Exa-DI Co-design Motif Workshop

Efficient discretisation for PDEs @ exascale

J-P. Vilotte (CNRS), PI Exa-DI project

NumPEx: Capable Exascale

NumPEx: 5 years national project (CNRS, CEA, INRIA, Universities) - 40,8 M€

Coordinators: J.Y. Berthoud (INRIA), J. Bobin (CEA), M. Krajecki (CNRS)

- Aggregate the French HPC/HPDA/AI community, foster new collaborations and synergies
- Co-develop, integrate, validate and deliver an expanded exascale software stack to accelerate Exascale applications productivity and sustainability
- Contribute and accelerate the emergence of a European sovereign exascale software stack and performant strategic exascale applications
- Build a multidisciplinary national workforce and develop training to improve CSE application development and software integration methodologies

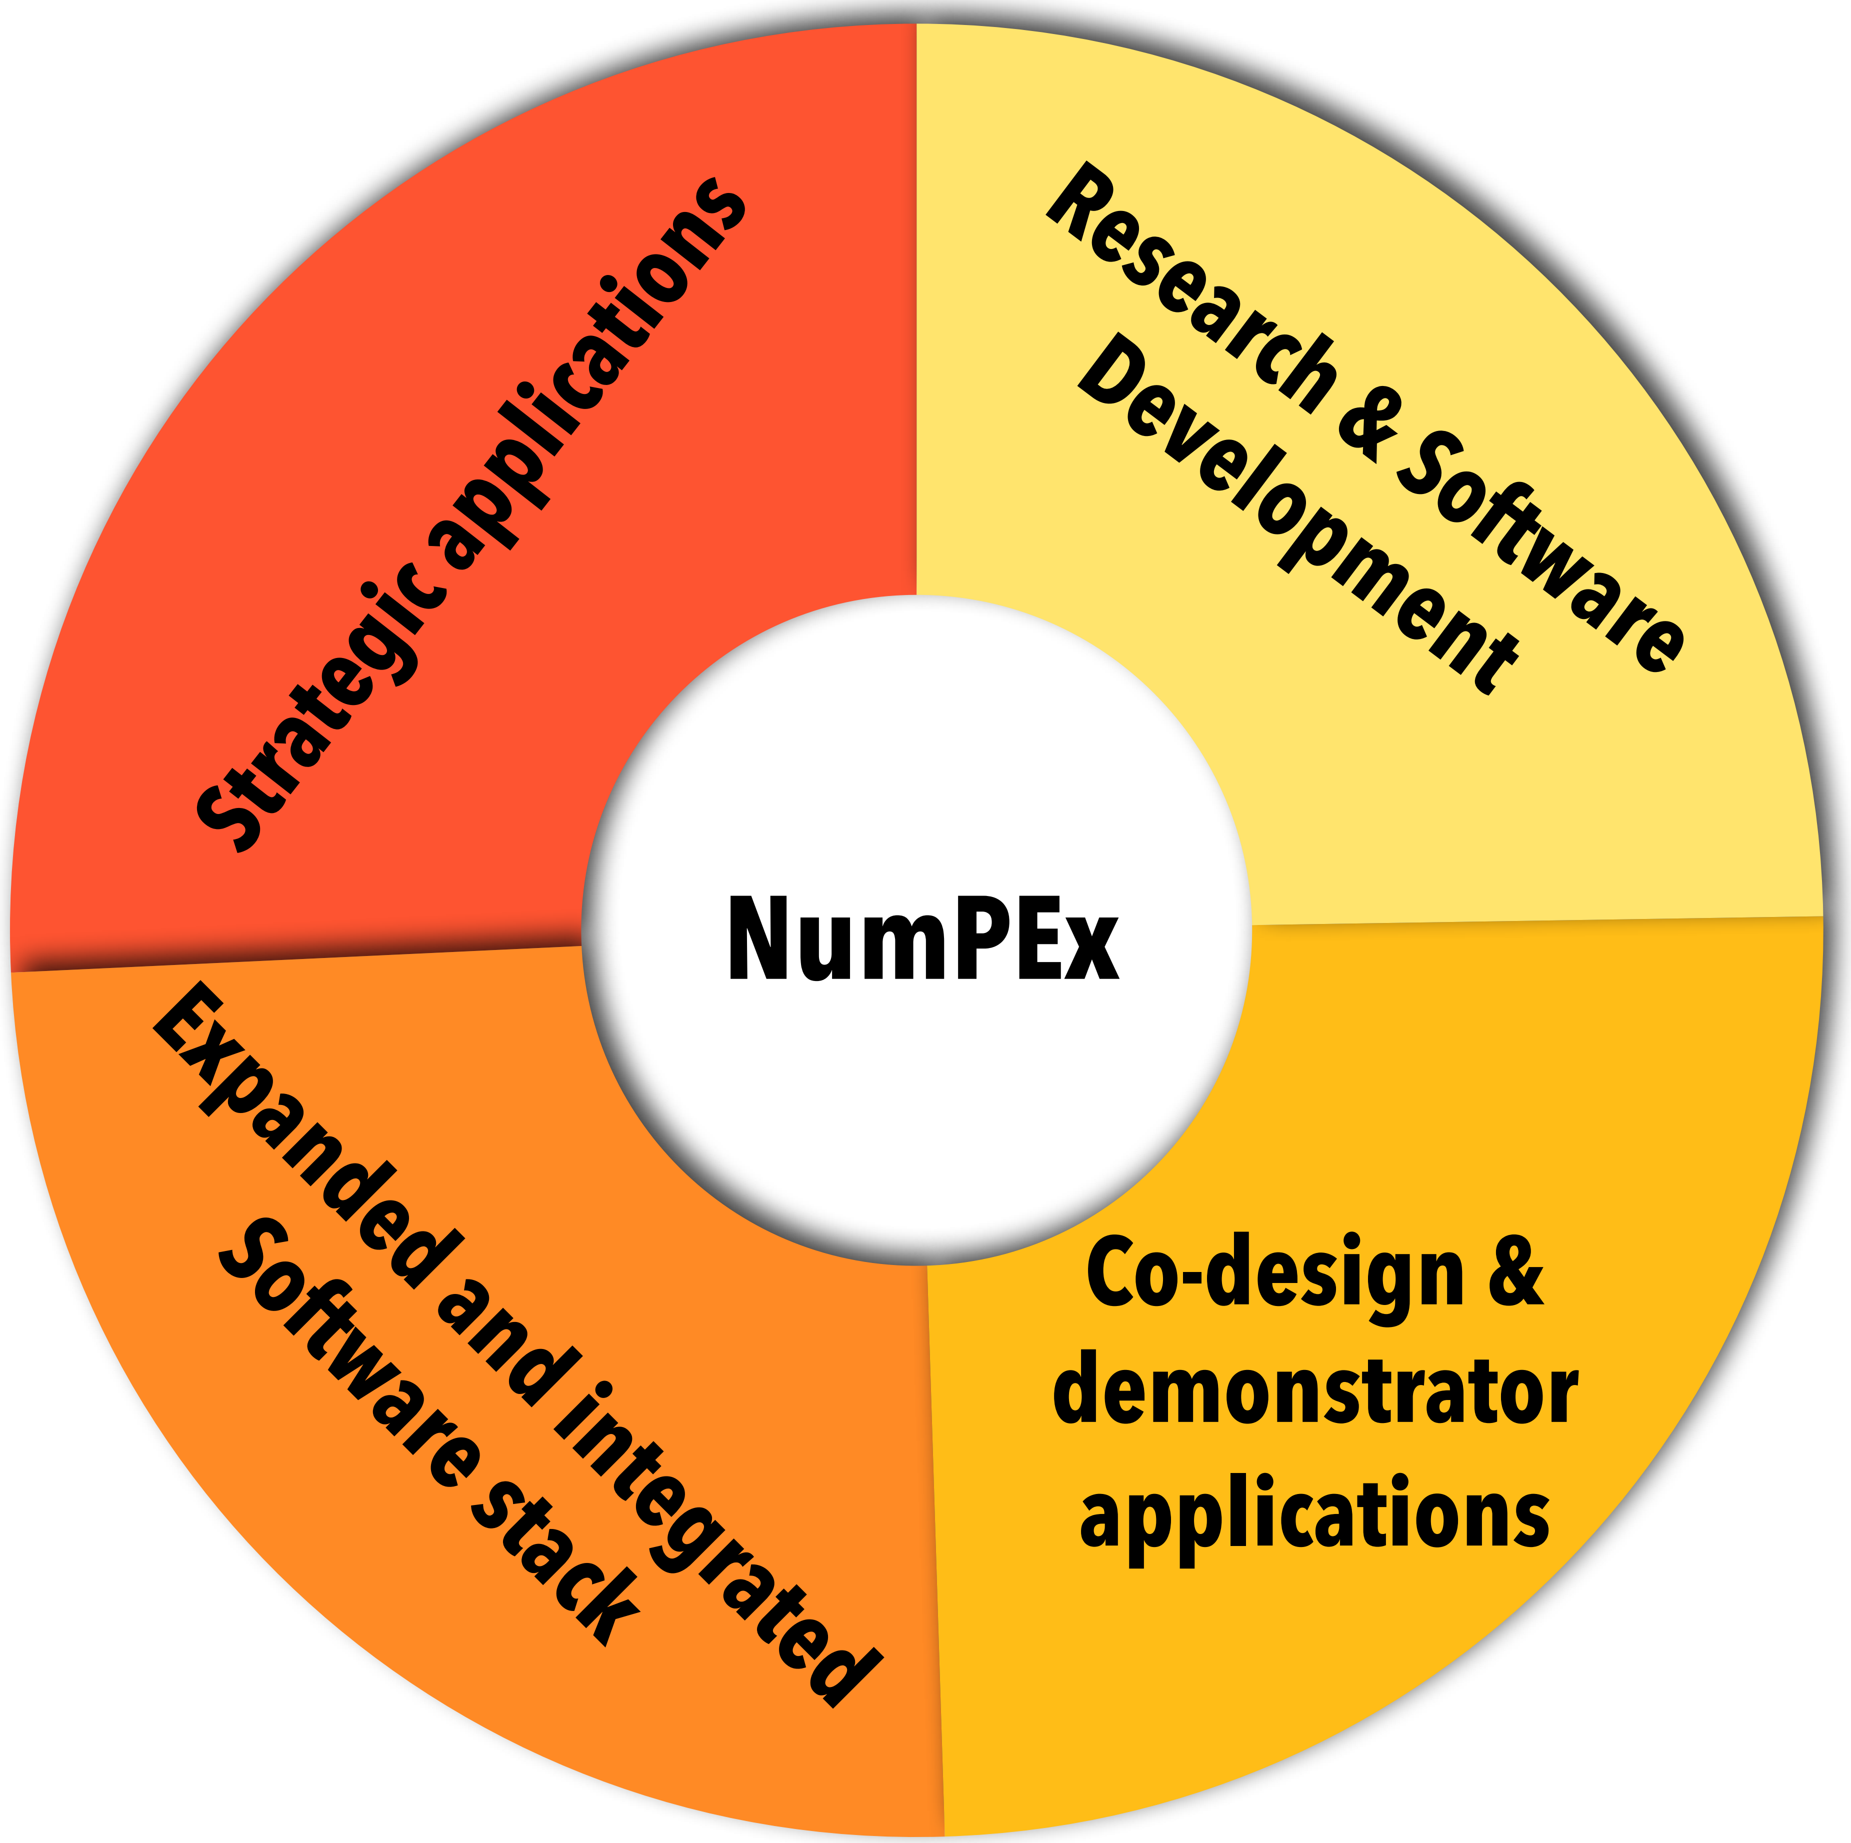

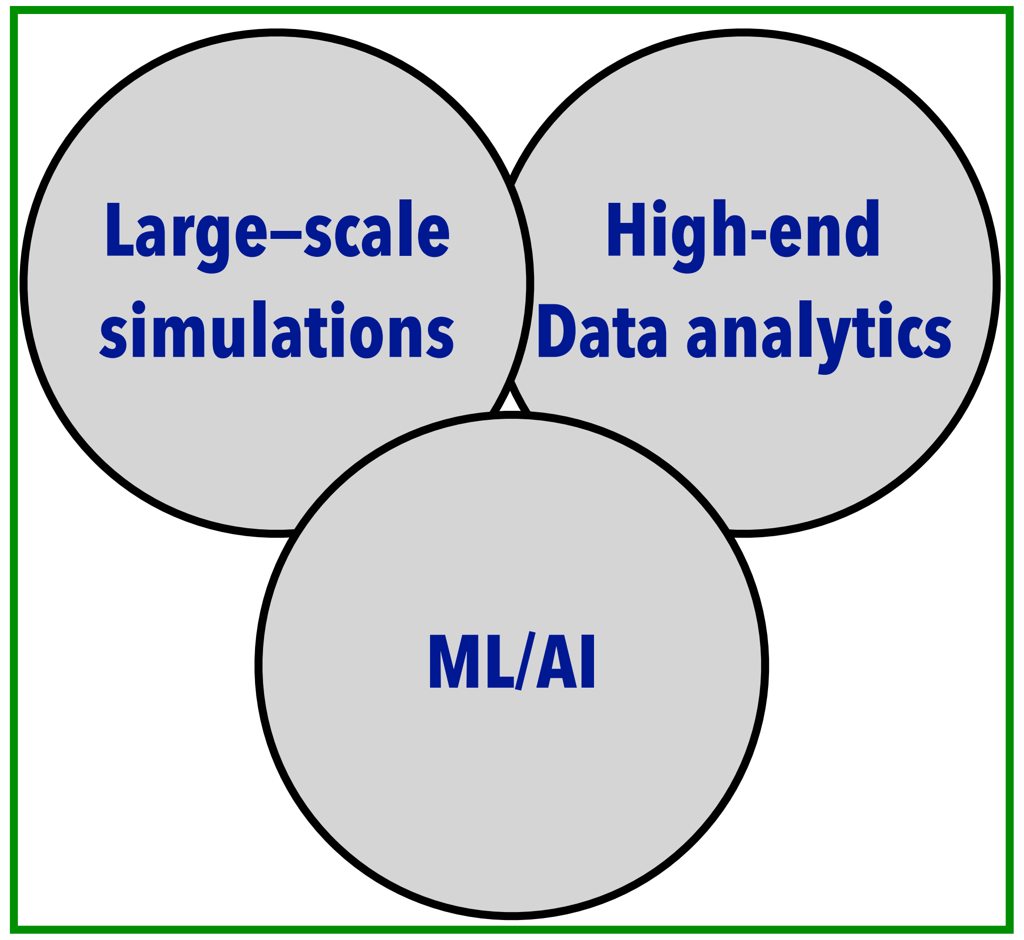

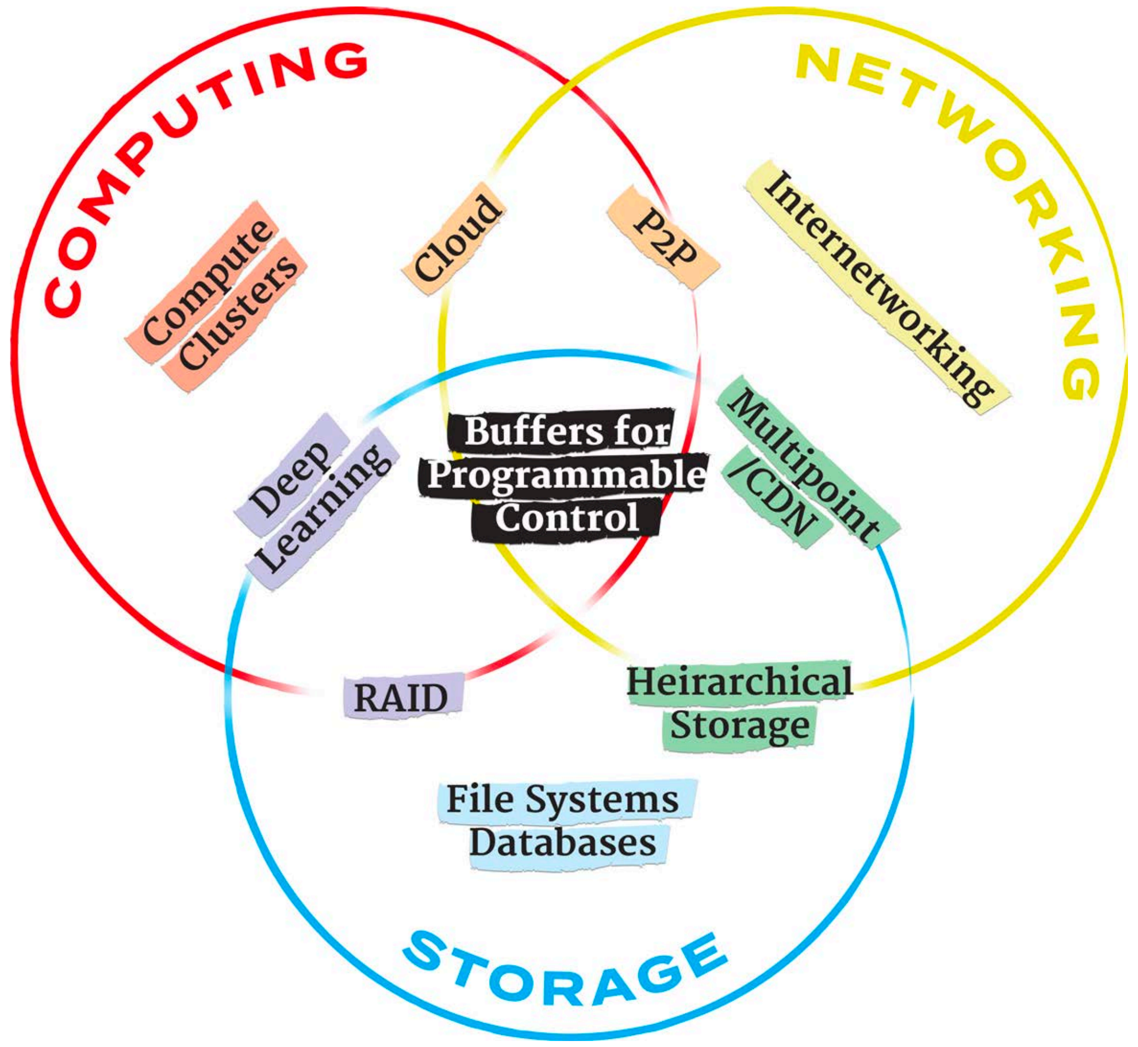

NumPEx
Exascale computing

NumPEx work plan

CSE Applications

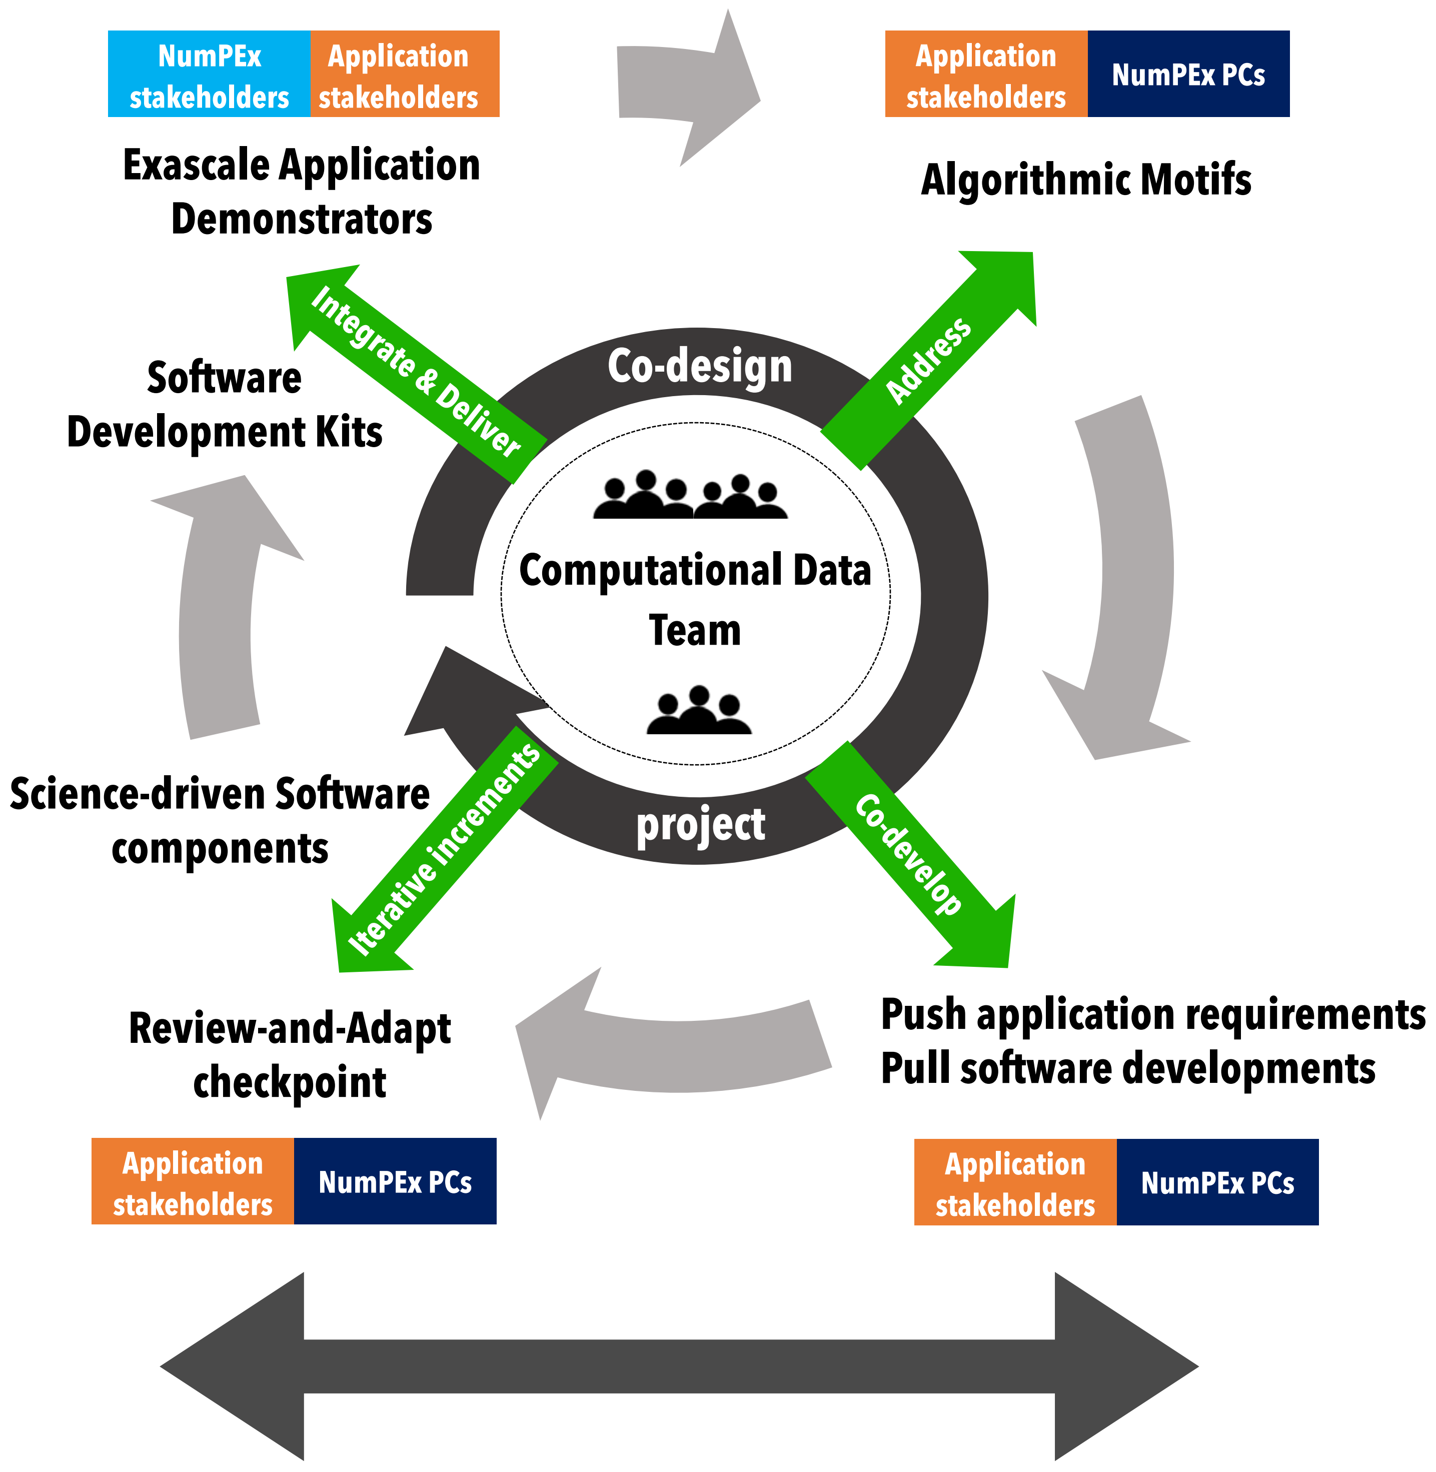

ExaDIP
software co-design & co-development, integration and delivery
J.-P. Vilotte (CNRS)
V. Brenner (CEA)

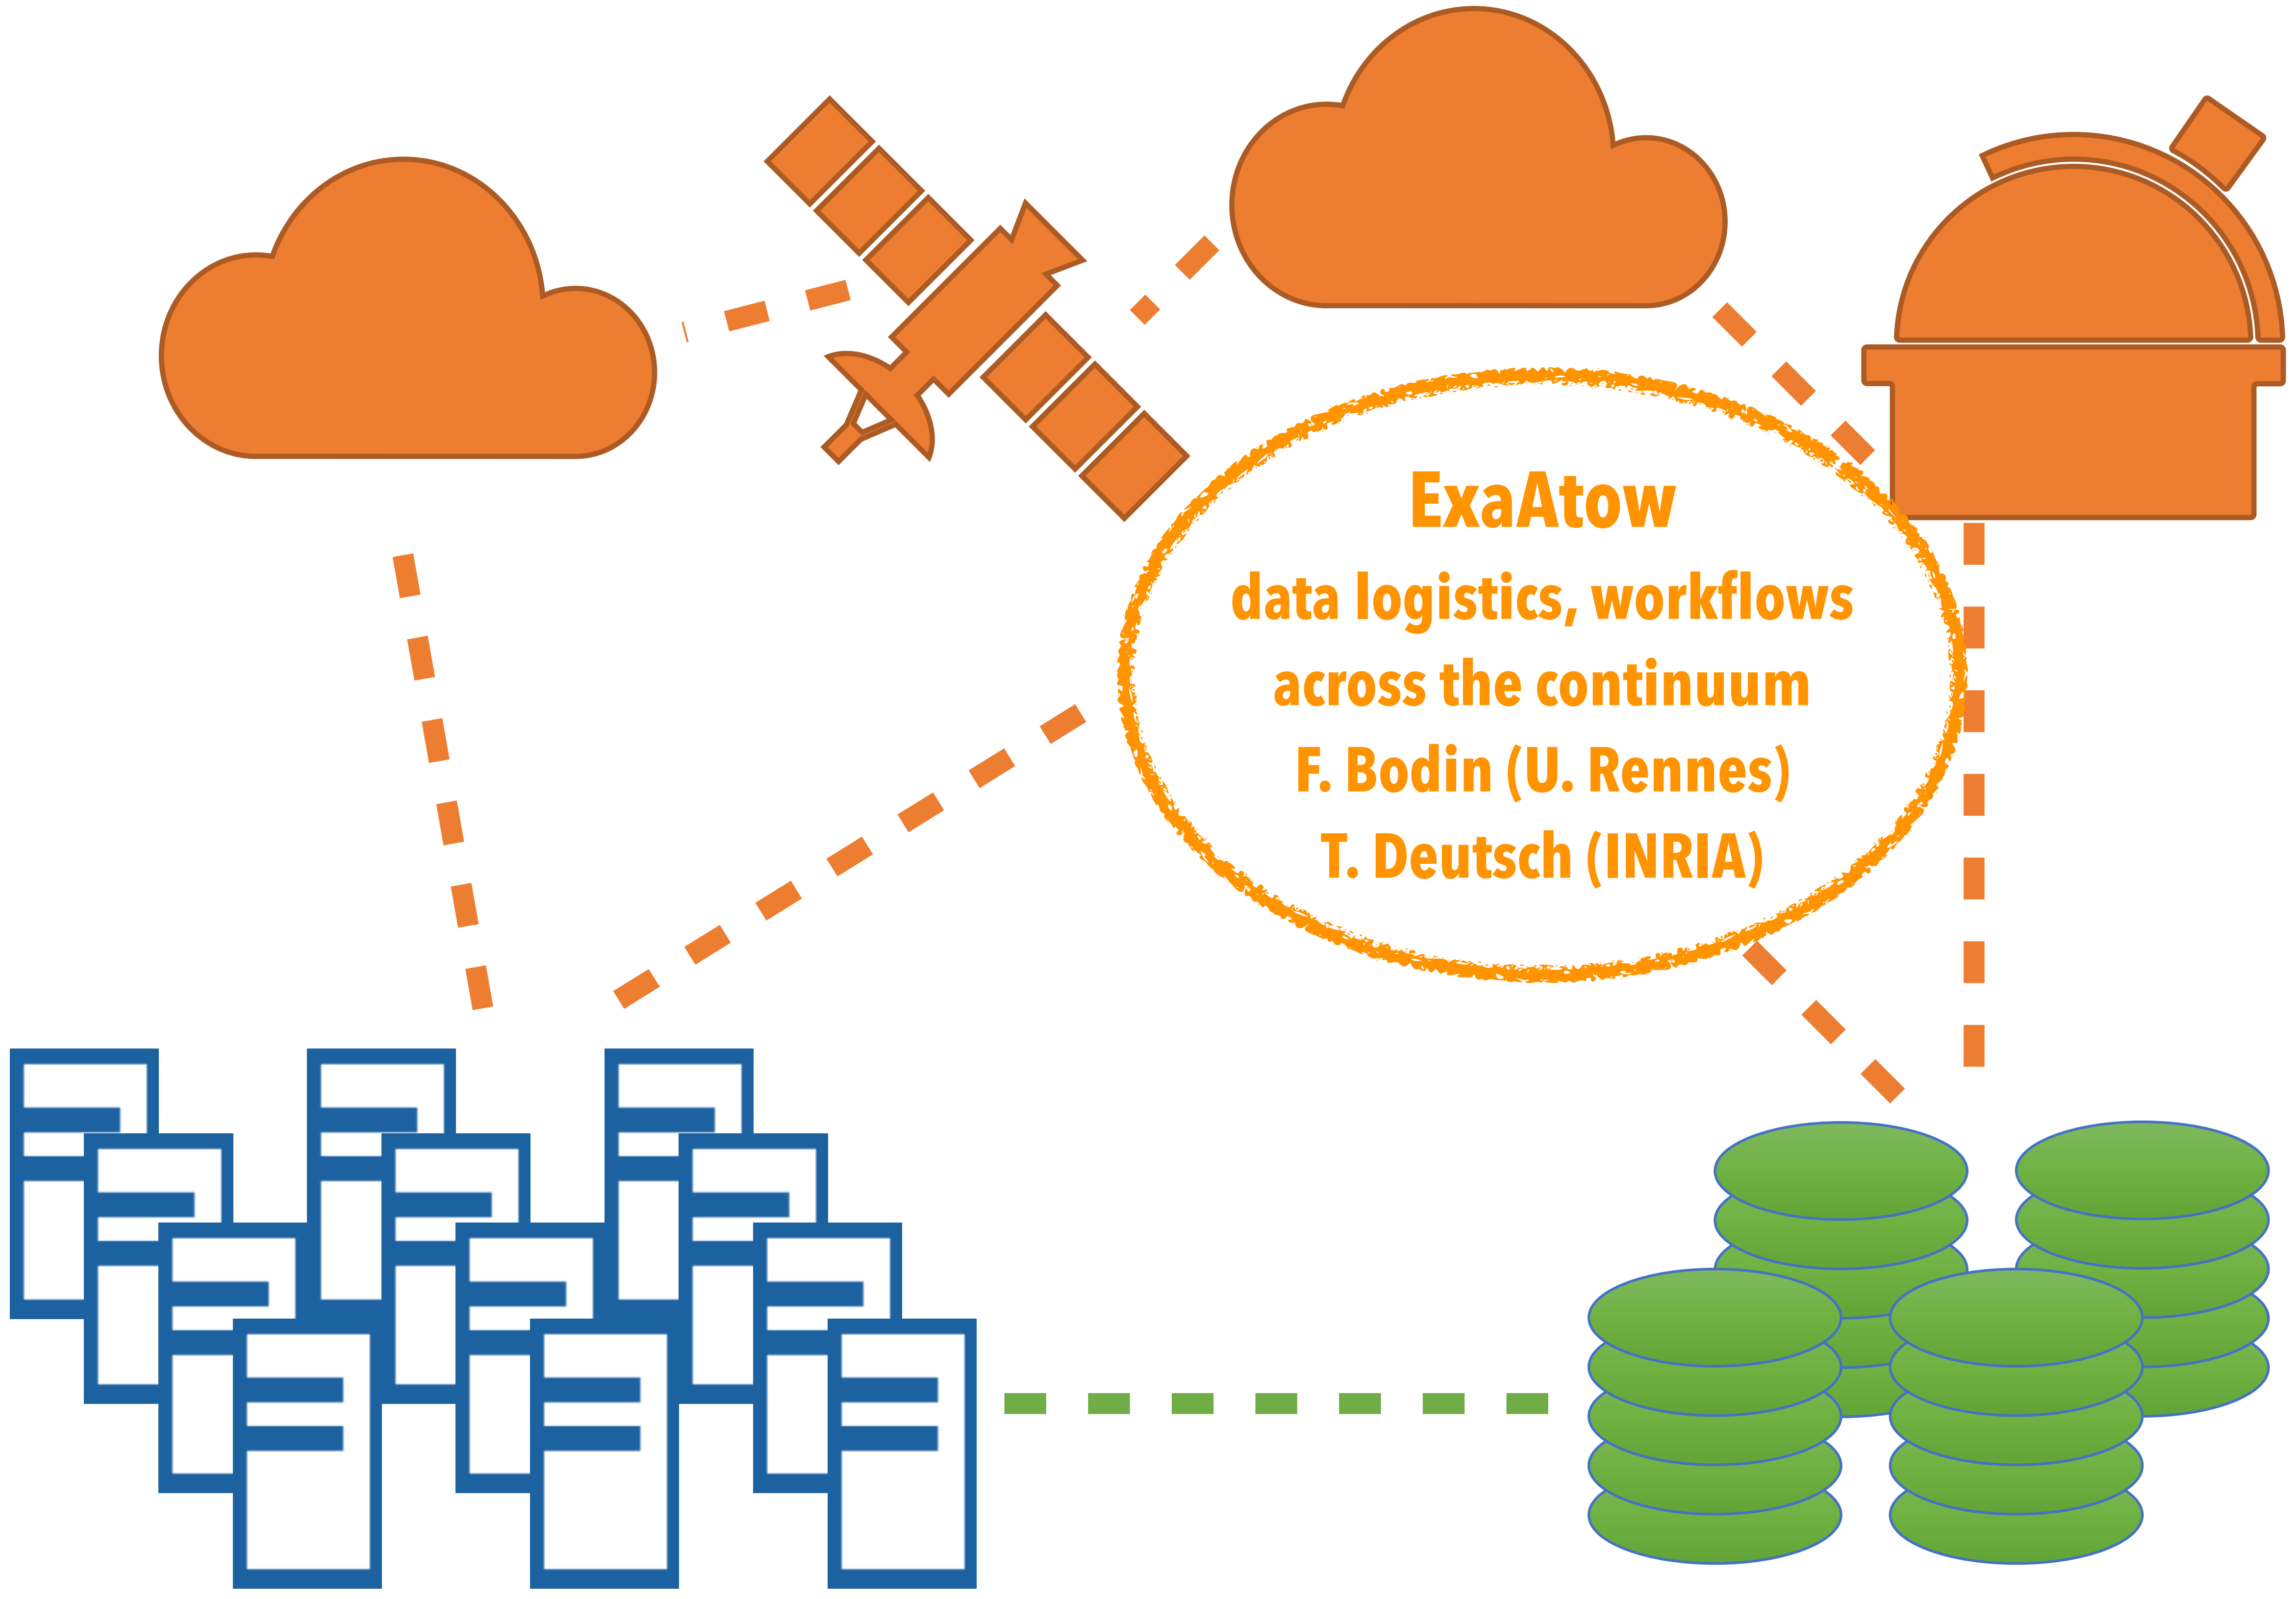

ExaMA
algorithms, math libraries
C. Prud'homme (UNISTRA)
H. Barucq (INRIA)

ExaSoft
parallel programming and execution environments
R. Namyst (INRIA)
A; Buttari (CNRS)

ExaDost
in-situ data reduction and analytics, storage, IO
G. Antoniu (INRIA)
J. Bigot (CEA)

NumPEx
Exascale computing

Exa-DI project

Challenges:

- *Exascale CS&E applications.*
- *Expanded Exascale science-driven software stack.*
- *Exascale software components integration, deployment and reuse.*
- *Exascale application performance portability and sustainability.*

Objectives:

- *Co-analyse CSE application demonstrators (ADs) across NumPEx*
- *Co-identify cross-cutting computational and communication motifs*
- *Co-develop libraries, frameworks, proxy and mini apps*
- *Co-develop Community Software Policies and sustainable software foundations*
- *Integrate and deliver interoperable SDKs as needed by the Exascale ADs*
- *Exercise, expand and harden sustainable core components for performance portable programming models (e.g. Kokkos, Raja)*
- *Establish a national computational data team to enable agile co-design projects*
- *Broaden awareness/use of robust software components in exascale application development methodologies*

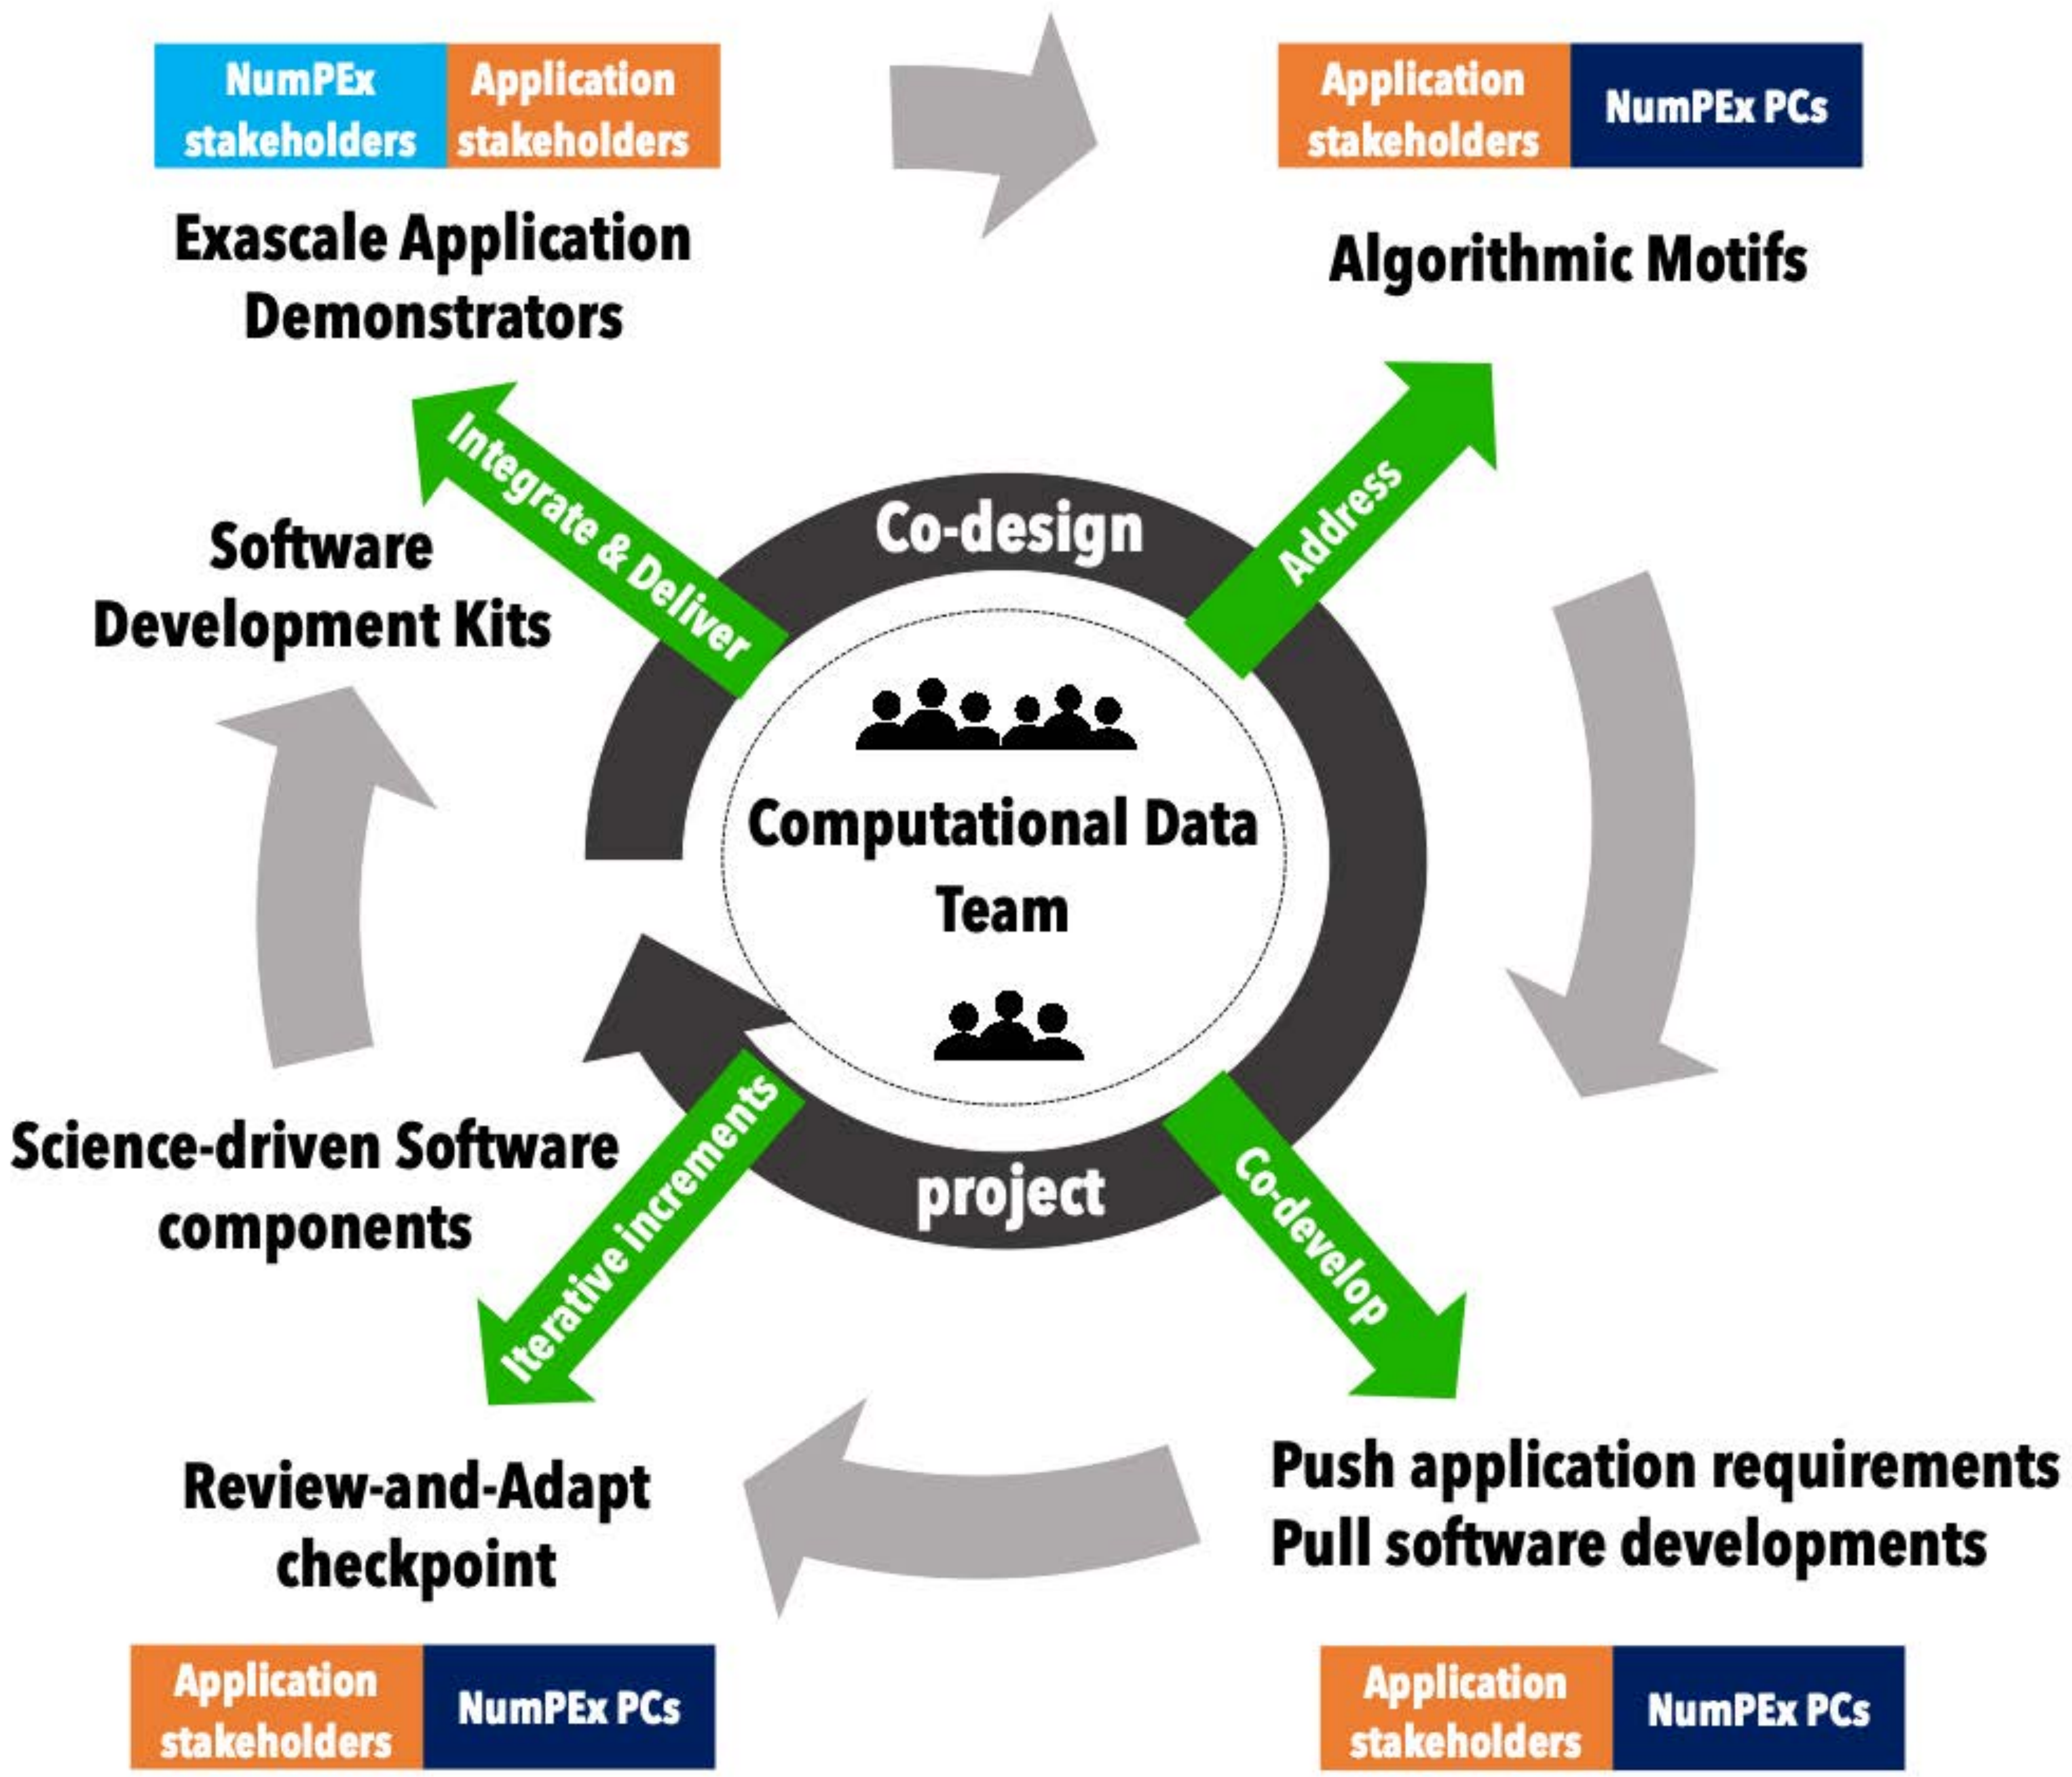

Team: CNRS, CEA, INRIA

Budget: 9,3 M€ (including CoEs co-funding)

Exa-DI International Context

Co-designed software portfolios, proxy-apps & benchmarking

Context:

- Exascale Computing Project (ECP): DOE, NSF
- Fugaku & Fugaku NEXT co-design projects
- Euro-HPC JU initiatives, ETP4HPC

Problematics:

- CSE applications development methodologies, accuracy & performance portability
- Co-designed application-driven software technology and logical software components portfolios
- Proxy-apps / many-apps suites
- Hardware & Integration, Testing & Profiling tools, Benchmarking specifications

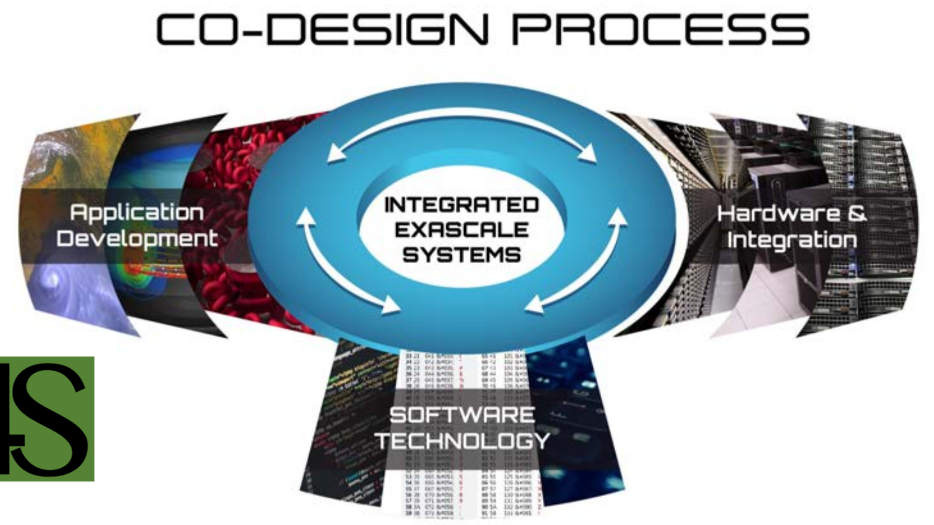

Fugaku Codesign Report

- FLAGSHIP 2020 Project Technical Report -

Fugaku

FLAGSHIP 2020 Project
RIKEN Center for Computational Science (R-CCS), RIKEN

March, 2022

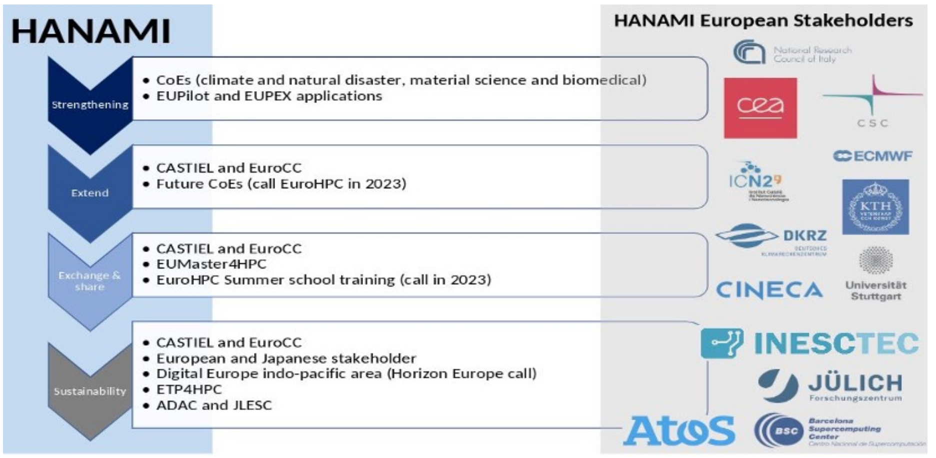

CSE application demonstrators

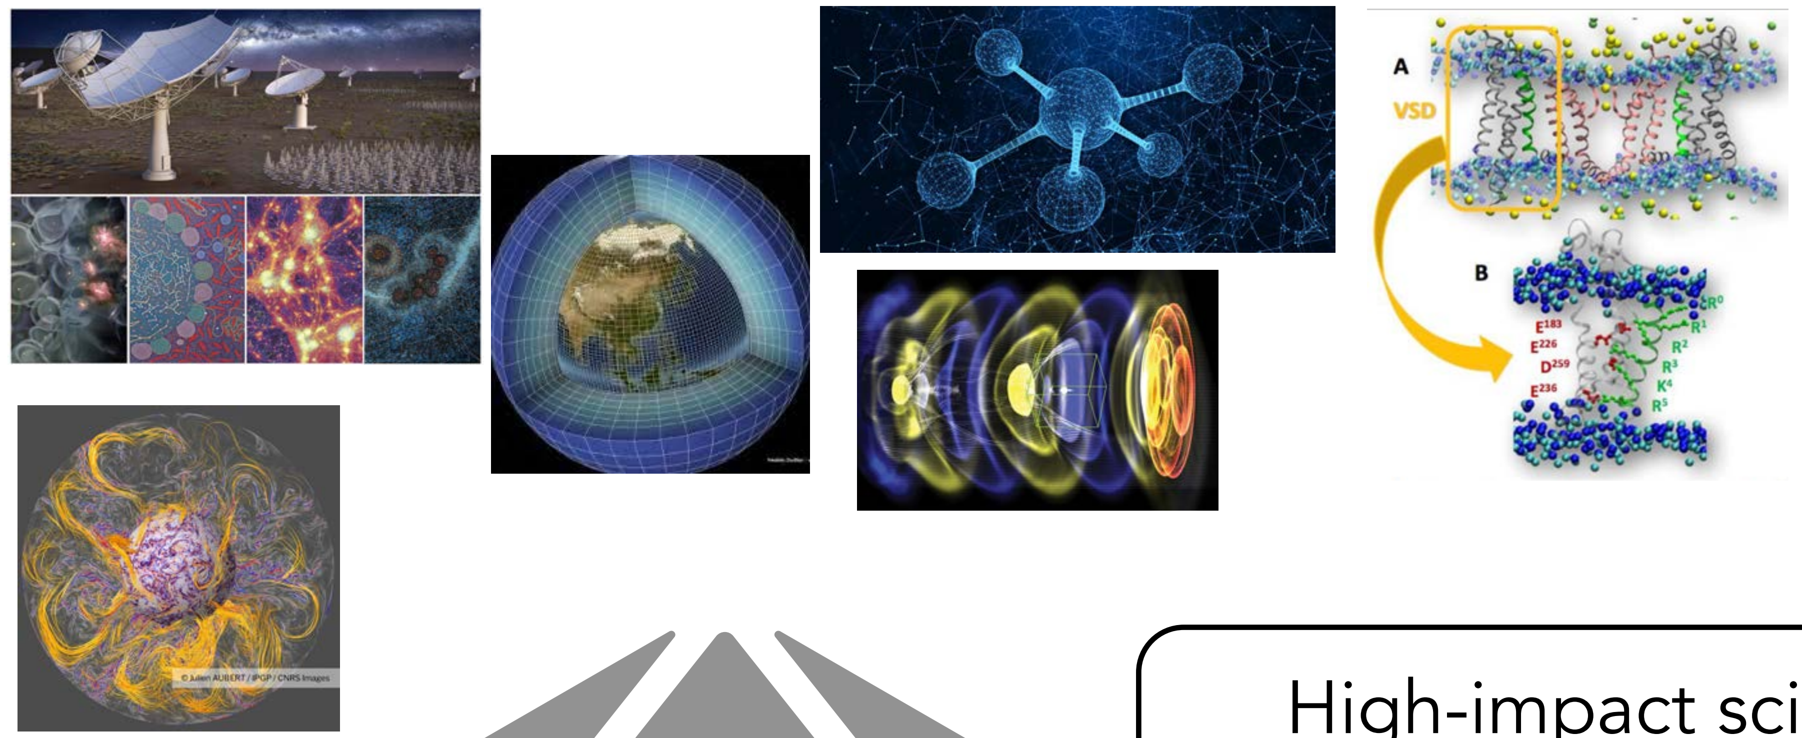

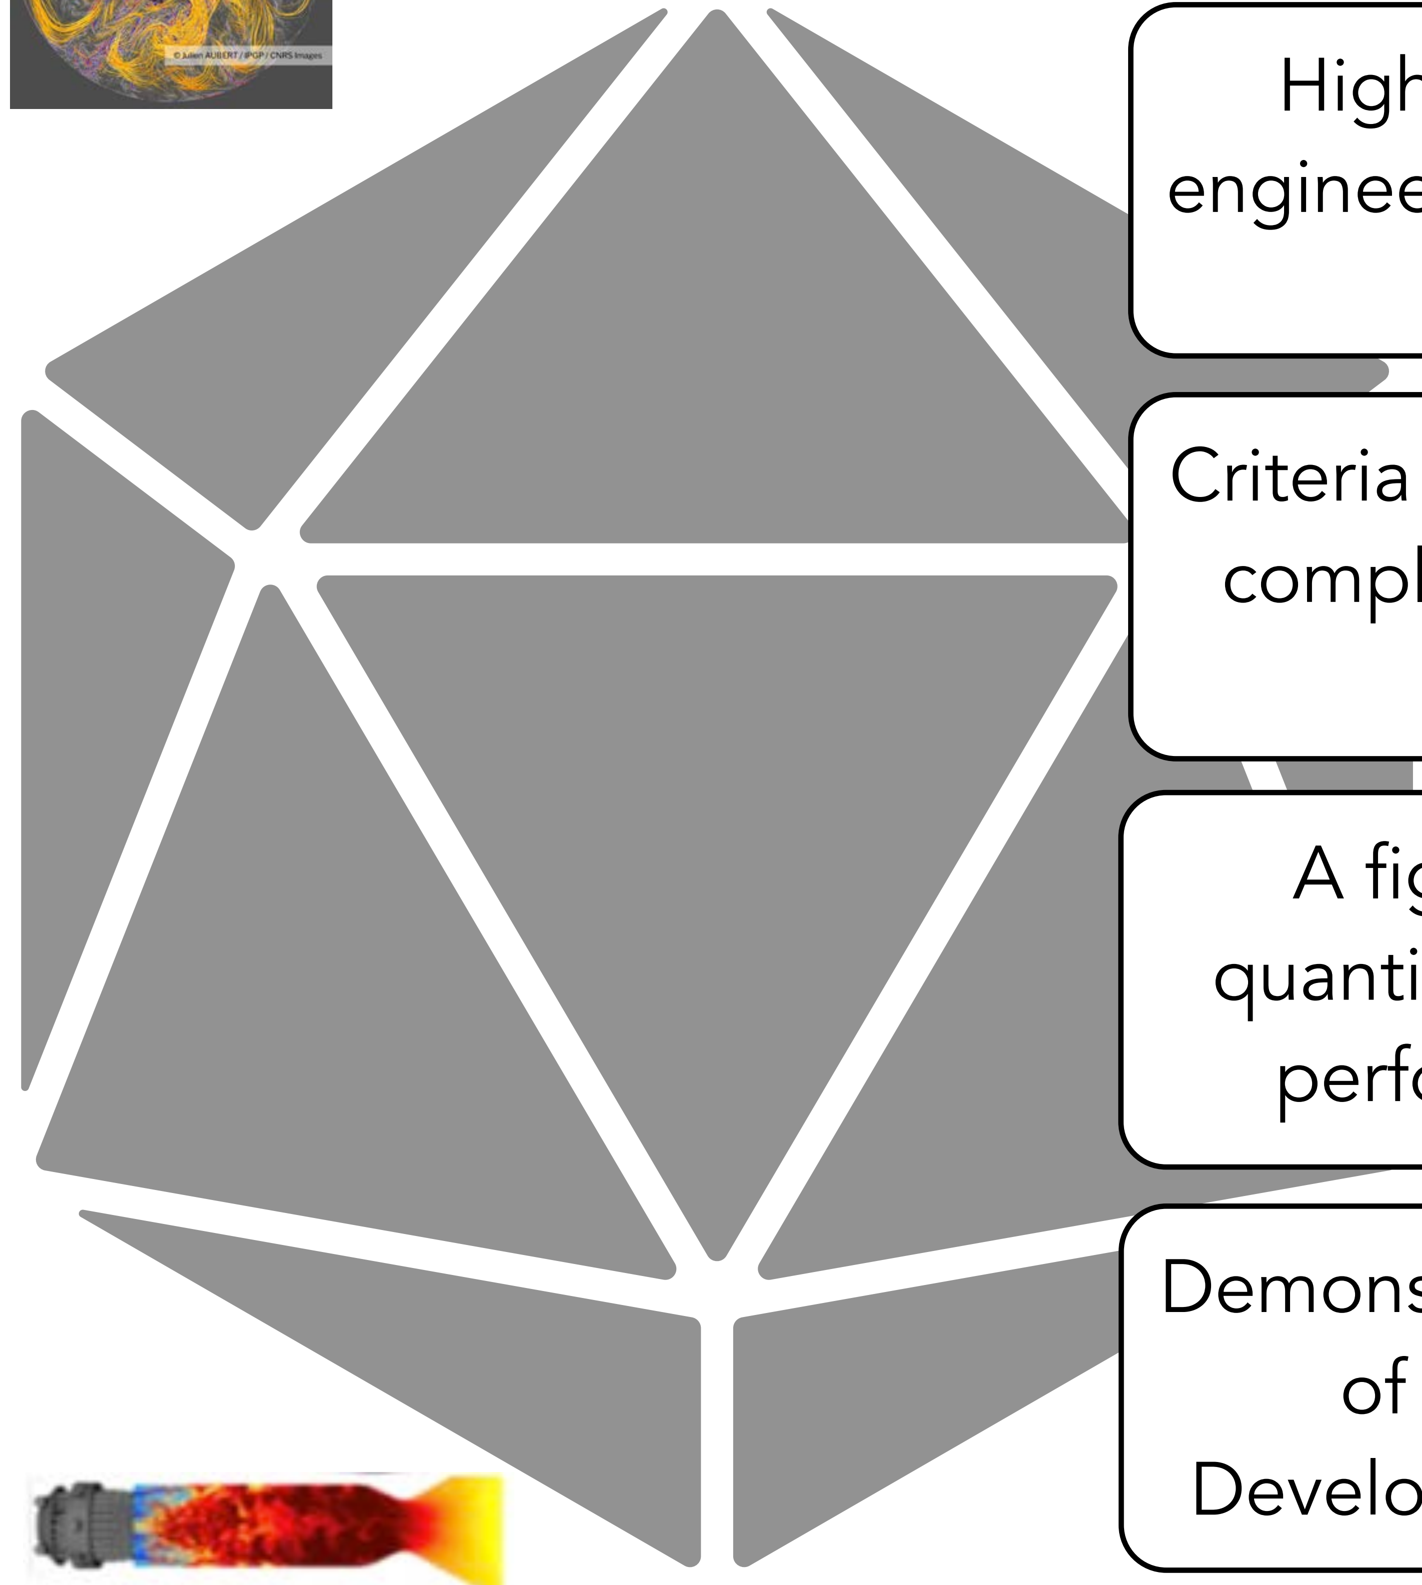

- High-impact science and engineering exascale challenge problem
- Criteria for assessing successful completion of the challenge problem
- A figure of merit (FOM) quantifying enhancement of performance and rate of science
- Demonstration and assessment of effective Software Development Kits integration

Astronomy & Astrophysics
Earth System Models & environment
Environmental extreme events
Computational biology & Life science
Laboratory laser-plasma physics
High-energy particle physics
Quantum chemistry and materials
Digital health
Environmental & societal risks
Urban systems planing
Magnetically confined fusion plasma (ITER)
Sustainable Transport & mobility
Energy production & transport
…

Exascale Challenges

- Heterogeneous exascale and post-exascale architectures
- New multi physics and multi-scale capability
- On-line streaming data analysis /reduction
- AI-enabled big data analytics and dimension reduction
- New mathematical approaches and model improvements
- Leveraging robust and accurate logical collection of interoperable software components (libraries, tools and frameworks)
- Improving performance portability by exercising new performance portable programming models
- New scalable task-based programming and execution models
- Fondations for a sustainable exascale scientific software stack

Co-designed computational and communication motifs

Objectives

- proceed and integrating logical suites of software components and tools with emerging hardware technologies
- Deliver software development kits easily deployable and instantiated into CSE application environments for testing, development acceleration and requirements feedback.
- Facilitate collaboration between software packages, enable integration and/or interoperability with overall NumPEx technologies stack, streamline developper and user workflows, and coordinate releases

Flexibility/Interoperability

- support a wide range of CSE application algorithmic requirements
- enable exploration of new algorithms and frameworks as need by the CSE applications.
- avoid unnecessary restrictions on programming models and how application developers construct their algorithms
- allow application developers to interact with software components at different levels of abstraction
- exercise and harden abstraction layers (e.g. Kokkos, Raja) for hardware-aware technologies

Proxy and mini app suites

- represent algorithms, data structures and layouts, optimisation and other computational characteristics
- allow to measure performance gain, scalability, performance portability, and software quality
- enable continuous integration, testing and benchmarking experiments with standardised performance tools to guide optimisations, understand bottlenecks and trade-off between software components

Cross-cutting computation and communication motifs

Efficient discretisation for PDEs @ Exascale

- *PDE-based multi-physics multi-scale simulations (FV, FEM, SEM, HDG), unstructured hexahedral and tetrahedral meshes, AMR*

Block-structured AMR @ Exascale

- *Multi-physics CFD, transport and particle-based simulations*

Particle-based methods @ Exascale

- *MD, QMD, AIMD, gravitational N-body, SPH, PIC simulations*

On-line analysis @ Exascale

- *Online analysis, reduction, aggregation, AI-based multi physics and multi-scale coupling*

Big data analytics @ Exascale

- *AI-based data analytics, scattering transform, detection, classification, reinforcement learning, manifold learning, DML training*

Artificial Intelligence @ Exascale

- *Physics-based machine learning, PINNs, surrogate models, DML learning, AI-based inverse, inference, control and design problems*

Combinatorial methods @ Exascale

- *Static and dynamic combinatorial (graph) kernels at scale*

Cross-cutting co-design algorithmic motifs

- Accelerate the development and improve performance portability of exascale ADs
- Address common patterns of communication and computation
- Improve scientific software development methodologies balancing trade-off between fine-grained control and modularity based on integrated/interoperable libraries

Agile co-development projects

- Co-analyse existing software components, e.g. from ECP co-design centres
- Co-identify gaps and missing functionalities as needed by the ADS
- Co-develop logical software components (libraries, programming and execution models, frameworks) with different interoperability levels and abstraction layers
- Integrate into re-usable SDKs that can be easily managed, deployed and instantiated as needed by ADs
- Co-develop meaningful proxy-apps and mini-apps to integrate, test and benchmark the developed software components (performance, trade-offs)

NumPEx
Exascale computing

Exa-DI flow of product delivery and deployment

NumPEx
Exascale computing

ExaDI Software Development Kits (SDKs)

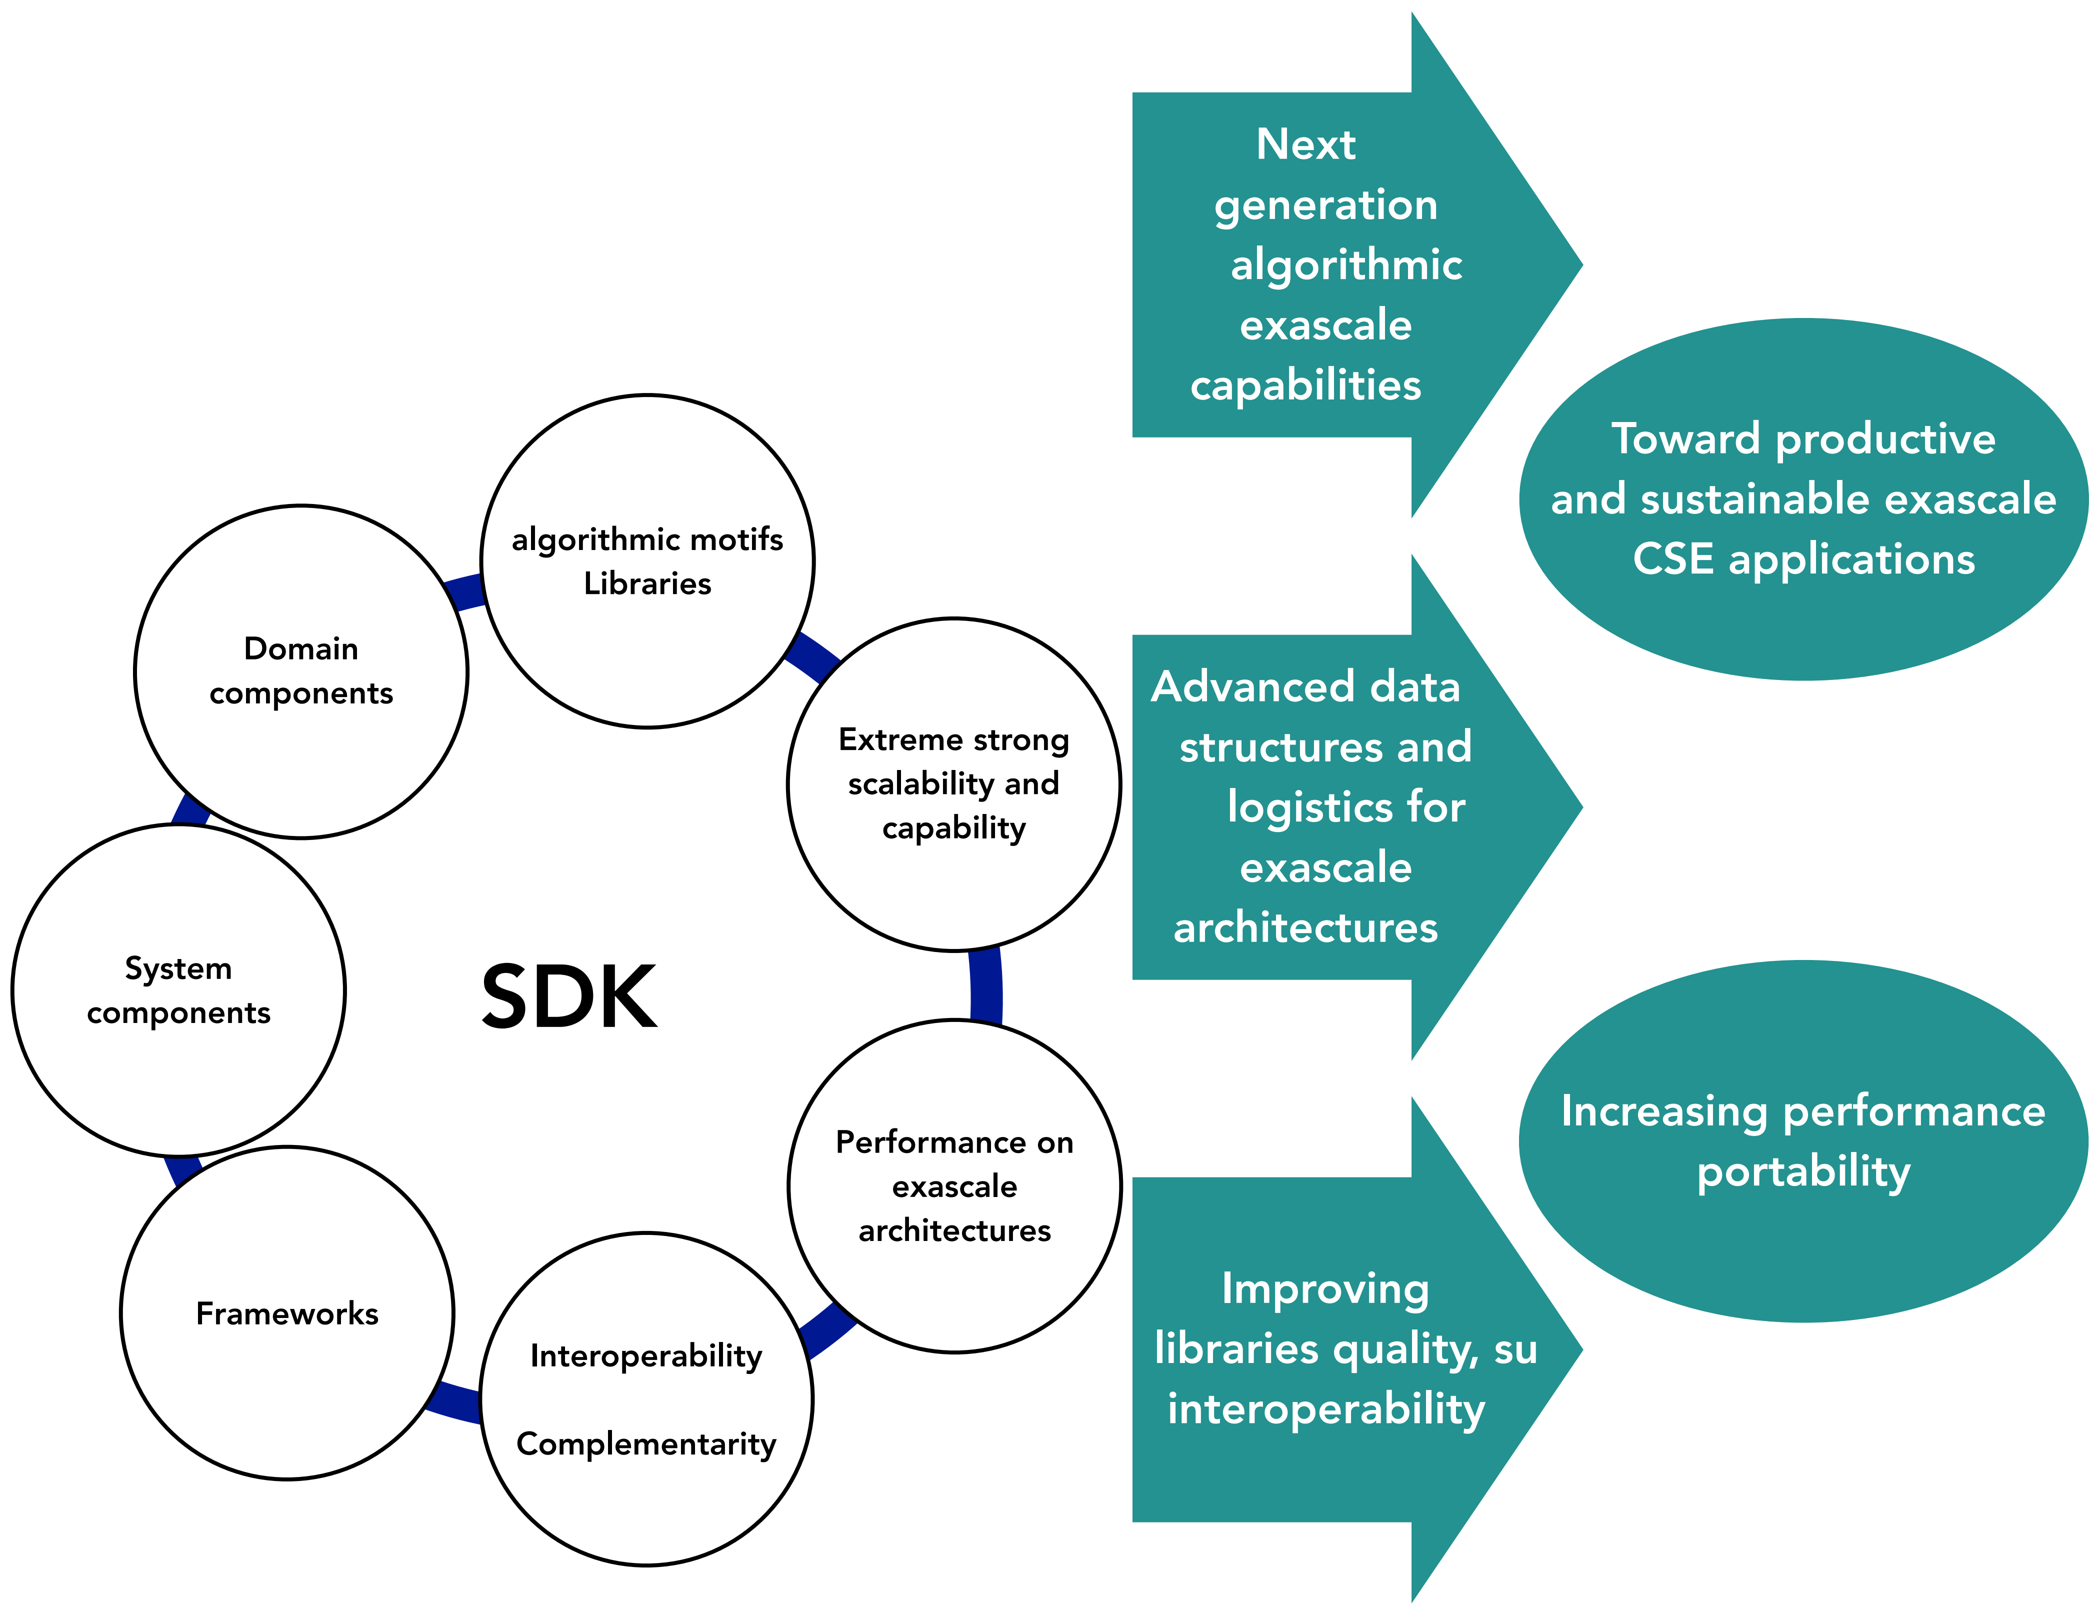

Timeline: SDK release 1 — SDK release 2 — .... — SDK release n

Software Development Kits

- Logical collections of value-added interoperable software components as needed by ADs,
- Integrated and delivered using meta-builder and container systems enabling a combined deployment on exascale systems and combination as needed by CSE applications

Community Software Policies

- Improve software quality, usability, access and sustainability;
- Provide foundation for deeper levels of interoperability
- Establish a certification process to label software (maturity, portability, compliance)

Software packaging and deployment technologies

- Promote common Meta-builder systems (e.g. SPACK, GUIX, NIX) and container technologies (e.g. Singularity)
- Extend/harden new capabilities enabling deployment on exascale systems and regression testing

Software Integration hub

- Enable access to externally managed software integration and testing platforms
- Synergetic collaborations with national computing facilities, vendors and other initiatives

Efficient Discretisation for PDEs @ Exascale

Ecosystem

- Unstructured PDE discretisation libraries
- High-order meshes libraries
- Unstructured adaptive mesh refinement
- Dynamic unstructured anisotropic meshing
- Efficient high-order operator format / Representation
- Batched dense tensor contractions
- Distributed linear algebra operations
- Scalable matrix-free solvers
- General interpolation
- Visualisation of high-order meshes and functions

NumPEx Application Demonstrators

- ***Climate/Ocean***
  - Croco, PEPR TRACCS
- ***Aeronautics/CFD, Combustion***
  - Sonics/Onera, Safran/INRIA, Yales2
- **Magnetically confined fusion plasma**
  - Gisela-X
- ***Earth & environment***
  - GEOXIM-IFPEN, GEOSX-Total
  - CoE ChEESE-2P: xSHELLS, SPECFEM3D
- ***Coupled Urban systems***
  - *CoE Hidalgo2, OpenFOAM*
- ***Industrial Risk & Safety evaluation***
  - *Manta*

ECP CEED SDK

LibParanumal

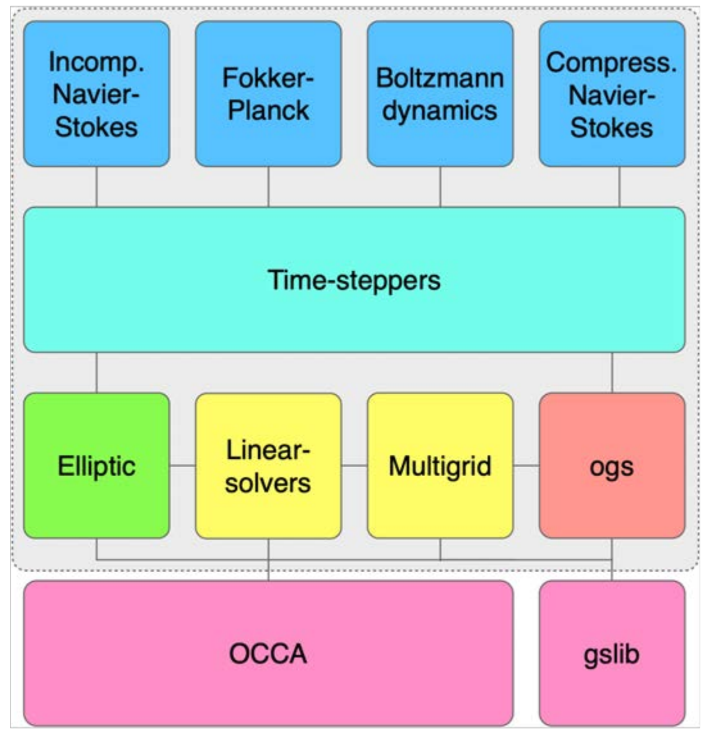

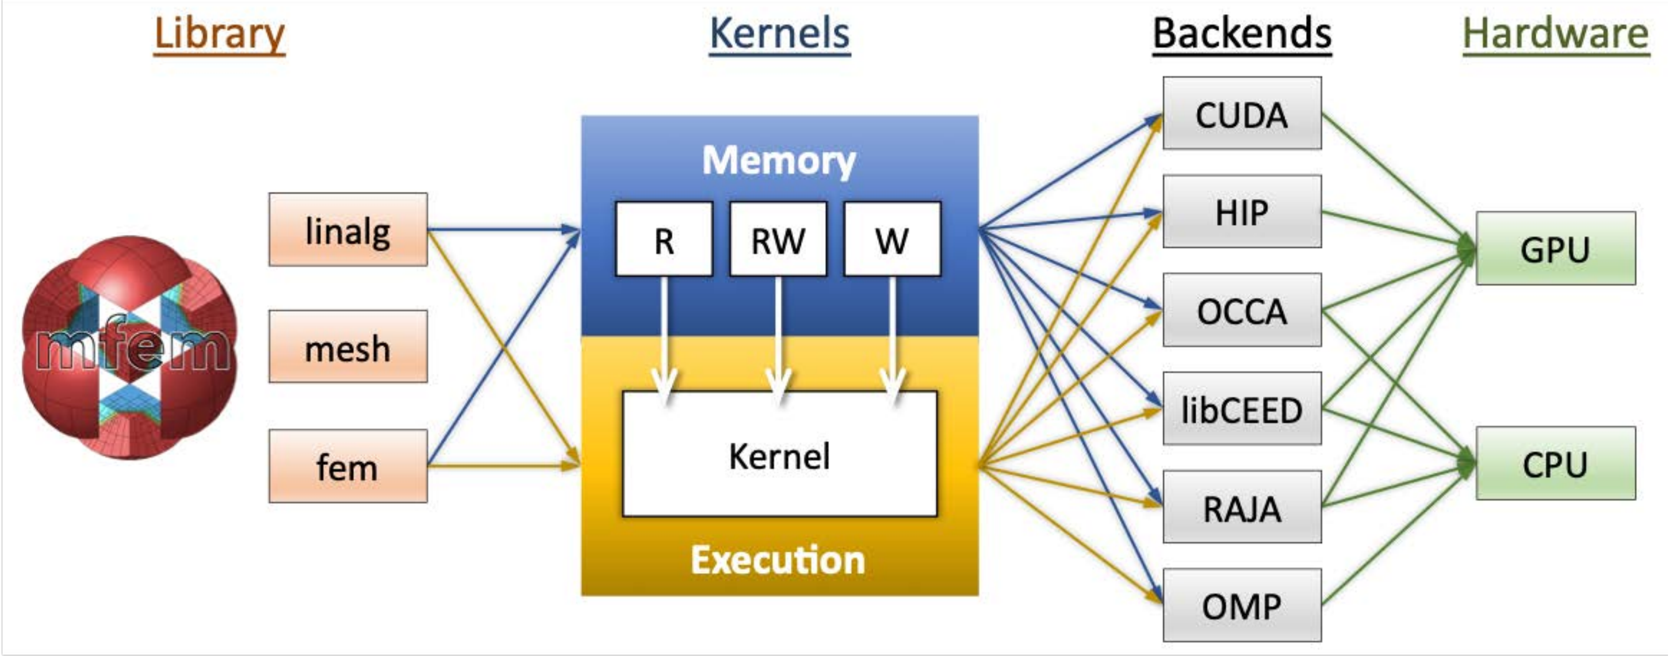

MFEM high-level API

OCCA API Performance portability

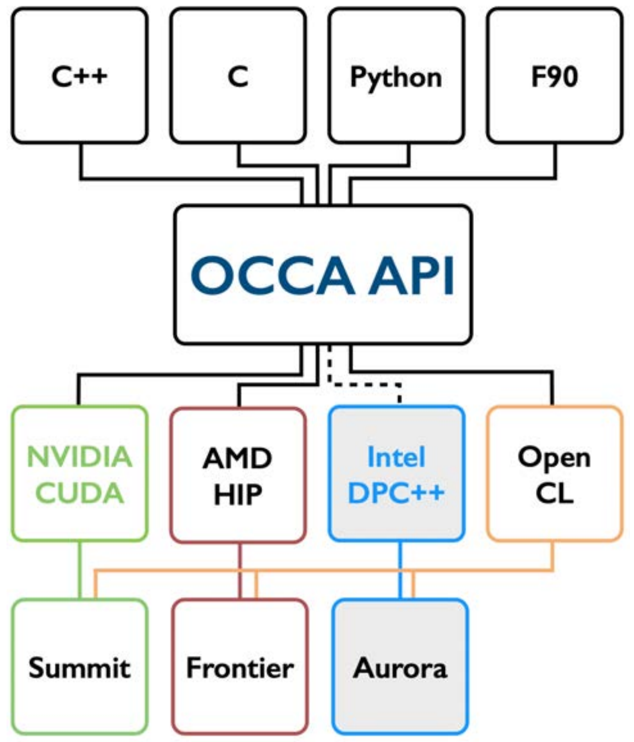

Lib-CEED low-level API

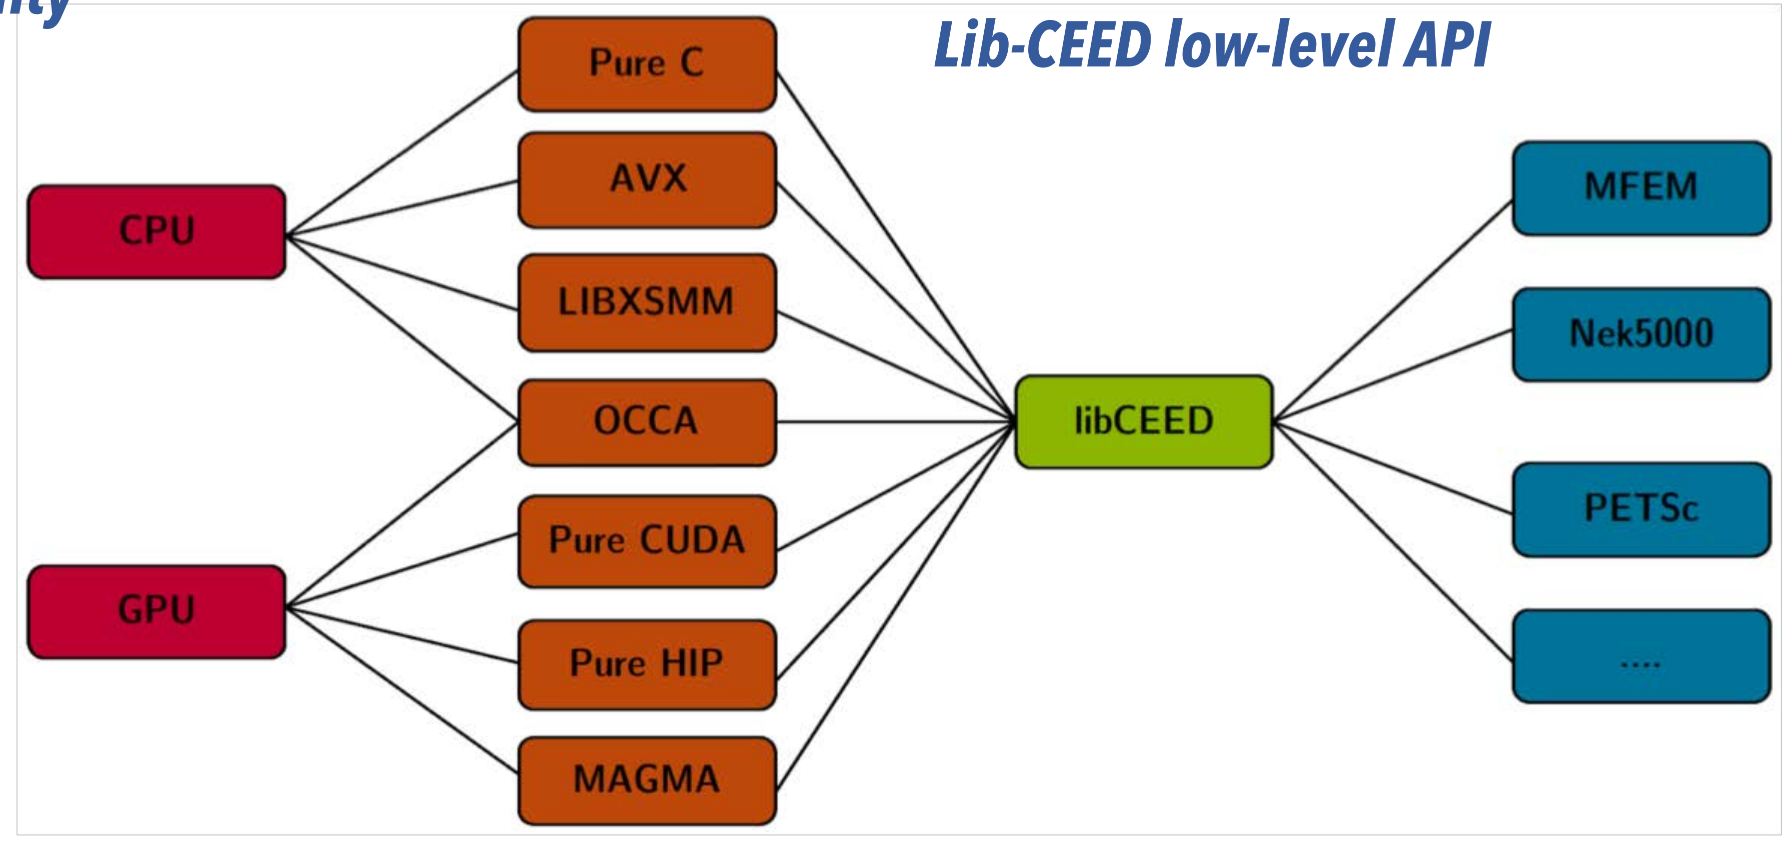

Software components, standards

- **ECP CEED**: MFEM, LibCEED, LibParanumal, PUMI, Omega_h, MAGMA, PETSc, SUNDIALS, Hypre, SLEPc, SuiteSparse, Gingko, METIS, GLVis, VisIt, ParaView, Umpire, OCCA, RAJA, HDF5/ADIOS, PyPI
- **NumPEx:** Feel++, Freefem++, GEOSX, MANTA, Mathphys++, Gyselalibx, PetSc, Hypre, MUMPS, PaStiX , Chameleon,HPDDM, Alien, Scotch,PARADIGM/CWPI, METIS, GAMMA's mesh, Mmg/ParMmg, DDC, CGGAL, MED/Coupling, AGIOS, HDF5, StarPU, Dask, Deisa, Kokkos-Kernels, PyPi

NumPEx
Exascale computing

Workshop objectives and return

Develop and share a common understanding of:

- Exa-DI co-design and co-development agile process driving collaboration between software packages, integration and/or interoperability with overall NumPEx technologies, streamline developper and user workflows, maintaining testing and benchmarking, and coordinate SDK releases
- role and responsibility in this process between the R&D teams in the different NumPEx PCs, the ADs' development teams, and the Exa-DI Computational and Data Team
- CDT application liaisons and one-to-many interactions based on the development of efficient discretisation for PDEs @ exascale

Co-identify across several ADs:

- existing exascale software issues and barriers,
- software components (libraries, frameworks, abstraction layers, programming and execution environments) to be developed and integrated for accelerating exascale ADs development and improving their performance portability

Co-analyse the software stack developed in the ECP CEED project:

- what software components can be leveraged,
- what gaps and missing functionalities need to be addressed
- what new complementary and/or alternative components need to be co-developed by agile teams (objectives, deliverables)
- what mini-apps and proxy-apps need to be co-developed with the ADs and shared with specifications for testing and benchmarking methodologies

Co-define agile co-development teams

➡ identified contributions of different R&D teams in-and-across the NumPEx PCs, and of the ADs,

➡ necessary resources and expertises to be gathered in the CDT to drive the agile co-development process

Organise follow-up with series of identified deep-dive sessions