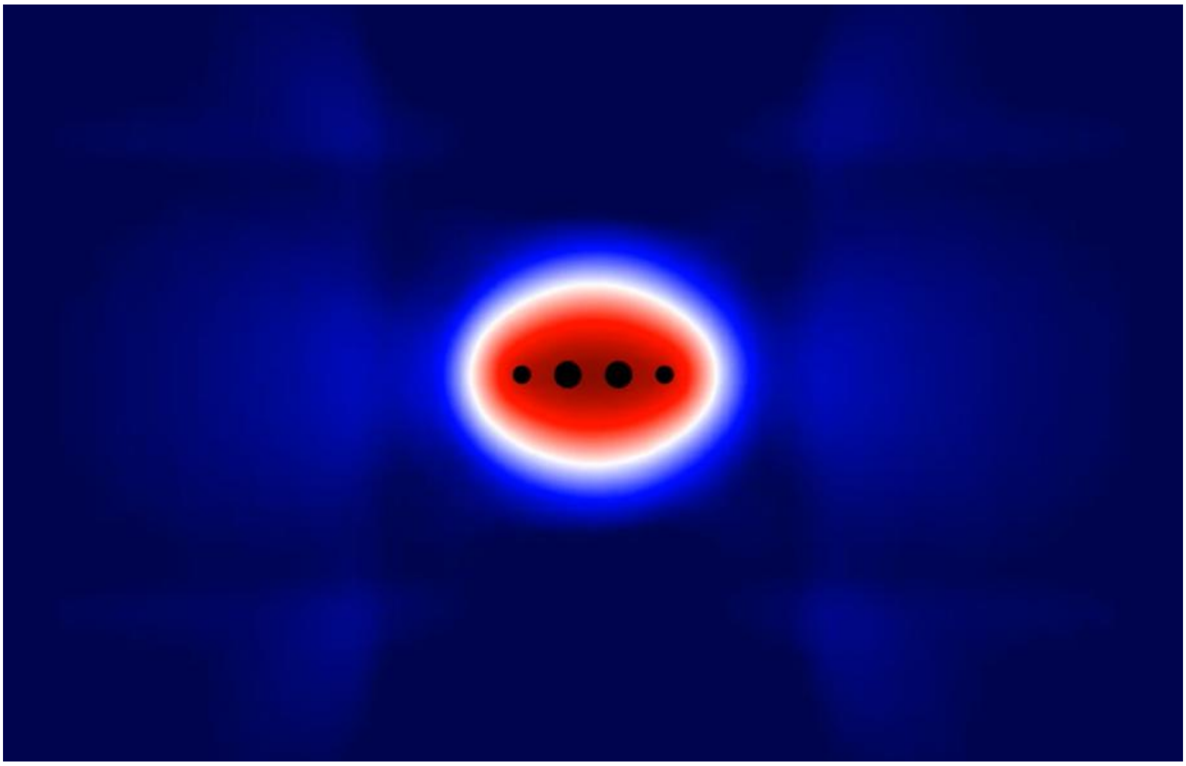

Source: https://www.archer2.ac.uk/about/gallery/2022-image-comp/

Figure 62 ARCHER2 animation depicting the moment that the dipole instability occurs, 100 fs after an Acetylene molecule is subjected to a 50 terawatt, 30 femtosecond(fs) ultraviolet laser pulse.

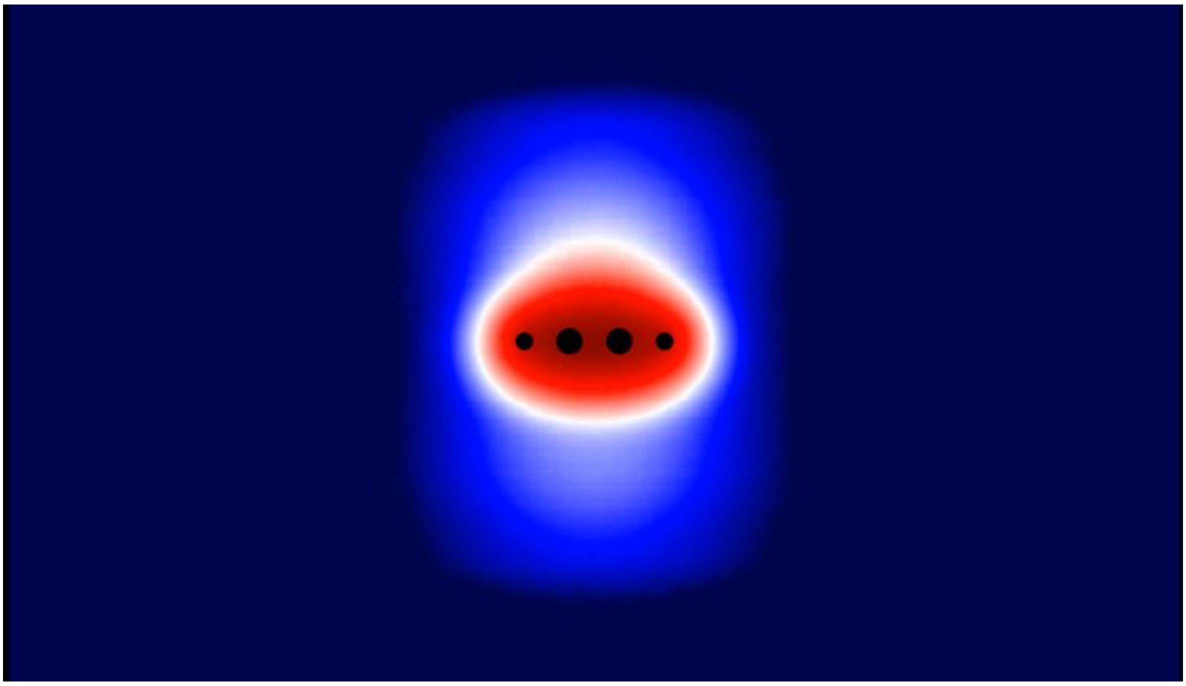

Source: https://www.archer2.ac.uk/about/gallery/2022-image-comp/

A5.4 Modelling extreme wave conditions for floating wind turbines

As the UK moves toward net zero, floating offshore wind is set to play a bigger role in delivering clean, reliable electricity[148]. Offshore wind farms use large turbines placed at sea, where the winds are stronger and steadier than on land. In shallow waters, turbines can be fixed to the seabed, but the best wind areas are often much deeper. Floating turbines, which sit on platforms held in place by cables, make it possible to reach these deeper waters and take advantage of the stronger winds found there. Designing these floating turbines in a way that can handle strong waves, powerful winds, and fierce storms is a major engineering challenge that demands accurate and efficient modelling.

Traditionally, engineers test floating wind turbines by simulating the ocean over a long time, with waves of all different sizes and directions, just like real seas. These mixed, changing conditions are called "random sea states." This helps them see how much force the turbine will have to handle. But these tests take a long time to run and use a lot of computer power, which slows down the early stages of designing new turbines.

The team at the University of Plymouth has tested a new approach to make this process faster and more efficient by using shorter, more focused wave simulations known as short design waves (SDWs). To support this, they used ARCHER 2 to run hundreds of simulations across a wide range of wave and wind conditions combining engineering-level numerical models with Computational Fluid Dynamics CFD simulations, which are advanced computer models that simulate how air and water

[148] See https://www.gov.uk/government/publications/clean-power-2030-action-plan/clean-power-2030-action-plan-a-new-era-of-clean-electricity-main-report?utm_source=chatgpt.com

move around the turbine. Thanks to ARCHER2's processing power, the team investigated a much larger set of conditions, known as a larger parameter space, in less time.

The team tested four different types of short, focused wave simulations. Some of these simulations zoomed in on a single, carefully chosen wave — like taking a snapshot of the biggest wave you might face. Others created a mini-sequence by adding a few waves before and after the peak, better mimicking real sea behaviour where the lead-up to the biggest wave can be just as important as the wave itself. These two styles — single waves and surrounding waves — capture different aspects of how a floating turbine moves and reacts. The team tested how turbines shift forward and backward, bob up and down, tilt, and how much strain is placed on the cables anchoring them to the seabed. They found that including the background wave history gave a clearer and more reliable picture, especially for predicting how the cables behave in rough seas.

This approach is different from traditional methods, where engineers usually run long, random simulations of the sea and wait for big waves to happen by chance. This is why the process can take weeks or even months. By focusing directly on the most critical wave events, the new method cuts down the time dramatically. In fact, simulations were up to 20 times faster. For developers, this means they can find the most dangerous conditions much earlier in the design process, saving both time and money.

The research team has been working closely with DNV[149], a leading organisation that sets international standards for offshore wind. The advanced simulations and faster modelling approaches developed in this project are now helping to build a stronger evidence base that DNV will use to update best practice guidelines for floating wind turbines. These improvements will make certification processes faster and more cost-effective for developers, reducing the number of expensive simulations they need to perform and helping new designs reach the market more quickly.

"By speeding up simulations, we're making it easier for developers to iterate quickly and bring safer designs to market." — Scott Brown from the University of Plymouth

A5.5 Personalising heart care with digital twins

Summary

Access to ARCHER2 enabled the Cardiac Electromechanics Research Group (CEMRG) at Imperial College London to advance personalised simulation models of the human heart. By using high-performance computing (HPC) to build digital twins, the team produced new clinical insights, supported regulatory guidance and contributed to securing major programme funding. Their work demonstrates the important role of computational modelling in improving cardiovascular care and reducing unnecessary procedures.

Research motivation

Cardiovascular disease (CVD) affects approximately seven million people in the UK and accounts for one in four premature deaths. The annual cost of cardiovascular disease in the UK was £29 billion in 2021/22, including £16.6 billion in direct costs to health services and £12.4 billion in indirect costs such as lost productivity.Reducing the costs of CVD could be significant for NHS and other worldwide

[149] https://www.dnv.com/focus-areas/floating-offshore-wind/

health providers. It is also a leading contributor to health inequalities, with individuals in the most deprived 10% of the population being twice as likely to die from CVD compared to those in the least deprived 10%.[150] A significant proportion of CVD cases are linked to high-risk conditions such as atrial fibrillation (AF), an irregular heartbeat that significantly increases the risk of stroke.[151]

The CEMRG at Imperial College London seeks to improve diagnosis and treatment of CVD by developing personalised computer models of the heart. These simulations, known as digital twins, integrate patient-specific imaging and pressure data to recreate the electrical and mechanical behaviour of individual hearts. The models enable researchers to explore how patients may respond to therapies in a risk-free, simulated environment.

Traditional clinical tools often fail to capture the complex, patient-specific variations that drive differential treatment responses. By contrast, digital twin approaches offer the potential to support more precise decision-making, reduce unnecessary interventions and improve outcomes.

EPSRC and ARCHER2 contribution

Access to ARCHER2, provided by EPSRC, was described by the research team as "vital for the success" of many CEMRG projects. For example, a sensitivity analysis by Strocchi et al. (2023) required approximately 500 individual simulations, each running for five hours on 512 cores, equivalent to over 20 desktop computers operating in parallel.

Moreover, the group’s ARCHER2-based work contributed to securing an EPSRC programme grant, CVD-Net, worth £8.8 million[152]. Led by Professor Steven Niederer, this initiative aims to develop digital twins for patients with pulmonary arterial hypertension in collaboration with partners such as Medtronic, the University of Heidelberg, GlaxoSmithKline, NHS England and the US Food and Drug Administration.

In the absence of ARCHER2, the research team would have had to consider different high-performance computing facilities, including Amazon Web Services or local high-performance computing facilities available within Imperial College London. However, these can be expensive, harder to access or limited in terms of the number of computing nodes available to them. This would have probably forced the research team to reduce the number of simulations they could perform and therefore leading to less ambitious research. For example, they would have needed to perform a less complex sensitivity analysis requiring fewer simulations per patient, or they would have reduced the number of patient models they generated, limiting the clinical relevance of their research.

[150] NHS England. (n.d.). Cardiovascular disease (CVD). Retrieved July 3, 2025, from https://www.england.nhs.uk/ourwork/clinical-policy/cvd/

[151] British Heart Foundation. (2025, January). UK Factsheet. https://www.bhf.org.uk/-/media/files/for-professionals/research/heart-statistics/bhf-cvd-statistics-uk-factsheet.pdf

[152] UKRI. (2023). Networks of Cardiovascular Digital Twins (CVD-Net). UK Research and Innovation. https://gtr.ukri.org/projects?ref=EP%2FZ531297%2F1#/tabOverview

Key outcomes and impact

Research conducted using ARCHER2 has led to the publication of two papers and one which is under review[153] focused on the application of cardiac simulation. These include studies on atrial fibrillation treatment strategies, global sensitivity analysis and advancements in whole-heart modelling.

CEMRG brings together expertise in cardiac modelling, simulation and data integration. The group is led by Professor Steven Niederer and includes Dr Marina Strocchi, who leads simulation work, and Dr Cristobal Rodero, a specialist in shape and whole-heart modelling. PhD researchers Rosie Barrows, Tiffany Baptiste and Shambhavi Malik contribute to specific areas such as left atrial mechanics and cardiac electrophysiology. All researchers rely on ARCHER2 for running high-resolution, patient-specific simulations.

CEMRG's work benefited from academic collaborations. Prof Gernot Plank and Dr Christoph Augustin (Medical University of Graz) developed the simulation software used on ARCHER2, bringing advanced expertise in cardiac mechanics. Statistical guidance was provided by Prof Chris Oates (University of Newcastle) and Prof Richard Wilkinson (University of Nottingham), supporting methodological development in model calibration and sensitivity analysis.

CEMRG's research is cited in the European Society of Cardiology guidelines (2021)[154], based on modelling work led by Dr Angela Lee and Professor Niederer. This study investigated sex-specific responses to CRT by simulating how differences in heart structure may affect outcomes. Such insights are significant given that 30-50% of CRT patients do not experience the intended improvement. Better identification of responders could reduce unnecessary procedures and associated healthcare costs.

The group's modelling work has also been referenced in regulatory science outputs by the US Food and Drug Administration (FDA), including guidance on model credibility and inclusion in an FDA regulatory science tool for assessing patient-specific modelling in medical device software.

CEMRG has demonstrated a strong commitment to public engagement. Team members have participated in a range of outreach activities aimed at increasing awareness of heart disease and the role of modelling. This includes:

- Community events on the impact of diabetes on heart health (e.g. with the Harrow Diabetes community and Latin American community)
- Youth and education events such as the Great Exhibition Road Festival, Great Exhibition Big Ideas, Imperial Lates, STEM for Britain, AI UK, and FameLab
- Presentations at New Scientist Live, reaching broader public audiences

[153] Strocchi, M., Longobardi, S., Augustin, C. M., Gsell, M. A. F., Petras, A., Rinaldi, C. A., Vigmond, E. J., Plank, G., Oates, C. J., Wilkinson, R. D., & Niederer, S. A. (2023). Cell to whole organ global sensitivity analysis on a four-chamber heart electromechanics model using Gaussian processes emulators. PLoS Computational Biology, 19(6), e1011257. https://doi.org/10.1371/journal.pcbi.1011257; Barrows, R. K., Augustin, C. M., Gsell, M. A. F., Roney, C. H., Solís Lemus, J. A., Xu, H., Young, A. A., Rajani, R., Whitaker, J., Vigmond, E. J., Bishop, M. J., Plank, G., Strocchi, M., & Niederer, S. A. (2025). Rhythm control benefits left ventricular function compared with rate control in patients with atrial fibrillation: A computational study. Heart Rhythm O2. https://doi.org/10.1016/j.hroo.2025.04.014; Strocchi, M., et al. (Under revision). Integrating imaging and invasive pressure data into a multi-scale whole-heart model. Journal of Biomechanical Engineering.

[154] European Society of Cardiology. (2021). ESC Guidelines. European Heart Journal, 00(1-94). https://doi.org/10.1093/eurheartj/ehab364

Their work has also featured in media outlets and university communications, including Barbados Today[155], The Independent[156], MIT Technology Review[157] and the University of South Wales[158].

Looking Ahead

CEMRG's work demonstrates how high-performance computing can support the development of virtual testing tools that improve diagnosis, reduce risk and inform clinical decision-making. As digital twin technologies mature and gain regulatory confidence, they could support more efficient healthcare delivery and improved patient outcomes across the NHS.

A5.6 NEMO (Nucleus for European Modelling of the Ocean)

Ocean modelling plays a vital role in predicting climate patterns, supporting marine conservation, improving disaster preparedness, and informing sustainable development. ARCHER2 provides the high-performance computing power needed to run complex ocean and climate models offering insights into the ocean's influence on global weather, ecosystems, and human activities.

One of the key modelling tools powered by ARCHER2 is the Nucleus for European Modelling of the Ocean (NEMO), a state-of-the-art framework used for both research and operational ocean forecasting.[159] In the UK, Professor Andrew Coward leads the ARCHER2 N01 "Ocean and Shelf Seas" consortium, which uses NEMO to run large-scale simulations to explore a wide range of issues, including how CO_2 can be stored in the deep ocean, what drives sea level changes, and how changing conditions affect fish stocks and food security.[160]

The group also produces high-resolution coastal models that are particularly relevant to UK industry, for example in helping to assess the stability and safety of offshore wind farms, especially as they face increasingly intense ocean conditions. The team's work is used to evaluate potential risks and improve the design and siting of offshore infrastructure. Because these coastal models depend on input from global-scale simulations, the ability to run these models on ARCHER2 is essential.

ARCHER 2 also supports the Joint Marine Modelling Programme (JMMP), a collaboration between the Met Office, British Antarctic Survey, National Oceanography Centre (NOC), and Centre for Polar Observation and Modelling. The JMMP applies the NEMO modelling framework to regional ocean modelling, developing tools that are crucial for accurate forecasting and environmental monitoring.[161] NOC's contributions to JMMP are coordinated by Professor Jeffrey Polton, who helps lead the co-development of the model configurations that underpin these regional applications.

With support from ARCHER2, the JMMP team developed a new and more detailed model of the UK's coastal seas. This development of NEMO on ARCHER2 has directly led to upgrades in the UK Met Office's operational numerical weather forecast system. The model shows smaller and more

[155] Brathwaite, S. (2024, April 9). Weather, healthcare 'could be AI test bed' in the Caribbean. Barbados Today. https://barbadostoday.bb/2024/04/09/weather-healthcare-could-be-ai-test-bed-in-the-caribbean/
[156] Massey, N. (2024, June 8). Scientists and artists create live art based on groundbreaking research. The Independent. https://www.independent.co.uk/news/science/scientists-imperial-college-london-london-south-kensington-b2558999.html
[157] Hamzelou, J. (2024, December 19). Digital twins of human organs are here. They're already changing medical treatment. MIT Technology Review. https://www.technologyreview.com/2024/12/19/1108447/digital-twins-human-organs-medical-treatment-drug-trials
[158] Niederer, S. (2024, October 16). "Wow, this is the Future!": Prof Steven Niederer visit stimulates new ideas. UNSW News. https://www.unsw.edu.au/news/2024/10/wow-this-is-the-future-prof-steven-niederer-visit-stimulates-new-ideas
[159] See https://www.nemo-ocean.eu/
[160] See https://docs.archer2.ac.uk/research-software/nemo/
[161] See https://www.metoffice.gov.uk/research/approach/collaboration/joint-marine-modelling-programme

dynamic ocean patterns such as swirling currents, tides, and temperature changes that were not visible in the older version. It provides a much more accurate and realistic picture of coastal ocean dynamics, replacing the old version for operational forecasting. It enables the Met Office's weather prediction system to incorporate dynamic ocean conditions for the first time, significantly improving forecast accuracy.

ARCHER2 has been key to running the complex simulations required for this work and for enabling close collaboration between academic researchers and operational partners at the Met Office. Without ARCHER2, it would not have been possible to carry out the testing and refinement of these high-resolution models, nor would the collaborative development of improved ocean and weather prediction systems have been feasible.

These upgrades have a direct impact on weather forecasting, especially in coastal areas, which are vulnerable to extreme events. The improvements contribute to better preparedness and response, supporting efforts to protect lives, infrastructure, and ecosystems as we face increasing risks from climate change.

A5.7 VAMPIRE (Versatile Atomistic Monte Carlo, Integrator, and RElaxation)

VAMPIRE is an open-source software which simulates atoms of magnetic materials. Atoms are found everywhere and are the building blocks of all matter. VAMPIRE helps model and predict the behaviour of atoms of magnetic materials, i.e., materials that are attracted to magnets or those that can be magnetised. Such simulations of atomistic spin dynamics reveal properties about materials such as interfaces, defects and roughness. These simulations have a wide range of applications: they can be used to design new materials, understand chemical reactions, study biological processes and explore properties at the atomic level.

Starting in 2009, the University of York developed VAMPIRE.[162] The initial phase of the project involved a trial using Archer1 and later transitioned to development on ARCHER2 (Archer is UK's national supercomputing service)[163]. Over time, the code has been significantly enhanced, enabling simulations up to a billion atoms. This capability provides unprecedent insights into the behaviour of complex materials.

ARCHER2: role, performance and policy

ARCHER2 has been instrumental in the development of VAMPIRE. According to the developer of Vampire, enabling atomic simulations of materials requires strong computational power and "would not have been possible without ARCHER2"[164]. The software support of ARCHER2 has also been described as "exceptional", playing a critical role in advancing research.

Archer2 is reliable, accessible, and sustainable. Standard simulations, including VAMPIRE, were run on Archer2 and two other high-performance computing (HPCs) systems, Viking2 and Oracle's E5 shape (a cloud computing service), to compare performance, energy use and cost. Archer2's performance is comparable to that of E5 and better than Viking2. It is also the most scalable option, designed for large simulations and workloads, with the most consistent performance.

[162] See https://vampire.york.ac.uk
[163] See http://www.archer.ac.uk/casestudies/Archer_magnet_casestudy.pdf
[164] Quote by Richard Evans during the interview.

Archer2 runs on sustainable energy and is very energy efficient. Despite ARCHER2 being five years older than Viking2, its alternative, the energy efficiency is similar in both HPUs. While sustainability has been built into the design of Viking2, energy efficiency in ARCHER2 has been improved over time through benchmarking and research.

For academics, ARCHER2 is the most affordable option compared to its alternatives, based solely on cost of computing and ignoring other costs such as storage. UKRI, owning 97.5% of ARCHER2, offers funding to use ARCHER2 which makes access to the supercomputer very affordable for successful academic applicants.[165] Outside of the UKRI funding, University of Edinburgh also allows commercial access to Archer2 at a comparatively higher price than that offered by UKRI. However, commercial access to Archer2 is still much cheaper: about 4 times less than access to OCI. Overall, the public HPC system, ARCHER2, was found to be reliable, cost-effective and accessible compared to other alternatives.

Funding policies linked to ARCHER2 access have also played a key role in the development of VAMPIRE. The software was funded by three Embedded Computational Science Engineering (eCSE) projects, a support that comes along with ARCHER to help its community and build up capabilities.[166] Such complementary policies, supported by EPSRC, are much-needed public investments in high-performance computing (HPC) to not only broaden the work being carried out by the Archer community but also achieve full potential of supercomputers.

Impact on science and industry

VAMPIRE serves a niche but growing set of scientific and industrial application. It currently has about 500-600 users across major UK research hubs, including Edinburgh, Sheffield, Manchester, Leeds, and Lancaster. Given that it allows better understanding of magnetic materials, it has lots of potential for innovation in components where magnets are key such as medical imaging devices, power generators, and data storage devices.

One of the most prominent applications of the code has been the development of HAMR (Heat-Assisted Magnetic Recording) by a consortium led by Seagate Technology. HAMR is a new technology that has increased data storage capacity of hard drives from 1TB to 30 TB. This was achieved by increasing the density of the disk platter, allowing more bits to be packed into each square inch. VAMPIRE played an important role in simulating how different materials behaved under intense heat and magnetic fields, and uncover innovative material solutions. Additionally, by simulating behaviour computationally, VAMPIRE reduced the need for costly materials and manufacturing processes.

Such advanced data storage technologies are vital in the Data Age. The volume of data created, captured, copied and consumed has grown exponentially over the years, along with increasing use of cloud computing. The rapidly growing need to store data means that cheaper, lighter and more energy efficient materials are needed. VAMPIRE and HAMR have contributed to the research of other storage technologies such as advanced media design (Advanced Storage Research Consortium) and Magnetic Random Access Memory (MRAM) (Samsung Semiconductor, IMEC), and the development of new permanent magnetic materials (Toyota).[167]

[165] See: https://www.ukri.org/wp-content/uploads/2022/07/EPSRC-050722-ImpactEPSRCInvestmentsHighPerformanceComputingInfrastructure.pdf
[166] See https://gtr.ukri.org/projects?ref=EP%2FP022006%2F1
[167] See https://results2021.ref.ac.uk/impact/66736a9d-e5c5-4dc8-ad82-03976ee2c1cf/pdf

VAMPIRE has also enabled significant scientific advancements, including the understanding and visualization of Van der Waals magnets and the thermal behaviour of skyrmions—both cutting-edge areas in materials science. Van der Waals magnets, which can be exfoliated into atomically thin layers, offer promising potential for advanced magnetic storage and electronic applications. Similarly, skyrmions—tiny, stable quasiparticles—are being explored for their use in data storage, logic devices, neuromorphic computing, and quantum computing. VAMPIRE has been adopted for device and materials design by several companies including Seagate Technology, Western Digital, Samsung Semiconductor, Intel, and IMEC.[167]

Along with IT sector, VAMPIRE is also unlocking benefits in health. Magnetic hyperthermia is a promising treatment for brain and prostate cancer and uses the code for large simulations. Previous simulations have been unable to predict nanoparticles due to their small size. VAMPIRE will help understand if and how nanoparticles in magnetic hyperthermia provide efficient and reliable heating to develop a new cancer therapy. The same technique can also potentially be used for biosensing or targeted drug delivery. Such developments have significant health and quality of life benefits for people in the UK and beyond. The pharmaceutical sector can also capitalise on the developments leading to job creation and economic benefits in the future.

The VAMPIRE software package is expected to deliver broader benefits by enhancing UK's appeal for international and EU research funding. The UK is a leader in developing scientific codes- such as those for plasma codes and 2D materials- and access to HPC is a critical enabler. ARCHER2 currently satisfies the essential needs of British research and, without it, the UK risks falling behind other countries. Anecdotal evidence suggests that the UK is losing talent to better funding opportunities overseas and eroding its competitive advantage. The development of codes such as VAMPIRE, supported by infrastructure such as ARCHER2, is important to build and maintain UK's lead in science and innovation.

Conclusion

Lastly, while ARCHER2 has been key for the development of VAMPIRE, software development is also beneficial for efficient use of supercomputers. Software such as VAMPIRE enable efficient and impactful use of the computational powers of supercomputers. Software is a critical part of the HPC ecosystem: it makes existing hardware usable and also enhances its efficiency and scientific output. For example, drawing on estimated time savings from software development using Archer1, and assuming similar benefits reported by academic users to the EPSRC, software development is estimated to have delivered approximately £53.4 million in value through time savings.

Looking ahead, more industry adoption of VAMPIRE is envisioned, particularly in sectors like pharmaceuticals where high-resolution magnetic modelling has potential. As VAMPIRE continues to demonstrate its scalability and performance across HPC platforms, including Archer2, increased industry engagement could make HPC more competitive and sustainable.

A5.8 Supercharging Futures: How ARCHER2 is Empowering Early Career Researchers at Imperial College London

Summary

Access to ARCHER2 has had a transformative effect on the research activities of Sarah Rouse's group in the Department of Life Sciences at Imperial College London. The team explores a range of topics in molecular biology - including mitochondrial dynamics and viral or pathogen invasion - using advanced simulations of biologically realistic systems made possible by the high-performance computing capabilities of ARCHER2.

Their work not only supports close collaboration the pharmaceutical industry in developing more efficient gene therapies and drug-targeting techniques but also plays a crucial role in training a new generation of researchers in computational biology, many of whom go on to careers in academia or the pharmaceutical industry.

The team's achievements highlight the growing significance of HPC and biologically informed simulations in accelerating progress in modern therapeutics.

Research background

Conducting research with the use of ARCHER2 allows the team to explore complex molecular systems critical to the development of pharmaceutical treatments. These computational simulations support deeper understanding of how biological systems behave, paving the way for more precise and effective therapies.

One standout example is a collaborative PhD project with AstraZeneca, which investigates adeno-associated virus (AAV) packaging mechanisms in gene therapies which are associated with the treatment of genetic diseases such as cystic fibrosis, adenosine deaminase deficiency or cancer[168]. AAVs are widely used as vectors due to their ability to deliver genetic material with minimal immune response[169]. However, incomplete packaging remains a major manufacturing challenge – often requiring higher doses and increasing the risk of triggering immune responses. By modelling the structural and dynamic factors that govern packaging efficiency, the project aims to support innovation in gene therapy production.

In another project, the team partners with Vertex Pharmaceuticals to study G protein-coupled receptors (GPCRs) in neuronal membranes. GPCRs are a key class of drug targets and understanding how membrane composition influences their behaviour could unlock new strategies for neurological and psychiatric treatments.[170] Together, these collaborations demonstrate how high-performance computing is accelerating medical innovation from molecular insight to therapeutic impact.

EPSRC and ARCHER 2 Contributions

Rouse highlights that such large-scale, industry-relevant simulations of biological systems would simply not be possible without the use of ARCHER2. Models such as those used in AAV packaging or GPCR drug testing projects cannot be run on institutional systems due to their scale and complexity.

Moreover, the computing power of ARCHER2 has made it possible to bridge cutting-edge experimental methods, such as tomography, with computational modelling. In other words, access to the supercomputer has enabled the creation of complex, large-scale simulations informed by experimental data and techniques from structural biology. Without access to ARCHER2, many of these projects would not be feasible from a computational standpoint. Others would require over-

[168] Harris, J., Stanton, C., Griffiths, A. J. F., & Rogers, K. (2025, July 18). *Gene therapy*. Encyclopaedia Britannica. https://www.britannica.com/science/gene-therapy

[169] See https://biologyinsights.com/aav-packaging-processes-roles-and-techniques/

[170] Baccouch, R., Rascol, E., Stoklosa, K., & Alves, I. D. (2022). The role of the lipid environment in the activity of G protein coupled receptors. *Biophysical chemistry*, *285*, 106794. https://doi.org/10.1016/j.bpc.2022.106794

simplification, limiting the impact of the simulations. Overall, without ARCHER2, the critical computational insights gained across research projects would be lost or severely reduced.

Key outcomes and impact

Rouse's team developed large-scale molecular simulations - particularly models of large membrane systems - that replicate structures derived from experimental data under realistic conditions. For example, in the AAV project, molecular simulations were used to identify the molecular determinants that influence how the virus enters host cells. Industry collaborations, such as those with AstraZeneca and Vertex Pharmaceuticals, support the translation of molecular research into therapeutic applications. Employing ARCHER2 for simulations improves both the effectiveness of therapeutic methods and the efficiency of manufacturing processes. As such, ARCHER2 has been instrumental not only in advancing scientific discovery but also in driving industrial innovation.

Beyond research outcomes, the use of ARCHER2 had a significant positive impact on training and skills development, particularly among early-career researchers. With ARCHER2 embedded in both undergraduate and postgraduate projects, students gain hands-on experience in building computational pipelines - often with no prior HPC experience - directly relevant to careers in both academia and industry.

Skills developed include:

- Running large-scale molecular simulations
- Managing and analysing large datasets
- Building cross-disciplinary understanding across biology, chemistry, and computation
- Presenting and communicating computational science effectively

This real-world experience in applied computational biology has prepared students for future roles in UK biotech, the pharmaceutical industry, and academia. Among the 2020 cohort of ARCHER2-trained students, career destinations included:

- 3 DPhils at the University of Oxford
- 3 PhDs at the University of Cambridge
- 3 PhDs at Imperial College London
- 2 PhDs at the Francis Crick Institute
- 1 Role at a UK biotech startup
- 1 PhD at the University of Southampton
- 1 PhD at the University of Melbourne

In general, students trained on ARCHER2 are pursuing careers at the intersection of science and biomedical industries, contributing to health, industry, and economic growth. ARCHER2 access is not only crucial for advancing research but also plays a pivotal role in cultivating talent across the UK's scientific landscape.

A5.9 Supporting zero-emission aviation through High-Performance Computing (HPC)

Summary

EPSRC's investment in High-Performance Computing (HPC) infrastructure has supported UK-led research into hydrogen combustion, an important technology for decarbonising aviation.

Through access to ARCHER2, the researchers from Newcastle University and the University of Cambridge have partnered with Rolls-Royce on a project to develop hydrogen-powered aeroplane engines. This case study demonstrates how national HPC resources have enabled high-fidelity simulations that would not be feasible otherwise, generating knowledge and capabilities that support industrial innovation and the UK's net-zero target.

Research motivation

Decarbonising aviation represents a challenge in the transition to net zero; hydrogen-fuelled gas turbines offer a potential route to zero-emission flight. However, the combustion of hydrogen in turbulent environments, such as those encountered in aeroengines, is not well understood. Industry requires robust physical models to support the safe and efficient design of hydrogen combustion systems. Bridging this knowledge gap requires detailed simulation of hydrogen combustion under realistic aeroengine conditions.

EPSRC and ARCHER2 contribution

The principal investigator, Dr Umair Ahmed, has led two Pioneer Projects, EPSRC-funded, high-risk, high-reward research initiatives requiring access to national high-performance computing resources. ARCHER2 has been essential to this work, as the size of the simulations, investigating flame-wall interactions, requires large-scale high-performance computing facilities.

Results generated on ARCHER2 demonstrated the feasibility and industrial relevance of the research, leading directly to the establishment of a collaboration with Rolls-Royce. The company recognised the value of these findings and joined the project as an industrial partner. The collaboration continues to rely on ARCHER2, as the required simulations remain computationally intensive. Access to ARCHER2 has enabled the research to progress in line with industrial needs, supporting sustained collaboration and knowledge exchange.

Key outcomes and impact to date

The research conducted using ARCHER2 with EPSRC support has informed six papers[171] published in journals, including the *International Journal of Hydrogen Energy* and *Physics of Fluids*. These papers

[171] Young, F. W., Ahmed, U., & Chakraborty, N. (2025). Influence of preferential diffusion on the distribution of species in lean H2-air laminar premixed flames at different equivalence ratios. *Fuel, 381*(Part B), 133363. https://doi.org/10.1016/j.fuel.2024.133363; Ghai, S. K., Ahmed, U., & Chakraborty, N. (2025). Effects of wall temperature on scalar and turbulence statistics during premixed flame–wall interaction within turbulent boundary layers. *Flow, Turbulence and Combustion, 114*, 421-448. https://doi.org/10.1007/s10494-024-00603-w; Mohan, V., Young, F. W., Ahmed, U., & Chakraborty, N. (2025). Influence of equivalence ratio on the statistics of the invariants of velocity gradient tensor and flow topologies in turbulent premixed lean H_2-air flames. *International Journal of Hydrogen Energy, 98*, 35-51. https://doi.org/10.1016/j.ijhydene.2024.11.477; Mohan, V., Ahmed, U., & Chakraborty, N. (2025). Distributions of wall heat flux and wall shear stress and their interrelation during head-on quenching of premixed flames within turbulent boundary layers. *Flow, Turbulence and Combustion, 114*, 1361-1376. https://doi.org/10.1007/s10494-024-00633-4; Ahmed, U., Chakraborty, N., & Klein, M. (2024). Effects of laminar burning velocity to friction velocity ratio on turbulent premixed flame-wall interaction within turbulent

help improve our understanding of how hydrogen behaves when it burns, particularly in low-emission settings. The results help explain how heat is transferred near surfaces, how hydrogen mixes with air before ignition and how flames interact with surrounding materials. This knowledge supports the design of cleaner, more efficient energy systems using hydrogen fuel across sectors such as transport and power generation.

In the short term, this project has contributed to the advancement of mathematical models for combustion, which have been integrated into the open-source simulation platform Code_Saturne, developed by EDF Energy. The open dissemination of these models enhances reproducibility and facilitates their adoption by industrial stakeholders. This, in turn, supports innovation in hydrogen combustion modelling and broadens the applicability of high-fidelity simulations to sectors such as energy and aerospace, particularly in the development of low-emission combustion technologies.

While the impact of the current work is still emerging, prior research by Ahmed and Prosser (2016)[172], conducted using the ARCHER, demonstrates the potential for broader uptake. Their study implemented the flame-turbulence interaction (FTI) transport model within Code_Saturne, validating it against experimental and large eddy simulation (LES) data, and benchmarking it against established models. This work provides a robust foundation for replication and further refinement. According to Dr Ahmed, these models were specifically developed for industrial applications and are currently in use within that context.

The project has introduced new mathematical modelling techniques that enhance the simulation of industrial flows, reducing the need for costly physical experiments such as combustion tests in a full-scale jet engine combustor rig. Through an industry collaboration with Rolls-Royce, supported by EPSRC's provision of computational time on ARCHER2 via the Pioneer Projects, these techniques are being applied to the early-stage design of new zero-emission aeroengines.

Additionally, the research activities have supported the training of two PhD students and two Postdoctoral Research Associates. One PhD student completed their doctorate in 2023 and is now a PDRA in quantum computing at Imperial College London, with a University Fellowship at the University of Manchester starting in late 2025. Another PhD graduate is working on hydrogen combustion as a PDRA at the University of Cambridge. One of the PDRAs who used ARCHER2 is now employed as a data scientist at CPI in Durham. These outcomes illustrate how the skills developed through simulation work and data analysis on ARCHER2 have enabled the researchers to continue contributing to computational research within the UK.

Stakeholder insights

Stakeholders highlight the industrial relevance of the research and its potential to accelerate hydrogen technology development:

“We [Rolls-Royce] believe that the proposed ‘Application of Direct Numerical Simulation (DNS) to gain physical understanding of the flame and to identify thermo-acoustic instabilities at relevant combustor operating conditions for hydrogen-air swirl flames’ can offer the benefit of filling in the

boundary layers. *Physical Review Fluids, 9*(11), 113201. https://doi.org/10.1103/PhysRevFluids.9.113201; Ghai, S. K., Ahmed, U., Chakraborty, N., & Klein, M. (2024). Multiscale analysis of Reynolds stresses and its dissipation rates for premixed flame–wall interaction. *Physics of Fluids, 36*(10), 105199. https://doi.org/10.1063/5.0232629

[172] Ahmed, U., & Prosser, R. (2016). Modelling flame turbulence interaction in RANS simulation of premixed turbulent combustion. *Combustion Theory and Modelling*, 00(00), 1-16. https://doi.org/10.1080/13647830.2015.1115130

knowledge-gap and speed up the development of hydrogen combustors, in line with the company's (and the UK Government's) net-zero vision." – Rolls-Royce

"I am sure that the outcomes of this project will eventually be translated into models for industrial use and this will be of direct benefit to the industrial community resulting in the widest impact possible." – Renuda

"This area of research has received limited attention in the existing literature and the use of high-fidelity simulations proposed in this proposal will be timely, and important for ensuring that flame stability issues are resolved for power generation and gas turbine combustors of the future." – Professor Ryoichi Kurose, Kyoto University

Looking ahead

Continued access to national HPC infrastructure is necessary to support this research. The models being developed aim to inform the design of combustor components, contributing to industrial development in low-emission aviation.

The project focuses on very lean hydrogen combustion in aircraft engines, where the fuel is mixed with more air than in conventional systems. This work supports the design of combustors that are stable and consistent under lean conditions, with the potential to reduce nitrous oxide and eliminate carbon emissions. In the longer term, the research contributes to the development of zero-emission aeroengines in collaboration with industry partners, aligning with the UK Government's net-zero targets.

By 2025, aviation emissions could reach 1.5 to 2.0 gigatonnes of CO_2 annually[173], with an estimated annual social cost of these CO_2 emissions £254 to £338 billion[174] ($278 to $370 billion)[175]. These figures exclude additional warming from non-CO_2 effects such as NO_x and water vapour, which further increase aviation's climate impact. The project addresses this challenge by developing hydrogen combustor technologies that eliminate in-flight CO_2 emissions and reduce NO_x output through lean combustion strategies. Current estimates suggest this approach could cut the total climate impact of flight by 50 to 75% compared to conventional jet fuel,[176] highlighting the environmental benefits of the combustor designs informed by this research.

[173] Clean Sky 2 JU & Fuel Cells and Hydrogen 2 JU. (2020). Hydrogen-powered aviation: A fact-based study of hydrogen technology, economics, and climate impact by 2050. Publications Office of the European Union. https://www.euractiv.com/wp-content/uploads/sites/2/2020/06/20200507_Hydrogen-Powered-Aviation-report_FINAL-web-ID-8706035.pdf

[174] USD values were converted to GBP using the Bank of England's average 2022 spot exchange rate (1 USD = 0.808930594 GBP), then uprated to 2025 prices using GDP deflators from HM Treasury (2024) (2022 = 0.88607225, 2025 = 1.00). Sources: Bank of England. (n.d.). Data viewer. Retrieved May 2025, from https://wwwtest.bankofengland.co.uk/boeapps/database/fromshowcolumns.asp?Travel=NIxAZxSUx&FromSeries=1&ToSeries=50&DAT=RNG&FD=1&FM=Jan&FY=2015&TD=19&TM=May&TY=2025&FNY=&CSVF=TT&html.x=176&html.y=34&C=DMY&Filter=N; HM Treasury. (2024). GDP deflators at market prices, and money GDP: March 2024 (Quarterly National Accounts). Retrieved from https://www.gov.uk/government/statistics/gdp-deflators-at-market-prices-and-money-gdp-march-2024-quarterly-national-accounts

[175] The estimated annual cost to society of $278 to $370 billion is derived by multiplying projected global aviation emissions of 1.5 to 2.0 gigatonnes of CO_2 in 2025 by the central estimate of the social cost of carbon, which is $185 per tonne (Rennert et al., 2022). This results in a range of $277.5 billion to $370 billion. Rennert, K., Errickson, F., Prest, B. C., Rennels, L., Newell, R. G., Pizer, W., Kingdon, C., Wingenroth, J., Cooke, R., Parthum, B., Smith, D., Cromar, K., Diaz, D., Moore, F. C., Müller, U. K., Plevin, R. J., Raftery, A. E., Ševčíková, H., Sheets, H., Stock, J. H., Tan, T., Watson, M., Wong, T. E., & Anthoff, D. (2022). Comprehensive evidence implies a higher social cost of CO_2. Nature, 610, 687–692. https://doi.org/10.1038/s41586-022-05224-9

[176] Clean Sky 2 JU & Fuel Cells and Hydrogen 2 JU. (2020). Hydrogen-powered aviation: A fact-based study of hydrogen technology, economics, and climate impact by 2050. Publications Office of the European Union. https://www.euractiv.com/wp-content/uploads/sites/2/2020/06/20200507_Hydrogen-Powered-Aviation-report_FINAL-web-ID-8706035.pdf

A5.10 Computer Modelling of Liquid Repellent Surfaces

Fibrous materials are found in both nature, for example cotton, and man-made products, such as nylon. A key property of fibrous materials is their wetting dynamics, which describe how liquids interact with and absorb into fibres. Wetting dynamics mostly depend on the type of liquid and the size, arrangement and textures of the fibres themselves, known as the surface geometry. Domestic products that require high absorption, such as paper towels, require very different surface geometries compared to products such as diapers, which require repellent surfaces that effectively direct the liquid towards more absorbent layers.

The effective design of surfaces has numerous industrial applications. Despite the ability of advanced manufacturing techniques to produce a wide range of surfaces, the fundamental mechanism of how surface geometries influence wetting properties is poorly understood. Companies such as P&G undertake rigorous experiments involving trial and error to identify suitable designs for their products.[177] These experiments are time-consuming and resource intensive, leading to challenges in optimising the performance of their products.

Researchers at Durham specialise in the field of wetting phenomena, surface geometries and computational modelling and engage regularly with industry to develop analytical tools for designing and testing their products. In 2017, they applied for a New Investigator grant from EPSRC, aiming to expand these analytical tools in collaboration with P&G.[178] With access to ARCHER2, they developed complex simulations that mimic wetting dynamics observed during laboratory experiments. The simulation is based on genetic algorithms that can optimise surface design inspired by the process of natural selection, up to 10,000 times faster compared with previous computational methods.

The researchers found that changing variables of surface geometry will often affect different wetting properties in opposite ways. This meant that all parameters of surface geometry and liquid composition needed to be considered simultaneously to create a tool effective in capturing the whole system.

Achieving realistic simulations of these highly complex surface geometries therefore required an extremely large amount of code to be run in parallel to capture the intricate interplay between surface geometries and wetting dynamics within a single simulation. Without the large number of cores and nodes available at ARCHER2, the parallel calculations required by the simulation tool would not have been possible.

Beyond its unique computational power, ARCHER2 was also capable of storing the vast amount of data generated from these simulations. This enabled researchers to interpret the results and generate useful insights for P&G. Access to ARCHER2 through this grant was therefore fundamental to enabling the project. The collaboration with P&G focused on two problems.

One project involved designing better liquid repellent surfaces for improved cleaning products. The researchers were inspired by the naturally-occurring ability of the insect, Spring tail, to repel many kinds of liquid, termed omniphobicity.[179] Using ARCHER2, they were able to simulate specific features of the Springtail surface fibres that give rise to its omniphobicity. The mechanistic

[177] See https://results2021.ref.ac.uk/impact/a2523c89-20ca-455d-a584-d919167a6208?page=1
[178] See https://gtr.ukri.org/projects?ref=EP%2FL00030X%2F1
[179] See https://results2021.ref.ac.uk/impact/a2523c89-20ca-455d-a584-d919167a6208?page=1

understanding they gained from these simulations enabled them to incorporate similar properties into P&G's range of cleaning products.

In March 2025, a UN report highlighted that approximately half the world's population are facing water scarcity at some point every year.[180] The simulation tool also led to work with P&G's Transformative Platform Technology Division in Cincinnati, helping them to reduce the environmental impact of their products. They developed new laundry products that require less water and environmentally damaging surfactant whilst removing soil from clothing more effectively.

The powerful modelling approaches have also transformed the speed and efficiency at which P&G's is able to optimise its product design. The tool has enabled them to incorporate underlying physical principles of how liquids and semi-solids can be efficiently removed from solid surfaces. The same work done experimentally in a lab would have taken years to complete. The simulations have significantly reduced the experimental time for trial and error and resources required by P&G, thereby improving their efficiency and economic sustainability.

The project forms part of a wider partnership between Durham and P&G that has attracted inward investment of over £20 million for the North East region.[181] The estimated value of the P&G products and processes improved through this collaboration exceeds £800 million per annum, in markets worth £11 billion.

"The collaboration with the Durham group has benefited P&G to accelerate our innovation. The modeling work from Durham reduces time and resources associated with experimental trial and error approach. For example, in the soil de-wetting projects, there are many parameters and multiple mechanisms to be tested, which require years of laboratory work to be done systematically and will cost P&G multiple full-time staff members. The simulation results have allowed us to concentrate our in-house experimental research efforts on the most promising regimes highlighted by the Durham group, thus reducing experimental time and resources by more than a half. Furthermore, the understanding and insights provided by the modeling and simulations have also been very valuable for discovering new cleaning mechanisms. As the TPT works cross categories and sector business units, the foundational knowledge gained from Durham can impact a wide range of products to better delight P&G consumers across many product categories now and generations to come." — Senior staff member at P&G Transformative Platform Technology division in Cincinnati

The UK Consortium on Mesoscale Engineering Science (UKCOMES) is a High-End Computing consortium supported by EPSRC through UKRI's central Digital Research Infrastructure.[182] It provided both a platform for disseminating the simulation techniques developed during this project as well as continued access to ARCHER2 beyond the project, which has enabled collaborations with additional industry partners using same tool developed at Durham.

A collaboration with ExxonMobil was supported through EPSRC's Impact Acceleration Awards to investigate the feasibility of a new carbon capture mechanism.[183] Liquid infused surfaces (LIS) are a novel class of surface that contain a thin layer of lubricant on a solid surface that repels all other liquids. They can absorb gas molecules such as CO_2 and have great potential in producing cheaper,

180 UNESCO. (2025). UN World Water Development Report 2025. Retrieved from https://www.unesco.de/assets/dokumente/UNESCO/01_Publikationen/WWDR_2025_ff_EN_WEB.pdf
181 See https://results2021.ref.ac.uk/impact/a2523c89-20ca-455d-a584-d919167a6208?page=1
182 See https://gtr.ukri.org/projects?ref=EP%2FR029598%2F1
183 See https://gtr.ukri.org/projects?ref=EP%2FV034154%2F1

more environmentally friendly alternatives to current carbon capture methods that have been used since the 1930s.

However, a key weakness of LIS is that they lack durability due to the solid surface quickly losing the lubricant and the ability to absorb CO_2. As with the P&G collaboration, the capabilities at ARCHER2 were integral to model the complex interactions of the lubricant with solid surface to form the LIS. The models explored the underlying mechanisms that cause liquids to leak from the substrate. This has enabled ExxonMobil to design new ways to immobilise the liquid onto the solid surface, thereby making them more durable and creating more viable alternative carbon capture solutions.

In 2024, the simulation model was used in collaboration with Infineon, a semiconductor manufacturing company, to design more robust coatings for their electronic devices. The main challenge was to prevent any contaminants contained within water getting into the electronics, which is the main cause of damage beyond water itself. This work is currently being funded by an EPSRC CASE Conversion Award and continues to benefit from ARCHER2 and Tier 2 resources funded by EPSRC such as Cirrus.[184] The findings will lead to more robust design of electronic devices and therefore less material usage caused by needing to replace semiconductor parts.

By investing in national infrastructure such as ARCHER2, supporting academia-industry partnerships and building community networks such as UKCOMES, EPSRC have ensured transformational discoveries are translated. This will lead to dissemination across a wide range of industries and broader impacts for society.

Figure 63 Water droplets on a textured surface, showing the interaction of liquid against a hydrophobic material at high-resolution

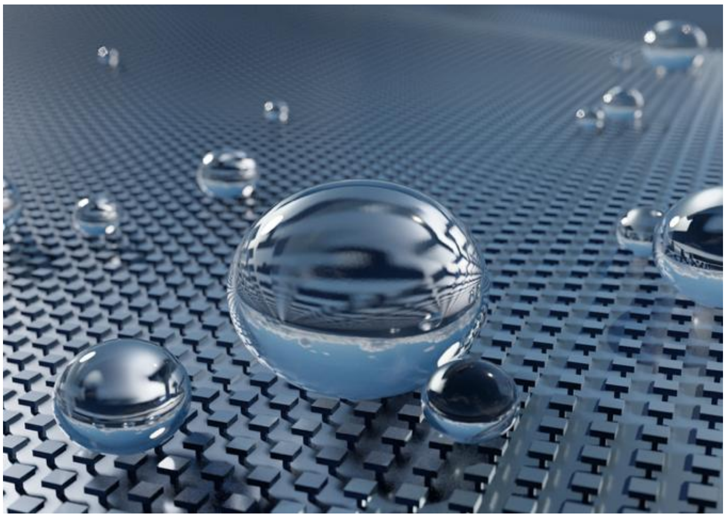

A5.11 Simulating accident scenarios in nuclear reactors to improve and accelerate the design process

Under normal operation, the systems in and around a nuclear reactor's core are subject to severe conditions, like intense radiation, high temperatures, mechanical stress and exposure to harsh

[184] See https://gtr.ukri.org/projects?ref=EP%2FP007139%2F1

chemicals. These factors are well understood and are routinely incorporated into reactor design. In contrast, conditions during accident scenarios are less well characterised, largely due to their rarity and the challenges involved in replicating them. Nonetheless, ensuring reactor safety under worst-case conditions is essential. To do this, engineers rely on advanced computer models to predict how parts of the reactor, like the coolant system, would behave during an accident. These models use a method called Computational Fluid Dynamics (CFD) to simulate how liquids and gases move.[185] But CFD needs a lot of computing power, especially when done at the high resolutions needed for design safety. That's why supercomputers like ARCHER2 are used to run these complex simulations.

Although access to supercomputers has improved thanks to cloud computing and national initiatives like ARCHER2, it remains limited due to high operating costs and the growing demand for these resources arising from the increased complexity of artificial intelligence and the Internet of Things. To reduce reliance on these powerful but costly resources, engineers often use simpler, less detailed modelling techniques during much of the reactor design process. These methods can run on standard desktop computers and are effective for simulating normal, steady-state reactor conditions. However, they fall short when it comes to accurately capturing the brief but intense events that occur during accident scenarios.[186]

To meet the need for more accurate yet less computationally intensive simulations of nuclear reactor conditions, EDF UK has been collaborating with researchers at STFC since 2017 to develop an intermediate modelling tool. This work, carried out over four years with support from the former BEIS (now DESNZ), led to the development of Subchannel CFD (SubChCFD). SubChCFD is a cost-effective thermal hydraulics simulation tool that significantly reduces computational demands compared to traditional CFD, while offering much improved simulation of extreme scenarios than existing engineering tools in large, complex reactor systems. [187]

Before EDF can fully integrate SubChCFD into its standard design processes, the tool's results must first be validated against those of conventional CFD, which it aims to emulate. To achieve this, EDF once again partnered with STFC in 2023 on an 18-month Pioneer Project. Using ARCHER2, the project ran high-resolution CFD simulations of High Temperature Gas-cooled Reactors (HTGRs) under various hypothetical accident scenarios, including a Loss of Flow Accident.

The project had two main goals: first, to publish the simulation outputs for use by academic and industry stakeholders in improving advanced nuclear reactor designs; and second, to use the data as a benchmark to validate the accuracy and performance of SubChCFD. The first six months of the project were spent setting up and testing the HTGR model in the ARCHER2 environment. The remainder of the project focused on running simulations and analysing the resulting data, with the work concluding in January 2025.[188]

Through the Pioneer Project, STFC was able to simulate multiple accident scenarios in high detail across the entire reactor system. They also validated SubChCFD's results against these high-resolution simulations, confirming its accuracy in replicating complex reactor behaviour.

Findings from the simulations will be used to optimise designs and accelerate the deployment of Advanced Modular Reactors (AMRs, of which HTGRs are just one type), which are safer and more efficient than previous reactor types. These findings also support EDF's goal of decarbonising local heavy industrial sites, such as Hartlepool, which currently emits 3.1 million tonnes of CO_2 annually,

[185] Zinkle, S.J., Was, G.S. (2013) Materials challenges in nuclear energy, Acta Materialia Volume 61, Issue 3, Pages 735-758.
[186] See https://world-nuclear.org/information-library/safety-and-security/safety-of-plants/safety-of-nuclear-power-reactors
[187] See https://sheffield.ac.uk/heft/subchannelcfd
[188] See https://www.okoone.com/spark/technology-innovation/awss-new-hpc-service-expands-access-to-supercomputing-power/

by deploying HTGRs. This aligns with the UK government's goal of using AMRs to increase nuclear capacity in the UK from around 6 GW to 24 GW by 2050.[189]

Figure 64 Depiction of an HTGR core

Source: Wikimedia

SubChCFD is built on open-source code developed by EDF, and validating its accuracy is key to building confidence in its use. By showing that SubChCFD provides a reliable middle ground between detailed CFD and traditional subchannel codes, STFC and EDF aim to encourage its broader adoption in nuclear R&D. Reducing reliance on supercomputing resources could, in turn, help accelerate the development and deployment of HTGRs.

EDF is a strategic collaborator and non-government funder for STFC, so this project not only supports the long-term advancement of the nuclear industry and benefits electricity consumers, but also strengthens the partnership between STFC and EDF, paving the way for future collaboration. EDF has indicated its intention to continue developing SubChCFD with STFC. Meanwhile, STFC is raising its profile in the nuclear sector by sharing project outcomes at high-profile events such as ARCHER2's annual Celebration of Science, the CCP-NTH annual technical meeting, and NURETH-21, the leading international conference on nuclear reactor thermal hydraulics.[190]

"This pioneer project has allowed STFC to produce validation data for coarse CFD modelling that can be deployed for the design, certification and consequently the adoption of High Temperature Gas-cooled Reactors in the UK"

A5.12 Improving combustion research with High-Performance Computing (HPC)

Summary

EPSRC's investment in High-Performance Computing (HPC) infrastructure has supported research into combustion processes relevant to transport and power generation. With access to ARCHER2, Professor Nedunchezhian Swaminathan and his team at the University of Cambridge have developed and tested high-fidelity models for turbulent combustion that are being used by industry to inform engine design. This case study outlines how national HPC resources have enabled simulations that would not be feasible otherwise, contributing to ongoing efforts to reduce emissions and improve engine performance in line with decarbonisation objectives.

[189] See https://researchbriefings.files.parliament.uk/documents/CDP-2024-0036/CDP-2024-0036.pdf
[190] See https://www.okoone.com/spark/technology-innovation/awss-new-hpc-service-expands-access-to-supercomputing-power/

Research motivation

Combustion is the process of burning fuel to produce heat and energy. It supports a wide range of activities across power generation, transport, manufacturing and construction. For example, it drives engines in vehicles and aircraft, produces electricity in power stations and supplies heat for industrial processes such as steel or cement production.

Although widely used, combustion can have unintended effects. One of these is sound, an unwanted but inevitable by-product that can damage engine components and, in some cases, contribute to engine failure. Another one is pollutants such as CO2 and oxides of nitrogen which have their effects on environment and human health.

By modelling combustion processes, engineers can explore new approaches, improve fuel efficiency, reduce emissions and develop quieter, safer and more effective systems without depending on costly or high-risk real-world testing.

EPSRC and ARCHER2 contribution

Professor Swaminathan's research has been supported by EPSRC through access to ARCHER and ARCHER2, including a Pioneer Project, which provides HPC access for early-stage, high-potential research. The simulations required for this work are computationally intensive and could not be carried out on standard university systems. ARCHER2 provided the capacity needed to run, test and further develop high-fidelity models that simulate combustion and its effects under realistic conditions.

This and earlier access helped attract industry support, with companies contributing to both research and personnel costs. The project has produced tools and datasets that are now used by industry and made available to the wider research community. The pioneer project has also helped six PhD students (three graduated and three in progress) and one senior research fellow with their studies and careers. Among the graduates, one is a Henslow Fellow at the University of Cambridge, supported by the Cambridge Philosophical Society, a highly prestigious and competitive fellowship with a success rate of just 1 in 100. Another is a Postdoctoral Fellow at the Cambridge University Engineering Department, collaborating with Rolls-Royce. A third will be joining Rolls-Royce in Bristol as an Aerothermal Analyst starting in September 2025.

According to MHI, ARCHER/ARCHER2 access has enabled 15 PhDs (12 completed, three ongoing) and trained seven postdoctoral researchers over the past decade. Many of these individuals now hold senior positions globally, ranging from Engineering Analyst to VP/CEO roles in engineering, data science and investment banking.

Without ARCHER2, this research would not have been possible. While Tier-2 systems, such as CSD3 at Cambridge, are also used, access is limited in capacity and costly. For example, the simulations conducted in the pioneer project used many thousands of cores, which is impossible with the CSD3 system; even if users force it to conduct these simulations, they will be nearly **6.5 times more expensive**. Additionally, there are other limitations on data size. These factors demand ARCHER2-type service on practical grounds. GPU-based systems and AI integration are being explored, but these approaches require significant redevelopment effort and are not yet suitable for the scale and type of simulations involved. Therefore, maintaining national HPC infrastructure is required for continuing this line of research and, as MHI notes, for preserving the UK's leadership in science and technology on a global scale.

Key outcomes and impact to date

Professor Nedunchezhian Swaminathan has written **16 papers** (six published[191], five in review[192] and seven in preparation[193]) that utilised ARCHER2 through the Pioneer Project. His research focuses on advancing understanding of hydrogen as a low-carbon fuel and exploring new computational techniques, including machine learning, to support combustion science.

Mitsubishi Heavy Industries (MHI) outlined that the supported projects have produced **36 papers in Q1 international journals** and attracted nearly **£5 million** in research funding to the UK. These funds helped to attract highly talented people to the UK and supported them through their PhD and postdoctoral training.

Hydrogen is increasingly recognised as a key component in future energy systems aiming to reduce carbon emissions. However, its safe and efficient use in engines and turbines remains a challenge due to its high reactivity and distinctive combustion behaviour. This body of work investigates how blending hydrogen with small amounts of methane could help address these issues. It examines how the fuel behaves when injected into a fast-moving air stream, replicating conditions found in advanced gas turbines.

Alongside conventional simulation techniques, the research uses machine learning methods to accelerate high-fidelity modelling. These approaches offer a way to generate detailed insights more efficiently, supporting the development of safer and more effective combustor designs.

In addition, the project has generated open-access datasets that are used by researchers internationally. Open access publishing enhances the visibility and dissemination of research by removing paywalls, leading to higher usage and citation rates. It also facilitates educational reuse

[191] Murugavel, A. B., Massey, J. C., Tanaka, Y., & Swaminathan, N. (2024). The effect of methane addition on reacting hydrogen jets in crossflow. *International Journal of Hydrogen Energy, 80*, 57-67;
Kumar, A. D., Massey, J. C., Murugavel, A. B., Chen, Z. X., & Swaminathan, N. (2025). Spectral characteristics of the heat release rate in confined turbulent flames. *Journal of Fluid Mechanics, 1007*, A29-1.

[192] Murugavel, A. B., Massey, J. C., & Swaminathan, N. (2025). Pressure effects on mixing and combustion mode of a hydrogen/helium jet in crossflow. *Journal of Fluid Mechanics*;
Kumar, A. D., Soundararajan, P. R., Li, C. J., Massey, J. C., & Swaminathan, N. (2025). Characteristics of heat release rate spectra in turbulent flames with different fuels. *Proceedings of the Combustion Institute*;
Li, C. J., Dere, R., Massey, J. C., Ezenwajiaku, C., Schumann, C. D. K., Talibi, M., Balachandran, R., Tanaka, Y., & Swaminathan, N. (2025). Hydrogen flame blow-off in lean direct injection burner. *Proceedings of the Combustion Institute*;
Yang, H., Massey, J. C., & Swaminathan, N. (2025). JPResUnet: A joint probability density function translation model in partially premixed flames. *Combustion and Flame*;
Schumann, C. D. K., Massey, J. C., Liu, Y., Li, C. J., & Swaminathan, N. (2025). Large eddy simulation of partially premixed hydrogen combustion with differential diffusion effects. *Combustion and Flame*;
Schumann, C. D. K., Massey, J. C., Li, C. J., & Swaminathan, N. (2025). Large eddy simulation of transient leading-edge propagation in a turbulent lifted hydrogen jet flame. *Proceedings of the Combustion Institute*;
Yang, H., Massey, J. C., & Swaminathan, N. (2025). Large eddy simulation of turbulent premixed combustion with a data-driven probability density function model. *Applied Energy Combustion Science*.

[193] Murugavel, A. B., Massey, J. C., & Swaminathan, N. (2025). The effect of hydrogen enrichment on the combustion characteristics of methane jet in crossflow at elevated pressure conditions. *Combustion Science and Technology*;
Liu, Y., Massey, J. C., Faldella, F., & Tanaka, Y. (2025). Lean premixed hydrogen/methane Bunsen flames. Part 1: Effects of hydrogen content and pressure on flame shape characteristics. *International Journal of Hydrogen Energy*;
Liu, Y., Massey, J. C., Schumann, C. D. K., & Tanaka, Y. (2025). Lean premixed hydrogen/methane Bunsen flames. Part 2: Influence of chemical mechanism and non-unity Lewis number formulation on LES. *International Journal of Hydrogen Energy*;
Li, C. J., Massey, J. C., Dere, R., Ezenwajiaku, C., Schumann, C. D. K., Rocha, N., Talibi, M., Balachandran, R., Tanaka, Y., & Swaminathan, N. (2025). Flow field characteristics of turbulent round jet with jets-in-crossflow entrainment under free and non-free conditions. *Physical Review Fluids*;
Han, Y., Li, Z., Massey, J. C., & Swaminathan, N. (2025). A priori analysis of the dilution equation and its application in large eddy simulation.;
Han, Y., Massey, J. C., & Swaminathan, N. (2025). Large eddy simulation of MILD combustion incorporating a perfectly stirred reactor-based dilution model.

and promotes public engagement, allowing professionals outside academia to benefit from the latest findings.[194]

The computationally efficient and accurate models developed by Professor Swaminathan, tested and validated using ARCHER2, are routinely used by the UK and international gas turbine industries. These tools support informed design decisions and deliver significant cost savings of approximately **£100,000** per year in computational costs, **£1 million** per year in testing and verification and a reduction of over one year in development time.

The models developed by this research are used by companies such as Rolls-Royce, Siemens and Mitsubishi Heavy Industries to support design decisions. MHI reported **savings of approximately £2 million and one year in engine development resources** by using these tools. The simulations help industry users identify performance issues before engines are built, which can reduce the need for physical prototypes and reduce costs.

Stakeholder insights

"It is not easy to achieve these without the access to ARCHER2 and hence the need for such a national facility is imperative and very much required for UK to keep its leadership position and edge in S&T."
– MHI

Looking ahead

The project aims to develop combustion models that are both computationally efficient and accurate, with a focus on future low- and zero-carbon fuels. These models are intended to support the design of lower-emission gas turbines for power generation and aviation, and can also be applied to the development of cleaner engines for road and marine transport.

An additional focus is on combustion-induced noise, which is expected to become a key factor in future aeroengine design as airports move toward implementing noise-related levies. High-fidelity simulations can help estimate noise levels during engine operation, providing valuable insights for designing quieter engines.

By enabling the development of engines that produce less noise and fewer emissions, these simulations have the potential to reduce both the environmental and social impacts of aviation. They also offer practical benefits for airlines, which may face increasing charges or taxes based on their noise and emissions profiles.

A5.13 HPC modelling for next-generation battery materials in electric vehicles

In 2024, transport from road vehicles was the largest emitting sector in the UK, accounting for 30% of the nation's CO_2 emissions. In 2025 alone, electric vehicle sales are expected to increase by 25% globally, with predictions that their market share will rise to 40% by 2030.[195] Development of affordable, reliable and long-range electric vehicles is therefore essential to meet this growing demand and achieve the UK's Net Zero commitments by 2050.[196] The UK Government's Battery

[194] Cambridge University Press. (n.d.). Benefits of open access. Retrieved June 17, 2025, from https://www.cambridge.org/core/open-research/open-access/benefits-of-open-access
[195] See https://www.gov.uk/government/speeches/phasing-out-the-sale-of-new-petrol-and-diesel-cars-from-2030-and-support-for-zero-emission-vehicle-zev-transition
[196] See https://gtr.ukri.org/projects?ref=EP%2FX035859%2F1&pn=8&fetchSize=25&selectedSortableField=date&selectedSortOrder=ASC

Strategy sets out an ambitious vision for the UK to become a world leader in sustainable battery design and manufacture by 2030.

The rechargeable lithium-ion battery has been instrumental in powering the global revolution in portable electronics. This success puts an ambitious challenge in sight: the electrification of road transport, a crucial ingredient for a low carbon future to deal with climate change. The performance of these battery technologies is crucially dependent on the fundamental properties of the component materials.[197]

Of the battery components, the cathode represents one of the major barriers to increasing the energy density of lithium-ion batteries. Current batteries use metal oxides such as cobalt and nickel to make cathodes, which perform well in electric vehicles. However, these materials are in limited supply globally. Therefore, there is an urgent need for the UK to develop scalable solutions for high performance batteries produced from more abundant materials. Lithium-rich materials (so called O-redox cathodes) are based on layered transition metal oxides and a new family of cobalt-free compounds called disordered rocksalts. Current cathodes store charge on the transition metal ions. The lithium-rich materials can store additional charge on the oxygen ions, thus increasing energy density by up to 50%. However, they are not yet used in electric vehicles because they are highly susceptible to degradation and voltage fade upon multiple cycles of charging and discharging.

The mechanism of voltage fade is not fully understood, and the atomic- to nanoscale picture remains incomplete, largely due to challenges in experiments and modelling to characterize these complex materials during cycling. Current experimental methods require X-ray synchrotron techniques, which are complex and expensive to run. Attempting to perform computational studies of voltage fade in lithium-ion battery materials is also highly challenging. The process involves subtle and irreversible changes to the structure and redox properties. This makes it challenging to identify viable structures to use in a simulation. Achieving realistic simulations is also extremely computationally intensive. The crystal structure of the battery material becomes disordered with each charging cycle. Therefore, simulations must run over large enough length scales and timescales to gain a true understanding of how the battery functions over time.

Researchers at the University of Oxford lead a Faraday Institution (FI) funded consortium on lithium-ion cathode materials called CATMAT, comprised of 6 universities, 3 research centres and 12 industrial partners, with the aim to understand and develop new high energy density cathodes. Using the computational power of ARCHER 2, they sought to understand the structural and electrochemical processes at the nanoscale level and how this leads to voltage fade in batteries. The CATMAT project was awarded £15.3m over five years as part of a wider programme of 10 projects under the Faraday Battery Challenge, which were supported by the FI, EPSRC, InnovateUK and UK Battery Industrialisation Centre (UKBIC).[198]

The Oxford materials modelling team, led by Professor Saiful Islam, accessed ARCHER 2 time through membership of the Materials Chemistry Consortium (MCC), who are also funded by EPSRC. Without ARCHER 2, they would not have been able to investigate such complex processes at the nanoscale level. Typically, researchers may be restricted to ab initio modelling of around 500 ions or atoms of the battery materials. New methodology in ARCHER 2 allowed researchers to study up to 50,000 ions, enabling larger scale nano structures to be investigated. Additionally, ARCHER 2 enables modelling of cathode materials through a complete ‘cycle’ of battery charging and then discharging, rather than time-bound, ‘snapshots’ of degradation processes. Performing similar molecular

[197] See https://www.faraday.ac.uk/2022-23-annual-report/
[198] See https://www.ukri.org/what-we-do/browse-our-areas-of-investment-and-support/faraday-battery-challenge/

dynamics simulations using in-house HPC capabilities would have taken weeks or months. Instead, ARCHER 2 can complete high-quality simulations of different material compositions in a few days, significantly reducing the time taken to gather insights. ARCHER 2 also enabled greater storage space as well for post-simulation analysis.

The simulations revealed the contribution of 'nanovoids' in the crystal structure to voltage fade in batteries during charging cycles. Nanovoids are gaps in the structure of the battery material that form irreversibly and contain O_2 molecules. With each charging cycle, the nanovoids increase in size, which leads to reduction in battery performance and lifetime. Further research has discovered that nanovoids forming closer to the surface of the cathode can lead to the material cracking and the voids opening up at the surface, thereby releasing O_2.[199]

The researchers have developed a model that predicts nanovoid formation and voltage fade, using a first principles approach based on the fundamental laws of physics. The modelling insights help to accelerate the identification of promising material compositions for batteries, compared to experimental approaches, as well as indicating materials design routes capable of mitigating voltage fade.

Work is now ongoing with academic collaborators and industry partners of CATMAT and other FI projects, such as UKBIC, BASF, Fortescue Zero and the Centre for Process Innovation, to develop scalable manufacturing processes for the most promising material compositions identified through ARCHER2 models.

The simulations enabled by ARCHER 2 have provided a crucial step towards the practical application of next generation batteries to electric vehicles. New battery technology will provide electric vehicles with longer range and durability, which will accelerate the adoption of electric vehicles, help the UK to reach its Net Zero target by 2050 and secure its position as a global competitor in battery technology. By convening UK researchers and industry, FI projects such as CATMAT will improve supply chain resilience in the UK by ensuring new battery technology is manufactured domestically, for example in new battery gigafactories.

A5.14 Advancing the development of compact electron accelerators through High-Performance Computing (HPC)

Summary

Researchers at the University of Oxford (led by Prof Simon Hooker and Emeritus Prof Roman Walczak) and Ludwig Maximilian University of Munich (LMU) (led by Prof Stefan Karsch) are working on a new type of electron accelerator; Plasma-Modulated Plasma Accelerator (P-MoPA), which could be smaller, cheaper and more energy-efficient than current systems.

An electron accelerator is a device that increases the speed and energy of electrons and guides them in a precise, controlled manner. These devices have a wide range of practical uses, for example:

- In healthcare, electron accelerators are used in radiotherapy to deliver controlled doses of radiation for the treatment of cancer.

[199] Marie, J. J., House, R. A., Rees, G. J., Robertson, A. W., Jenkins, M., Chen, J., ... & Bruce, P. G. (2024). Trapped O2 and the origin of voltage fade in layered Li-rich cathodes. *Nature Materials*, *23*(6), 818-825.

- In manufacturing and engineering, they support non-destructive testing by enabling internal inspection of components without causing damage.
- In the food and medical supply sectors, they are applied in sterilisation processes to reduce microbial contamination in medical equipment and packaged food products.[200]

Accelerators are typically large or very large, spanning from meters up to kilometres. Researchers from Oxford University and LMU seek to develop smaller and more efficient designs to make the technology more accessible and environmentally friendly.

To design and test ideas, the research team used ARCHER2, allowing them to run complex simulations. The project has led to new research papers, partnerships and funding for experimental verifications of simulations and is now also exploring commercial opportunities.

Research motivation

The use and accessibility of traditional particle accelerator systems are often limited by their size, cost and energy use. Researchers from Oxford University and LMU address this by investigating the design of smaller accelerators to minimise environmental impact and be deployed in a wider range of settings, in areas such as medical diagnostics.

EPSRC and ARCHER2 contribution

This project was supported through the EPSRC Pioneer Projects, which provides access to national High-Performance Computing (HPC) resources for early-stage, high-potential research. According to Professor Walczak, who was leading the work on simulations, without access to ARCHER and ARCHER2, the research would not have occurred. No other computer available in the UK was powerful enough for the required simulations. Another reason was time and personnel constraints (size of the research group, time available for graduate students to complete their thesis, other research commitments, etc).

In the absence of Tier-1 computer, the researchers would have used one of Tier-2 computers. Alternative HPC systems include earlier versions of the Isambard supercomputer in Bristol. However, these computers are often not of the same potential as ARCHER2. Hence, another solution would be to get funding for using similar supercomputers abroad (e.g. EuroHPC supercomputer). But they would be facing similar personnel constraints. The research team is small including about 10 researchers, so engaging in the process of applying for funding and access to supercomputers abroad would significantly delay the research. Another issue is Brexit, which limits the access to the EuroHPC supercomputer to UK researchers. Hence, the researchers guess that, without ARCHER2 they would only have been able to publish two papers out of the four cited in the text. Without these papers, it would have been more difficult to convince other researchers and reviewers to give them funding and probably impossible.

The team used ARCHER and ARCHER2 to run large-scale simulations that supported the design and planning of a single experiment involving the Gemini laser system, where one aspect of the P-MoPA concept was tested. While Gemini was not the primary facility for the broader experimental program, which will take place in Oxford and at LMU, where the appropriate lasers are available, the simulations helped prepare for that specific experiment. Due to the computational work, the team was able to optimise the setup in advance, making it possible to complete the experiment

[200] Kutsaev, S. V. (2021). Novel technologies for compact electron linear accelerators (Review). Instruments and Experimental Techniques, 64(5), 641–656. https://doi.org/10.1134/S0020441221050079

within a single six-week session. Given that a session costs approximately **£500,000**, this preparation significantly improved cost-efficiency and enabled the team to gather good results and publish a paper. Looking ahead, the team expects similar benefits from ARCHER2 simulations in upcoming experiments, where computational predictions will guide and refine experimental setups before any physical trials are conducted.

Key outcomes and impact to date

The project has resulted in a range of outcomes. These include the publication of **4 research papers**[201] and collaborative partnerships established with the Central Laser Facility and the commercial laser manufacturer TRUMPF Scientific Lasers.

The P-MoPA project has secured several grants, including from the United States Air Force Office of Scientific Research (AFOSR) (**£778,000**), a **£363,000** share of a €10 million EU Plasma Accelerator systems for Compact Research Infrastructures (PACRI) grant and a **£2.16 million** EPSRC grant. In April 2024, an STFC Impact Acceleration Account (IAA) grant of **£27,000** was awarded to support an initial study into the commercial potential of P-MoPA.

Intellectual property generated through the project has been formally protected. A PCT patent application (PCT/GB2022/051390), entitled "Driver for Particle Accelerator", has been filed.[202] The intellectual property is managed by Oxford University Innovation Limited, ensuring professional oversight of its commercialisation potential.

Looking ahead

The patent application and the grant to access market potential will provide a route to commercialisation, with opportunities to develop new products or services based on the technology. The research team is currently working with Oxford University Innovation (OUI) and participating in programmes such as Oxford Venture Scouts to explore these possibilities and better understand the spinout process. The team plans to pursue the formation of a spinout company once the results from upcoming experiments, which are intended to validate their simulation findings, become available. These results are expected within the next 12 to 24 months.

While many plasma accelerators offer size and cost benefits over traditional systems, P-MoPA stands out for its potential performance improvements. It is designed to be around **1,000 times** more energy-efficient than existing plasma accelerator models and can fire electron pulses at least **100 times** more frequently. This enables it to produce stronger and more continuous beams of particles or radiation. P-MoPA is also expected to operate with built-in stability by using Machine Learning and Artificial Intelligence to automatically monitor and adjust its performance. This reduces the need for expert operators and could make the system more accessible to a broader range of users.

---

[201] Jakobsson, O., Hooker, S. M., & Walczak, R. (2021). Gev-Scale Accelerators Driven by Plasma-Modulated Pulses from Kilohertz Lasers. Physical Review Letters, 127(18), 184801. https://doi.org/10.1103/PhysRevLett.127.184801; van de Wetering, J. J., Hooker, S. M., & Walczak, R. (2023). Stability of the modulator in a plasma-modulated plasma accelerator. Physical Review E, 108(1), 015204. https://doi.org/10.1103/PhysRevE.108.015204; van de Wetering, J. J., Hooker, S. M., & Walczak, R. (2024). Multi-GeV wakefield acceleration in a plasma-modulated plasma accelerator. Physical Review E, 109(2), 025206. https://doi.org/10.1103/PhysRevE.109.025206; Ross, A. J., Chappell, J., van de Wetering, J. J., Cowley, J., Archer, E., Bourgeois, N., Corner, L., Emerson, D. R., Feder, L., et al. (2024). Resonant excitation of plasma waves in a plasma channel. Physical Review Research, 6(L022001). https://doi.org/10.1103/PhysRevResearch.6.L022001

[202] Oxford University Innovation Limited. (2022). Driver for a particle accelerator (WO2022254207). World Intellectual Property Organization. https://patentscope.wipo.int/search/en/detail.jsf?docId=WO2022254207

Preliminary findings suggest that P-MoPA is less likely to be adopted in routine medical or industrial settings in the near term, as its complexity and energy levels are beyond what those applications currently require. However, it shows strong potential in advanced imaging and scientific use. These include compact X-ray sources for security scanning, industrial inspection and research. In future, it could also support the development of Free Electron Lasers that generate highly detailed images of biological samples and materials using precise wavelengths of light.

A5.15 Driving transition to electrification of flight

Aviation accounts for approximately 2.5% of global CO_2 emissions. As average global income and quality of life rises, so too will the demand on air travel, highlighting the need for cleaner alternatives. FlightPath 2050 is a strategic vision set by the European Commission to achieve more environmentally friendly aviation in Europe through the electrification of flight.[203]

The development of a new generation of turbine parts, for example turbine blades, is crucial for the transition to electrically powered flight. Key physical characteristics such as lightness and tolerance to extreme temperatures must be incorporated into the design of future turbines. These desirable characteristics rely heavily on the production of metals exhibiting a property known as anisotropy. A material that is anisotropic contains atoms that are arranged in one direction, otherwise known as a single grain, which means the material will have different physical properties depending on the direction in which you measure them. One of the main properties is strength, for example cutting a block of wood in the direction of the grain is easier than cutting it against the grain.

Turbine parts are made from a type of metal that is highly anisotropic and are therefore very strong and durable against the intense forces they face within an aeroengine. The anisotropy of a turbine part is highly dependent on the production method used, which determines the way in which the metal solidifies during the production process. Furnace technologies can precisely control the temperature to enable the metal to solidify gradually, which ensures a single grain is produced. However, during the solidification process, the metal can also form secondary grains that are arranged in a different direction, which make the turbine part unfit for function in aeroengines. These secondary grains are found during detailed inspection, and the parts are scrapped.

The physical processes that lead to these unwanted features are poorly understood. Rolls-Royce have spent decades trying to perfect the production of turbine parts that are highly anisotropic. However, a central challenge within the aviation industry is that the causes of unwanted solidification features, such as the formation of secondary grains, are difficult to investigate. This is partly due to the final product being dependent on so many conditions within the furnace, but also because anything that you use to look at what is going on during the production process will either be destroyed (by the extreme temperatures in the furnace) or change the production process itself.

In 2023, researchers from Greenwich began to tackle this challenge initially by using the university's High Performance Computing (HPC) facilities to develop computer simulations of the turbine production process at a smaller scale to ARCHER 2.[204] In 2024, Greenwich secured £9M funding from Research England, of which £1M is being invested into upgrading the university's HPC infrastructure. This new investment will act as an additional stepping stone that will be useful in the development and validation of their modelling approach, before using ARCHER 2 within the next year.

[203] See https://ourworldindata.org/global-aviation-emissions
[204] See https://gtr.ukri.org/projects?ref=EP%2FX025454%2F1#/tabOverview

The primary aim at Greenwich is to apply the simulations to ARCHER 2 during this project to understand the complex mechanisms that lead to the formation of secondary grains and how these defects can get worse over time. A simulation capable of modelling how an entire turbine blade solidifies at the microscopic level would be a World first.

A whole-systems view is crucial to understanding the causes and effects of secondary grain formation. The use of ARCHER 2 will create huge opportunities for Rolls-Royce to optimise the production of their turbines. Modelling across different scales and in 3D requires hundreds of billions of computational cells to simulate both the physics of turbine solidification and the complete set of conditions within a furnace. It would be highly expensive for the project team to run a simulation of this scale and store the terabytes of resulting data within their own organisation. ARCHER 2 offers around 750,000 cores at max capacity, representing step change in capacity in orders of magnitude compared to what is available to the team at their organisation. In the absence of ARCHER 2, the team at Greenwich would have relied on their university's internal infrastructure, which would limit the scope and size of their simulations, reducing the pace in which interesting findings could be investigated and inform new ideas. Therefore, the lack of ARCHER 2 would also constrain the team in terms of exploring new areas of research and the ability to re-focus objectives during the project, which they identified as an important aspect for the success of a large-scale project such as ARCANE.

Government regulations around IP also prohibit the export of company data to be used on a supercomputer overseas. Therefore, other facilities considered by the researchers would need to be based in the UK. Viable alternatives for ARCHER 2 include the new Isambard 3 supercomputer based in Bristol. However, constraints such as available memory required to store the resulting simulation data would remain a limiting factor to the scale of the simulations. These facts highlight how essential it is to have ARCHER 2 capabilities based in the UK.

A key feature of the Prosperity Partnerships is the requirement for match funding from an industry collaborator. The match funding contribution from Rolls-Royce in this project has enabled the team to carry out expensive but valuable experimental trials that will validate modelling approach being developed by the researchers. Furthermore, Rolls-Royce will help the researchers test their modelling approaches in cases that do not involve aircraft to discover alternative uses for the model. This demonstrates a key benefit of the Prosperity Partnership award in enabling mutually beneficial relationships between academia and industry. The researchers also recognised the critical importance of their access to ARCHER 2 in terms of gaining interest from Rolls-Royce in forming this partnership. Without ARCHER 2, the opportunity to partner with Rolls-Royce may not have existed.

The researchers have already identified uses for the simulations across similar industrial process. For example, the same model has been used to study 3D printing in microgravity conditions with potential for use in space. Additionally, the model has relevance to improving battery performance, where failure mechanisms are often caused by the same defects seen in the production of turbine parts.

The UK Government's recent Compute Roadmap announced up to £750 million in funding for a new national supercomputer to replace ARCHER 2. This investment has created new opportunities for the researchers at Greenwich to collaborate with industry – both within this project and across other initiatives.

In the near term, the simulations will enable Rolls-Royce to improve their manufacturing processes to produce strong and more durable turbines. This will reduce the cost and environmental impact of air travel whilst also reducing wastage and carbon emissions by producing fewer engine parts.[205]

"We have saved tens of millions of pounds by using process simulation to increase yield over the years. This value increases every year with the increasing complexity of components that we're manufacturing"

- Collaborator at Rolls-Royce

The simulations will also inform future generations of turbine blades, which will involve more lightweight, efficient designs necessary for electrically powered engines. This will enable Rolls-Royce to meet the growing demand for sustainable air travel and achieve the emissions targets set by FlightPath 2050. New innovations will also boost national security through self-reliance and expertise developed within UK. They will help meet public demand for environmentally friendly alternatives, not just for travel, but lowering the carbon footprint of products.

Figure 65 Large scale 40 mm cube, 8 billion cell simulation of polycrystalline dendritic directional solidification using HPC, capturing component scales at a microscopic resolution.

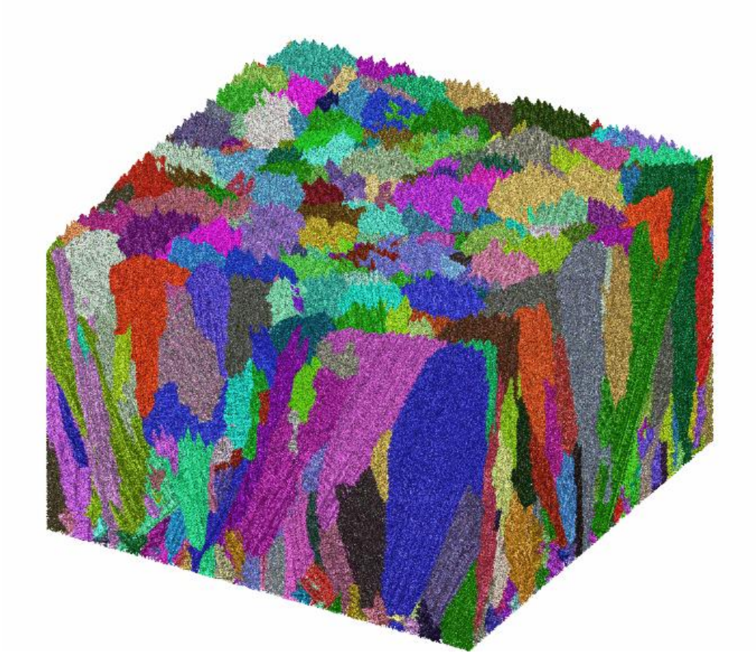

The different colours represent different orientations of dendrites.

[205] See https://www.ukri.org/what-we-do/browse-our-areas-of-investment-and-support/prosperity-partnerships-business-and-academia-collaborations/

LE
London
Economics

Somerset House, New Wing, Strand,
London, WC2R 1LA, United Kingdom
info@londoneconomics.co.uk
londoneconomics.co.uk
𝕏 @LondonEconomics
+44 (0)20 3701 7700