UK Atomic Energy Authority

Andrew Davis

Enabling Exascale Nuclear Science & Engineering

Computation in High Energy Physics 2019 - November 4th-8th 2019 - Adelaide

Introduction

- This talk is a little un-usual - there will be lot of framing regarding the background to why I’m discussing this bizarre mismatch of topics, but the main drive here is answering questions around:

  “What am I going to model?”

  “How am I going to get my exascale?”

  “What are the benefits to me for aiming at the exascale?”

  “What am I going to do with my exascale?”

- This framing of questions, I think helps frame what “Enabling Exascale Nuclear Science” gets us

What do I need to model?

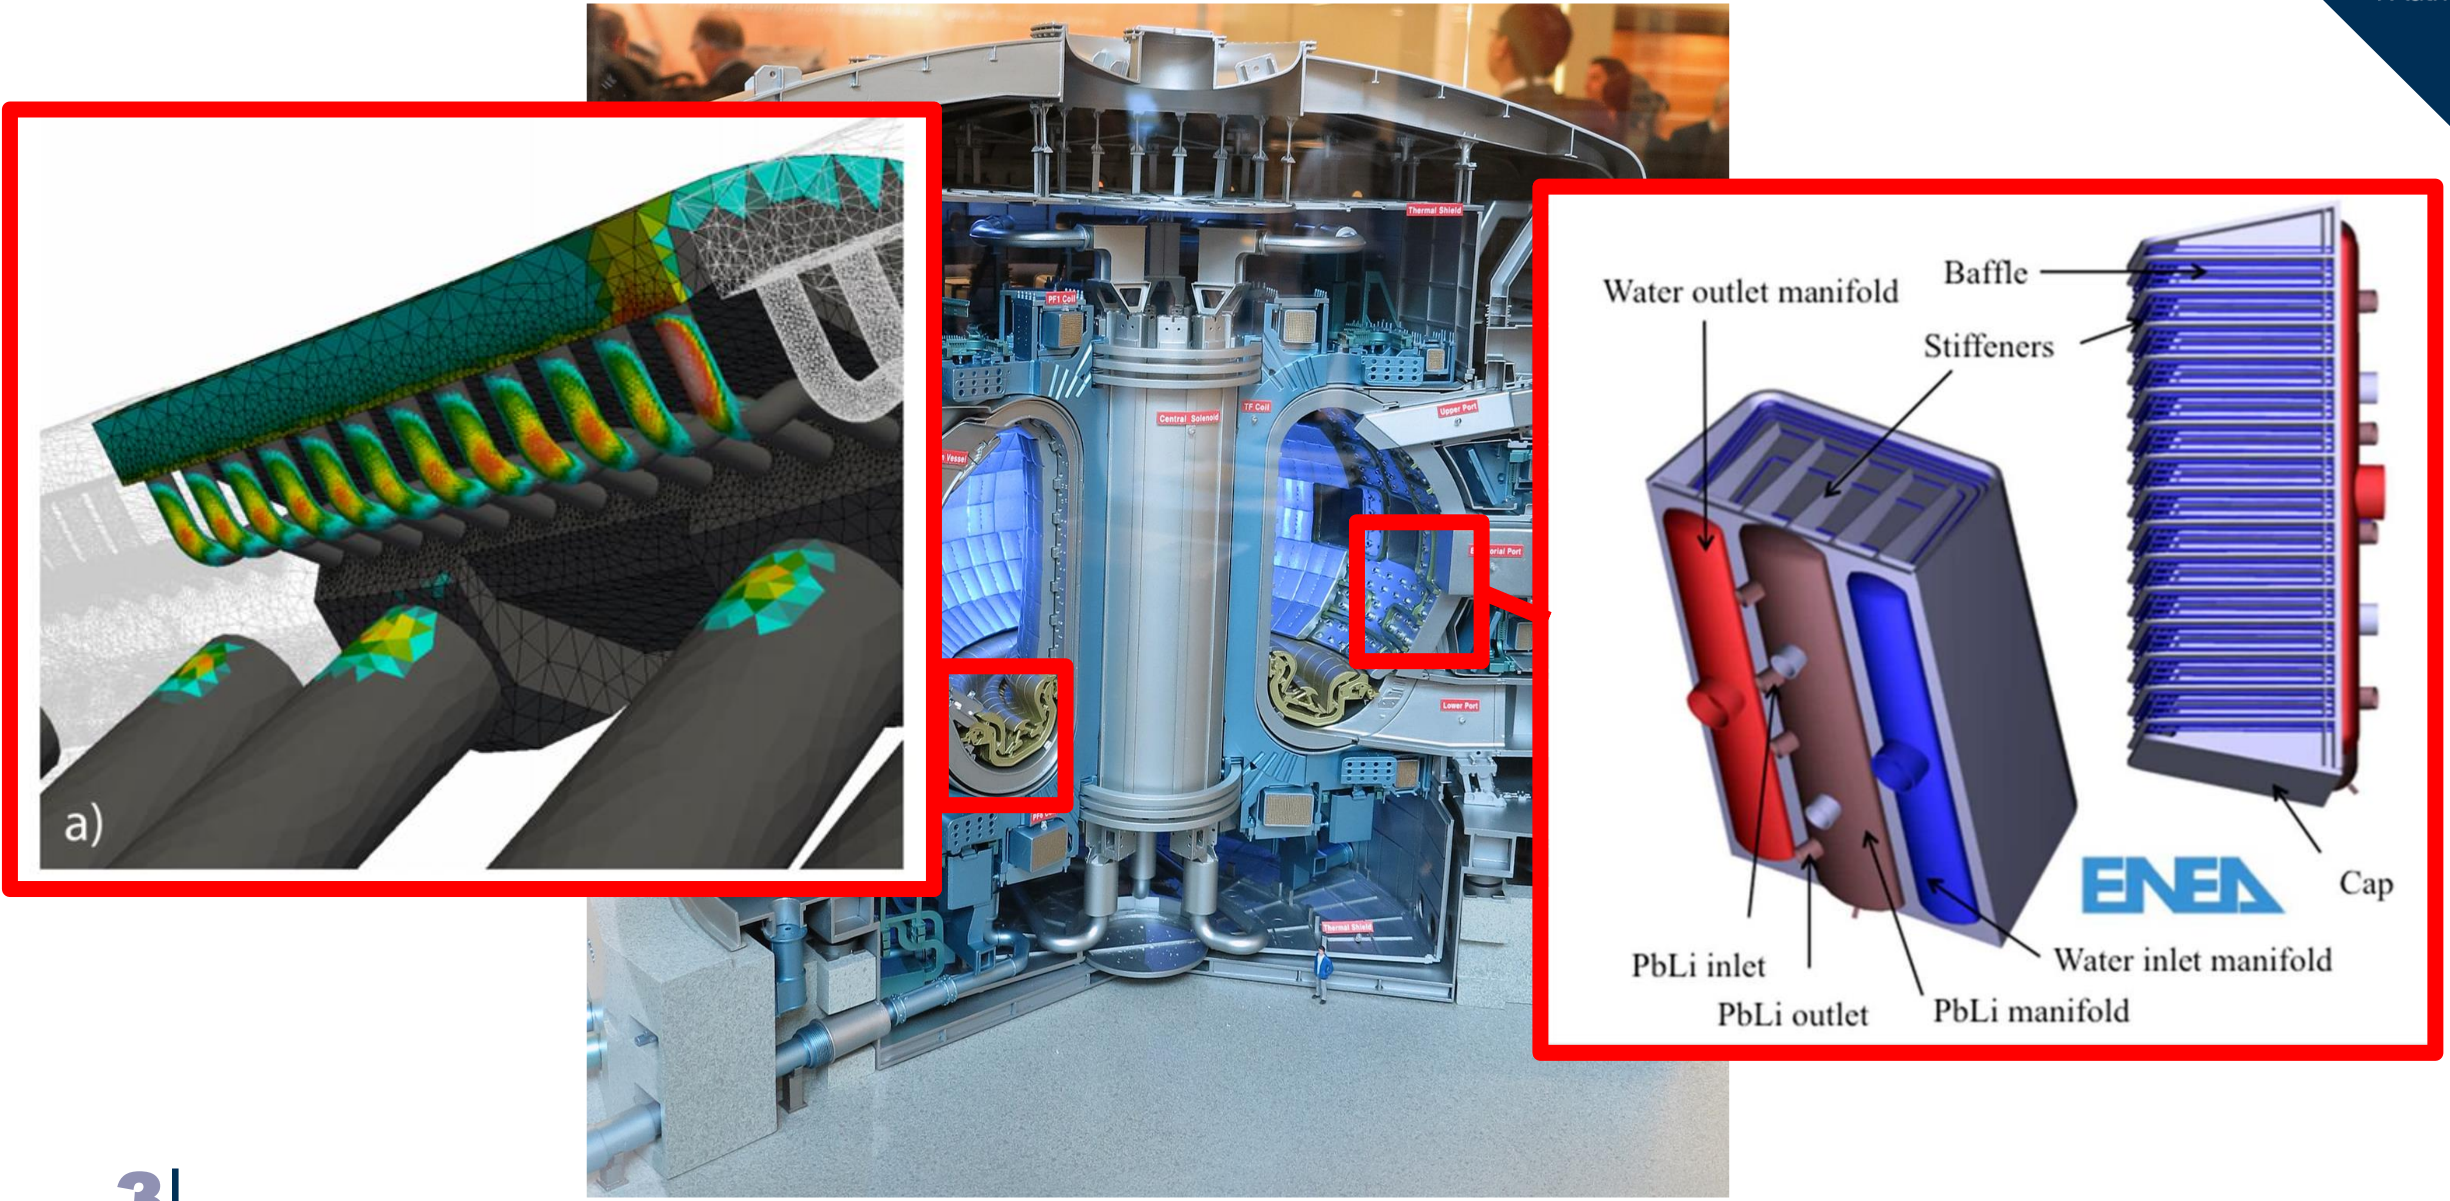

What do I need to model?

- To support fusion facilities I need nuclear digital twins to for predictive computation
  - A digital twin is a the digital mimic of a real world system that can accurately reproduce the results of experiment
  - I also need the real world system!
    - Well instrumented
    - low uncertainty
- I really really need a nuclear digital twin
  - EM, Thermal, CFD, MHD, Radiation Transport, Chemistry
  - Need to use computation as discovery
  - validated tools

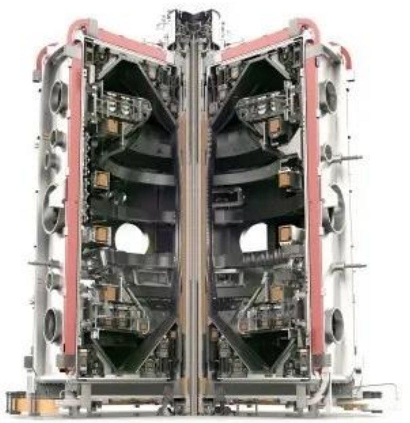

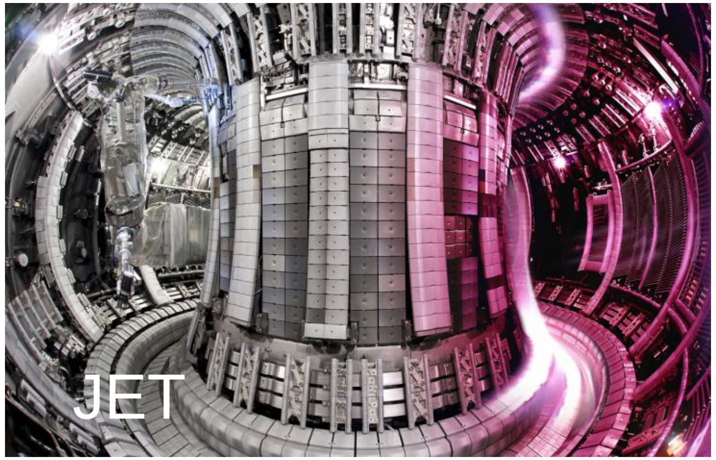

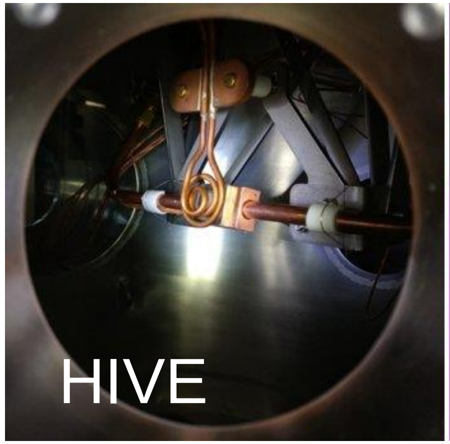

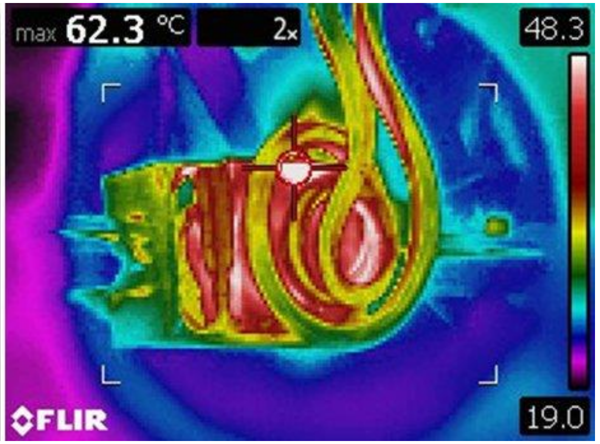

Where is my exascale?

UK Atomic Energy Authority

- Don’t feel sorry for the large US national labs (ORNL, LANL, LLNL, ANL, INL) they will get their exascale (as will labs in China, Japan)
- According to EuroHPC the EU hopes to have 2 top 5 exascale machines
- Three pillars:
  - 1) PRACE, 2) ETP4HPC, 3) Centres of Excellence

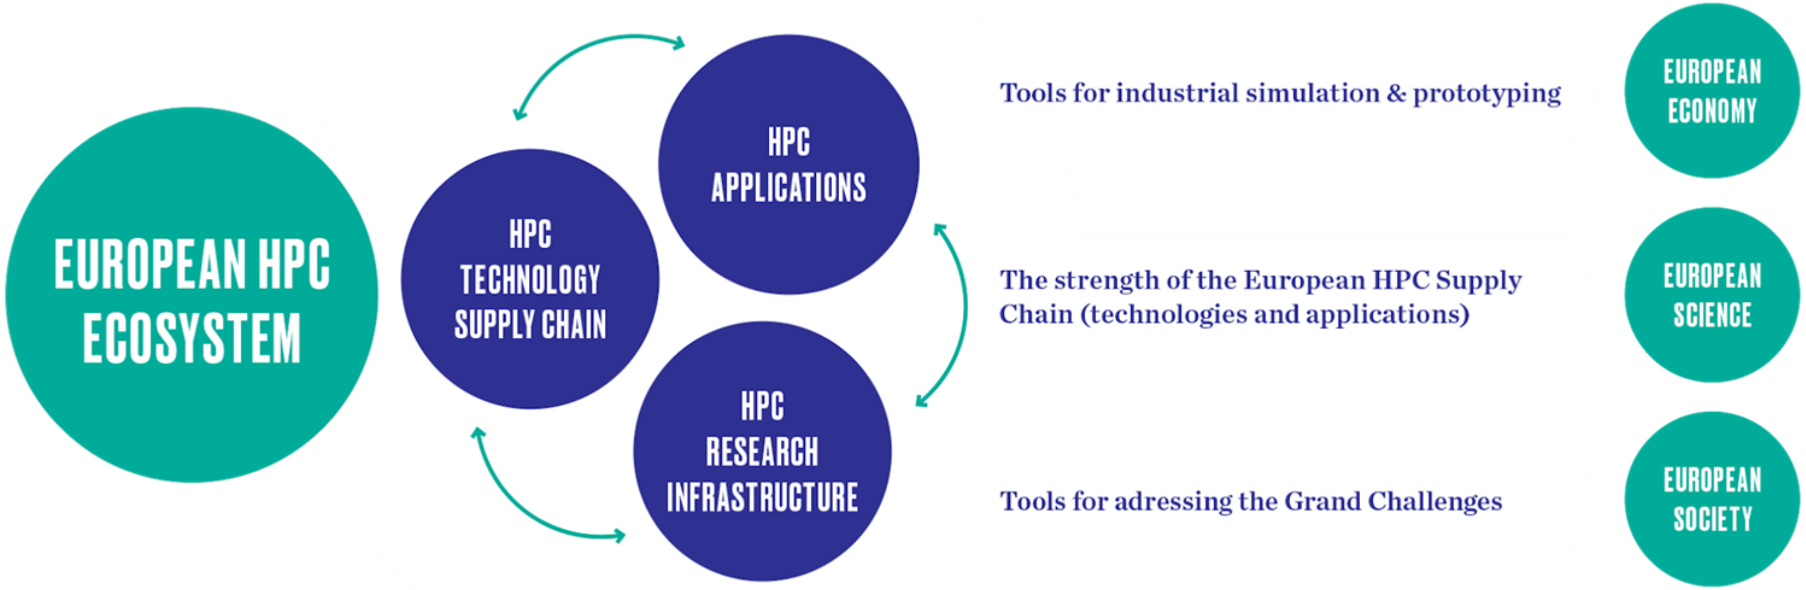

https://www.etp4hpc.eu/euexascale.html

UK Atomic Energy Authority

What will I get from exascale?

- What do I get by focussing my efforts on the exascale?
  - Faster and more optimised code - trickle down
  - better algorithm choices
  - more maintainable code
  - access to accelerators (GPU, FPGA)
  - chance to modernise the software stack

ExCALIBUR
E-TASC
EPC

- 10^{18} flops (what am I going to with all them flops?)
- GPU Co-compute
  - Maybe FPGA co-compute - FPGA is hard?
- Fast IO - can’t be hanging around waiting (30 MW ~ $5k/minute)
  - object stores?
  - Is serial IO & MPI_BCAST enough?
  - Where does RAM based storage fit?

csg2csg - enabling change

- Anytime one wants to trial another MC code, we take our favourite geometry which encapsulates the problem you're interested in, but aaaah aghast, the syntax isn't the same, oh and the materials are defined differently, one can't specify isotopes…… I haven't got time
- `csg2csg` is a minimal dependency Python3 code that ingests MC CSG geometry and exports it to another format (MCNP->FLUKA/OpenMC/PHITS/Serpent2)

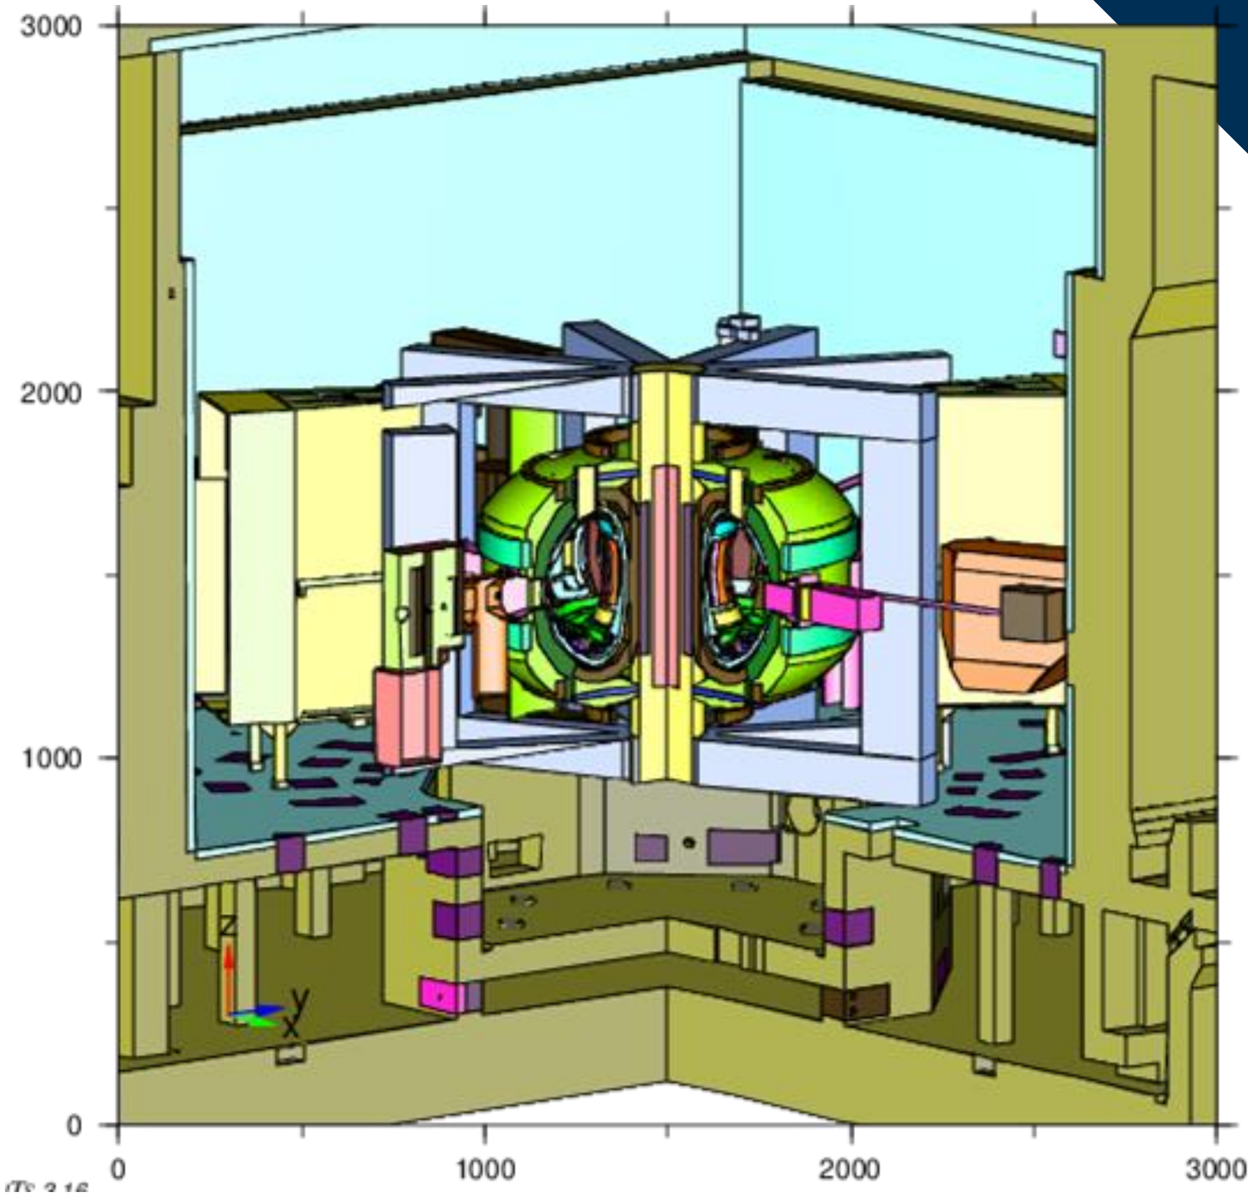

OpenMC

UK Atomic Energy Authority

- Started life as a PhD project at MIT, cleared the US export control barriers with the help of a few lawyers and is now professionally developed by Argonne National Lab (ANL) and the United Kingdom Atomic Energy Authority (UKAEA)
- It’s not feature complete, and is very much a work in progress but offers an open route to fusion specific problems; improvements are coming rapidly

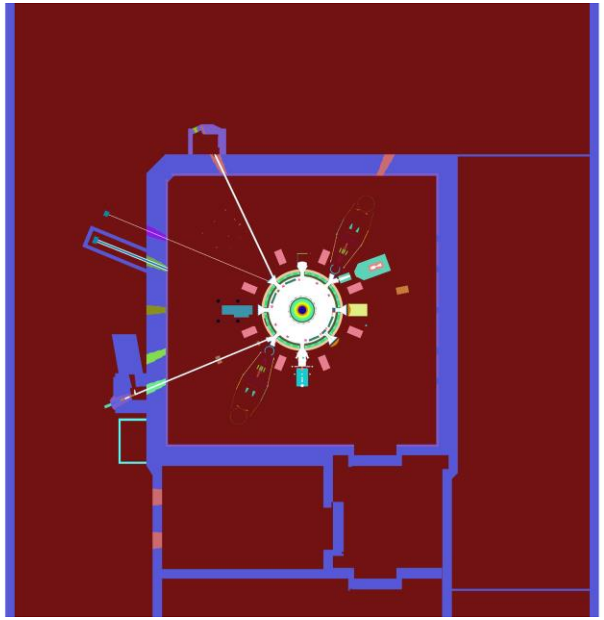

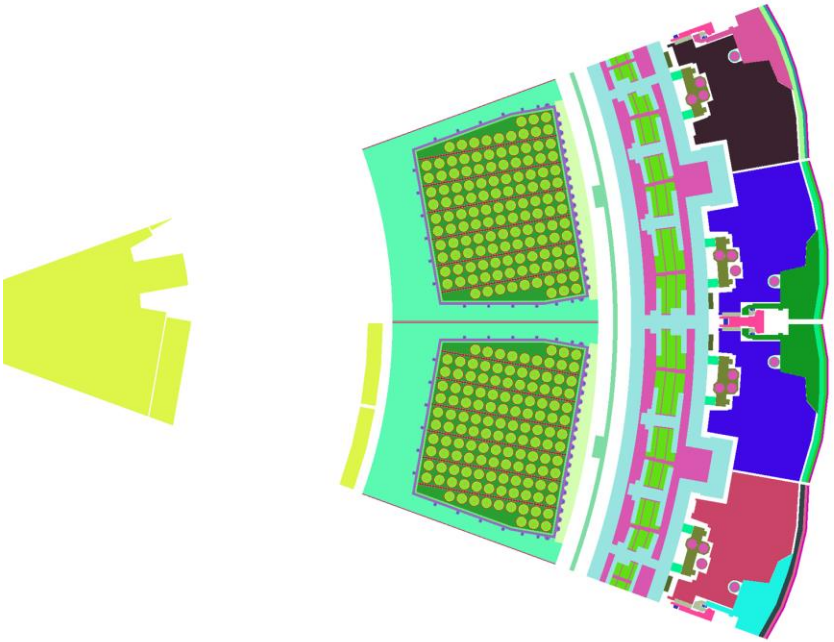

OpenMC

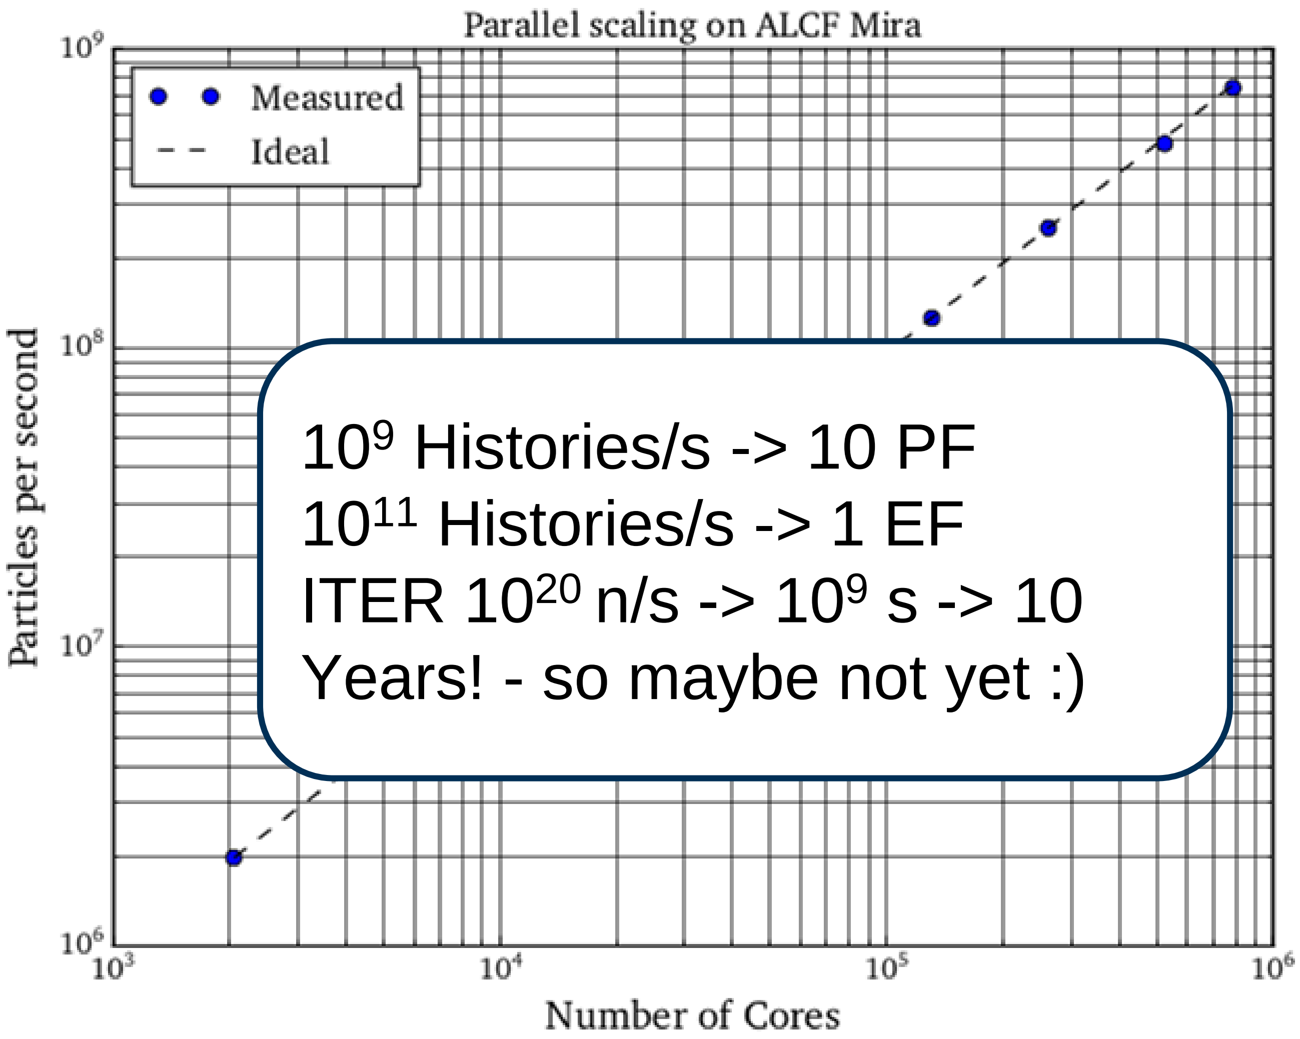

10^9 Histories/s -> 10 PF
10^{11} Histories/s -> 1 EF
ITER 10^{20} n/s -> 10^9 s -> 10 Years! - so maybe not yet :)

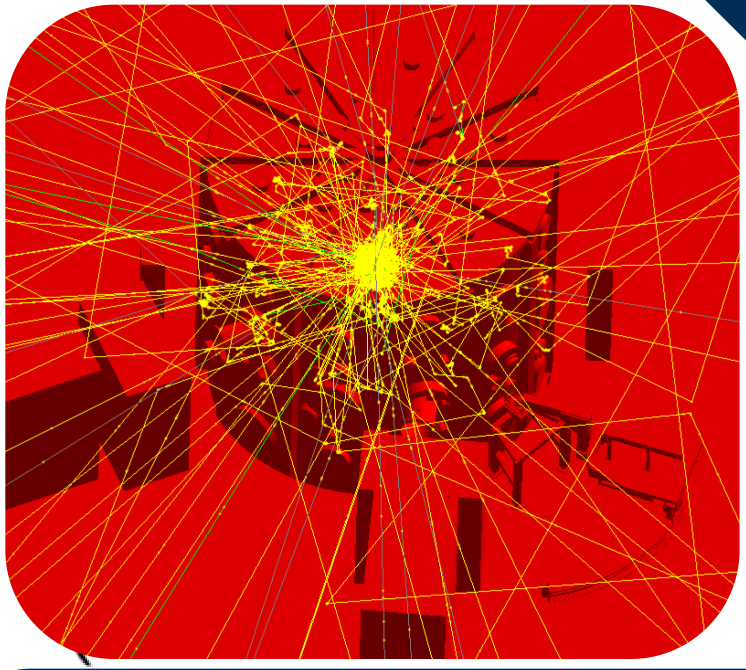

MAST-U - 10^{16} n/s
10^5 s -> 3 days

MIRA 8.59 (10.06 theoretical) PFlops - 3.9 MW

UK Atomic Energy Authority

Multiphysics

- The detailed design of ‘burning plasma’ fusion reactors is driven by multiphysics
    - High heat load from the plasma
        - Heat load on divertor 10 MWm^2>Q>75MWm^2
    - High volumetric heat load from neutrons
        - 10 MWm^3> Q > 100 MWm^3
    - High magnetic field
        - 10 T on plasma axis, 30 T at magnet
    - High radiation damage
        - 20 50 FPY > DPA > 50 FPY
    - Cooling channels maybe with liquid lead - MHD or liquid metal divertors
- Innovative engineering solutions needed
- Computation is challenging
    - Many strongly coupled effects
    - Geometrically complex

R Lewis - University of Swansea

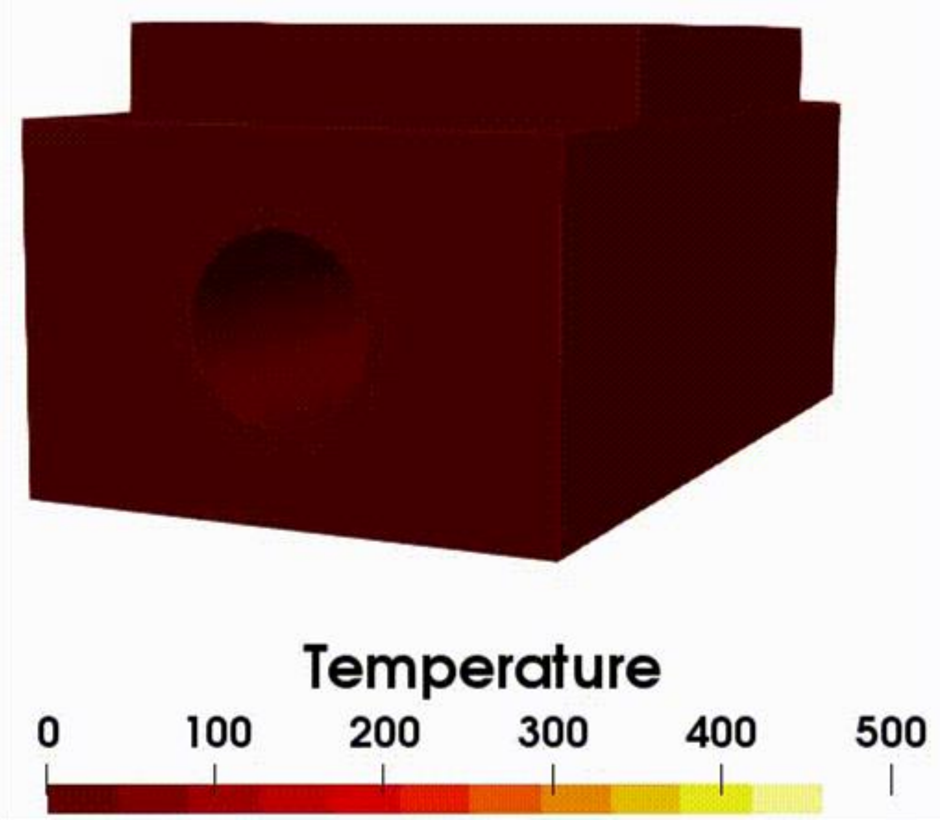

FEM Solvers

- A number of promising exascale FEM solvers exist:
  - MFEM looks very promising
    - Hybrid MPI+OMP & CUDA co-compute
    - Range of element types, including high order

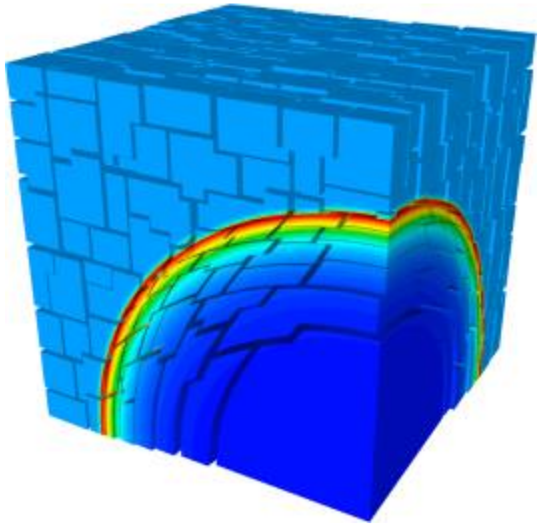

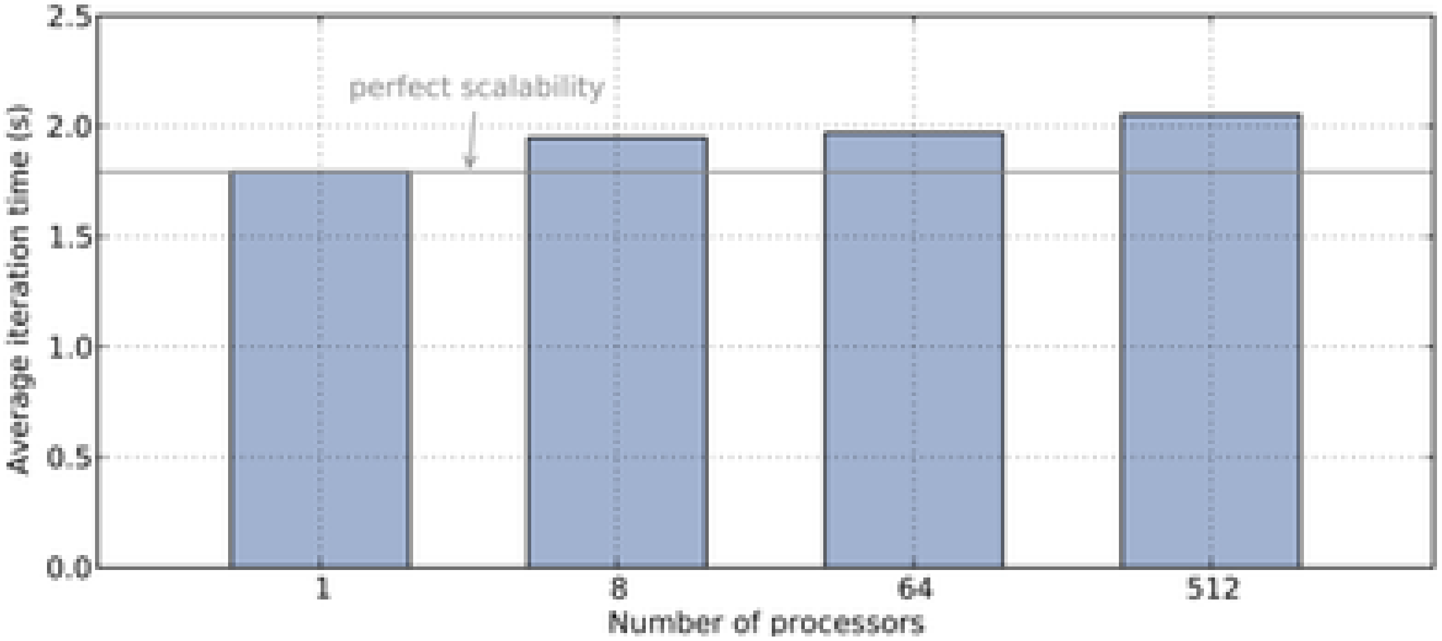

Integrated Workflows

UK Atomic Energy Authority

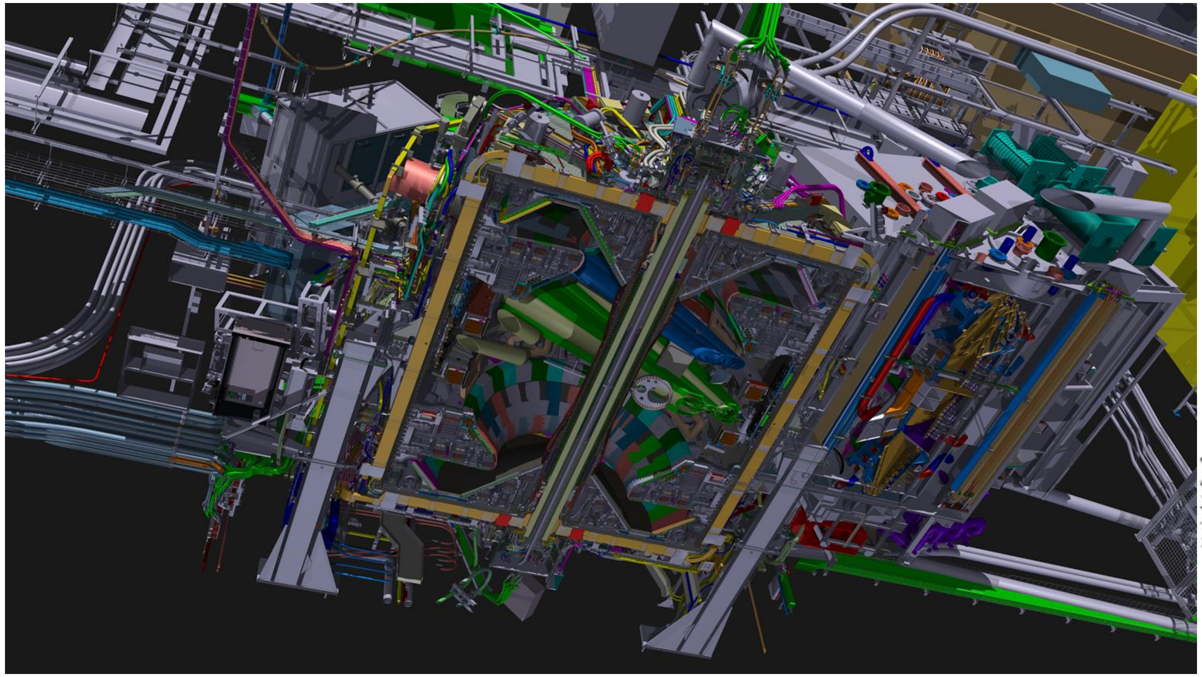

- Enabling holistic nuclear design
  - CAD as source of truth
  - Multiphysics driven
  - Use as 3rd mode of discovery

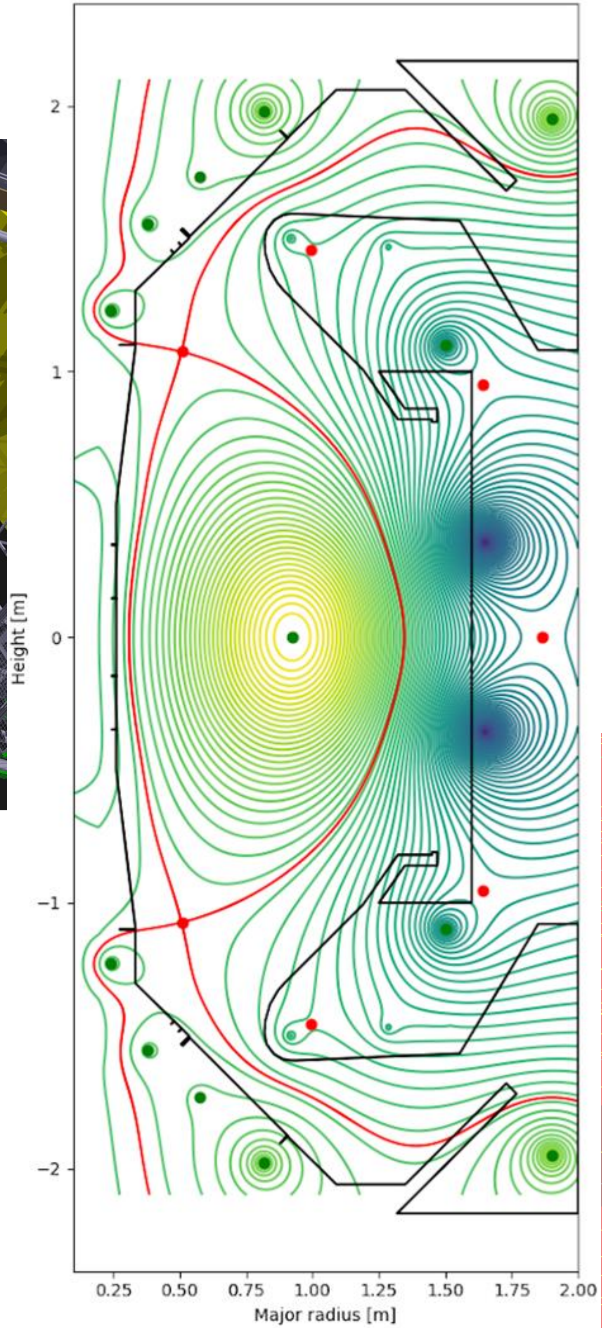

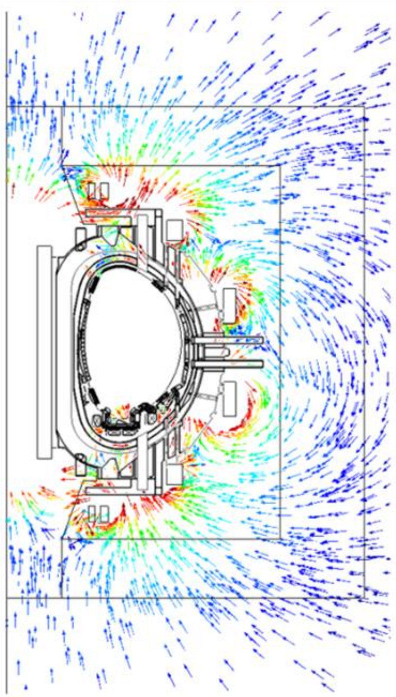

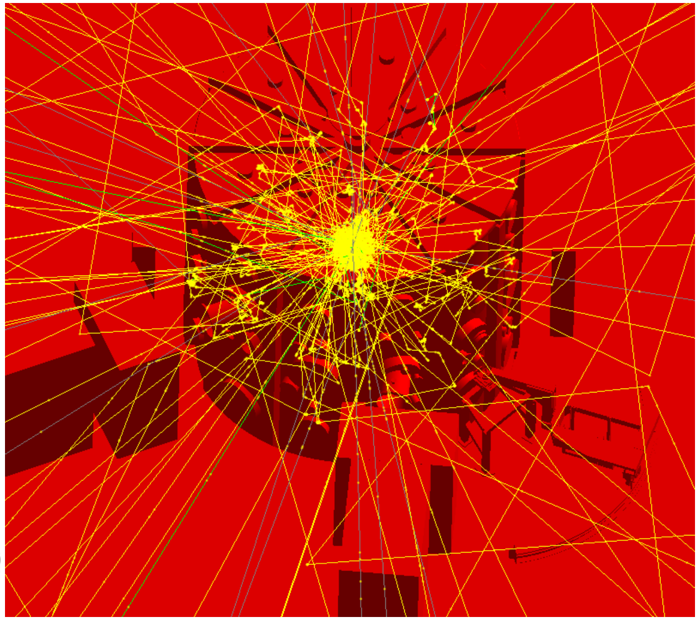

Thanks to Chris Wilson & Ruben Otin

Design Under Uncertainty

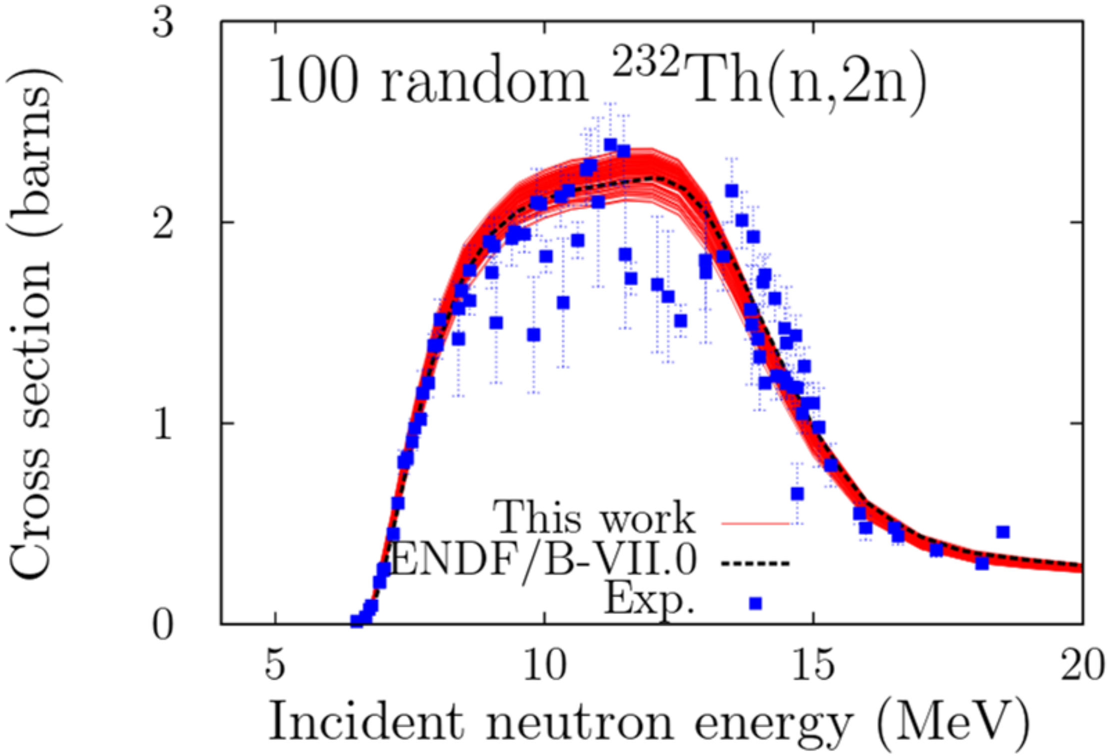

D Rochman e al., “Nuclear data uncertainty propagation: Total Monte Carlo vs. covariances”

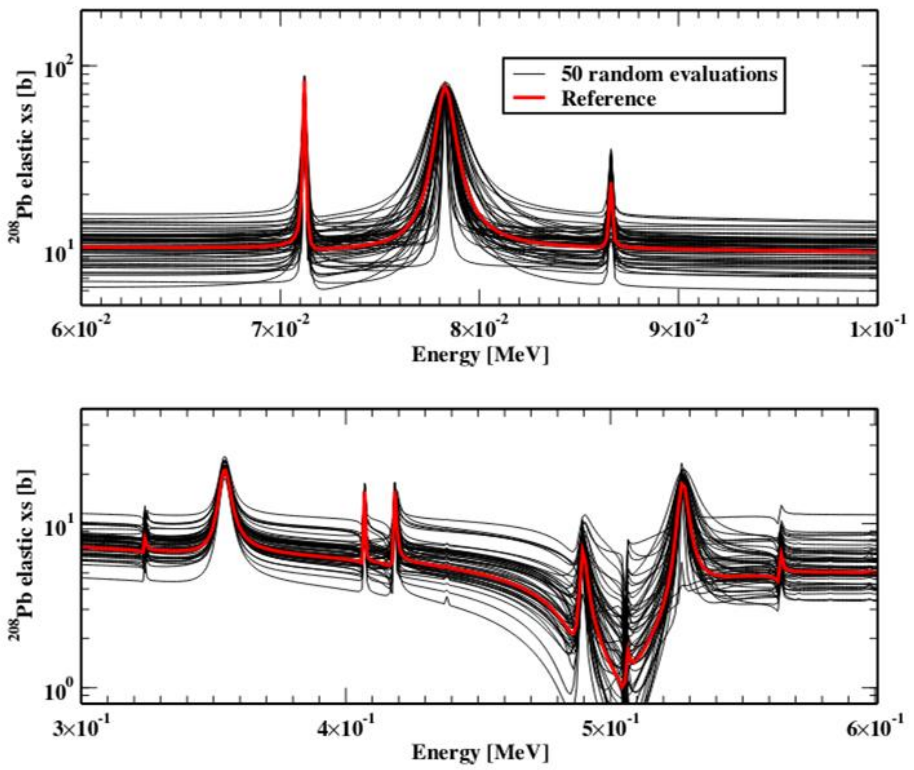

M Aufiero et al., “XGPT: Extending Monte Carlo Generalized Perturbation Theory capabilities to continuous-energy sensitivity functions”

Highest standard uncertainty propagation, but need nuclide-wise iterations e.g. 100 calculations per nuclide

Reliable when uncertainties are small and distribution of uncertainties is “well behaved” e.g. no bifurcated distributions

Design Under Uncertainty

UK Atomic Energy Authority

- Calculation of ITER neutron flux distribution with nuclear data uncertainty
  - > 400 unique materials (including "tally materials"
  - 108 Nuclides
- Perturbing only the nuclide data with 100 samples per nuclide
  - 10k individual calculations
  - Target MC Uncertainty < 1%
  - ~ 1000 CPUs ~ 1 Month
  - 10 MCPUs ~ 1 Month
    - 1 calculation 3000 s on 1 MCPU
- If we also perturb density x400
  - definitely an exascale problem!
  - Gets more challenging when including multiphysics

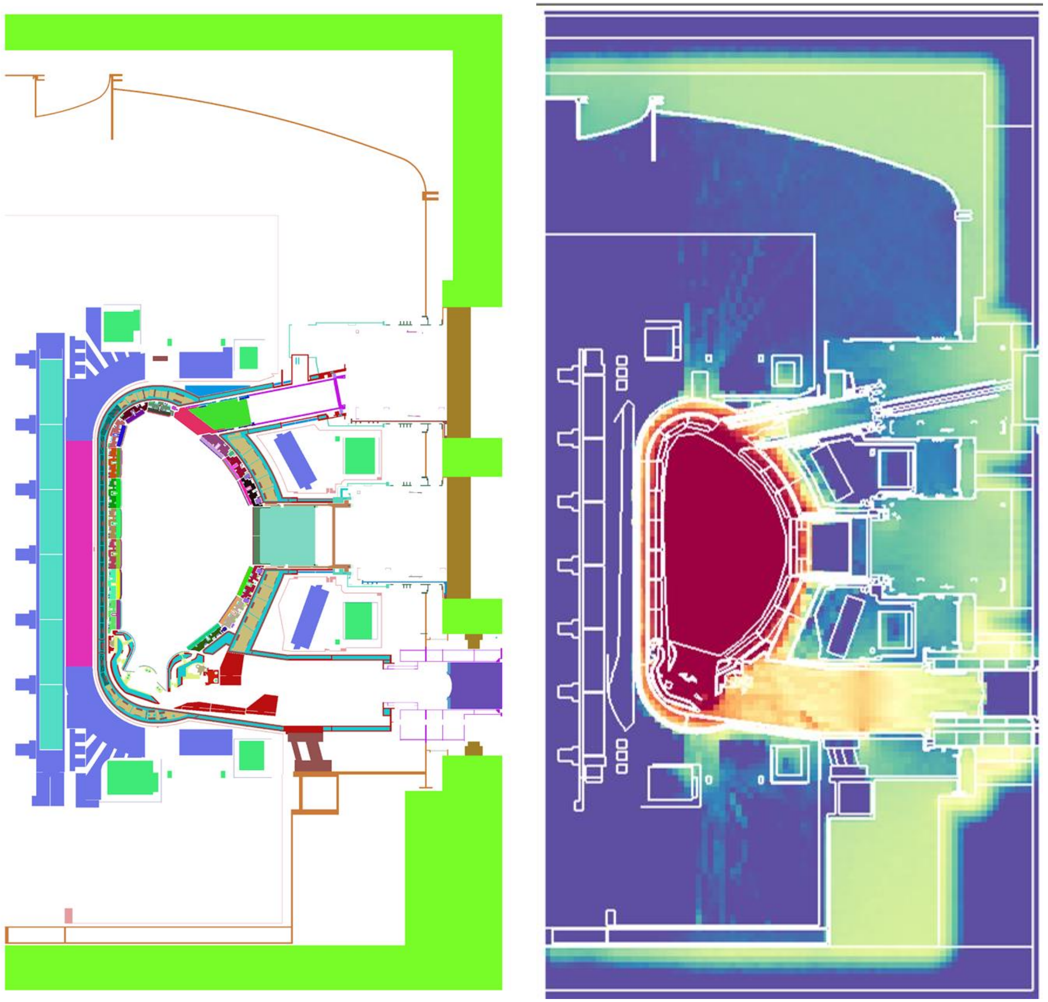

ITER C-Model thermal neutron flux distribution

Exascale Simulation Issues

UK Atomic Energy Authority

- Startup time - $5k per minute compute cost
- Stochasticity of the problem
  - Variance reduction tends to make MC scaling worse
  - Can lead to long histories - worst possible scaling - cannot be tolerated at exascale!
- Cost of tallying
  - Writing huge amount of tally data to disk (mesh tallies)
- US ECP Program
  - OpenMC is part of US ECP Program
  - Geant is involved in US ECP via collaboration with HEP
- CAD based FE Meshing is very slow - not parallel - massive bottleneck
- Visualisation - how do we visualise FE meshes with 10^{12} elements?

What are the applications?

- Fusion reactors
  - Blankets
  - Divertors
- Fission reactors
- Accelerators & targets
  - Magnet design
  - spallation neutron sources
- Fusion & HEP are solving very similar issues in very similar domains, we should work to together more
  - Particular on training for graduate - we find the same problems as you regarding programming & computational skills

Questions

UK Atomic Energy Authority

Thanks for listening!

andrew.davis@ukaea.uk