Structure & Timetabling of Online HPC Courses

Pros and Cons of Different Approaches

David Henty
EPCC Director of Teaching
d.henty@epcc.ed.ac.uk

epcc

THE UNIVERSITY OF EDINBURGH

Reusing this material

epcc

THE UNIVERSITY OF EDINBURGH

Short survey

- go to: www.socrative.com
- Login -> Student Login
- Room Name: **HPCQUIZ**

EPCC

- Founded in 1990 as Edinburgh Parallel Computing Centre
  - A Centre in The University of Edinburgh, Scotland, UK
  - self-funded with over 100 full-time staff
  - relocated to the Bayes Centre in 2018

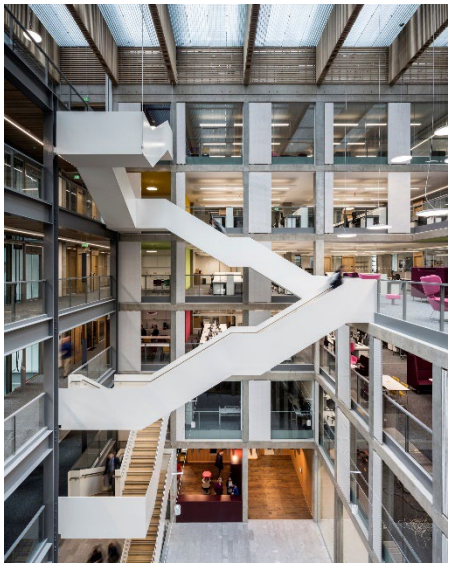

epcc

THE UNIVERSITY OF EDINBURGH

Activities

- Projects in HPC, Data Science & Software Development
  - supporting UK academic research
  - European funding
  - industrial partners
- National Supercomputing services
  - ARCHER2 and Cirrus
- Education and Training
  - Postgraduate Masters Programmes in HPC
  - HPC training courses for national services

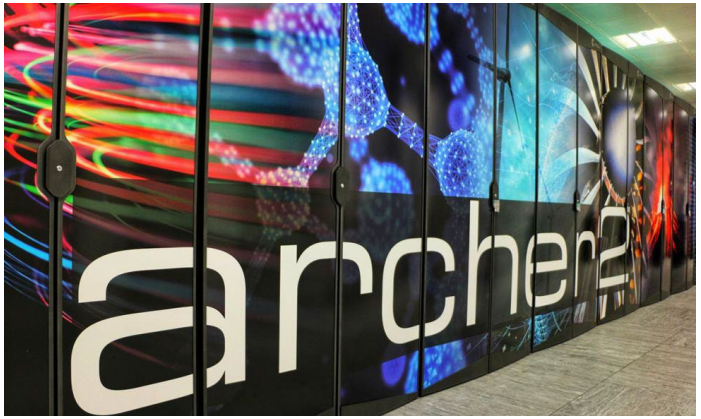

epcc

THE UNIVERSITY OF EDINBURGH

From: Department for Science, Innovation and Technology, UK Research and Innovation, The Rt Hon Michelle Donelan MP, and The Rt Hon Alister Jack MP

Published 9 October 2023

- Edinburgh nominated to host next-generation compute system, 50 times more powerful than our current top-end system
- national facility – one of the world's most powerful – will help unlock major advances in AI, medical research, climate science and clean energy innovation, boosting economic growth
- new exascale system follows AI supercomputer in Bristol in transforming the future of UK science and tech and providing high-skilled jobs

Edinburgh is poised to host a next-generation compute system amongst the fastest in the world, with the potential to revolutionise breakthroughs in artificial intelligence, medicine, and clean low-carbon energy.

The city has been named as the preferred choice to host the new national exascale facility, as the UK government continues to invest in the country's world-leading computing capacity – crucial to the running of modern economies and cutting-edge scientific research.

EPCC's Training Portfolio

- MSc in HPC / HPC with Data Science
  - University of Edinburgh degree since 2001 (Data Science from 2014)
  - currently around 50 on-campus students (12 months full time)
  - 60 students on online versions (3 - 6 years part time since 2020)
  - www.epcc.ed.ac.uk/msc/
- Courses for UK National HPC Services since 1994
  - also a PRACE Training Centre 2012 - 2022
  - around 65 days of training per year (mostly face-to-face, some online)
  - www.archer2.ac.uk/training/
- Online self-study courses
  - PRACE MOOC: www.futurelearn.com/courses/supercomputing/
  - self-service MPI and OpenMP: www.archer2.ac.uk/training/online/

ARCHER2

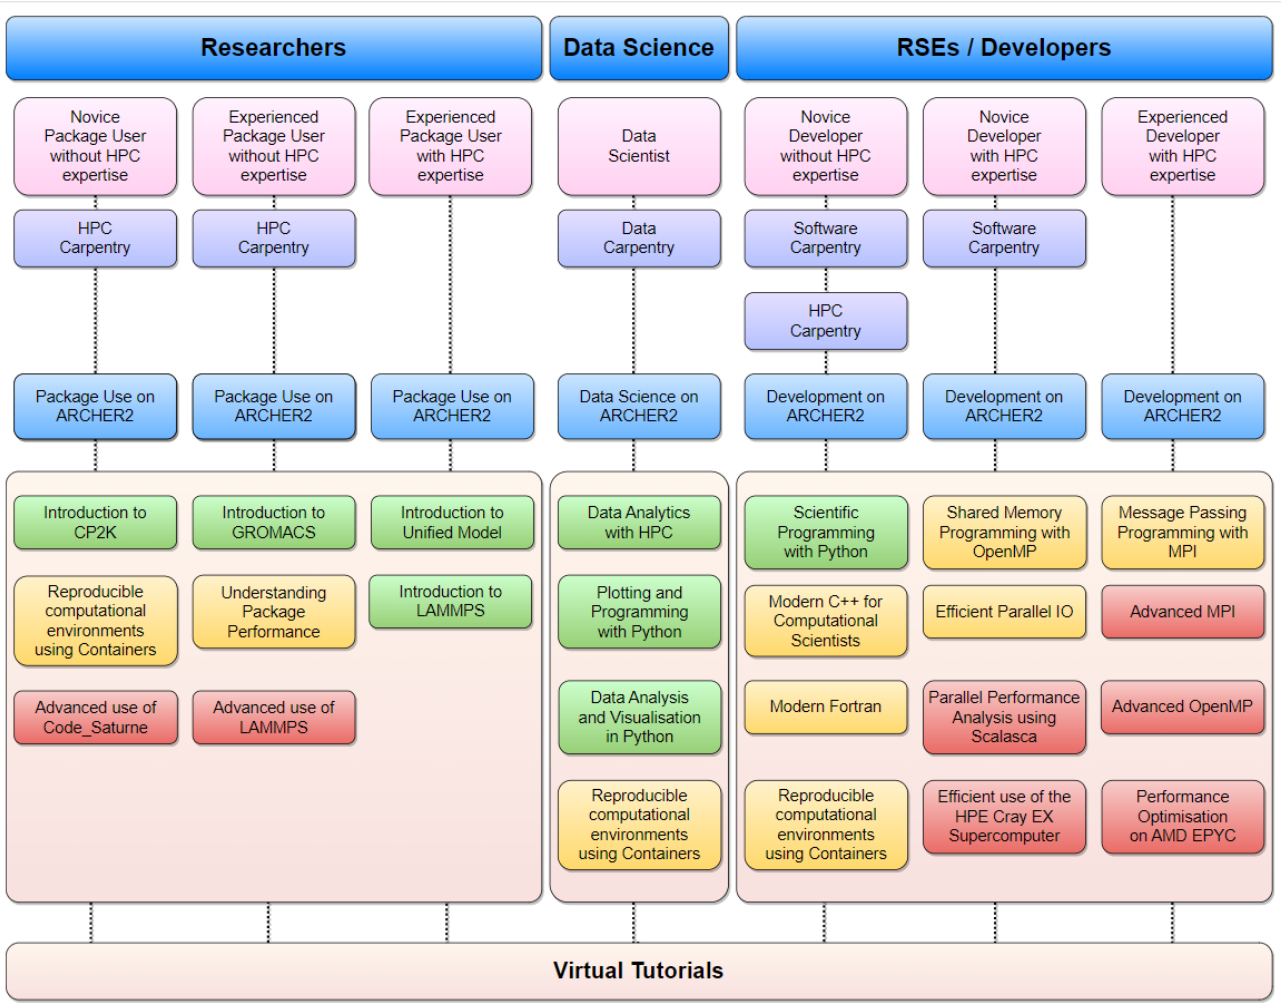
Researchers
Data Science
RSEs / Developers
Novice Package User without HPC expertise
Experienced Package User without HPC expertise
Experienced Package User with HPC expertise
Data Scientist
Novice Developer without HPC expertise
Novice Developer with HPC expertise
Experienced Developer with HPC expertise
HPC Carpentry
HPC Carpentry
Data Carpentry
Software Carpentry
Software Carpentry
HPC Carpentry
Package Use on ARCHER2
Package Use on ARCHER2
Package Use on ARCHER2
Data Science on ARCHER2
Development on ARCHER2
Development on ARCHER2
Development on ARCHER2
Introduction to CP2K
Introduction to GROMACS
Introduction to Unified Model
Reproducible computational environments using Containers
Understanding Package Performance
Introduction to LAMMPS
Advanced use of Code_Saturne
Advanced use of LAMMPS
Data Analytics with HPC
Plotting and Programming with Python
Data Analysis and Visualisation in Python
Reproducible computational environments using Containers
Scientific Programming with Python
Modern C++ for Computational Scientists
Modern Fortran
Reproducible computational environments using Containers
Shared Memory Programming with OpenMP
Efficient Parallel IO
Parallel Performance Analysis using Scalasca
Efficient use of the HPE Cray EX Supercomputer
Message Passing Programming with MPI
Advanced MPI
Advanced OpenMP
Performance Optimisation on AMD EPYC
Virtual Tutorials

EPCC History

- Previous experience of online training
  - MPI On-Line (1997)
  - ARCHER Virtual Tutorials / Webinars (2013-2020)
  - Supercomputing MOOC (2017-2019)
  - Online MPI course (2018)

Minty and Westhead (EPCC,1997)

MPI On-line: A Teaching Environment for MPI

- MPI On-line now provides a specially developed multimedia training environment. The course is presented as a slide show accompanied by a streamed audio sound track.
- Tied into this presentation are animations illustrating the concepts described, such as the different communication modes provided by MPI.
- Understanding of these concepts is reinforced by simulations for the student to study at intervals during the presentation and by the programming exercises which accompany the main sections of the course.

- Supported by programming exercises from a Web browser
  - "like Netscape or Mosaic"
  - used EPIC system
    - similar model to Jupyter Notebooks

ARCHER Virtual Tutorials

- Around 100 webinars on HPC topics for ARCHER service
  - delivered live using Blackboard Collaborate
    - local or remote presenters
  - recorded for ARCHER YouTube channel

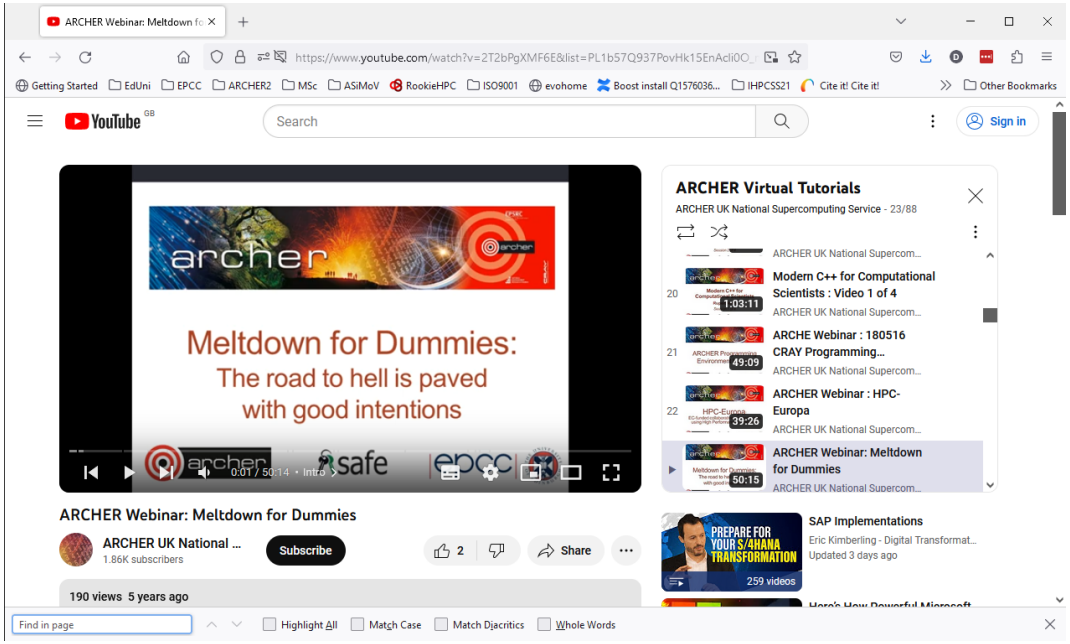

futurelearn.com/courses/supercomputing/

- Free introductory course
  - videos, articles, quizzes, discussion boards etc.
  - certificate upon completion of multiple-choice test (+fee!)
  - largely conceptual
  - tutors contribute to discussions
  - “ask an expert” sessions

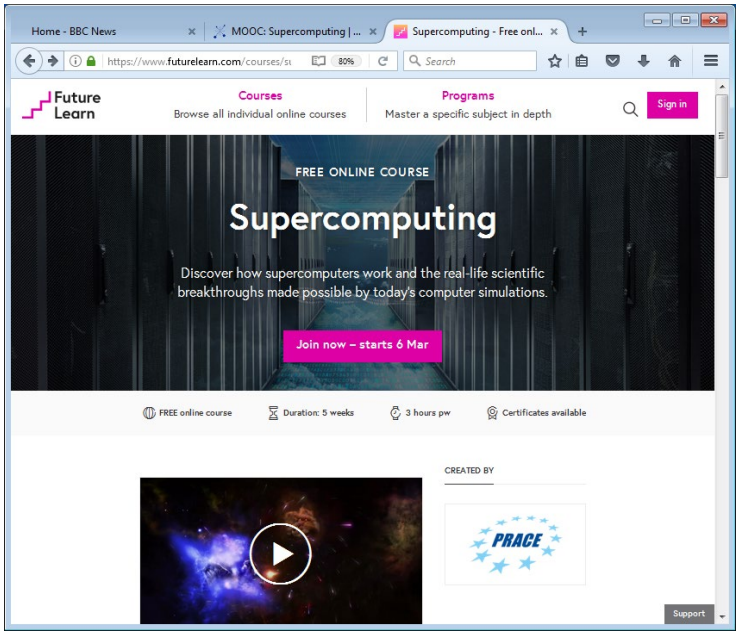

- Five runs from 2017 - 2019
  - typically around 3000 “joiners”
  - typically around 300 “completers”

- Still running in “unfacilitated” mode

Online MPI (based on US XSEDE project)

- XSEDE format
  - around 25 dedicated sites
  - HD video + high quality audio
  - groups attend at each site
  - locals provide basic IT support
  - presenter does technical support
- ARCHER trial over Collaborate (c.f. Virtual Tutorials) in 2018
  - MPI over four Wednesday afternoons
  - ~50 attendees at start, ~25 at the end
  - a few groups at sites but mainly individuals

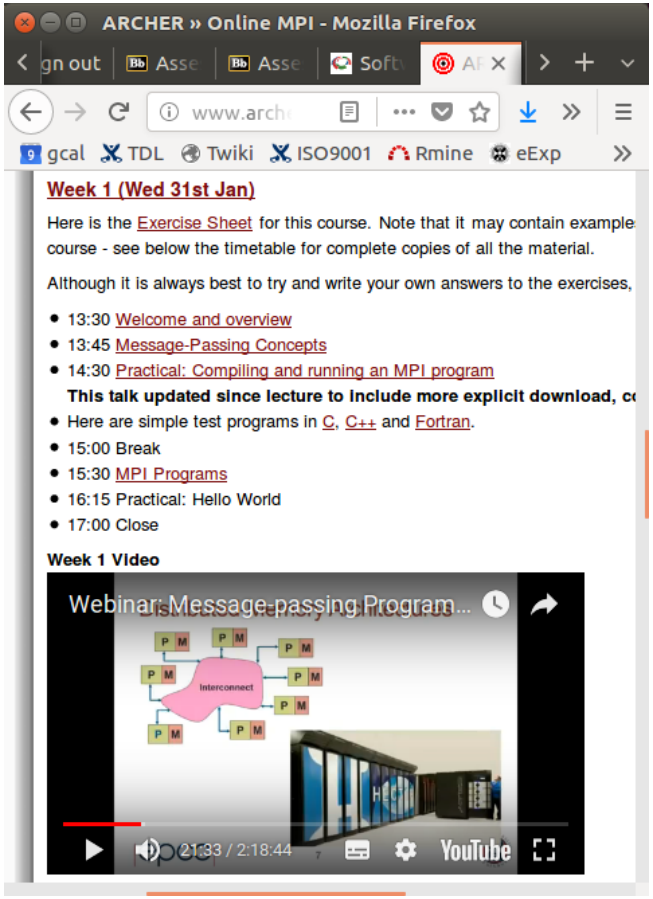

Pre-COVID19

- Already had a range of delivery mechanisms
  - but still mainly face-to-face
- ARCHER or PRACE course
  - two or three days intensive delivery in a training room
  - lecture/practical/lecture/practical ...

epcc

THE UNIVERSITY OF EDINBURGH

EPCC Ecosystem (pre-COVID)

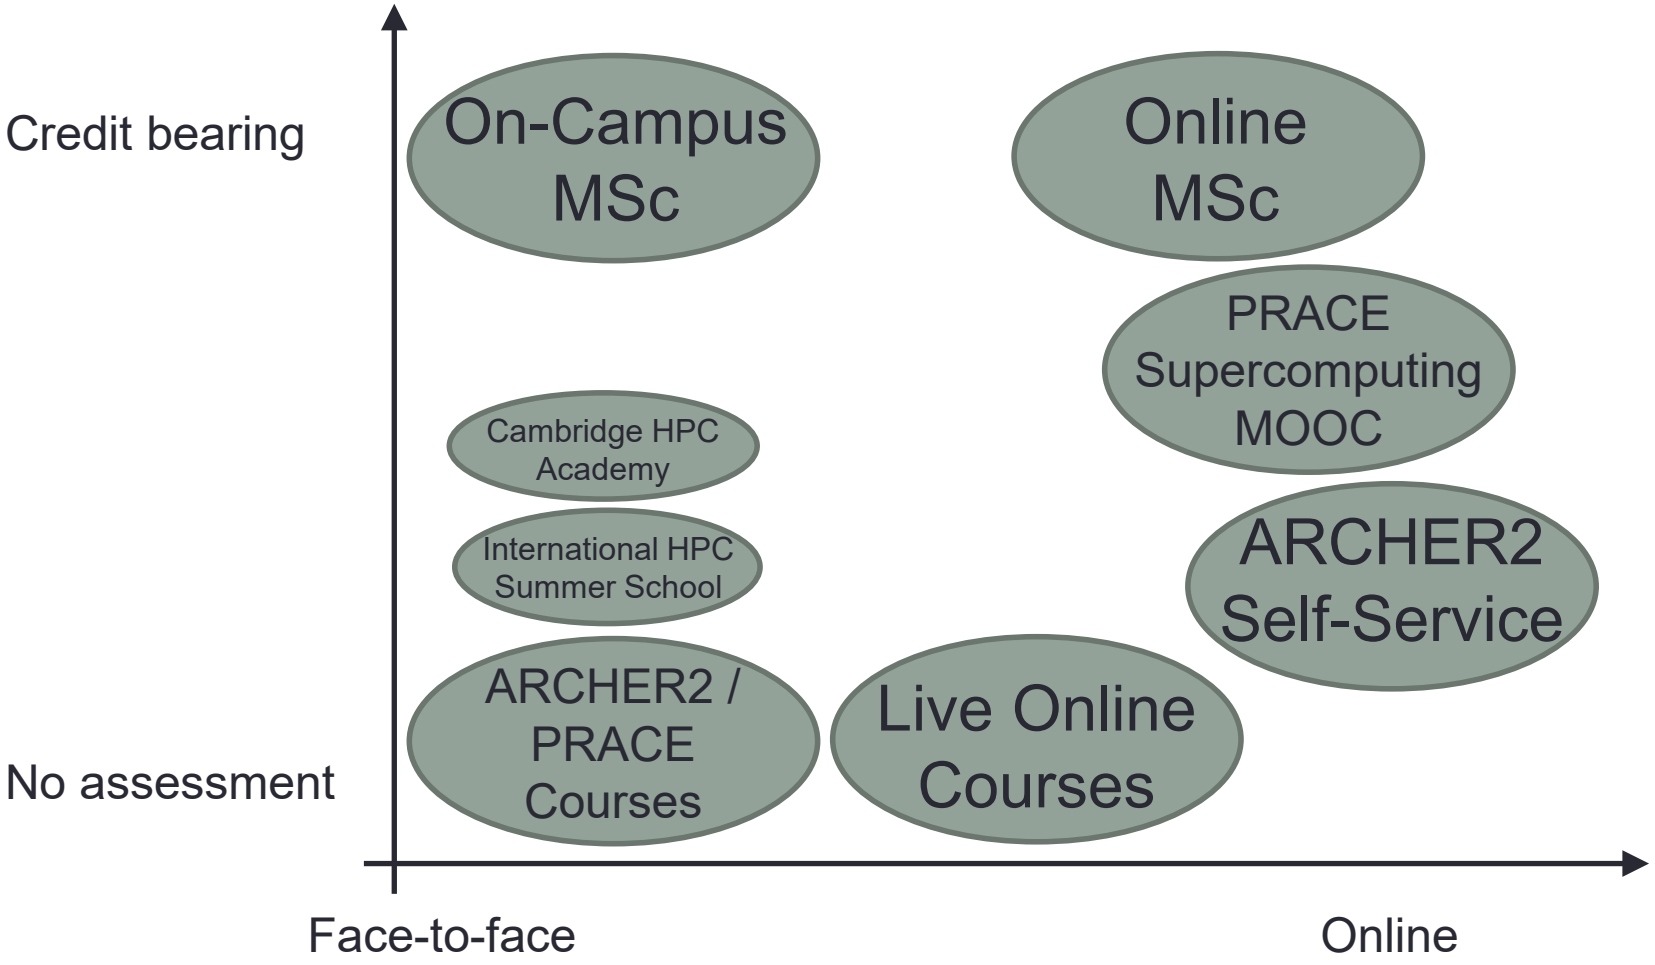

epcc

THE UNIVERSITY OF EDINBURGH

EPCC Ecosystem (during COVID)

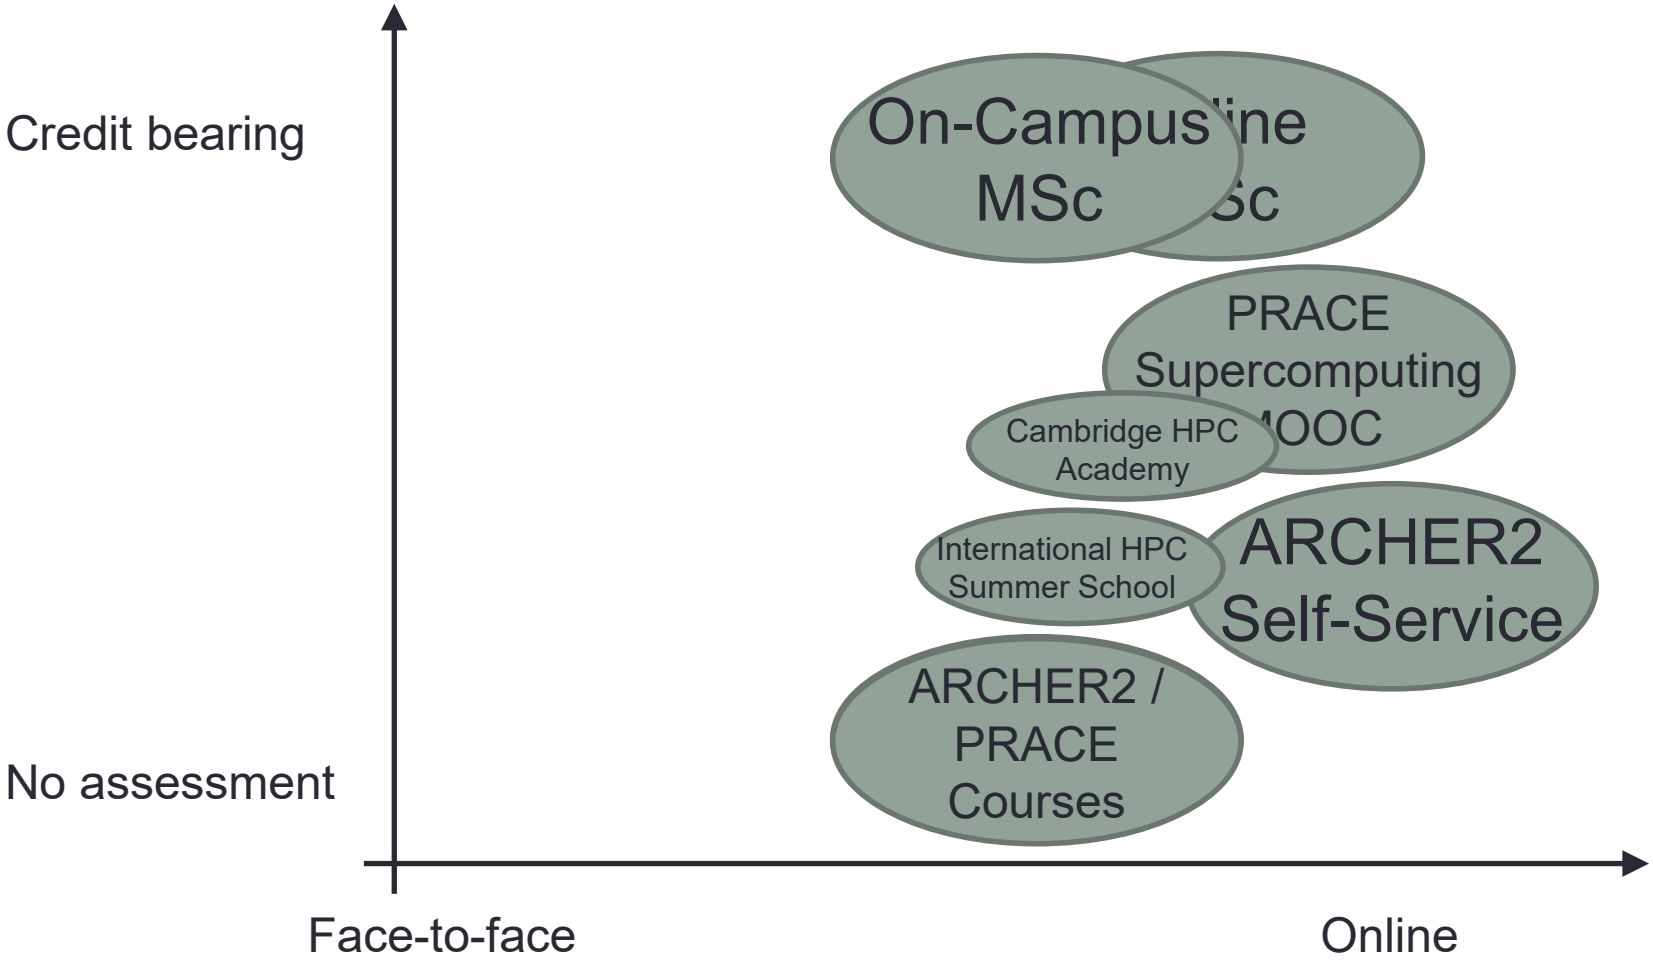

epcc

THE UNIVERSITY OF EDINBURGH

During COVID-19

- All training moves online
  - ARCHER2 and MSc
- Technology makes this simple in practice
  - live two-way audio and video
  - chat sessions and breakout rooms
  - computing practicals were already 100% online!
- But is the intensive multi-day course the right format?

Aside: is online HPC training special?

- Typical format:
  - lecture / hands-on / ...
  - programming exercises trivial to port
- Need to provide access to, say, ARCHER2
  - a well-oiled process
  - each course is a separate project
- Access well documented
  - some users need personal help
  - breakout sessions / screen sharing essential

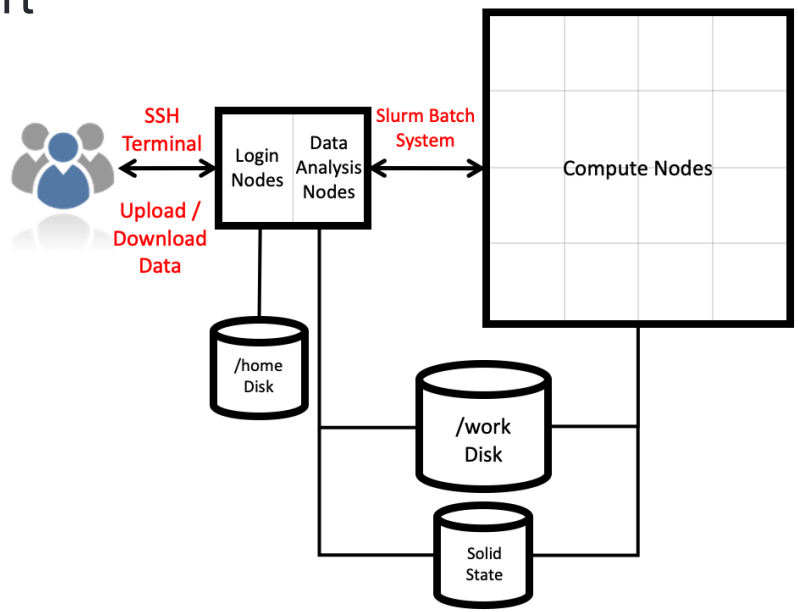

Post-COVID

- Resuming face-to-face training
  - but more online than pre-COVID
- Was a lot of discussion about best technologies
  - Teams
  - Zoom
  - Collaborate
- I found them all rather similar
- What is the right balance?
  - what have we learned about the best delivery formats?

Options

- Maintain intensive delivery
  - consecutive days
- My experience
  - people (attendees and lecturers) experience burn-out
  - no “down-time” chatting to other attendees at coffee / lunch
  - interaction in practical sessions is often poor
- Intensive delivery model **entirely based on physical travel**
  - true for both lectures and hands-on practical sessions
  - minimises travel & accommodation costs and time away from home
  - only way to enable non-local attendees
  - we don’t use it for most University courses

Synchronous Delivery (i)

- Live as opposed to recorded lectures
  - Pros:
    - attendees feel more engaged
    - allows for some interaction, e.g. quizzes and polls
  - Cons
    - more effort from lecturer
    - seems pointless unless you actively encourage questions during lectures

- I prefer live lectures
  - people don't usually watch recorded lectures in advance
  - I find timetabled slots for viewing very awkward and unnatural

Synchronous Delivery (ii)

- "long-thin" vs "short fat"
  - e.g. every Tuesday for 3 weeks
  - or every Monday and Wednesday afternoon for three weeks
- Pros
  - more time for attendees to assimilate info or review material
  - more time for attendees to attempt practicals in between sessions
  - much more relaxed with less burnout for all concerned
- Cons
  - harder to provide practical support (online chatboards help)
    - only limited success with online "office-hours"
  - attendees may not be available for every session (recordings help)

Synchronous Delivery (iii)

- Scheduling of lectures and practicals

- Face-to-face
  - I typically do lecture / practical / lecture / practical/ ...
  - each session around 45 minutes

- Online
  - I prefer lecture / lecture / practical / practical / ...
  - more flexibility for attendees
  - easier to schedule lecturers and demonstrators

Synchronous Delivery (iv)

- Standard format is just slide-based material done online
- How about using a completely new format?
- Software / Data / HPC Carpentry
  - all material on the web with no classical slides
  - entire course is a hands-on walkthrough with live coding etc.
- Online
  - Pros
    - material is ideally suited to online delivery
    - live coding keeps people engaged: a good idea for **all** practical sessions
  - Cons
    - development costs for new material
    - hard to replicate the one-to-one assistance of face-to-face format

Asynchronous Delivery

- e.g. fully recorded lectures or tailor-made MOOC
- Pros
  - very cost-efficient once material is developed
  - can teach large numbers of students
  - very accessible
- Cons
  - difficult to maintain (recordings hard to edit)
  - limited interaction between attendees and staff
    - only limited success with online “office-hours”
- Solutions
  - for online MSc we have chatboards and regular tutorials
  - encourage attendees to study as a cohort (e.g. FutureLearn MOOC)

ARCHER2 Self-Service Courses

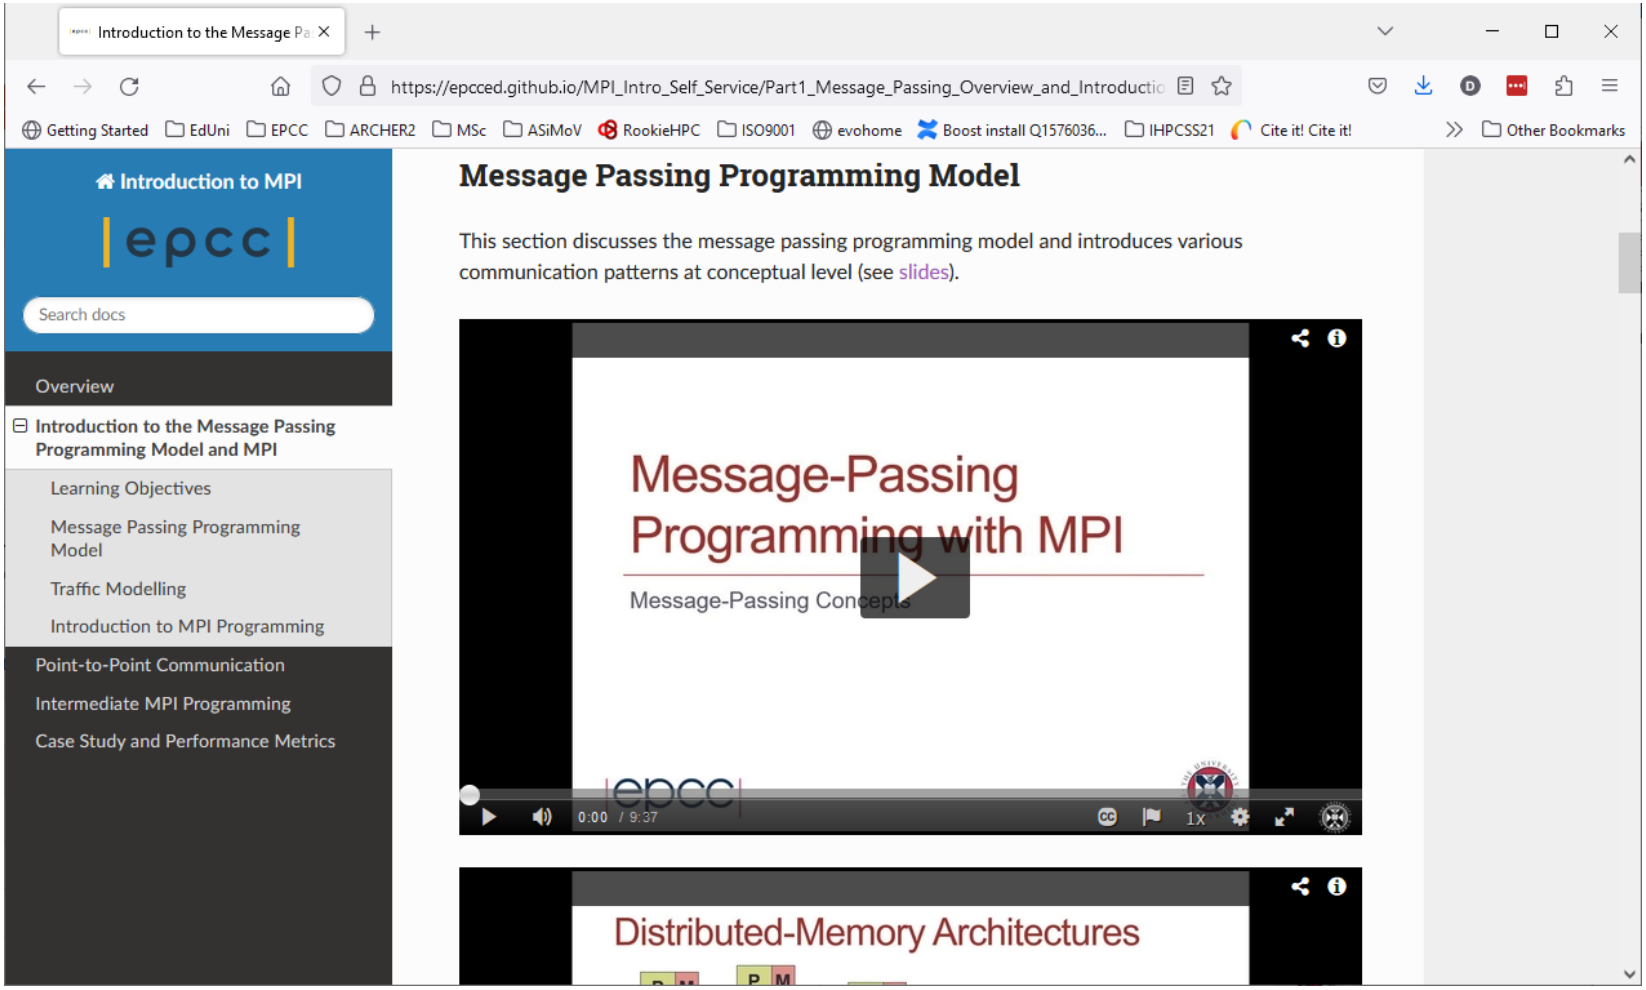

Diversity and Inclusion

- Don’t be worried about declining attendance
  - online makes attendance more diverse in background etc.
  - it is a huge commitment to attend a physical course
    - online, can attend a few lectures to see if it’s useful with little cost
    - widens access to the courses: what matters is **how many finish** not start!

- People have other commitments (life, childcare etc)
  - recording lectures guards against occasional clashes
  - automatic closed captions increase in quality every year

epcc

THE UNIVERSITY OF EDINBURGH

Others have different views

- One partner in a large project mandated for online courses:
  - live delivery but no questions answered by lecturer until the end
  - mandatory attendance (?)

Conclusions

- Online enables a wide range of delivery styles
  - we must exploit these
  - I am surprised at how little discussion there is of high-level issues
    - e.g. schedule within a day / timetabling across days
    - all focussed on details of particular technologies (Teams vs Zoom vs ...)
- Good material + good trainers even more important online
  - trainer must be engaged, open and enthusiastic
  - don't worry about informality and reacting to situation
    - if you're just reading a script you might as well pre-record
- All online and no face-to-face is not a good situation
  - but post-COVID I think we should do a lot of both