FinnFusion Annual Seminar, Helsinki, Finland, 27-28 May 2024

Advanced Computing Hubs within EUROfusion E-TASC programme

Mervi Mantsinen

ICREA and Barcelona Supercomputing Center (BSC)

This work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Programme (Grant Agreement No 101052200 — EUROfusion). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

Few words on me & my journey

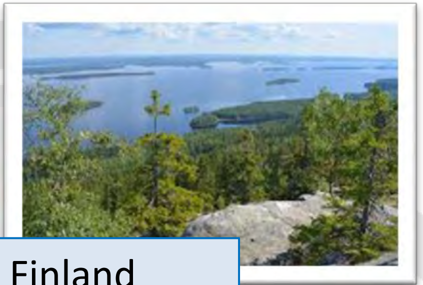

Roots in Northern Karelia, Finland

HUT (now Aalto):
- 1992 MSc in Technical Physics
- 1994 Lic. Tech
- 1999 PhD

1995-2004 at JET (UK):
- 1995-1999 PhD research
- 2000-2002 Scientific Coordination
- 2003-2004 Dep. Task Force Leader

2006-2008 IPP Garching (Germany)
- RF expert

2013-now at BSC (Spain) :
- ICREA Research Professor and Fusion Group Leader
- Since 2021 PI of Advanced Computing Hub (ACH) at BSC

Outline of this talk

Context

Fusion modelling: status and prospects in exascale and beyond

EUROfusion Theory and Advanced Simulation Coordination (E-TASC)

Advanced Computing Hubs with E-TASC

Conclusions

World's energy consumption is growing

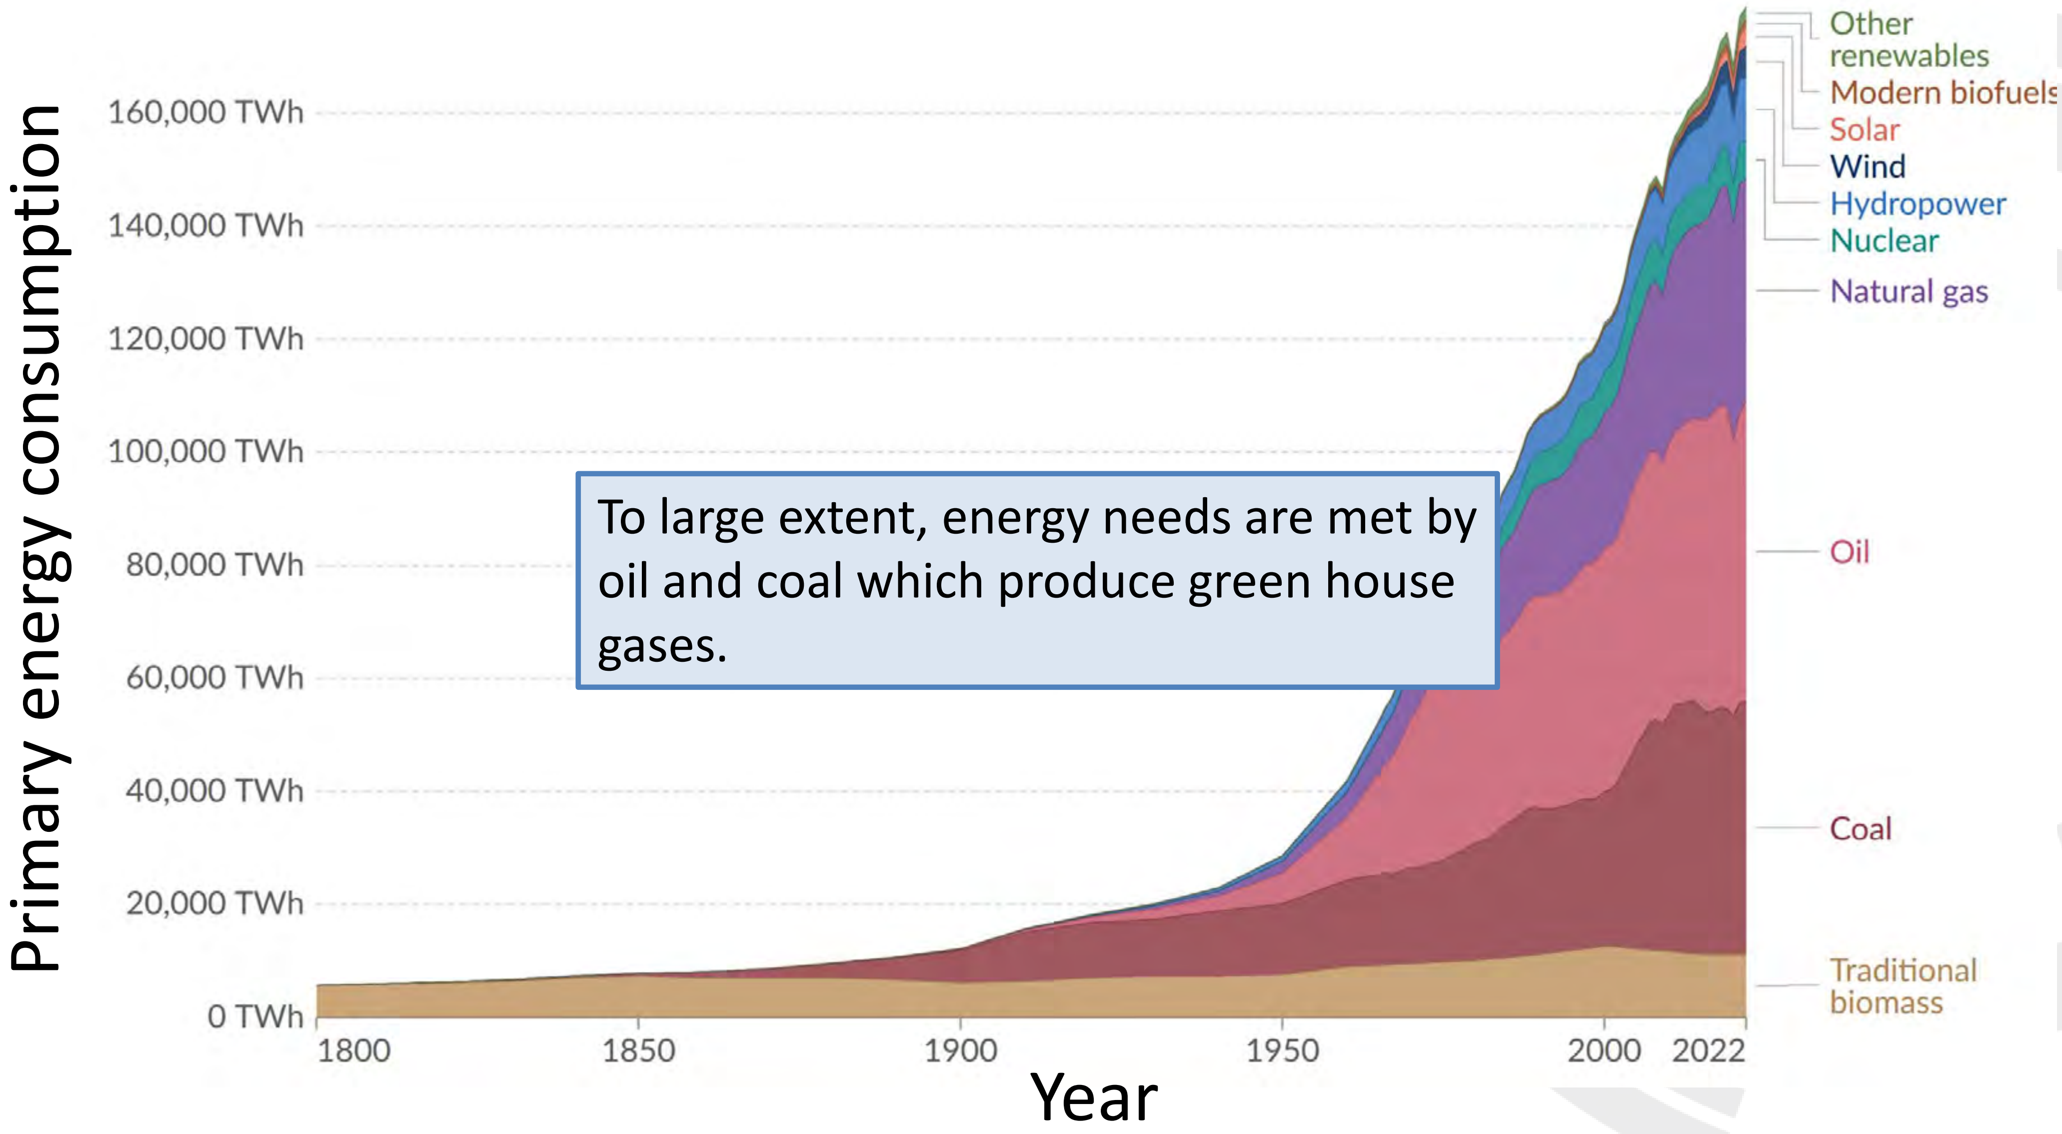

To large extent, energy needs are met by oil and coal which produce green house gases.

Global average surface temperature will continue to increase due to green house gases

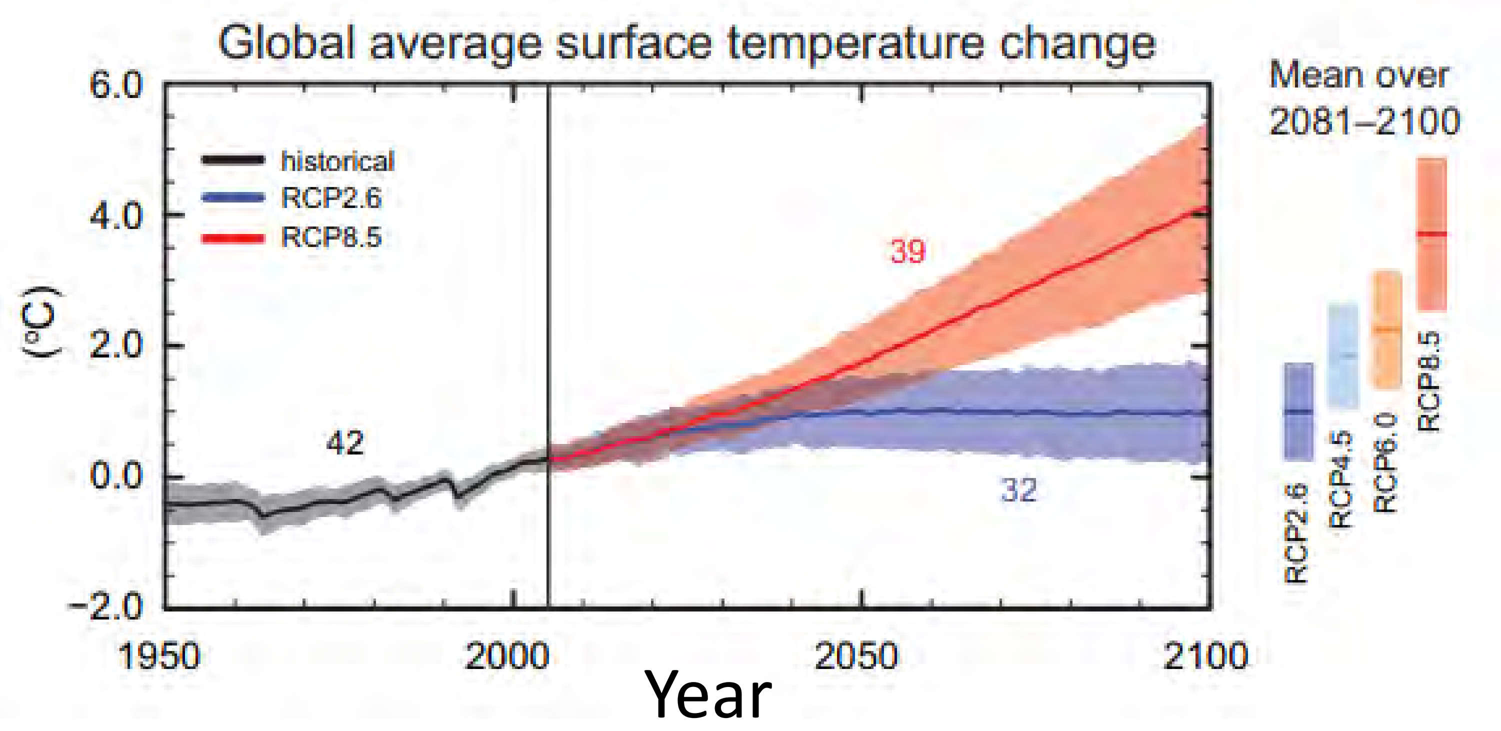

Finland by 2100 (RCP8.5):

- winters will be 4-10 °C and summers 2-7 °C hotter.
- winter rain fall will increase by 4–30 %.

- To reduce the damage, **action is needed**.
- There is **room and need for nuclear fusion**, which is carbon-free, safe and potentially abundant energy source, **within our future energy mix**.

Fusion research is on its path towards burning plasmas

Plasma “burns” when plasma self-heating compensates energy losses

Key requirements:

- Large central density n and temperature T
- Long enough energy confinement time $\tau = E_{plasma}/P_{loss}$

Main current limitations:

- Onset of large-scale instabilities
- Small scale instabilities, i.e. turbulence

Magnetic confinement fusion

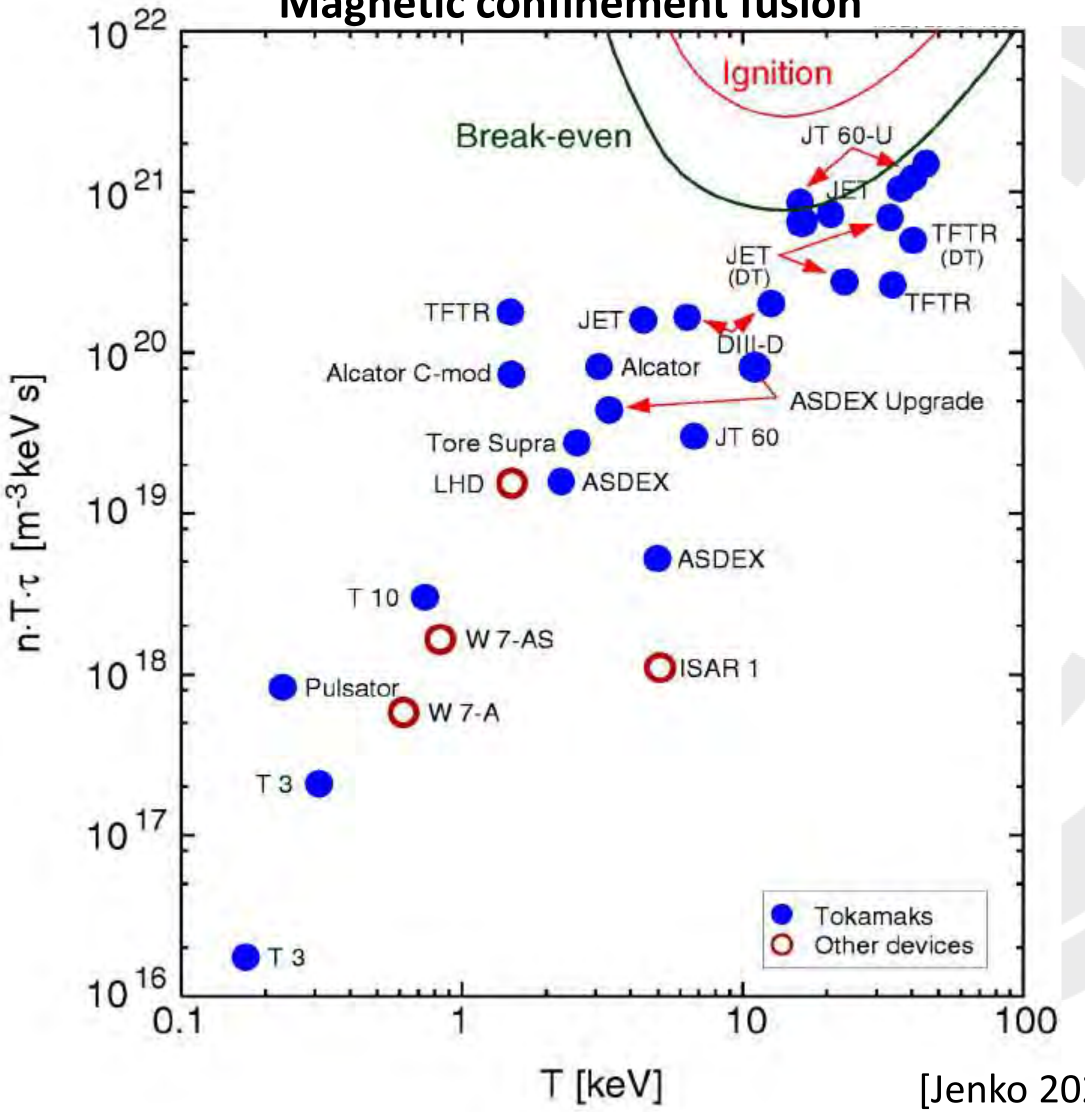

[Jenko 2024]

Recent progress has renewed confidence in fusion power realization

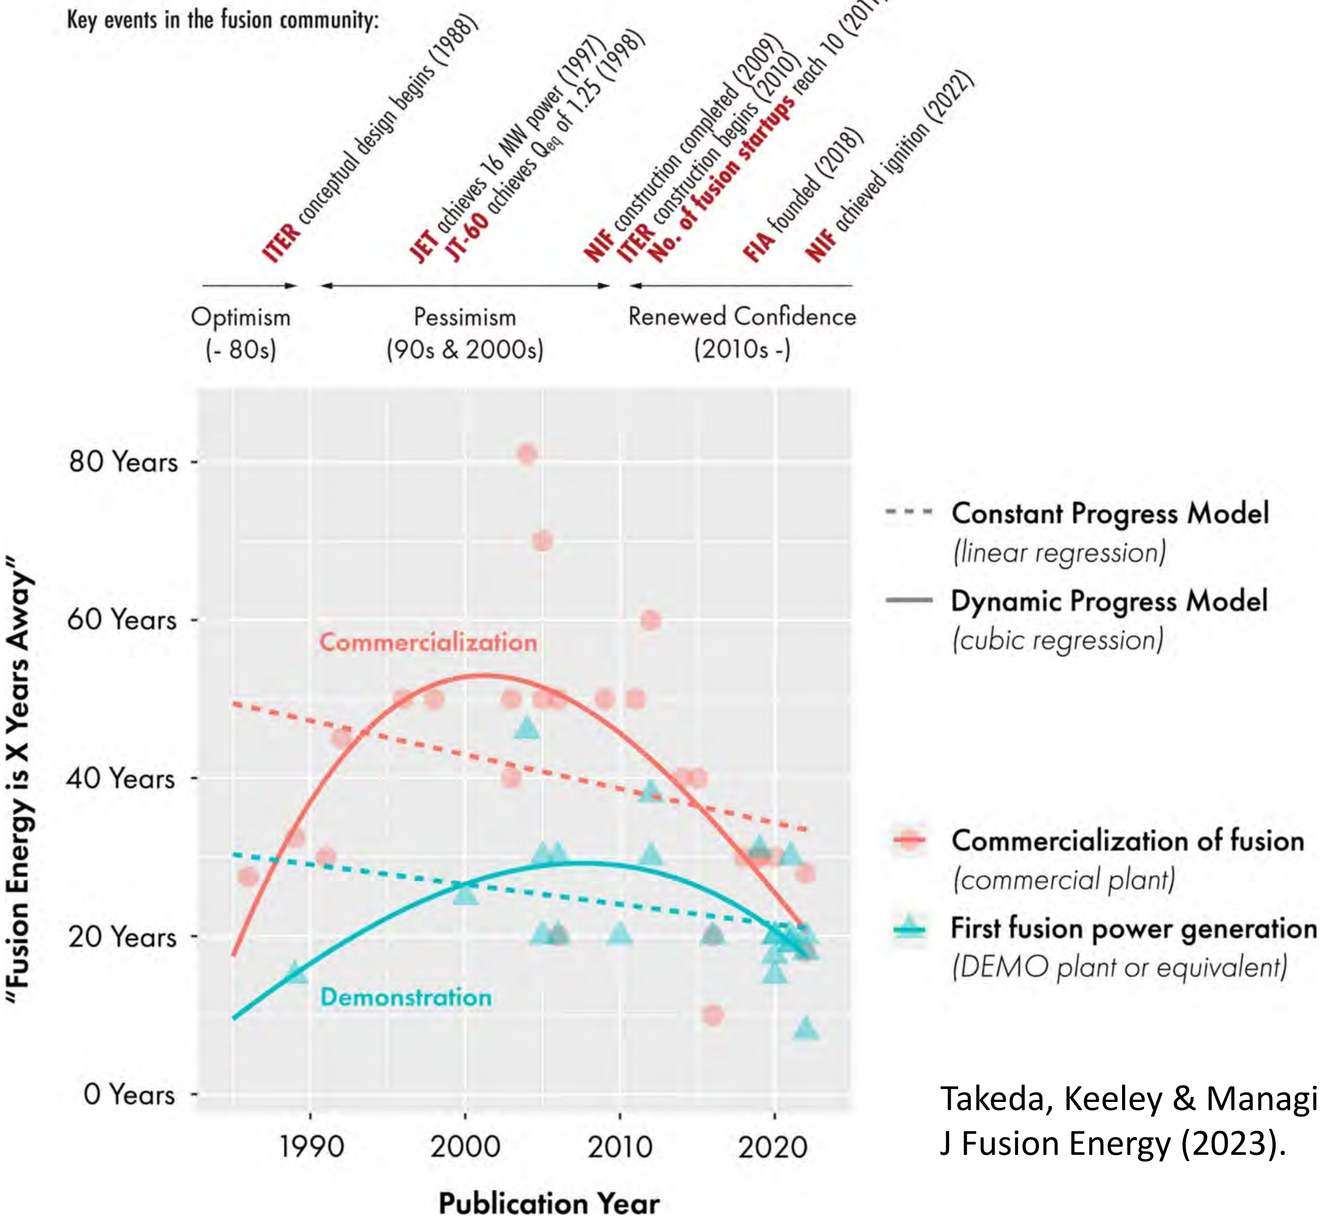

Takeda, Keeley & Managi
J Fusion Energy (2023).

Update of ITER baseline expected in 2024
Next phase: ITER operation

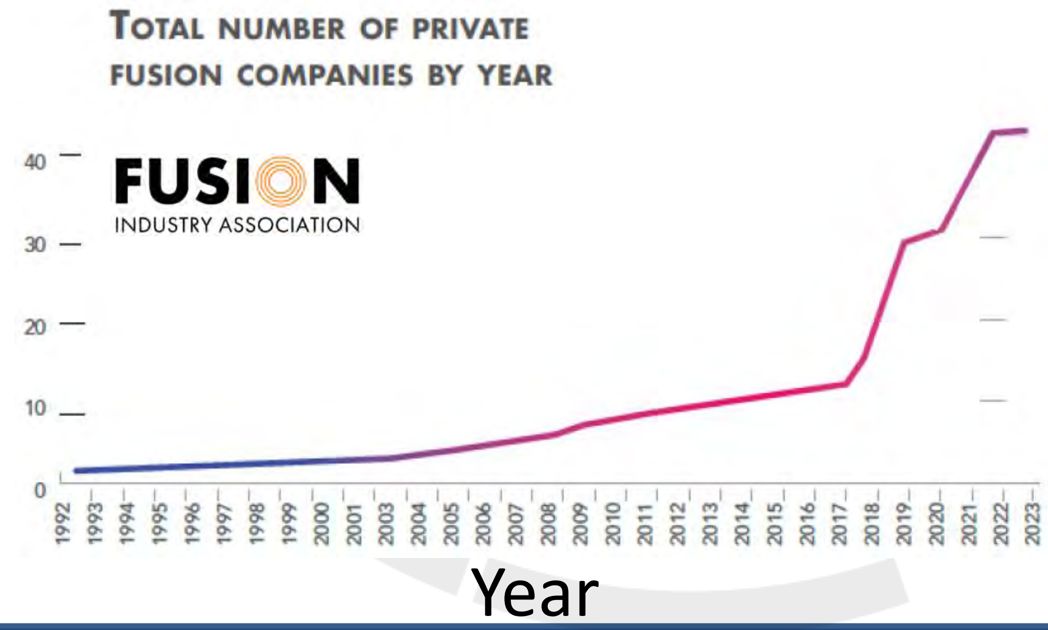

European Research Roadmap

- Structured way forward to commercial energy from fusion.
- Basis for the programmes of EUROfusion, F4E and IFMIF-DONES
- DEMO project to design a commercial fusion power plant.

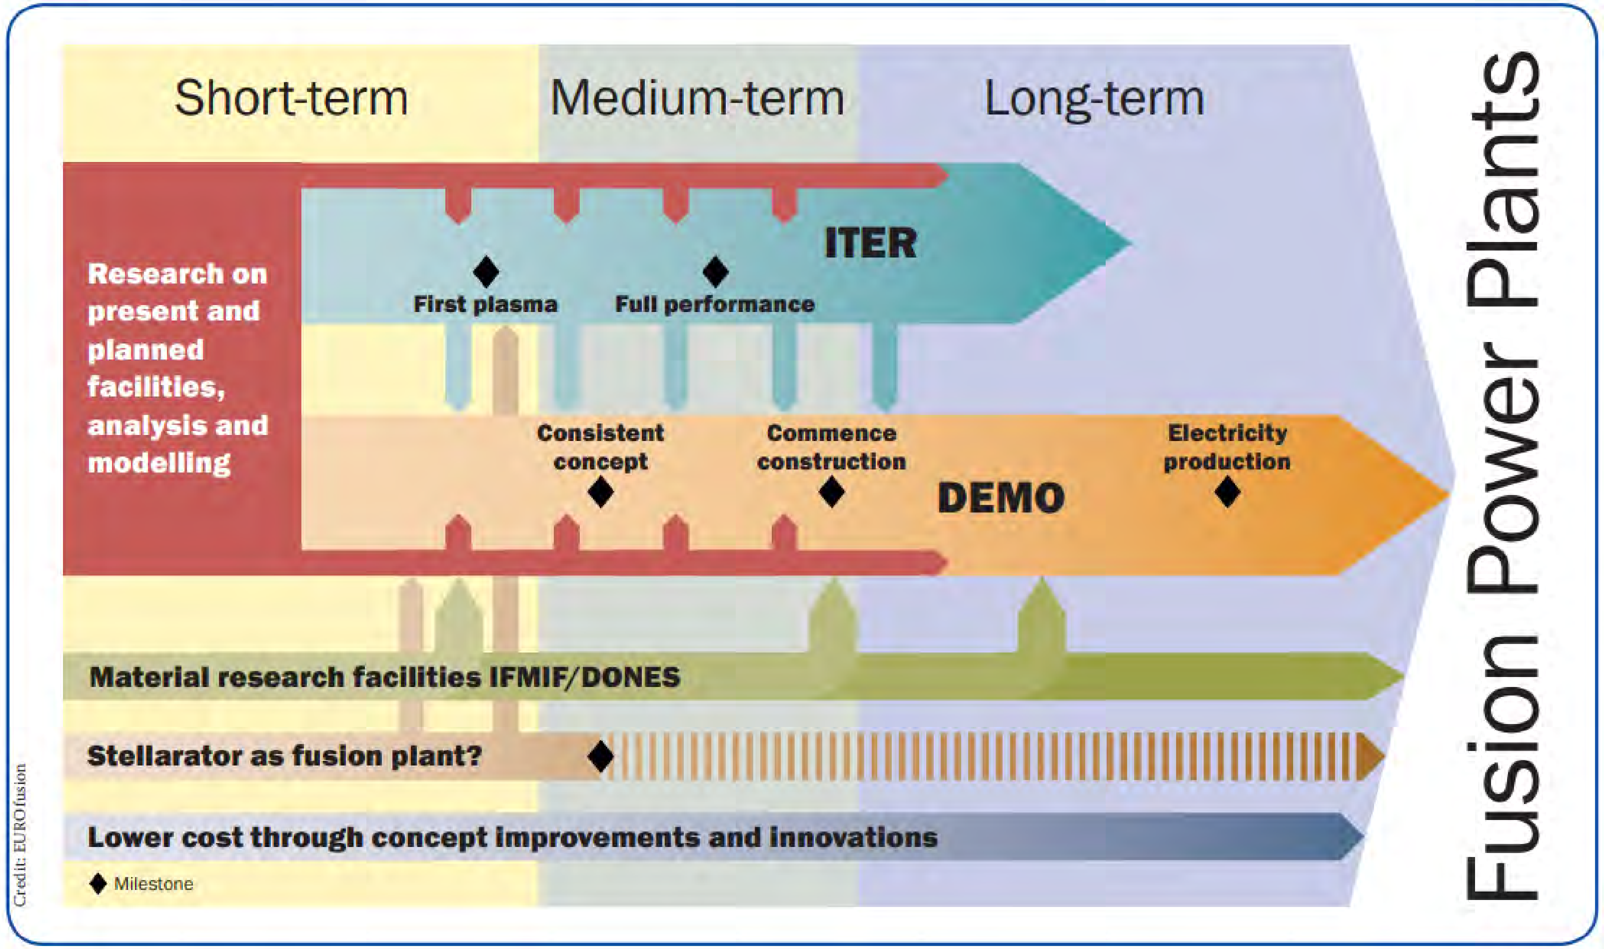

Need for modelling and theory effort in fusion

- Moving from ITER to DEMO is challenging.
- Experimental data from ITER and IFMIF-DONES will not be sufficient to design DEMO with confidence.
- Need for **high-quality suite of research codes** to model data from existing facilities and to extrapolate to future devices to accelerate the development of fusion energy

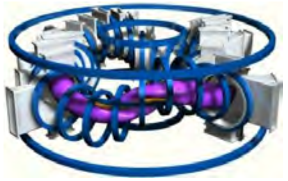

WEST

ASDEX Upgrade

TCV etc.

iter

&

Future fusion reactor

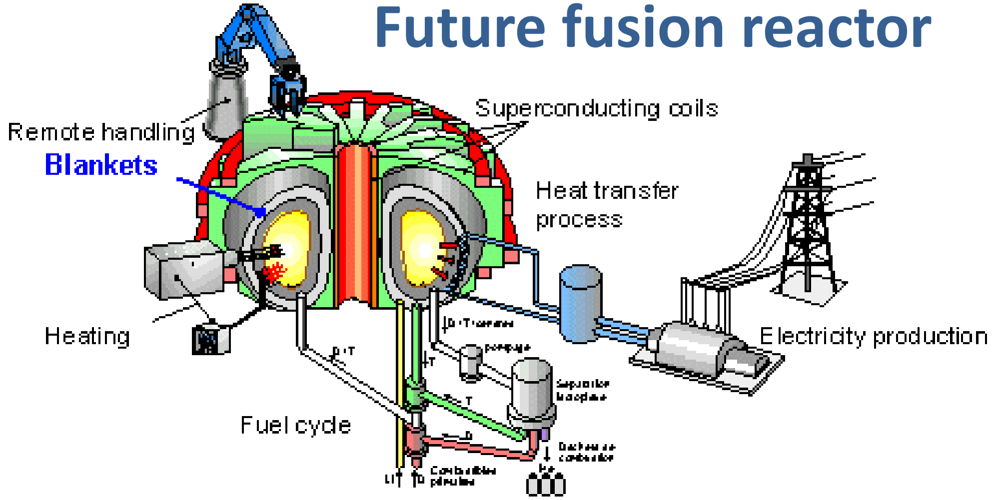

EUROfusion prepares for this transition with Theory and Advanced Simulation Coordination (E-TASC) programme, to maximize the benefits of facility investments.

Outline of this talk

Context

Fusion modelling codes: status and prospects in exascale and beyond

EUROfusion Theory and Advanced Simulation Coordination (E-TASC)

Advanced Computing Hubs with E-TASC

Conclusions

Fusion reactor modelling involves complex plasma, materials and engineering problems that are coupled

It is multi-scale, multi-physics challenge.

- Tremendous progress in recent years.
- However, all these phenomena cannot be yet modelled in a single simulation in a reasonable time.

Typical numerical methods and example codes for selected modelling problems in fusion

Purple: 12 codes supported by BSC- ACH

Blue: example HPC codes we use in our own fusion R&D at BSC

For fusion materials modelling codes, see e.g. [Nordlund; Knaster]

Role of High-Performance Computing (HPC) in fusion modelling

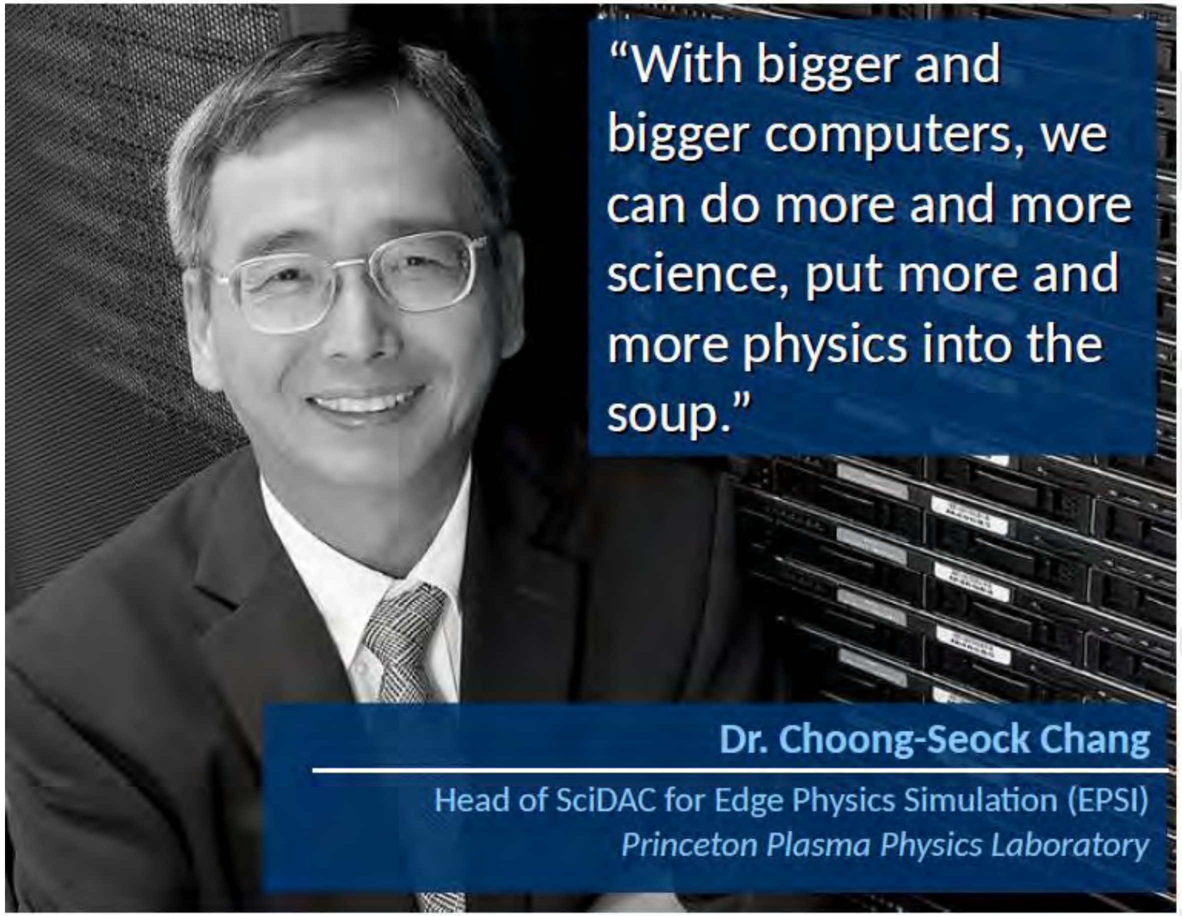

HPC has entered the exascale era

- **2 exascale supercomputers (in USA)** with > 10^{18} Floating point operations / second (FLOPS, FLOP/s)
- **8 pre-exascale supercomputers (4 in EU)** with > 10^{17} FLOPS
- Unprecedented level of **heterogeneity in architecture**
- **GPU are dominating**
- Large **variability in CPU/GPU combination**

Rank	System	Cores	Rmax (PFlop/s)	Rpeak (PFlop/s)	Power (kW)
1	Frontier - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE DOE/SC/Oak Ridge National Laboratory **United States**	8,699,904	**1,206.00**	1,714.81	23
2	Aurora - HPE Cray EX - Intel Exascale Compute Blade, Xeon CPU Max 9470 52C 2.4GHz, Intel Data Center GPU Max, Slingshot-11, Intel DOE/SC/Argonne National Laboratory **United States**	9,264,128	**1,012.00**	1,980.01	39
3	Eagle - Microsoft NDv5, Xeon Platinum 8480C 48C 2GHz, NVIDIA H100, NVIDIA Infiniband NDR, Microsoft Azure Microsoft Azure United States	2,073,600	561.20	846.84	
4	Supercomputer Fugaku - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,630,848	442.01	537.21	30
5	LUMI - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE EuroHPC/CSC **Finland**	2,752,704	379.70	531.51	
6	Alps - HPE Cray EX254n, NVIDIA Grace 72C 3.1GHz, NVIDIA GH200 Superchip, Slingshot-11, HPE Swiss National Supercomputing Centre (CSCS) **Switzerland**	1,305,600	270.00	353.75	5
7	Leonardo - BullSequana XH2000, Xeon Platinum 8358 32C 2.6GHz, NVIDIA A100 SXM4 64 GB, Quad-rail NVIDIA HDR100 Infiniband, EVIDEN EuroHPC/CINECA **Italy**	1,824,768	241.20	306.31	7
8	MareNostrum 5 ACC - BullSequana XH3000, Xeon Platinum 8460Y+ 32C 2.3GHz, NVIDIA H100 64GB, Infiniband NDR, EVIDEN EuroHPC/BSC **Spain**	663,040	175.30	249.44	4
9	Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148.60	200.79	10
10	Eos NVIDIA DGX SuperPOD - NVIDIA DGX H100, Xeon Platinum 8480C 56C 3.8GHz, NVIDIA H100, Infiniband NDR400, Nvidia NVIDIA Corporation United States	485,888	121.40	188.65	

With more computing power we can simulate bigger systems

TCV@EPFL	JT-60: 100x TCV	ITER: 500x TCV	DEMO: 5000x TCV
	Q = 1.25	Q = 10, 500MW	Q = 25, 2000MW
100 Teraflops (10^{14} FLOPS)	10 Petaflops (10^{16} FLOPS)	100 Petaflops (10^{17} FLOPS)	1 Exaflops (10^{18} FLOPS)

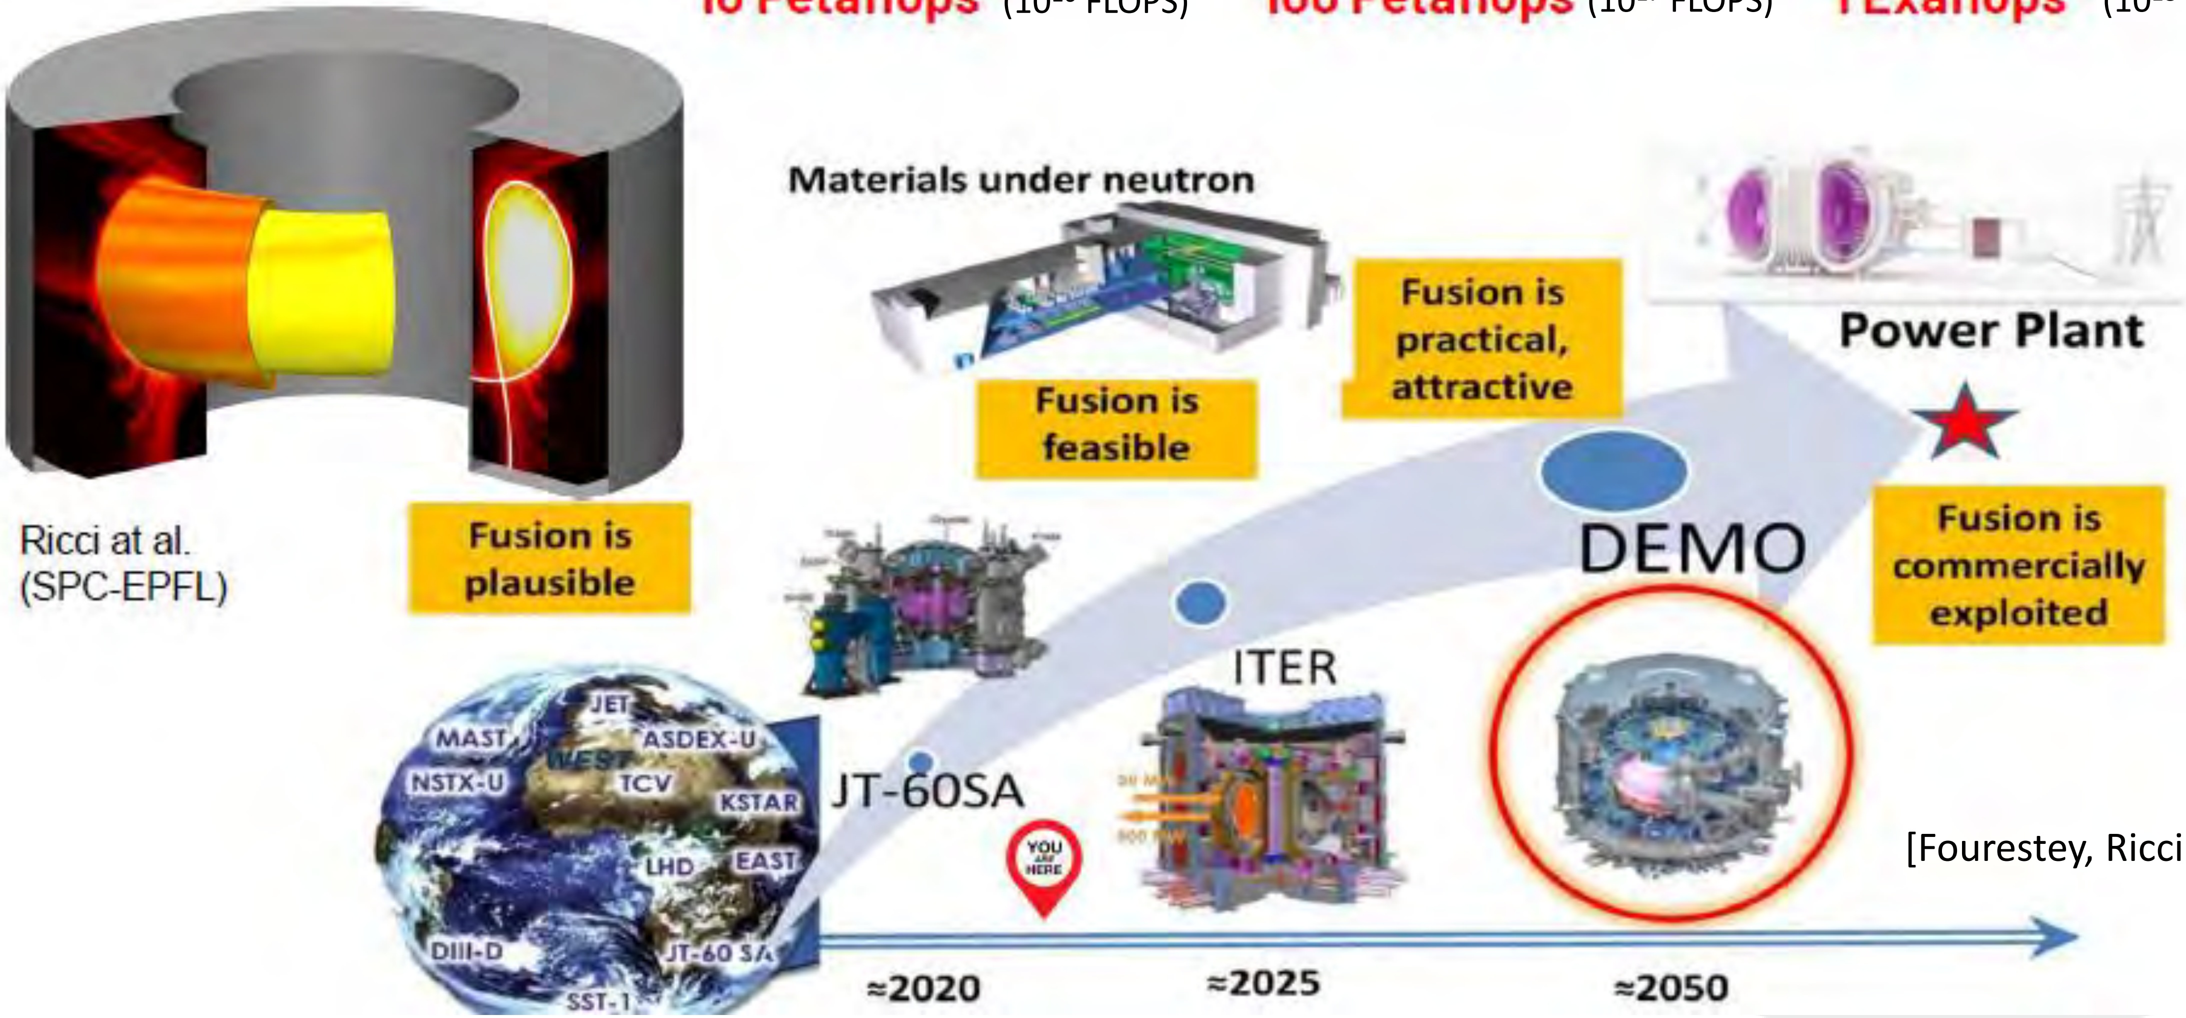

[Fourestey, Ricci 2023]

Fusion reactor modelling capability increases with computing power too

By bridging spatio-temporal scales and connecting different physical processes, we are moving towards a **fusion reactor digital twin** and **a virtual fusion plasma**.

Exascale brings new opportunities in fusion modelling

Exascale allows us to make predictions of **larger systems with more fidelity**.

To take full advantage of new resources, we need to adapt fusion codes to the emerging HPC platforms.

There are also many opportunities to **improve the computational efficiency** of the codes.

Outline

Context

Fusion modelling: status and prospects in exascale and beyond

EUROfusion Theory and Advanced Simulation Coordination (E-TASC)

Advanced Computing Hubs with E-TASC

Conclusions

Guiding principles of EUROfusion E-TASC

Innovative theory and simulation research takes typically place when it is driven by scientists and engineers themselves.

However, **production of a new portfolio of EUROfusion standard software requires** a more directed approach.

To accommodate both, two inter-linked structures have been created with E-TASC:

(1) **Theory-Simulation-Verification-Validation (TSVV) Tasks**

(2) **Advanced Computing Hubs (ACHs)**

These are in addition to **JET data center**, at DTU (Denmark), and **dedicated supercomputer** at CINECA (Italy)

[Litaudon, 2021]

TSVVs and ACHs work together to produce EUROfusion standard software

TSVVs perform research and channel science and engineering into scientific codes

ACHs modernize and industrialize research codes into HPC standards

EUROfusion's Software Standards:

- **SOFTWARE ENGINEERING**: Version control, coding standards, test-driven.
- **CODE INTERFACES**: Graphical User Interface, post-processing and visualisation tools, interfaces to the IMAS Data Dictionary.
- **VVUQ**: Code Verification and Validation, uncertainty quantification, reports/papers available, validation against experimental results.
- **CODE DISSEMINATION**: Up-to-date release version of the source code available, trainings provided.
- **CODE DOCUMENTATION**: High-quality technical documentation and user manual available.
- **USER SUPPORT**: Responsive support team in place, tools for managing support requests.

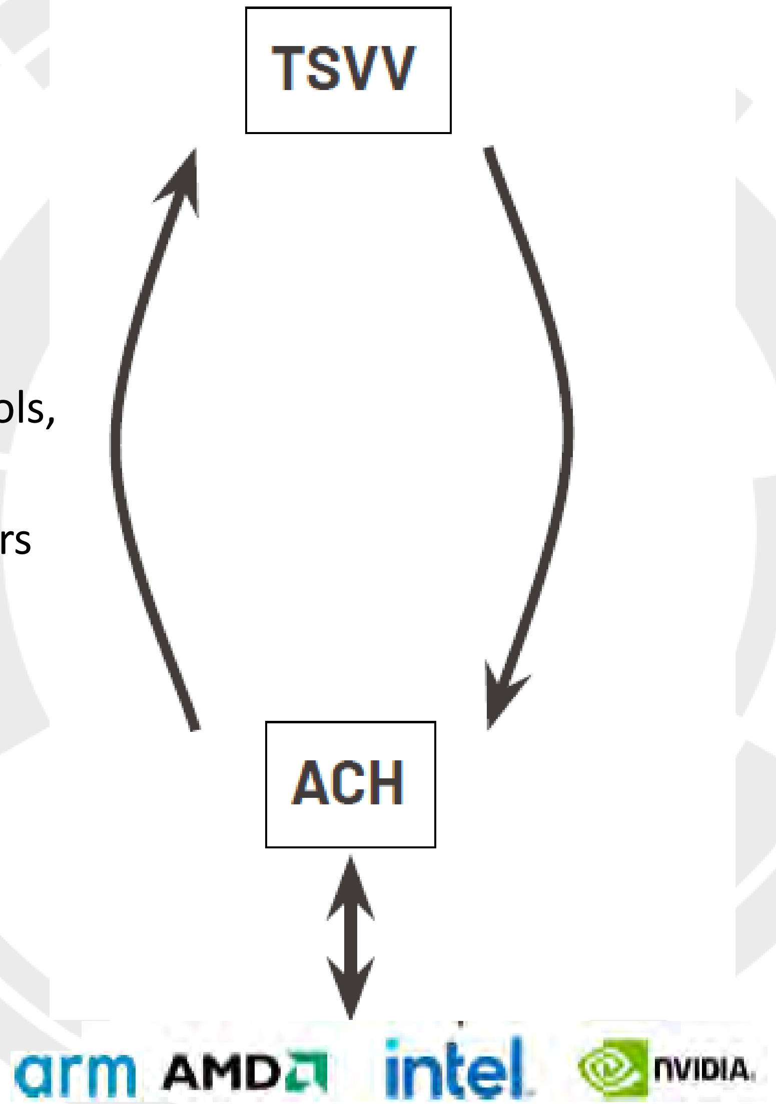

Theory, Simulation, Verification and Validation (TSVV) Tasks

TSVV tasks aligned with the European Fusion Roadmap priorities and missions

Aim to maximize the benefit from EUROfusion facility investments

Address high-priority issues for ITER operation and DEMO design

- Advance understanding and predictive capabilities

Develop a high-quality suite of European standard software

- To model data from EUROfusion facilities
- To extrapolate to future facilities: ITER and DEMO

[Litaudon, 2021]

E-TASC key challenges during Horizon Europe (2020-2025/2027)

Validated predictive capability of the **L-H transition and pedestal physics** in ITER and DEMO

Validated predictive capability for **heat exhaust** in ITER and DEMO

Integrated scientific work for ITER and DEMO on

- **plasma-wall interactions**
- **disruptions** including their prediction, mitigation, and avoidance
- **burning plasmas**

[Litaudon, 2021]

15 TSVV tasks to address high-priority issues

1	Physics of the L-H Transition and Pedestals
2	Physics Properties of Strongly Shaped Configurations
3	Plasma Particle/Heat Exhaust: Fluid/Gyrofluid Edge Codes
4	Plasma Particle/Heat Exhaust: Gyrokinetic/Kinetic Edge Codes
5	Neutral Gas Dynamics in the Edge
6	Impurity Sources, Transport, and Screening
7	Plasma-Wall Interaction in DEMO
8	Integrated Modelling of transient MHD Events
9	Dynamics of Runaway Electrons in Disruptions and Start-Up
10	Physics of Burning Plasmas + fast ions physics W7X
11	Validated Frameworks for the Reliable Prediction of Plasma Performance and Operational Limits in Tokamaks
12	Stellarator Optimization
13	Stellarator Turbulence Simulation
14	Multi-Fidelity Systems Code for DEMO

15 Pulse Design Tool, new for 2024-2025

[Litaudon, 2021]

Overarching goals and scope of E-TASC

Provide validated predictive capabilities

- Convert existing research codes into professional and widely used tools
- Develop new codes to fill any gaps

Include:

- Multi-fidelity modelling
- Validation, verification, and uncertainty quantification (VVUQ)

Involve range of existing and future fusion devices for validation and applications

Maintain flexibility to launch new tasks and to close completed tasks

Advanced Computing Hubs in support of entire EUROfusion simulation programme

[Litaudon, 2021]

Outline

Context

Fusion modelling: status and prospects in exascale and beyond

EUROfusion Theory and Advanced Simulation Coordination (E-TASC)

Advanced Computing Hubs with E-TASC

Conclusions

EUROfusion Advanced Computing Hubs (ACHs)

ACHs **bring together fusion physicists, mathematicians, materials scientists, engineers and computer scientists** to make efficient use of high-performance computers (HPC) and, thereby, to accelerate the development of fusion energy.

5 ACHs (2021-2025)

- **3 ACHs in HPC:** Scalable algorithms, code parallelization, performance optimization, code refactoring, GPU enabling etc.
- **1 ACH in Integrated Modelling and Control:** Code adaptation to ITER Modelling and Analysis Suite (IMAS), IMAS framework development, code integration etc.
- **1 ACH in Data management:** Open access, data management, data analysis tools, aspects of AI & Validation, verification, and uncertainty quantification (VVUQ) etc.

[Fourestey, 2023]

ACH at CIEMAT / Barcelona Supercomputing Center (BSC), Spain

- BSC, affiliated to EUROfusion via CIEMAT, hosts one of ACHs **in HPC**
- Recently described in [Saez PPCF 2024].
- Involves **three groups** at BSC:
  - 7 people in Fusion Group
  - 2 people in Operations
  - 2 people in Best Practices for Performance and Programmability (BePPP)
- **Total effort**: 328 pm ≈ **27 ppy in 2021-2025**
- From 2024 on, at **full size** of ~100 pm/year ≈ **8 ppy/year**.

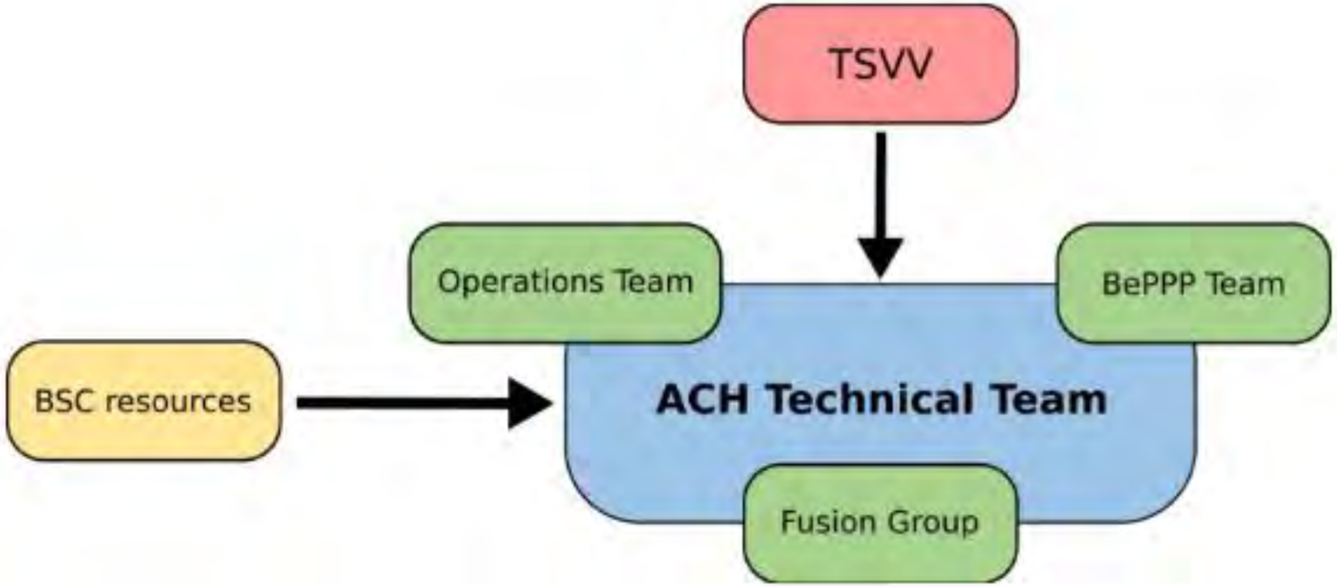

BSC ACH codes and tasks (2021-2024)

Up to now, **12 codes supported.**

Code characteristics:

- Mainly plasma physics codes
- Mostly CPU-only
- Written in C, C++ and Fortran
- MPI and/or OpenMP
- Under active development by physicists, mathematicians etc.

ACH tasks mainly on

- performance optimization
- GPU-enabling
- scalable algorithms
- code parallelization

Our general principles:

- Maximize performance
- Aim for good portability between architectures
- Aim for least modification to the code/no rewrite

ACH task highlight: reduce memory requirements by matrix compression in JOREK

Challenges

Huge matrices that do not fit into a single compute node

Objectives

- Reduce memory consumption and improve performance
- Apply factorization techniques to decompose the matrices
- Compress the matrices
- Write code to read, factorize, compress and compute matrix-matrix and matrix-vector operations

Issues

- ScaLAPACK library offers a routine to compute the SVD of a given matrix.
- ScaLAPACK requires that the matrix distribution format is 2d block cyclic, which is not implemented in JOREK
- Switching from one to the other is not trivial

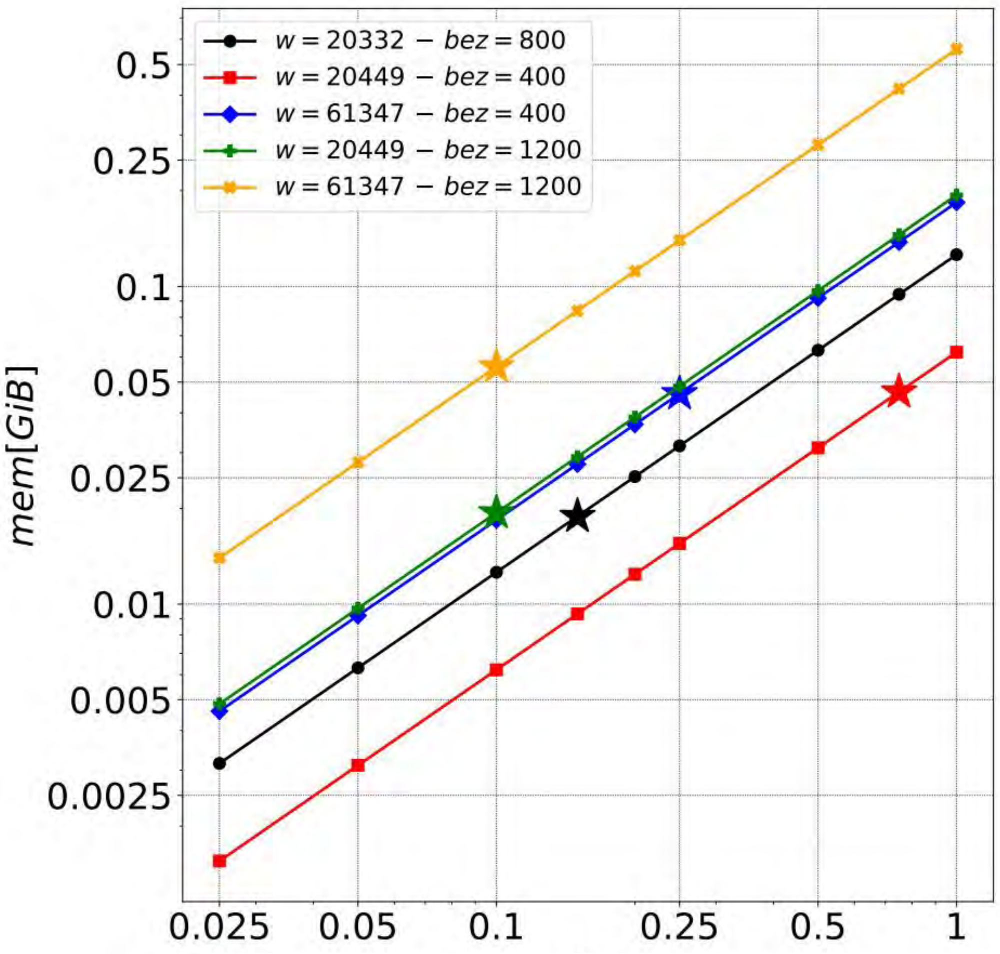

Rate of retained singular values

[Cipolletta, submitted PPCF]

ACH task highlight: Implement Parallel Poisson solver in SPICE2

Challenges

All internal routines are parallel (MPI) except the Poisson solver, which is serial.

Objectives

Implement parallel solver required to perform detailed simulations:

The number of cores in simulations can be increased to at least 128 (current practical limit is ~32) and the grid size can be increased (current limit is 4000 cells in one dimension)

Issues

– Conjugated gradient solver implementation was not as good as expected.

– PETSc library solver (KSP linear solver)

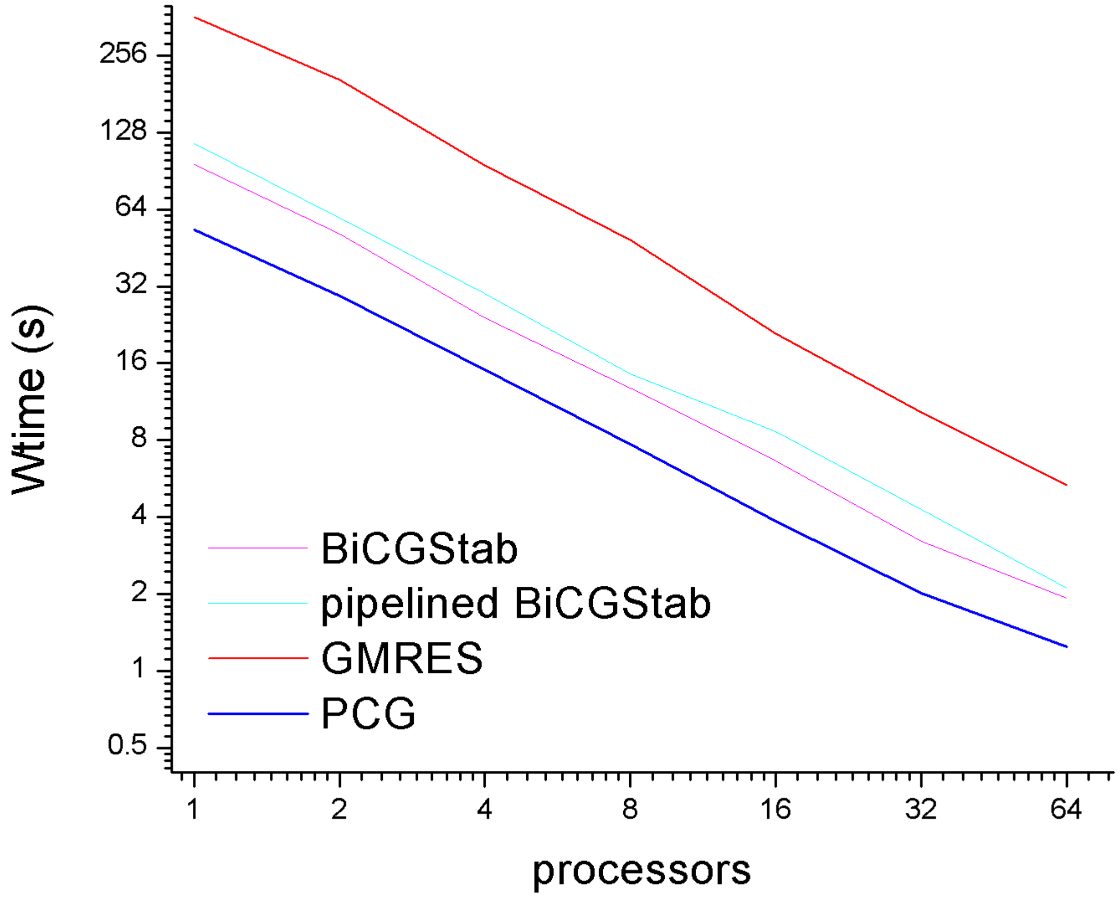

[Saez PPCF 2024]

ACH task highlight: ERO2

Challenges

Reduce time execution and port the code to GPUs

Objectives

- Perform a performance analysis
- Implement first version of OpenACC

Issues

- Reduce load imbalance on particle distribution.
- Implement iterative BFS (Breadth-First Search) octree traversal for GPUs

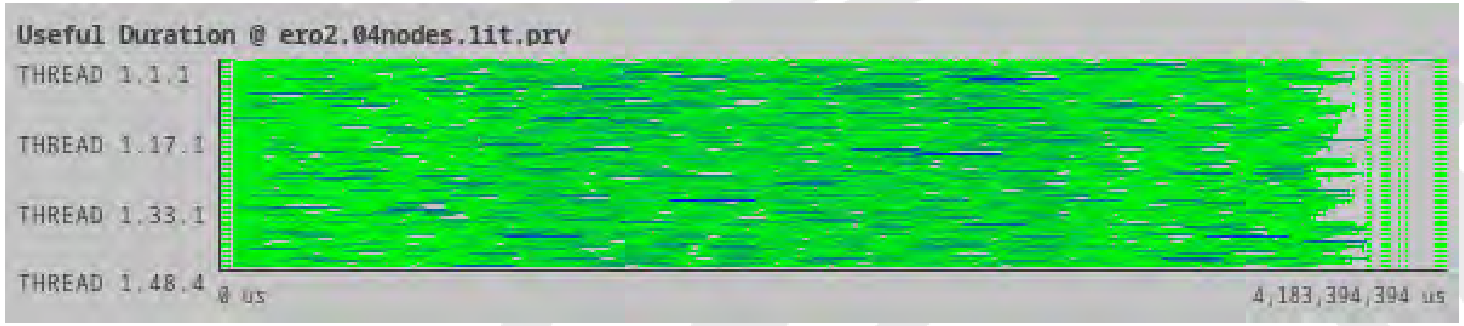

	192(48x4)[1]	384(96x4)[2]	768(192x4)[3]	1536(384x4)[4]
-- Hybrid Parallel efficiency	80.20	70.59	58.33	40.43
-- MPI Parallel efficiency	88.95	79.11	66.42	47.71
-- MPI Load balance	95.46	90.80	84.69	68.30
-- MPI Communication efficiency	93.18	87.13	78.42	69.86
-- Serialization efficiency	93.20	87.18	78.53	70.03
-- Transfer efficiency	99.98	99.94	99.86	99.75
-- OpenMP Parallel efficiency	90.17	89.23	87.83	84.74
-- OpenMP Load Balance	95.87	95.53	90.39	91.17
-- OpenMP Communication efficiency	94.05	93.41	97.17	92.94

Percentage(%) colour scale: 0, 20, 40, 60, 80, 100

[Vinyals, PASC 2024, Saez PPCF 2024]

ACH task highlight: GPU porting

Approach to GPU porting	Codes (ACH)	Observations
Library encapsulation (e.g. Kokkos, PETSc, AmgX, BLAS/Lapack…)	• BOUT++ (BSC) • XTOR-K (BSC) • GBS (EPFL) • GENE (EPFL)	• easy to use, good out-of-the-box performance • not applicable to all codes and might require some rewrite (e.g. data structures)
Pragma directives (OpenMP offload / OpenACC)	• ERO2 (BSC) • XTOR-K (BSC) • BIT1 (BSC) • ASCOT5 (EPFL) • CAS3D (EPFL) • GyselaX (EPFL) • ORB5 (EPFL) • Soledge3X (EPFL)	• In general, sound approach • OpenMP offload is probably the best choice but it is not mature • small effort to port with OpenACC but it is not supported by all
Cuda/ROCm/SYCL	• ERO2 (BSC) • FELTOR (EPFL) • GBS (EPFL) • GRILLIX (EPFL)	• portable, relatively easy to use • some code rewrite and possible algorithmic modifications

Whatever approach(es) one chooses, one can **expect at least some code rewrite.** The higher the quality of the code structure is, the easier it typically is.

[Saez, 2024; Fourestey, 2023]

Outline

Context

Fusion modelling: status and prospects in exascale and beyond

EUROfusion Theory and Advanced Simulation Coordination (E-TASC)

Advanced Computing Hubs with E-TASC

Conclusions

Conclusions

- **Fusion theory and advanced simulations**, developed within **EUROfusion new E-TASC programme**, **are essential for progress in fusion R&D**. This is because empirically based predictions are uncertain in unexplored environments like ITER and DEMO.
- Via a synergy between **Theory-Simulation-Verification-Validation (TSVV) Tasks** and **Advanced Computing Hubs, E-TASC** develops validated **predictive capabilities** for key challenges in fusion research **for ITER and DEMO**.
- Creation of **EUROfusion standard software requires close collaboration** between TSVVs and ACHs.

Conclusions

- **Examples of ACH work** show different degrees of success to build trustable software capable of running efficiently in HPC systems.
- In some cases, **trade-offs had to be sought**, and the solution that was reachable does not fully meet the original task objectives.
- Overall, **E-TASC's mission** is highly **ambitious**.
- To large extent, E-TASC progress **relies on -** and needs to adapt to - the **evolution of computing resources** for large scale simulation and integrated modeling.
- Securing the **continuation** of E-TASC beyond 2025 would be highly **important to guarantee the competitiveness of fusion R&D in Europe**.

Invitation to the 5th Fusion HPC Workshop

Fusion HPC Workshop

November 21-22, 2024

Where: Online

More details soon on hpcfusion.bsc.es

Thank you for your attention

mervi.mantsinen@bsc.es

@Fusion_BSC fusion.bsc.es fusion@bsc.es Fusion Group BSC