From Concept to CoSeC: A CCP Tale

A brief history of collaborative computational science and HPC in the UK

Prof. David Emerson,
STFC Scientific Computing,
Daresbury Laboratory

Science and Technology Facilities Council

CoSeC
Computational Science Centre for Research Communities

https://www.ccp.ac.uk
UK Collaborative Computational Projects

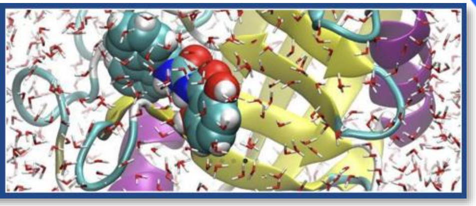

CoSeC Conference 2023 – Manchester Central Convention Complex – 6th December

Timeline of key moments

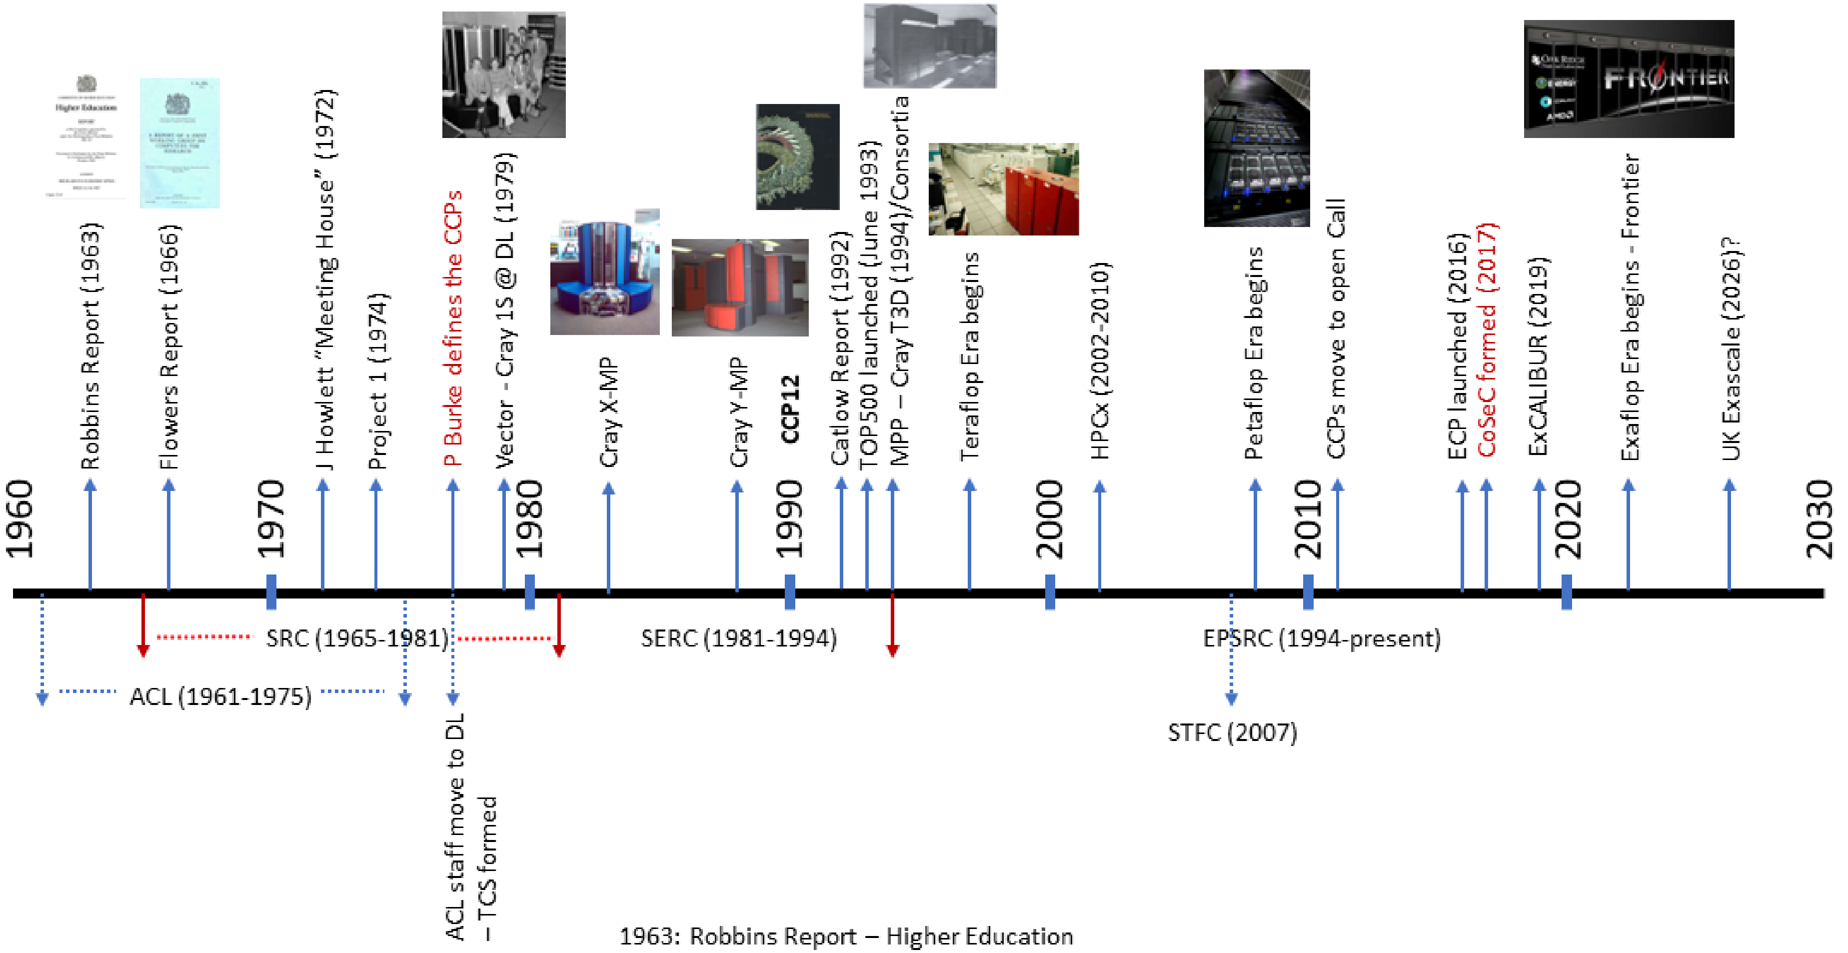

1963: Robbins Report – Higher Education

1966: Flowers Report – A Report of a Joint Working Group on Computers for Research

1974: Project 1: Correlation in Molecular Wave Functions

1992: Catlow Report - Research Requirements for High Performance Computing (SERC, September)

Sc
Te
Fa
Scientific Cor

What are the CCPs?

A Collaborative Computational Project (CCP) provides a focus that brings together like-minded researchers to tackle a significant computational challenge – typically too big for a single person/institution. An excerpt from a brief paper by Prof Paul Durham[1] (1996) highlights the thinking

Typically, CCPs:

- Implement flagship code development projects
- Maintain and distribute code libraries
- Organise training in the use of codes
- Hold meetings and workshops
- Invite overseas researchers for lecture tours and collaborative visits
- Issue regular newsletters

[1]The UK Collaborative Computational Projects, ERCIM News, #26, July 1996

What are the CCPs?

From Paul's paper, active CCPs at the time (1996) were:

- **CCP1** - The Electronic Structure of Molecules
- **CCP2** - Continuum States of Atoms and Molecules
- **CCP3** - Computational Studies of Surfaces
- **CCP4** - Protein Crystallography
- **CCP5** - Computer Simulation of Condensed Phases
- **CCP6** - Heavy Particle Dynamics
- **CCP7** - Analysis of Astronomical Spectra
- **CCP9** - Electronic Structure of Solids
- **CCP11** - Biosequence and Structure Analysis
- **CCP12** - High Performance Computational Engineering
- **CCP13** - Fibre Diffraction
- **CCP14** - Powder and Single Crystal Diffraction.

Science and Technology Facilities Council
Scientific Computing

Each CCP had a Chair and a Working Group which sets the scientific agenda. Note, CCP8 (Nuclear Structure) and CCP10 (Plasma Physics) were no longer operational at this point.

When do the CCPs begin?

So, what contributed to the creation of the first CCP?

To answer this question, we need to look further back and understand how things were changing with computers, society, and what was happening with universities,……..

Basically, we need to look at key societal changes starting in the 1960s…..and the corresponding technological advances in computers

Just over 60 years ago, an important report came out that led to a major change in the university landscape. The Robbins Report[2] (author Sir Lionel Robbins), published late 1963

Another significant report was the Flowers Report[3] (1966) discussing computers for research in the UK – this led to the creation of regional computer centres

Science and Technology Facilities Council

Scientific Computing

[2]*Committee on Higher Education (23 September 1963), Higher education: report of the Committee appointed by the Prime Minister under the Chairmanship of Lord Robbins 1961–63*

[3]*Report* of a *joint working group* on *computers for research, 1966*

Timeline of key moments

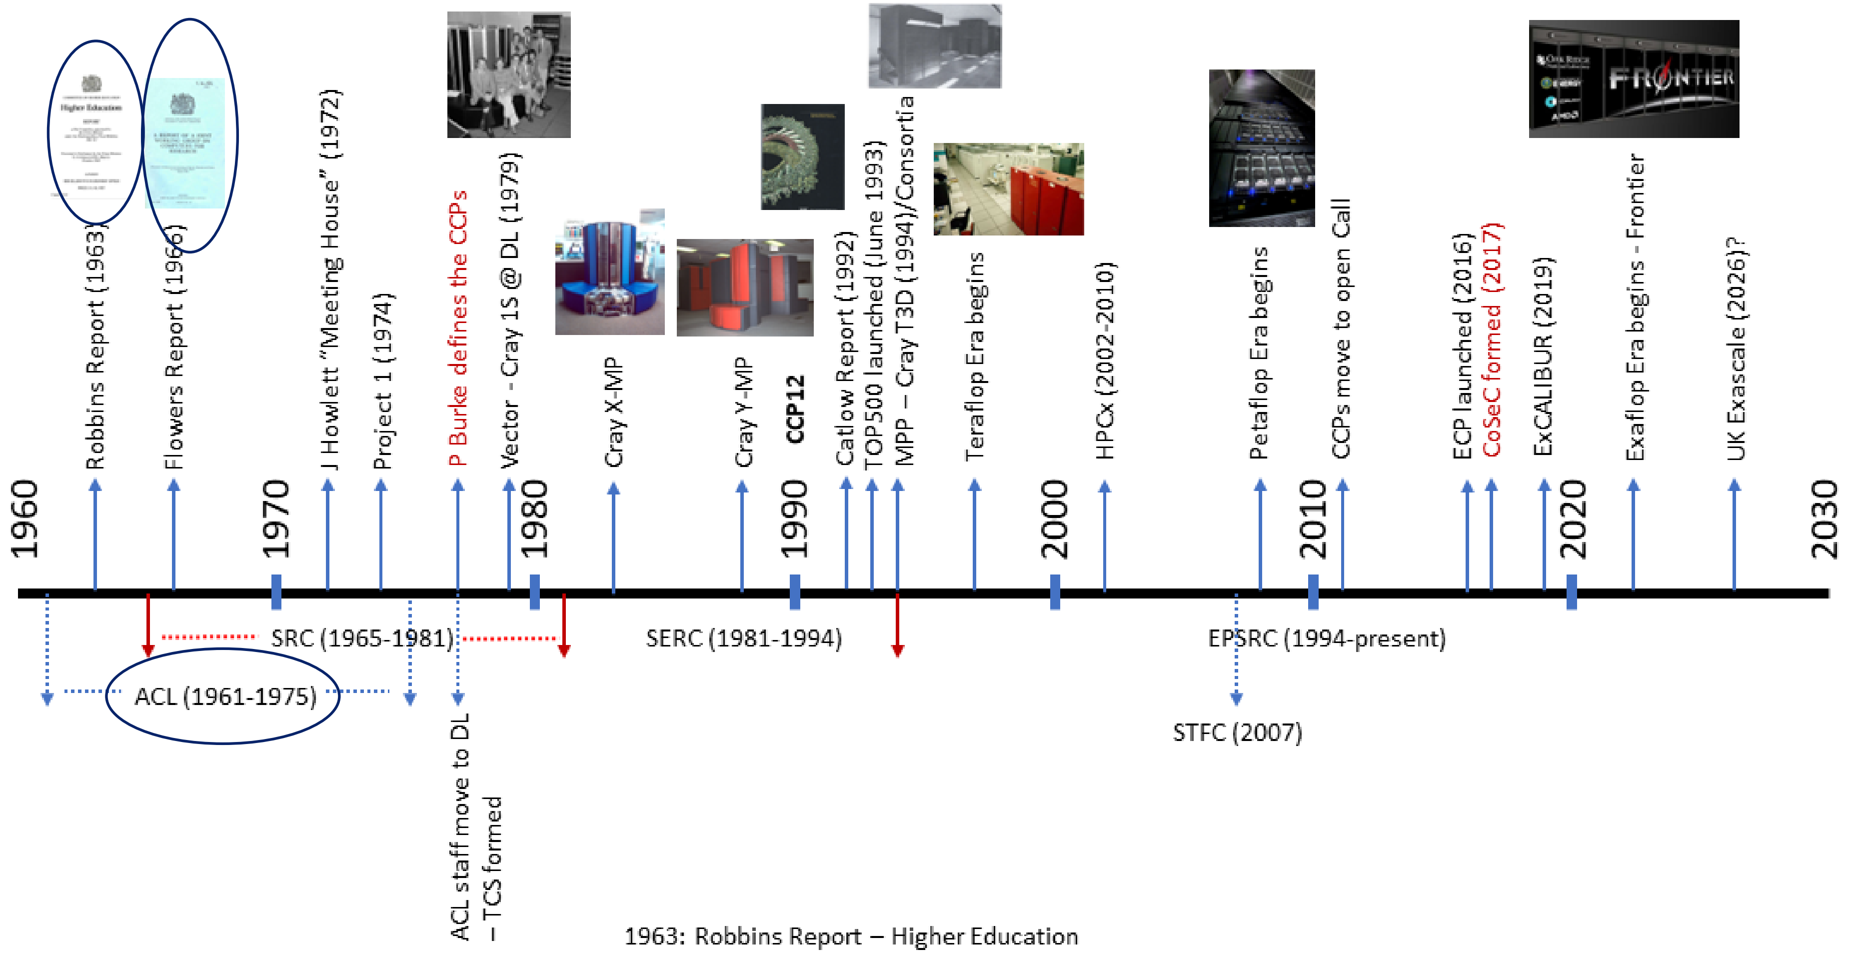

1963: Robbins Report – Higher Education

1966: Flowers Report – A Report of a Joint Working Group on Computers for Research

1974: Project 1: Correlation in Molecular Wave Functions

1992: Catlow Report - Research Requirements for High Performance Computing (SERC, September)

Scientific Cor

The Robbins Report: a social change

The Robbins Report is perhaps most famous for ‘the Robbins principle’ which underscored all its recommendations. This states: ‘courses of higher education should be available for all those who are qualified by ability and attainment to pursue them and who wish to do so’, Hillman (2023)[4]

Full-time higher education for young people was very different in the 1960s. Far fewer universities (~25) and a small range of other higher education providers. Student numbers were around 238,000 (63/64), with very few “working class” students, and very few female students….as noted by Barr[5]

“Only about four in every 100 young people entered full-time courses at university. Only 1 per cent of working-class girls and 3 per cent of working-class boys went on to full-time degree level courses.”

The report projected an increase to 558,000 (80/81) based on expansion of the universities (already underway when the report came out).

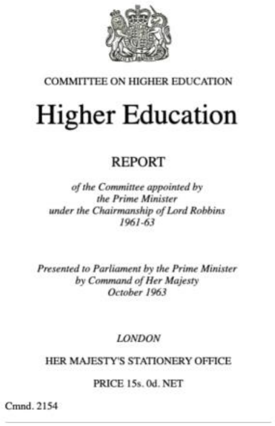

COMMITTEE ON HIGHER EDUCATION

Higher Education

REPORT

of the Committee appointed by the Prime Minister under the Chairmanship of Lord Robbins 1961-63

Presented to Parliament by the Prime Minister by Command of Her Majesty October 1963

LONDON

HER MAJESTY'S STATIONERY OFFICE

PRICE 15s. 0d. NET

Cmnd. 2154

The report also attracted criticism with letters being sent to The Times, e.g.

Many in universities were convinced that they were already scraping the bottom of the barrel – “more means worse”, to quote a notion popular at the time – and that any further expansion would spell disaster.

UKRI Science and Technology Facilities Council

Scientific Computing

[4]Nick Hillman, The Robbins Report at 60: Essential facts for policymakers today, (2023)

[5]Nicholas Barr, Shaping higher education: 50 years after Robbins (2013)

The Flowers Report

The use of computers was increasing in many areas, but the Flowers Report (1966) was looking at their impact in science. To address the growing demand for access to computers, one recommendation was to create a small number of “Regional Computer Centres” for UK universities.

For example, Manchester established its Regional Computer Centre in 1969 in direct response to the report.

The Working Group proposed buying 10 – 15 large US machines about 3 – 100 times more powerful than the IBM 7090 (Stretch) to meet expected demand.

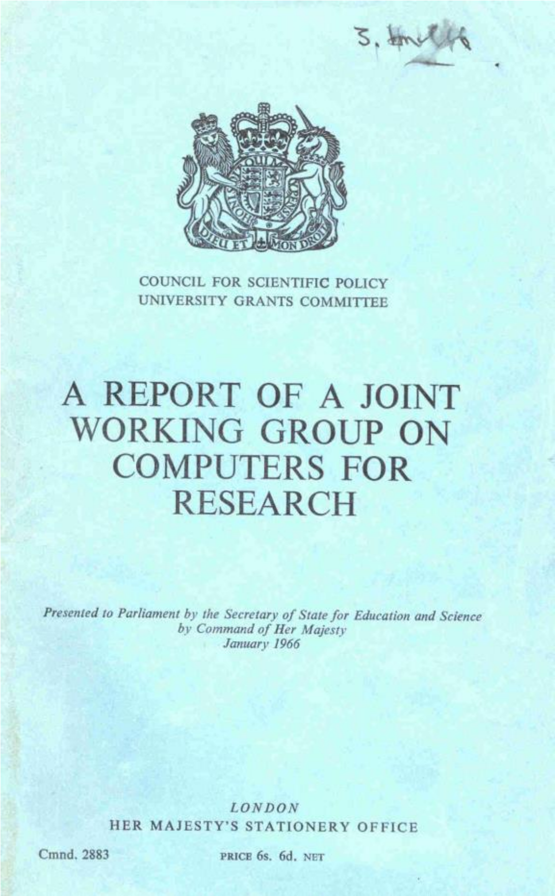

The Atlas Computing Laboratory

Atlas Computing Laboratory (ACL, 1961 – 1975) – a national computing centre, operational by 1964, contributing to OS, compilers, networking etc. and also software development for mathematics, chemistry…..

It housed, and was named after, Ferranti's Atlas computer, designed at Manchester University, and was a British machine, probably the most powerful in the world for a short time…..

Computational chemistry was a major driver of software developments, and some key names start to appear, namely Vic Saunders, who wrote code for Gaussian integral evaluations in Fortran IV, joined ACL in 1970. However, as noted by Smith and Sutcliffe[6], by 1970s several other groups, at Manchester, Oxford, Cambridge etc., were active using modest university resources. They also used facilities at the ACL, although numbers involved were actually fairly small and efforts quite fragmented. Another chemistry stalwart and important CCP figure joined ACL in 1972, Martyn Guest.

In 1975, the ACL was closed with some staff moving to Daresbury to join TCS…..

[6] SJ Smith and BT Sutcliffe, "The Development of Computational Chemistry in the United Kingdom", Reviews in Computational Chemistry, V10, 1997

Timeline of key moments

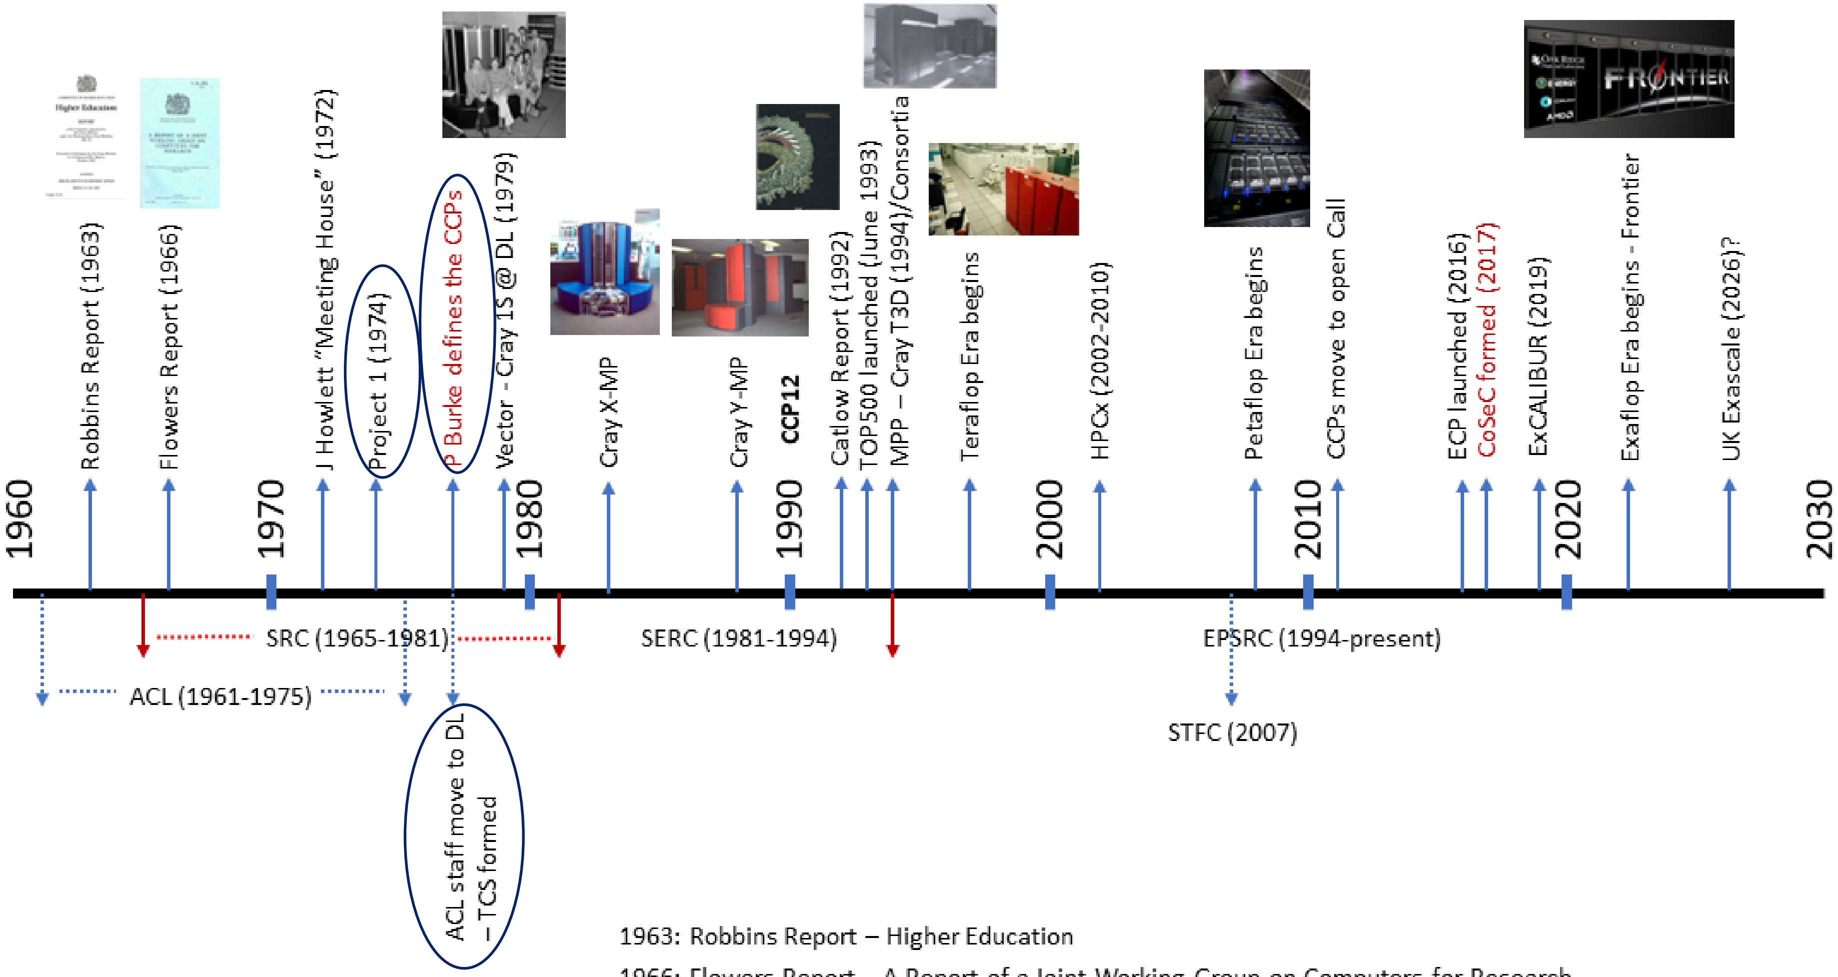

1963: Robbins Report – Higher Education

1966: Flowers Report – A Report of a Joint Working Group on Computers for Research

1974: Project 1: Correlation in Molecular Wave Functions

1992: Catlow Report - Research Requirements for High Performance Computing (SERC, September)

Scientific Cor

The 1970s: The Concept Emerges…

The primary funding body responsible for publicly funded scientific and engineering research at the time was the Science Research Council (SRC, 1965-1981).

High-level talks between ACL and its Science Board led to the idea of holding a "Meeting House", a term attributed to Prof Ron Mason, to bring together a small number (~4) of distinguished scientists to discuss a particular topic.

The idea was to meet at the ACL – Jack Howlett (ACL Director) proposed to Science Board that the first meeting focus on "molecular correlation errors in theories which surpass the Hartree-Fock theory in accuracy". Four leading professors were invited: Bransden (Durham), Burke (QUB), Coulson (Oxford), and McWeeny (Sheffield). Due to ill health, Coulson was replaced by Murrel (Sussex).

From this meeting, held Feb 6th 1974, came Project 1 on "Correlations in Molecular Wave Functions"

The 1970s: The Concept Emerges…

Project 1 had the go-ahead. A Working Group (renamed as a Steering Panel) met in March 1974 – agreed a PDRA be supported.

A further meeting in May 1974 discussed the work the PDRA would be involved in, and William Rodwell was appointed in summer of 1974.

The Working Group next met in October 1975 to discuss outputs. A key topic related to codes and their machine dependence. Porting of the s/w was therefore important. This was also motivated by the impact of the Flowers Report and creation of Regional Computing Centres.

A further meeting, held Dec 4th 1975, the next phase of work was decided which was on Valence Bond Techniques – this was recommended to Science Board and approved for 3 years at a spring 1976 meeting.

BUT, external changes were happening which shifted computational chemistry from ACL to Daresbury.

The 1970s: The Concept Takes Shape

The move to DL coincided with Phil Burke being director of the Theory and Computational Science Division. At the time, lots happening experimentally at DL with the Nuclear Structure Facility (NSF) and SRS.

Phil was highly regarded and very influential and organised a meeting of the Steering Panel at DL in October 1977. From this, the Collaborative Computational Projects were proposed!

Following ACL staff moving to DL, at a Daresbury Study Weekend in Dec 1977, Phil Burke introduces and defines the CCP programme

As a result, Project 1 was renamed as CCP1 and others soon followed, notably CCP2 (Continuum States in Atoms and Molecules, 1978), CCP3 (Surface Science, 1978), CCP4 (Protein Crystallography, 1979), CCP5 (Monte Carlo and Molecular Dynamics, 1979), and CCP6 (Heavy Particle Dynamics, 1979).

Timeline of key moments

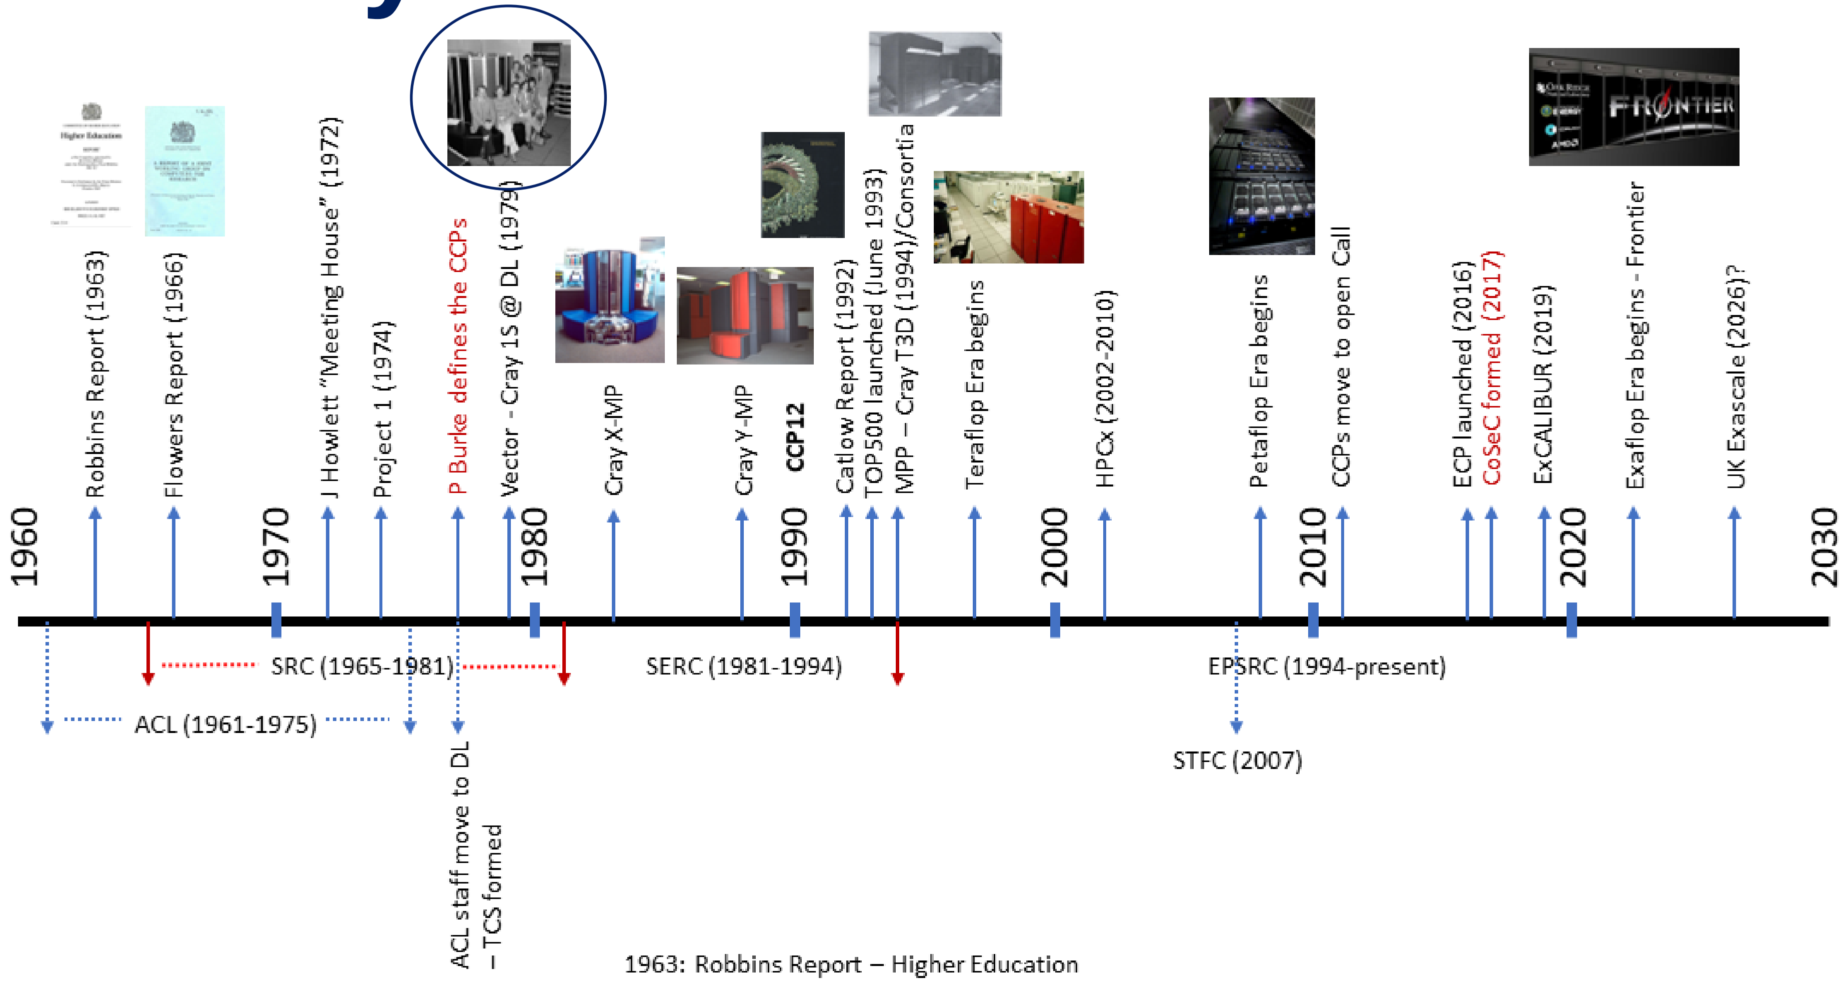

1963: Robbins Report – Higher Education

1966: Flowers Report – A Report of a Joint Working Group on Computers for Research

1974: Project 1: Correlation in Molecular Wave Functions

1992: Catlow Report - Research Requirements for High Performance Computing (SERC, September)

Scientific Cor

Enter the Cray 1

In 1979, the Cray 1S was delivered to Daresbury. It can be argued that this was a pivotal point in computational strategy and the use/benefit of HPC and supercomputers.

Chemistry led the way with Vic Saunders and Martyn Guest porting all ATMOL codes within a month. A speed-up of 16.2, compared to the CDC 7600 (at Manchester), was reported.

CCP1 was in an ideal position to quickly embrace this new technology with the community reaping the benefits very quickly. Sadly, not true for engineering…..

Enter the 1980s – the vector era

Vector supercomputers dominated in the 1980s. In 1983, the Cray X-MP was released and, in 1988, this was followed by the Cray Y-MP.

Computational science was making great progress with advances in hardware and algorithms (a great overview of the advances is given by Kaufmann & Smarr).

Although vector computers dominated this decade, the use of parallel (distributed memory) hardware was taking place e.g. Intel iPSC/1 (1985), iPSC/2 (1987), and the Intel iPSC/860 (1990). Despite Amdahl's Law, there was growing interest in parallel computing….

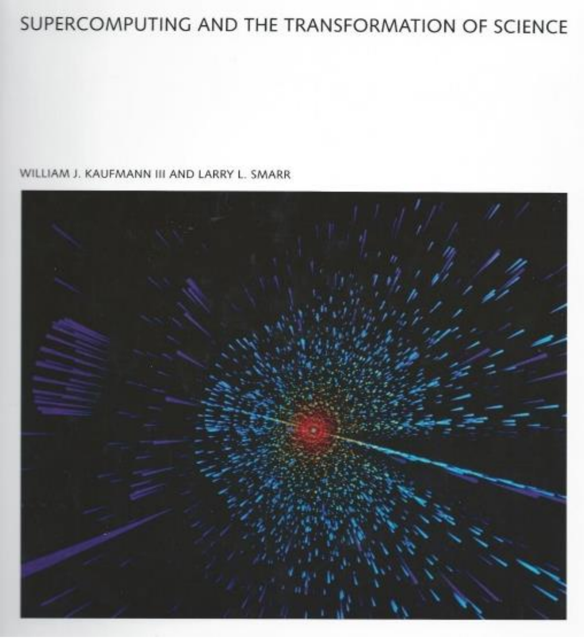

Supercomputing and the Transformation of Science
William J Kaufmann & Larry L Smarr

UKRI Science and Technology Facilities Council
Scientific Computing

A Brief Digression - Emergence of Scientific Computing

Computational Science and Engineering (CSE) has now emerged as the *science of simulation* - taking the “laws” of physics (e.g. Schrodinger, Maxwell, Newton) and creating **mathematical programs** to solve the governing equations of these laws.

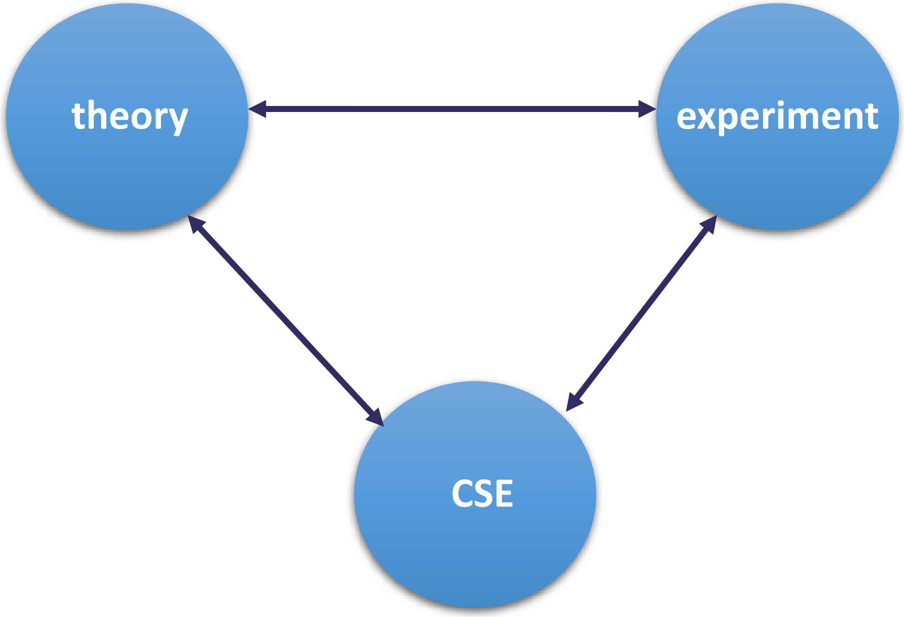

It is now regarded as the **third pillar** of research and its importance is increasing all the time due to developments in algorithms and computer power.

It was NOT the case in the early years where conflict with experimentalists and theory was common

UKRI Science and Technology Facilities Council

Scientific Computing

A Brief Digression - Emergence of Scientific Computing

Computing was still rapidly evolving, along with the programs being developed.

For some, their use was clear…Kawaguti (1953) indicates what it was like in the 1950s….. “*The numerical integration in this study took about one year and a half with twenty working hours every week, with a considerable amount of labour and endurance*”

In the paper by Smith et al., they report that *One well-known theoretical physicist is reported to have referred to computational scientists as “hairy-handed mahouts operating their elephantine computers”*

In CFD, there was some barbed banter between experimentalists and modellers, the latter joking that wind tunnels could be used to store their printouts. This was later turned on its head with an article appearing in 2096 entitled: Will the Wind Tunnel Replace the Computer? attributed to John Coopersmith. It jokes that the most accurate code available, Flo-1234.5, is so complex and expensive, it has never been run…..

Enough to say that computing today is an indispensable part of most studies.

Timeline of key moments

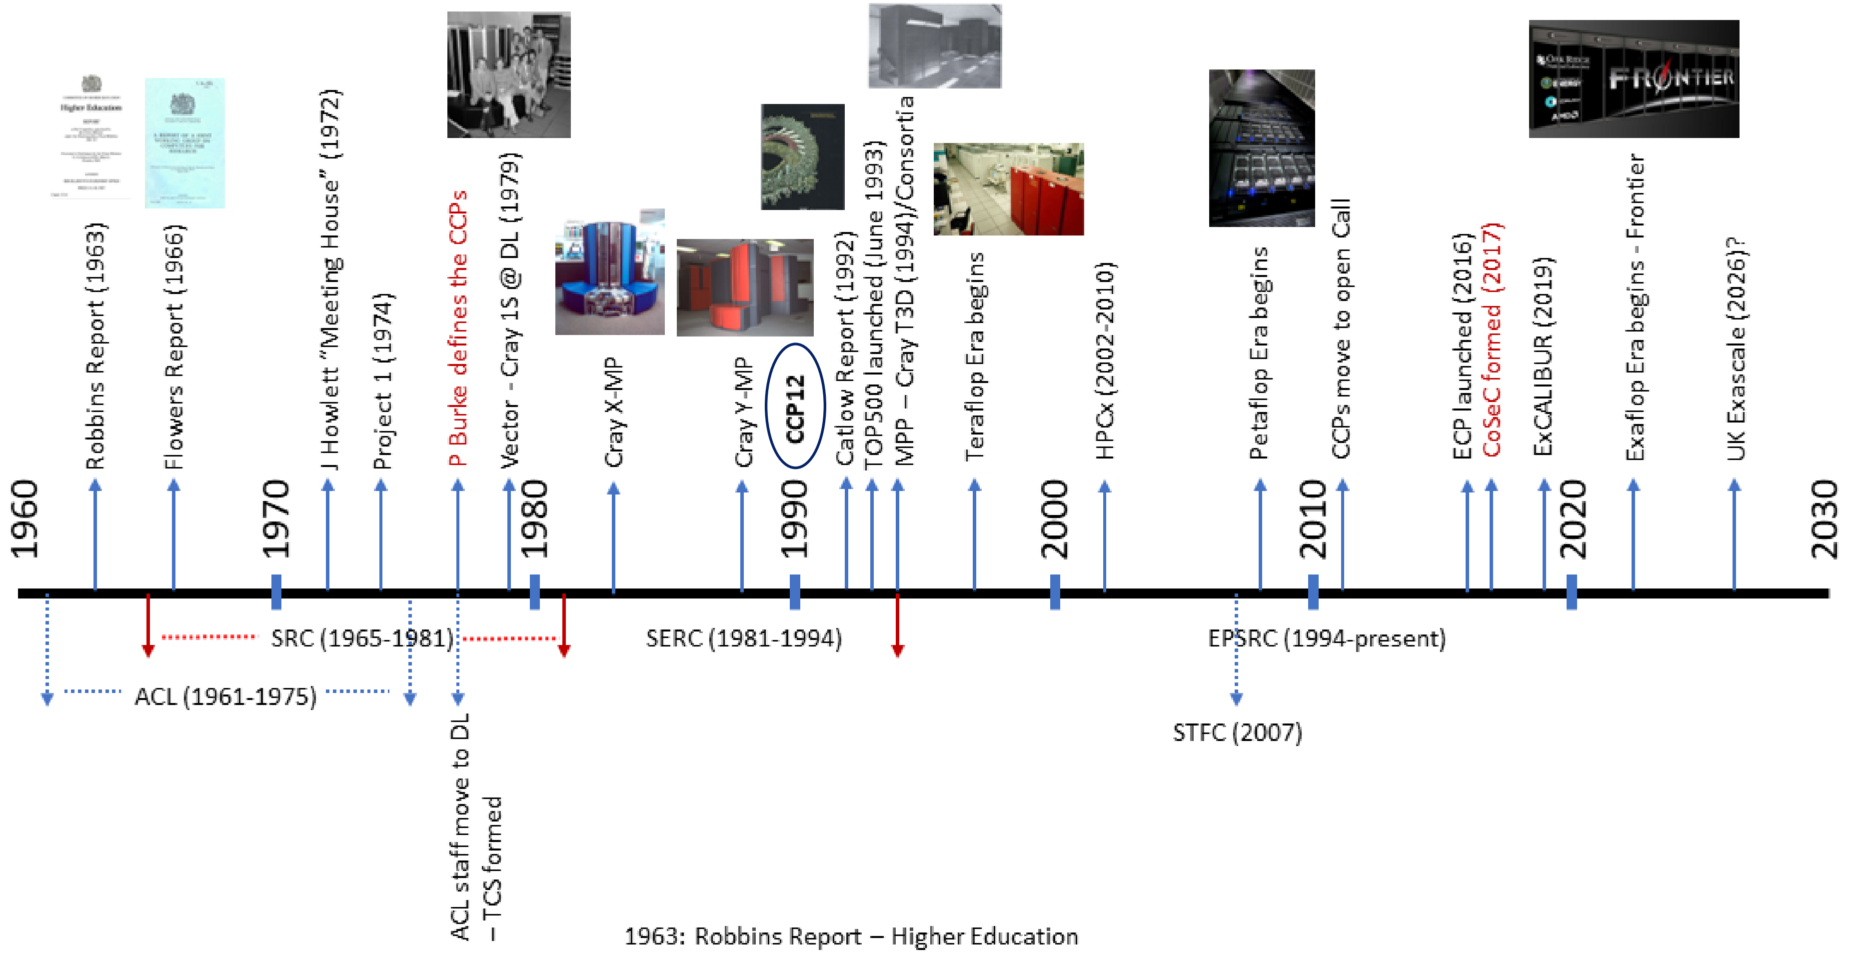

1963: Robbins Report – Higher Education

1966: Flowers Report – A Report of a Joint Working Group on Computers for Research

1974: Project 1: Correlation in Molecular Wave Functions

1992: Catlow Report - Research Requirements for High Performance Computing (SERC, September)

Sc
Te
Fa
Scientific Cor

Enter the 1990s – the emergence of CCP12

In 1989/1990, Engineering Board (EB) recognised the poor take-up of HPC by its engineering researchers (~4% Y-MP), especially CFD

The success of the CCPs provided the route to follow and CCP12 started in 1990 to investigate parallel and novel architecture computing in CFD. The initial desire was to create a flagship code for the community that could take full advantage of national HPC facilities.

To further encourage engineering to make use of modern computers, EB acquired a number of computers to distribute to HEIs: 2 x Meiko i860 (Bristol, Bath), 2 x HP/Apollo shared-memory systems (Brunel, Oxford), and 1 x Alliant FX/2808 (UMIST). These were allocated through the Parallel Computing Hardware Initiative, coordinated through DL

Science and Technology Facilities Council

Scientific Computing

Was CCP12 a success? In 2007, engineering represented 37% of HPCx

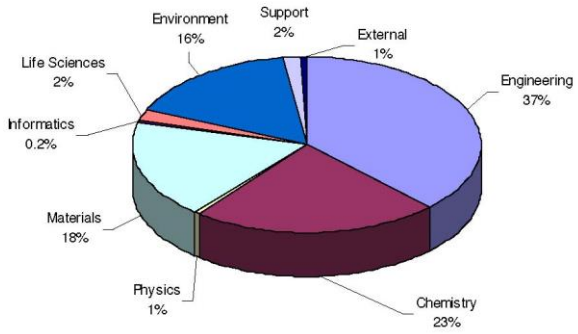

Enter the 1990s – the MPP era

In 1992, Prof R Catlow produced a report that opened the way to massively parallel computing and the procurement of a 256-processor Cray T3D in 1994.

Many of the CCPs were getting familiar with distributed memory systems making use of the Meiko, Intel iPSC/2, and Intel iPSC/860 hypercube at DL.

In reality, moving to distributed memory parallel machines was challenging with limited access to standards (MPI 1.0 was released June 1994). However, most codes associated with the CCPs were soon up and running. To manage access to the T3D, "consortia" were established (today's HEC are derived from this model)

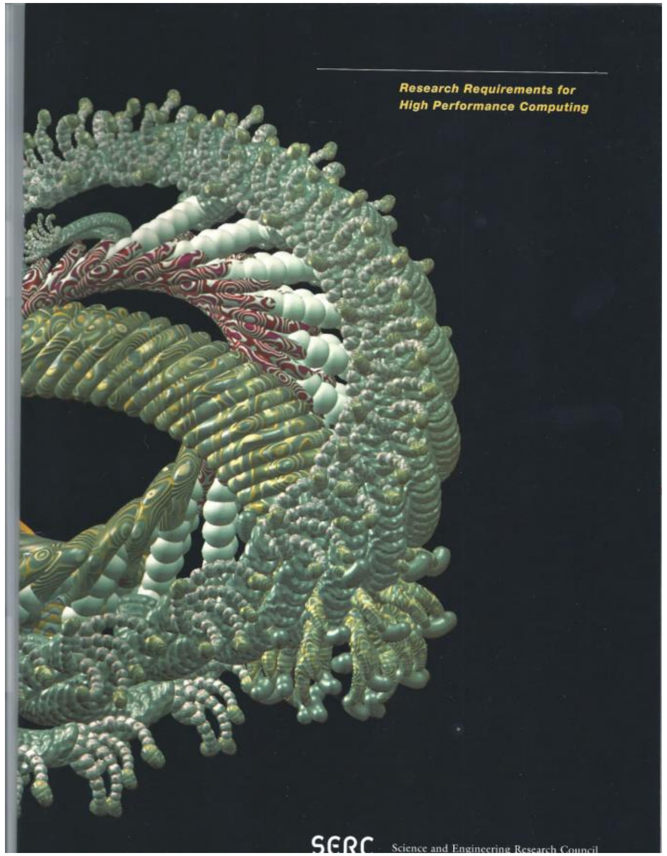

Research Requirement for High Performance Computing
Prof Richard Catlow FRS

Science and Technology Facilities Council
Scientific Computing

Enter the 1990s – the MPP era

The 1990s saw some significant developments. For example, the first Teraflop computer at Sandia (ASCI Red) and the first Gigahertz processor

The Accelerated Strategic Computing Initiative (ASCI) , established in 1995, saw major investment in HPC

With ASCI, the role that distributed memory parallel machines was to play was dominant and vector machines faded into the past….

Timeline of key moments

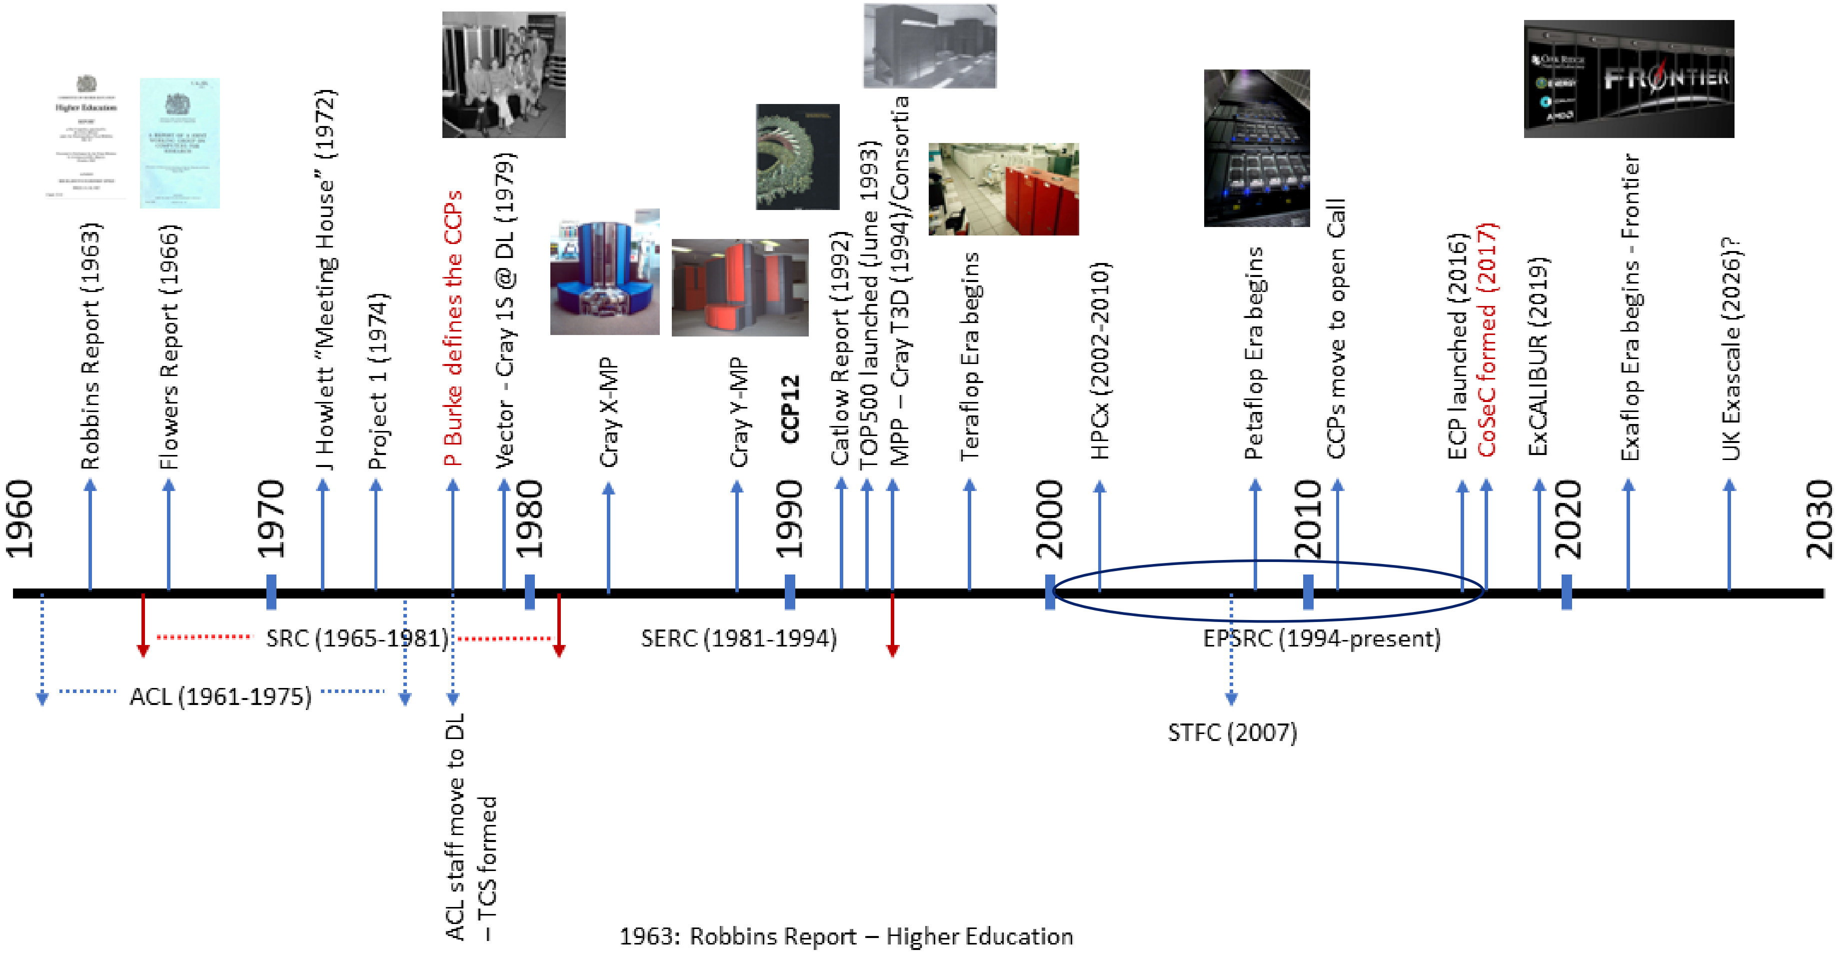

1963: Robbins Report – Higher Education

1966: Flowers Report – A Report of a Joint Working Group on Computers for Research

1974: Project 1: Correlation in Molecular Wave Functions

1992: Catlow Report - Research Requirements for High Performance Computing (SERC, September)

Sc
Te
Fa
Scientific Cor

Enter the 2000s….

The 2000s saw major developments leading to the first petaflop calculation by IBM's Roadrunner. Although a great milestone, the machine proved difficult to program.

During this period, the CCPs continued to evolve to work with the changing demands of the communities. Within engineering, the consortia model thrived e.g. the UK Turbulence Consortium has been funded since 1994, working closely with CCP12 and CoSeC, and continues to grow its membership with codes tackling turbulence on Tier 2, Tier 1 (ARCHER2), and Tier 0 (EU/US facilities).

Many of the CCPs are involved in exascale initiatives (ExCALIBUR) and the challenge of programming hybrid architectures with complex memory hierarchies remains formidable. With the consolidation of CCP support into CoSeC in 2017, there is huge potential to make a significant impact with UK science.

Science and Technology Facilities Council

CoSeC: Now and the Future

https://cosec.stfc.ac.uk

https://www.linkedin.com/in/cosec-ukri

@CoSeC_community

Dr Barbara Montanari,
Dr Stephen Longshaw,
STFC Scientific Computing,
Daresbury Laboratory

CoSeC
Computational Science Centre for Research Communities

Origins of CoSeC

CoSeC
Computational Science Centre for Research Communities
www.cosec.stfc.ac.uk

The **Co**mputational **S**cience C**e**ntre for Research **C**ommunities (CoSeC) was formed in 2017 to bring together the computational science support for CCPs and High-End Computing Consortia into a unified programme across four research councils to:

Scientific Software

- Development and maintenance
- Ensuring relevance for evolving scientific challenges
- Keeping up to date with changing computational hardware

Widen Participation

- Increase exploitation of scientific methods and codes
- Provide training
- Develop and maintain close scientific collaboration

Support Collaboration

- Help to support the coordination of research communities
- Broaden and strengthen UK science

Career Paths

- Provide and create career paths for computational scientists
- Create career development opportunities for research technology professionals

Widen Engagement

- Work with UK and international communities
- Ensuring a joined-up approach for UK computational science

UKRI Science and Technology Facilities Council
Scientific Computing

UKRI Engineering and Physical Sciences Research Council

UKRI Biotechnology and Biological Sciences Research Council

UKRI Medical Research Council

Evolution of CoSeC

CoSeC's programme continuously evolves as the communities supported change through periodic calls for renewal. Over the last 15 years, the evolution has been towards supporting:

- More communities
- More software packages
- More research areas
- Research across more length scales
- More stages in the research pipeline

CoSeC Communities Today

CoSeC Computational Science Centre for Research Communities
www.cosec.stfc.ac.uk

13 CCPs and **6 HECs** covering:

- Materials Science
- Tomographic Imaging
- Computational Engineering
- Biological Science
- Quantum computing

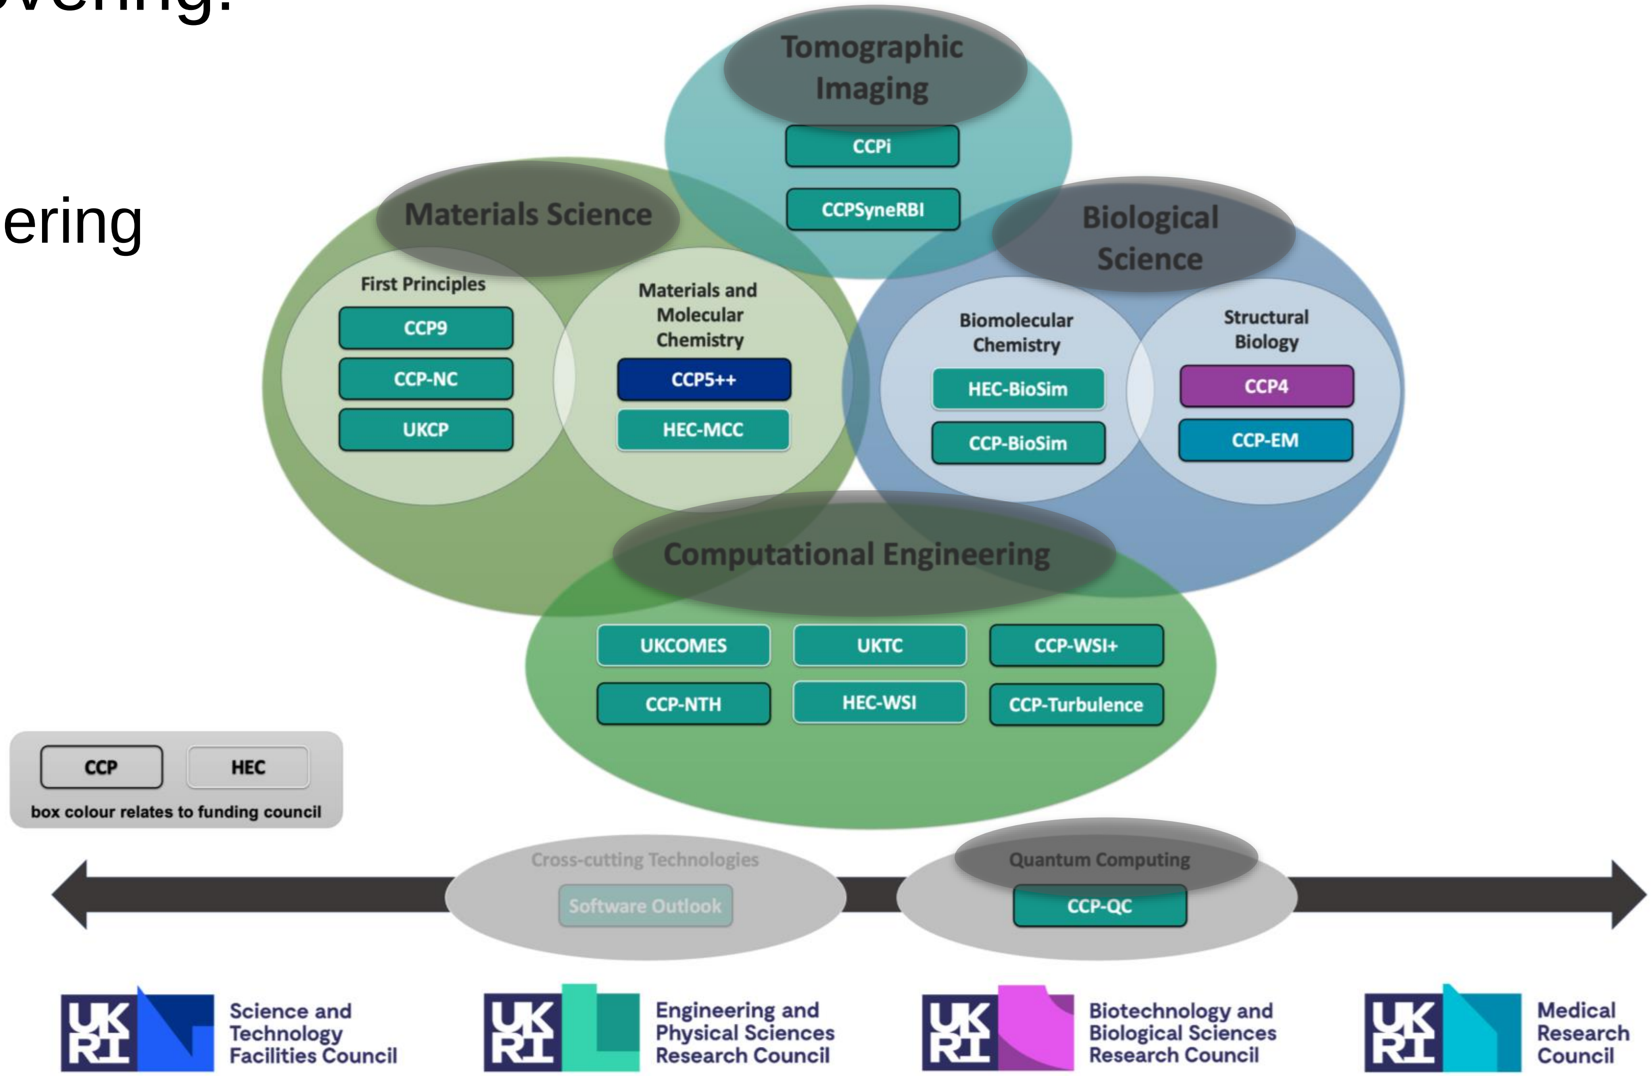

UKRI Science and Technology Facilities Council
Scientific Computing

Case Study: Structural Biology Within the CCPs and CoSeC

- Protein Data Bank (PDB): from 7 structures in 1971 to 200,000 in 2023!
- PDB has led to the AlphaFold Database and Meta AI's ESMFold Atlas
- This has been facilitated by
  - Continually evolving suites of software (CCP4, CCPN, CCP-EM, and others)
  - By a sense of a common goal for the overall community, fostered largely by the CCPs

UKRI Science and Technology Facilities Council

Scientific Computing

UKRI Biotechnology and Biological Sciences Research Council

https://www.deepmind.com/research/highlighted-research/alphafold/timeline-of-a-breakthrough

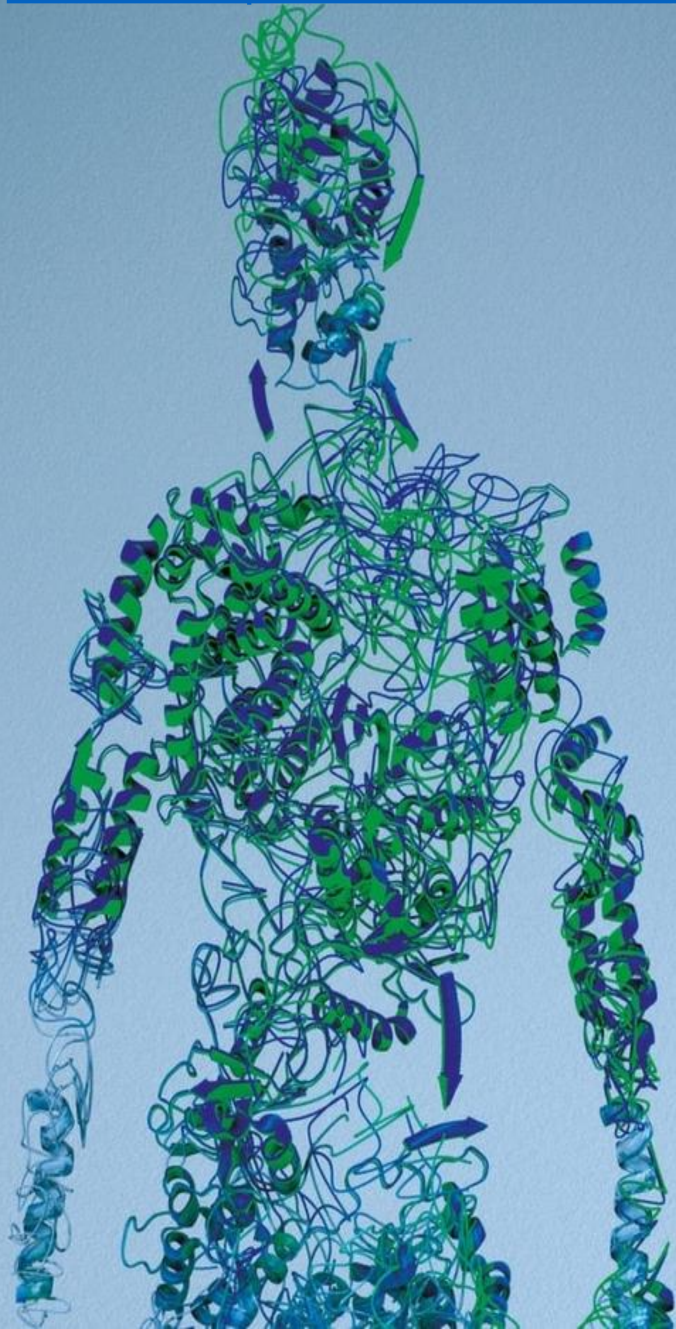

Impact Activities

CoSeC
Computational Science Centre for Research Communities
www.cosec.stfc.ac.uk

- CoSeC’s scientists and research technology professionals both enable and perform transformative research
- The programme celebrates science in a number of ways

Case Studies

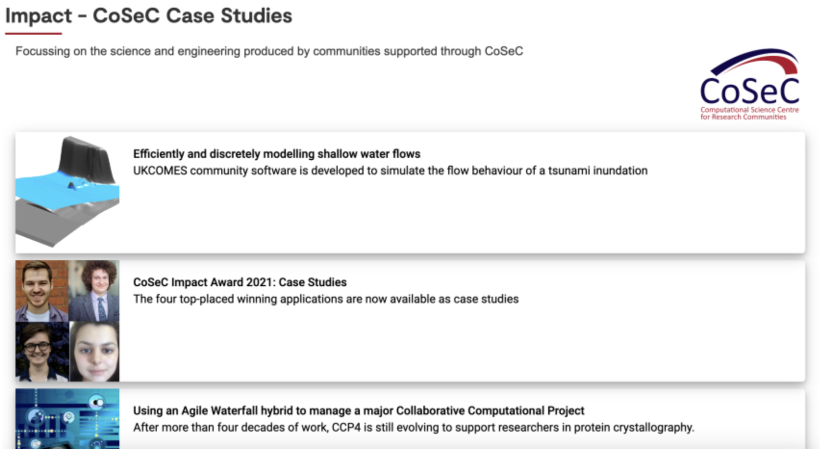
Impact - CoSeC Case Studies

Focussing on the science and engineering produced by communities supported through CoSeC

CoSeC Computational Science Centre for Research Communities

Efficiently and discretely modelling shallow water flows
UKCOMES community software is developed to simulate the flow behaviour of a tsunami inundation

CoSeC Impact Award 2021: Case Studies
The four top-placed winning applications are now available as case studies

Using an Agile Waterfall hybrid to manage a major Collaborative Computational Project
After more than four decades of work, CCP4 is still evolving to support researchers in protein crystallography.

CoSeC Impact Award


Dr Natalie Tatum
2020 Award Winner


Ryan Warr
2021 Award Winner


Claire Delplancke and Catherine Disney
2022 Joint Award Winners

CoSeC Conference

UKRI Science and Technology Facilities Council
Scientific Computing

CoSeC
Computational Science Centre for Research Communities
www.cosec.stfc.ac.uk

UK Research and Innovation

The UKRI Digital Research Infrastructure and CoSeC

Science and Technology Facilities Council
Scientific Computing

The UKRI DRI Landscape

www.cosec.stfc.ac.uk

- UKRI's Digital Research Infrastructure (DRI) programme underpins research and innovation in the UK:
  - The DRI has its own funding to enable **cross-council** activities
  - It is core to delivering key UK infrastructure and associated people, such as the 2023 £1.5bn investment in supercomputing resources around AI and exascale

- Councils within UKRI like EPSRC and NERC also maintain their own DRI programmes complimentary to the UKRI effort:
  - Platform grants for long-term training efforts of research technology professionals
  - Software and infrastructure like the ExCALIBUR exascale programme

Scientific Computing

UKRI DRI's Infrastructure

CoSeC
Computational Science Centre for Research Communities
www.cosec.stfc.ac.uk

Following the Future of Compute review, a £1.5bn investment in a **national AI resource** and an **exascale** supercomputer was announced by the chancellor in the 2023 spring budget:

- The AI Research Resource (AIRR) announced in November
- Federated supercomputers across the UK
- Dawn (Intel) at Cambridge
- Isambard-AI (NVIDIA) at Bristol's National Composite Centre
- Exascale – 1st phase investment in a 250PF machine hosted at Edinburgh.

These investments now need targeted computational science support for UKRI's communities

Science and Technology Facilities Council
Scientific Computing

UK Research and Innovation

Dawn at Cambridge University

The National Composites Centre

CoSeC and the UKRI DRI

CoSeC
Computational Science Centre for Research Communities
www.cosec.stfc.ac.uk

- CoSeC is well-placed to help the DRI achieve some of its goals, including:
  - Development and maintenance of cross-cutting scientific software
  - Collaborative computational science support and domain-specific enablement
  - Enabling cross-cutting training around aspects of computational science
  - A career-focussed development environment for research technology professionals
  - A central hub for cross-cutting activities relating to technologies like AI and quantum

Science and Technology Facilities Council
Scientific Computing
UK Research and Innovation

What's Next for CoSeC?

CoSeC
Computational Science Centre for Research Communities
www.cosec.stfc.ac.uk

- CoSeC is part of STFC and currently delivers strategic elements of EPSRC, BBSRC and MRC's programmes
- It is already **cross-council** and will maintain and expand these activities while looking to grow into new UKRI DRI roles

Maintain → **Create** → **Grow**

Maintain

Work with other councils to support, maintain and develop CCP and HEC like activities

Create

Interact with UKRI DRI to:

- Create new cross-council activities
- Leveraging existing domain-specific knowledge and expertise
- Helping to support new national infrastructure

Grow

Leverage new funding to develop the strategic expertise of CoSeC in areas like:

- AI and Quantum for computational science and engineering
- Environmentally sustainable computational science

UKRI Science and Technology Facilities Council
Scientific Computing

Conclusions

CoSeC
Computational Science Centre for Research Communities
www.cosec.stfc.ac.uk

- The CCP and HEC activities represent one of the UK's longest standing and successful programme of work around computational science.
- The CoSeC programme has been an integral part of this since the start.
- It provides custom-made training, software, algorithmic and methodological developments for the EPSRC, BBSRC and MRC CCP and HEC communities.
- With the advent of the UKRI Digital Research Infrastructure strategy, CoSeC is well placed to expand its remit, drawing on its existing cross-council working practices to identify and enable this successful model across the sciences.

UKRI Science and Technology Facilities Council
Scientific Computing

UKRI Engineering and Physical Sciences Research Council

UKRI Biotechnology and Biological Sciences Research Council

UKRI Medical Research Council

Science and Technology Facilities Council

Questions and Discussion