Trends in HPC; UKRI and European Context for Environmental HPC

Bryan Lawrence

NCAS & University of Reading

London, 27 Nov 18

University of Reading

NERC SCIENCE OF THE ENVIRONMENT

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

Aims

NERC HPC Status Quo

- Science objectives are built on decades of increasing HPC capability and capacity.
- Capital investment no longer buys massive improvements in capability or capacity.

  (and is dependent on "Benefits Realisation" and clear requirements which expose differences from other computing investment, e.g. cloud).
- NERC HPC recurrent budget is (currently) fixed and constrains "our share" of capacity.

Aims

NERC HPC Status Quo

- Science objectives are built on decades of increasing HPC capability and capacity.
- Capital investment no longer buys massive improvements in capability or capacity.
  (and is dependent on "Benefits Realisation" and clear requirements which expose differences from other computing investment, e.g. cloud).
- NERC HPC recurrent budget is (currently) fixed and constrains "our share" of capacity.

A future for NERC HPC

- Science objectives need to reflect the changing nature of HPC.
- Capital investment depends on consolidated science case with clear evidence of potential impacts. Expect this to need regular updates.
- If NERC HPC budget is to increased, NERC will need evidence as to why HPC should be funded at the expense of other activities - back to science case and impact metrics!

Expanding Communities and Expanding Usage within Communities

Interaction: A view from climate science:

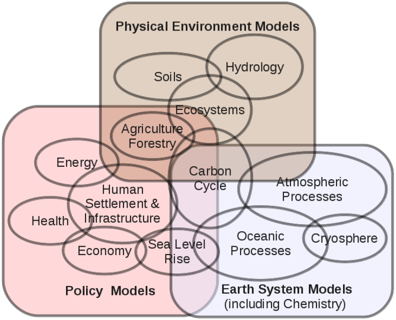

Plenty of other domains of NERC science as well
(e.g. Geosciences, Pollution, Waste Resources, and more).

Growing use of simulation in science:

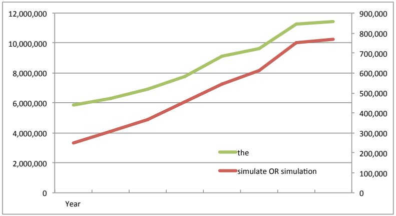

In the fifteen years (2001-2015), the number of papers published doubled, but the number of papers with the word simulation in the abstract tripled.

(Source: Analysis of Google Scholar Searches)

Trends

Our Experience:

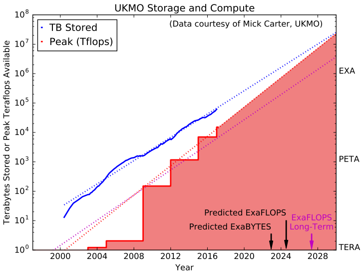

- Exponential Growth in Compute
- Exponential Growth in Storage
- You might think “Exa” in the near future …

Trends

Our Experience:

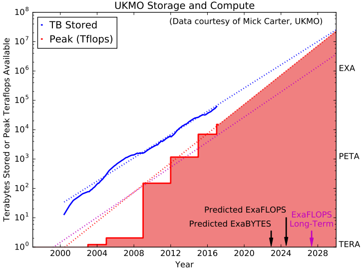

- Exponential Growth in Compute
- Exponential Growth in Storage
- You might think “Exa” in the near future ...

But; Impending Reality:

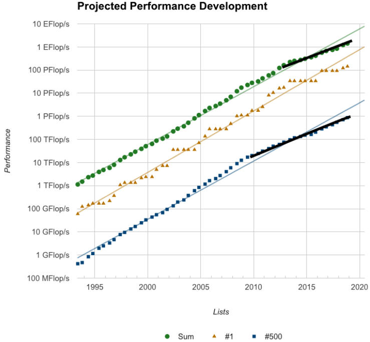

- Performance increases are slowing.
- Staying on the curve at the top end is becoming harder (& more costly).
- (Actual application performance is getting harder to extract with each generation of computing.)

What is driving this? The end of an era?

- **End of Dennard Scaling:** Clock frequency used to increase as transistor size shrunk, leading to increased performance per watt.
  - Since about 2006 this has not been true, with two practical consequences: clock frequencies are not increasing much (if at all, some are going down) & system power consumption is going up.
  - The consequence is that speed ups require more parallelism, and all other things being equal, consume more power (caveat GPU).
- **Moore's Law under pressure**. The number of transistors in a dense integrated circuit has doubled roughly every two years for a long time.
  - It's now longer than three years for a doubling, and the physics of making transistors suggests a limit is near.
- **Future computing may be slower, more parallel, more expensive (up front, and to run) and more dependent on customised hardware.**

What is driving this? The end of an era?

- **End of Dennard Scaling:** Clock fr… to increase as transistor size shrunk… mance per watt.
  - Since about 2006 t… practical consequences: clo… g much (if at all, some are going dov… n is going up.
  - The consequence is… arallelism, and all other things bein… caveat GPU).
- **Moore's Law under pr…** …sistors in a dense integrated circui… two years for a long time.
  - It's now longer than t… …sing, and the physics of making transistors suggests a limit is near.
- **Future computing may be slower, more parallel, more expensive (up front, and to run) and more dependent on customised hardware.**

Breakout Q1:

- How does our science face a future where computers do not run faster?
- Are there known algorithmic routes to improve time to solution?
- What else can we do?

It’s not just compute: declining Kryder rate too!

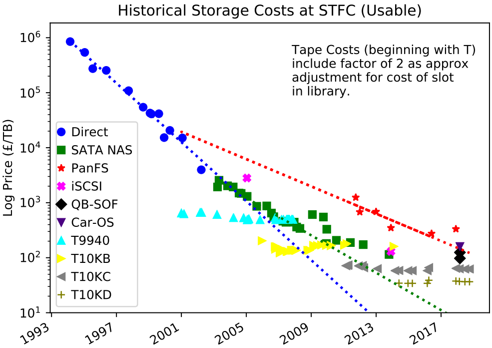

- Relative cost of storage going up: each new generation of disk has a “shallower Kryder rate” than the previous one.
- Note that each new generation of tape is cheaper, and doesn’t tend to change much in cost over the lifetime.
- (Expect tape to keep competitive cost advantage over disk for the foreseeable future. NERC investing in tape systems at JASMIN!)

It’s not just compute: declining Kryder rate too!

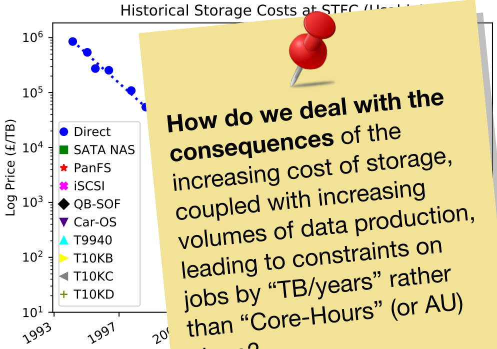

How do we deal with the consequences of the increasing cost of storage, coupled with increasing volumes of data production, leading to constraints on jobs by “TB/years” rather than “Core-Hours” (or AU) alone?

- Relative cost of storage going up: each new generation of disk has a “shallower Kryder rate” than the previous one.
- Note that each new generation of tape is cheaper, and doesn’t tend to change much in cost over the lifetime.
- (Expect tape to keep competitive cost advantage over disk for the foreseeable future. NERC investing in tape systems at JASMIN!)

How we use more compute - Large Scale Simulation

Weather and Climate

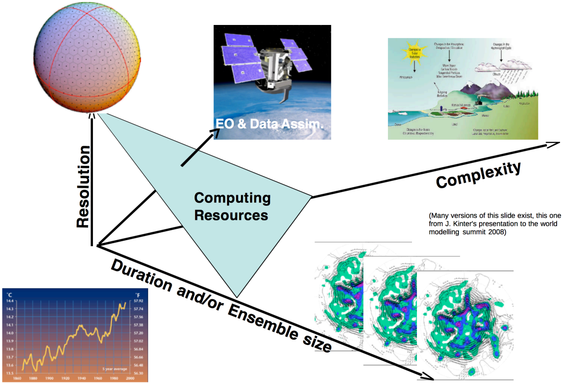

For *simulation* the pictures change with other disciplines, but the notions do not change much (even when we stray as far as astronomy and crystallography).

How we use more compute - Large Scale Simulation

Weather and Climate

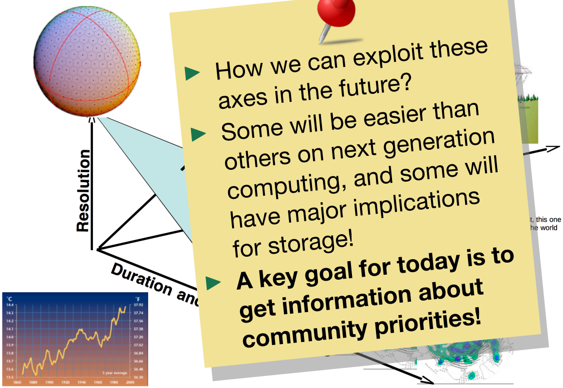

- How we can exploit these axes in the future?
- Some will be easier than others on next generation computing, and some will have major implications for storage!
- **A key goal for today is to get information about community priorities!**

For *simulation* the pictures change with other disciplines, but the notions do not change much (even when we stray as far as astronomy and crystallography).

Growing impact of Machine Learning and Artificial Intelligence

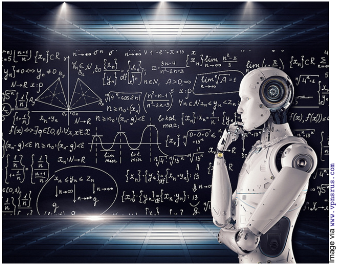

Image via www.vpnsrus.com

Gratuitous “robots are coming” image

Expect ML and AI to have major implications for both

- HPC architectures, and
- Algorithms, in use before, during, and after simulation (analytics)!

Growing impact of Machine Learning and Artificial Intelligence

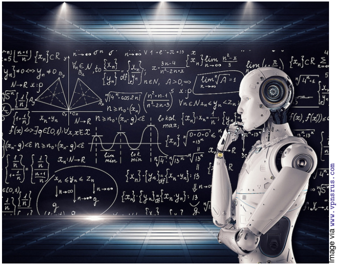

Image via www.vpnsrus.com

Gratuitous “robots are coming” image

Expect ML and AI to have major implications for both

- HPC architectures, and
- Algorithms, in use before, during, and after simulation (analytics)!

Initial emphasis on climate services, parameter estimation (for parameterisations) and emulation (potentially avoiding avoid long spin-up runs).

Two interesting examples contributed to the Gordon Bell competition this year:

- Preconditioning implicit solvers using artificial intelligence — ground breaking (!) simulations of earthquakes and building response : Ichimura et al 2018.

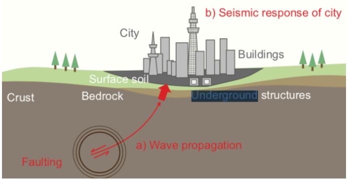

- Exascale Deep Learning for Climate Analytics - Extracting weather patterns from climate simulations: Kurth et al 2018, co-winner of 2018 Gordon Bell prize.

NERC HPC

NERC Supercomputing

Three Simulation Platforms:

- ARCHER (EPCC in Edinburgh, roughly quarter of the machine)
- Monsoon2 and NEXCS (UKMO in Exeter, similar size resource to ARCHER, much bigger platform).

One Analysis and Archive Platform:

- JASMIN (44 PB of spinning disk plus 12K cores plus tape)

Fast *and* reliable *and* fat network links!

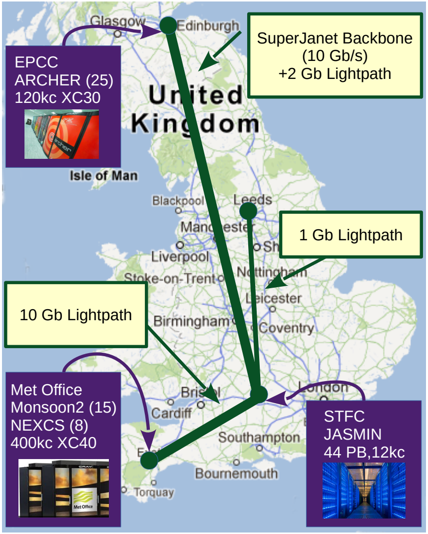

NERC HPC

NERC Supercomputing

Three Simulation Platforms:

- ARCHER (EPCC in Edinbu[…] roughly quarter of the mac[…]
- Monsoon2 and NEXCS (UK[…] in Exeter, similar size resour[…] to ARCHER, much bigger platform).

One Analysis and Archive Platform:

- JASMIN (44 PB of spinning disk plus 12K cores plus tape)

Fast *and* reliable *and* fat network links!

- ARCHER2 due late 2019.
- Expected to be roughly 5x throughput of ARCHER, and up to twice as fast in time to solution - roughly 2.5 x more nodes.
- NERC will have roughly a fifth of the machine (impact of fixed HPC budget).

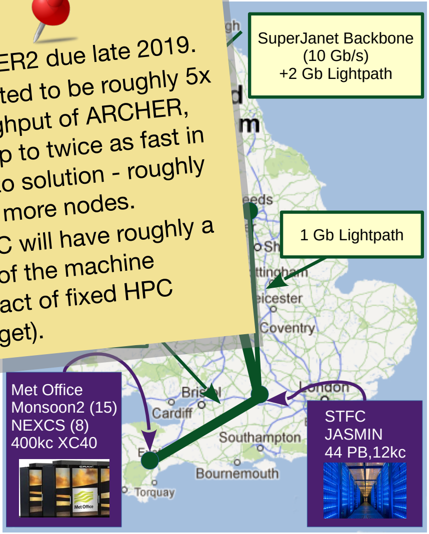

NERC HPC: JASMIN — The Data Commons

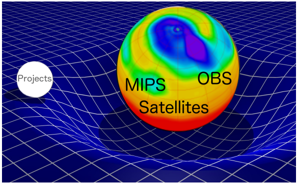

- Provide a state-of-the art storage and computational environment
- Provide and populate a managed data environment with key datasets (the “archive”).
- Encourage and facilitate the bringing of data and/or computation alongside/to the archive!
- Provide FLEXIBLE methods of exploiting the computational environment.

NERC HPC: JASMIN — The Data Commons

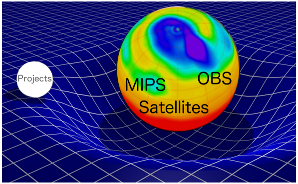

- Provide a state-of-the art storage and computational environment
- Provide and populate a managed data environment with key datasets (the “archive”).
- Encourage and facilitate the bringing of data and/or computation alongside/to the archive!
- Provide FLEXIBLE methods of exploiting the computational environment.

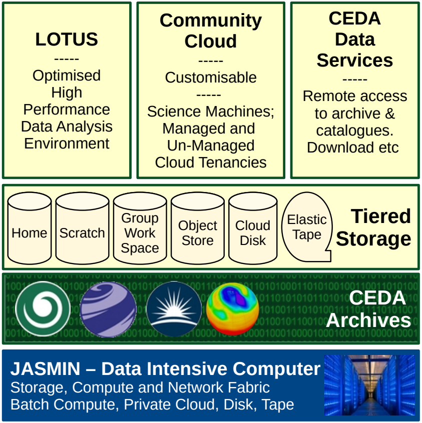

NERC HPC: JASMIN — The Data Commons

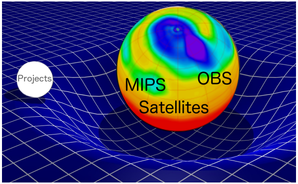

- Provide a state-of-the art storage and computational environment
- Provide and populate a managed data environment with key datasets (the “archive”).
- Encourage and facilitate the bringing of data and/or computation alongside/to the archive!
- Provide FLEXIBLE methods of exploiting the computational environment.

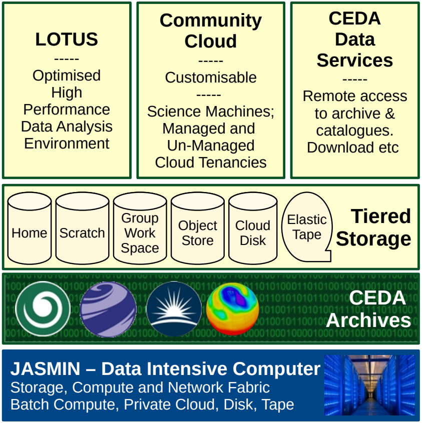

- NERC will not have access to the RDF after the advent of ARCHER2; all RDF data will need to migrate to JASMIN.
- Growing need for tape access to be met with new services in 2020 (develoment funding willing).
- Growing role for Cloud (within JASMIN and elsewhere!)

NERC HPC: JASMIN — The Data Commons

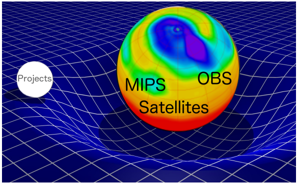

- Provide a state-of-the art storage and computational environment
- Provide and populate a managed data environment with key datasets (the “archive”).
- Encourage and facilitate the bringing of data and/or computation alongside/to the archive!
- Provide FLEXIBLE methods of exploiting the computational environment.

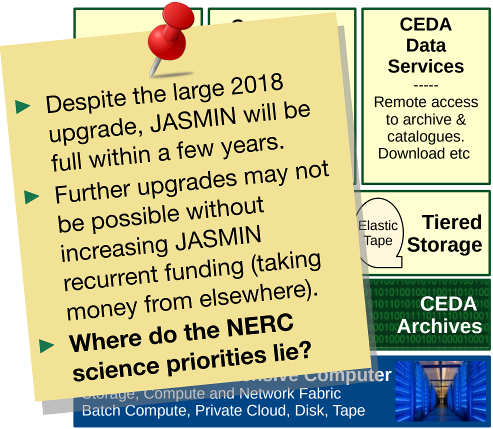

- Despite the large 2018 upgrade, JASMIN will be full within a few years.
- Further upgrades may not be possible without increasing JASMIN recurrent funding (taking money from elsewhere).
- **Where do the NERC science priorities lie?**

- NERC will not have access to the RDF after the advent of ARCHER2; all RDF data will need to migrate to JASMIN.
- Growing need for tape access to be met with new services in 2020 (develoment funding willing).
- Growing role for Cloud (within JASMIN and elsewhere!)

UKRI e-infrastructure roadmap - Supercomputing Component

…draft under development …

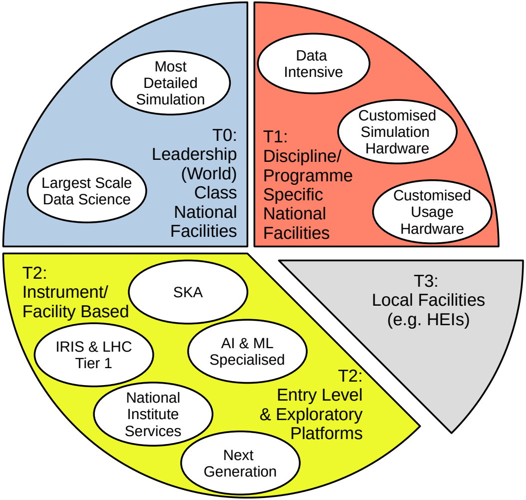

Tier 0: Leadership (World) Class National Facilities:

- ARCHER2

Tier 1: Discipline or Programme Specific National Facilities.

- JASMIN

Tier 2: Entry Level and Exploratory Platforms:

- ISAMBARD (EPSRC)

(Met Office supercomputing will likely appear on a national version of the UKR roadmap.)

HPC, NERC, and UKRI programmes

How does HPC funding work in a cross-disciplinary world?

- NERC science contributes to the arguments which justify capital expenditure.
- NERC recurrent funding quantifies the proportion of the resources which can be used by NERC scientists.
- NERC HPC resources are over-allocated supporting existng NERC science programmes (including Discovery Science).
- These statements lead to the following questions:
  1. Do we expect any large HPC requirements from non-core NERC and/or cross-council activity? (e.g. ODA Science, or ISCF, or SPF)? If so, who pays for HPC resources?
  2. What about non-UKRI (EC? Industry?)
  3. In all cases, even if the funds are found, from which council's allocation is the resource obtained?

HPC, NERC, and UKRI programmes

How does HPC funding wor

- NERC science contribut … fy capital expenditure.
- NERC recurrent funding … resources which can be used by NI
- NERC HPC resources are … tng NERC science programmes (inc
- These statements lead to
  1. Do we expect any lar … -core NERC and/or cross-council … Science, or ISCF, or SPF)? If so, who pays for HPC resources?
  2. What about non-UKRI (EC? Industry?)
  3. In all cases, even if the funds are found, from which council's allocation is the resource obtained?

Breakout 3:

- What are the implications of the new UKRI funding routes (e.g. Industrial Strategy Challenge Fund) for environmental HPC resource requirements (software, hardware, people)?

Benefits Realisation

How does NERC develop and report on capital investments?

If we want to sustain our hardware, we need to continually request updates and upgrades.

1. We need a science case, which we can keep updated and current, reflecting near-, medium-, and long-term ambition
   - We are used to this, but maybe not the idea of having it updated annually.
2. We need a business case? How is the facility to be operated? What will be the expected **economic** benefits to the country?
   - Large capital expenditure needs credible economic outcomes. How do we *evidence* such claims? This is now a larger and larger part of the over-all pass/fail for our HPC investments.
3. We need a "Benefits Realisation Plan". How do we ensure we deliver **both** the scientific and economic benefits?
   - Increased emphasis on realising the economic benefits, and on the likelihood that successful benefits realisation will be necessary for the next upgrade/update.

Benefits Realisation

How does NERC develop and report on capital investments?

If we want to sustain our hardware, we nee… …equest updates and upgrades.

1. We need a science case … …nd current, reflecting near-, medium…
   - We are used to this … …ng it updated annually.
2. We need a business case … …ated? What will be the expected **econom**…
   - Large capital expenc… …outcomes. How do we *evidence*… …ger and larger part of the over-all pa… …
3. We need a "Benefits Realis… …ow do we ensure we deliver **both** the scientific and economic benefits?
   - Increased emphasis on realising the economic benefits, and on the likelihood that successful benefits realisation will be necessary for the next upgrade/update.

Breakout 4:

- How do we gather information about the impact of our HPC related research?
- How do we demonstrate economic impact?
- We (Swindon and NCAS in support for HPC) cannot do this our own!

PRACE and EuroHPC

PRACE

- Piz Daint (2017) - Switzerland - #5
- SuperMUC (2018) - Germany - #8
- Marconi (2016) - Italy #19
- MareNostrum4 (2017)- Spain - #25
- JUWULS (2018) - Germany - #26
- Hazel Hen (2015) - Germany #30
- Joliet Curie (2017) - France #40

cf UK and other “climate compute”:

- UKMO (2016) - #23
- Cheyenne (2016) - NCAR, U.S. - #36
- Mistral (2015) - DKRZ, Germany - #62
- Cumulus (2018) - U.Cambridge - #87
- Scafell Pike (2017) - STFC, UK - #107
- ARCHER (2014) - EPCC, UK - #186

EuroHPC

- Massive upgrades due in 2020: Two “pre-exascale” machines (300K euros each)!
- Aimng for two “exascale machines” a few years later.

PRACE and EuroHPC

PRACE not standing still:

- Nov 2018, HLRS (Stuttgart) announces collaboration with HPE to build Hawk, 3.5 times faster than Hazel Hen (using AMD EPYC).
- Aims to be the world's fastest for industrial production, computational engineering and research …

PRACE

- …t (2017) - Switzerland - #5
- …UC (2018) - Germany - #8
- …(2016) - Italy #19
- …strum4 (2017)- Spain - #25
- …(2018) - Germany - #26
- …n (2015) - Germany #30
- …ie (2017) - France #40

…l other "climate …:

- …16) - #23
- …(2016) - NCAR, U.S. - #36
- …5) - DKRZ, Germany - #62
- …018) - U.Cambridge - #87
- …Pike (2017) - STFC, UK - #107
- ARCHER (2014) - EPCC, UK - #186

EuroHPC

- Massive upgrades due in 2020: Two "pre-exascale" machines (300K euros each)!
- Aimng for two "exascale machines" a few years later.

PRACE and EuroHPC

PRACE

- t (2017) - Switzerland - #5
- UC (2018) - Germany - #8
- (2016) - Italy #19
- 017)- Spain - #25
- Germany - #26
- Germany #30
- France #40

"climate

- ICAR, U.S. - #36
- Z, Germany - #62
- Cambridge - #87
- STFC, UK - #107
- CC, UK - #186

PRACE not standing still.

- No
  an
  wit
  tim
  (us
- Air
  fas
  pro
  en

- PRACE has been good to UK science!
- Top500 doesn't measure science throughput, but does measure ambition.
- Top machines easily deliver **more** sci, harder to deliver **better** sci.
- Better science needs ££ for s/w **and** h/w!

EuroHPC

- Massive upgrades due in 2020: Two "pre-exascale" machines (300K euros each)!
- Aimng for two "exascale machines" a few years later.

ExtremeEarth

European Commission investments in HPC go beyond HPC platforms

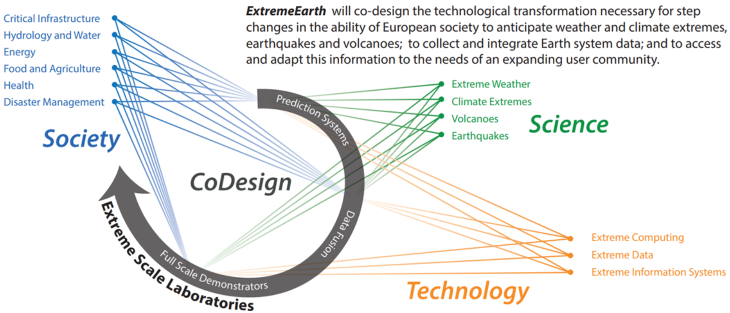

- Status: Stage 2 Pre Proposal under review: aim of that work would be to fully define an ambitious 10-year ExtremeEarth project to deliver the methods required for a step-change in predictive capabilities for Earth-system extremes.
- (Bid into what was the "Future and Emerging Technologies" programme. Not sure where it will land since that programme has been cancelled.)

European Strategy

- (PRACE, 2018): …rectify the under-investment in computational infrastructure covering hardware, technology, and software, but as these investments also compete with funding for academic research projects, …essential …cost-efficient, …(deliver) …a long-term vision and strategy how to deliver resources academia and industry depend upon.
- (ENES, 2017) Support common development and sharing of software and accelerate the preparation for exascale computing by exploiting next generation hardware and developing appropriate algorithms, software infrastructures, and workflows.
- (ENES, 2017) Exploit a blend of national and European high-performance facilities to support current and next generation science and work toward obtaining sustained access to world-class resources and next generation architectures.

(ENES: European Network for Earth Simulation, `https://enes.org`.)

European Strategy

- (PRACE, 2018): …rectify the under-investment in computational infrastructure covering hardware, technology, and software, but as these investments also compete with funding for academic research projects, …essential …cost-efficient, …(deliver) …a long-term vision and strategy how to deliver resources academia and industry depend upon.
- (ENES, 2017) Support common development and sharing of software and accelerate the preparation for exascale computing by exploiting next generation hardware and developing appropriate algorithms, software infrastructures, and workflows.
- (ENES, 2017) Exploit a blend of national and European high-performance facilities to support current and next generation science and work toward obtaining sustained access to world-class resources and next generation architectures.

(ENES: European Network for Earth Simulation, `https://enes.org`.)

Summary

Budget Fixed, Community Growing:

- Fixed recurrent budget, but recurrent costs going up.
- More science dependent on HPC.

Compute Growth Slowing, Costs Increasing:

- Speed $= f$(parallel, electricity).
- Harder in software to harness the speedup there is!

HPC is changing; new platforms, new science:

- Need to balance compute against storage.
- Greater role for analysis compute and machine learning systems.

Summary

Budget Fixed, Community Growing:

- Fixed recurrent budget, but recurrent costs going up.
- More science dependent on HPC.

Compute Growth Slowing, Costs Increasing:

- Speed $= f(\text{parallel}, \text{electricity})$.
- Harder in software to harness the speedup there is!

HPC is changing; new platforms, new science:

- Need to balance compute against storage.
- Greater role for analysis compute and machine learning systems.

Need Plans and Evidence:

- Need to understand and prioritise NERC science given constrained budgets.
- Case for computing depends even more on understanding benefit and impacts.

Ambitious Scientific Competitors:

- Delivering world class science depends on using world class tools!
- How can we sustain & exploit collaboration opportunities?

NERC needs a strategy:

- Begin with clear understanding of scientific goals (starting TODAY).
- Important role for participants (not just speakers) to contribute today, and in the future!