Exa-DI Co-design Motif Workshop

Block-structured AMR @ exascale

J-P. Vilotte (CNRS), V. Brenner(CEA)
PI and co-PI Exa-DI project

February 6-7, 2024 - Grand Amphi, Institut de Physique du Globe de Paris, CNRS-INSU

NumPEx: Capable Exascale

NumPEx
Exascale computing

NumPEx: 6 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 to and accelerate the emergence of a European sovereign exascale software stack and productive strategic exascale applications

Establish a multidisciplinary national workforce and develop training to improve CSE application development and software integration methodologies

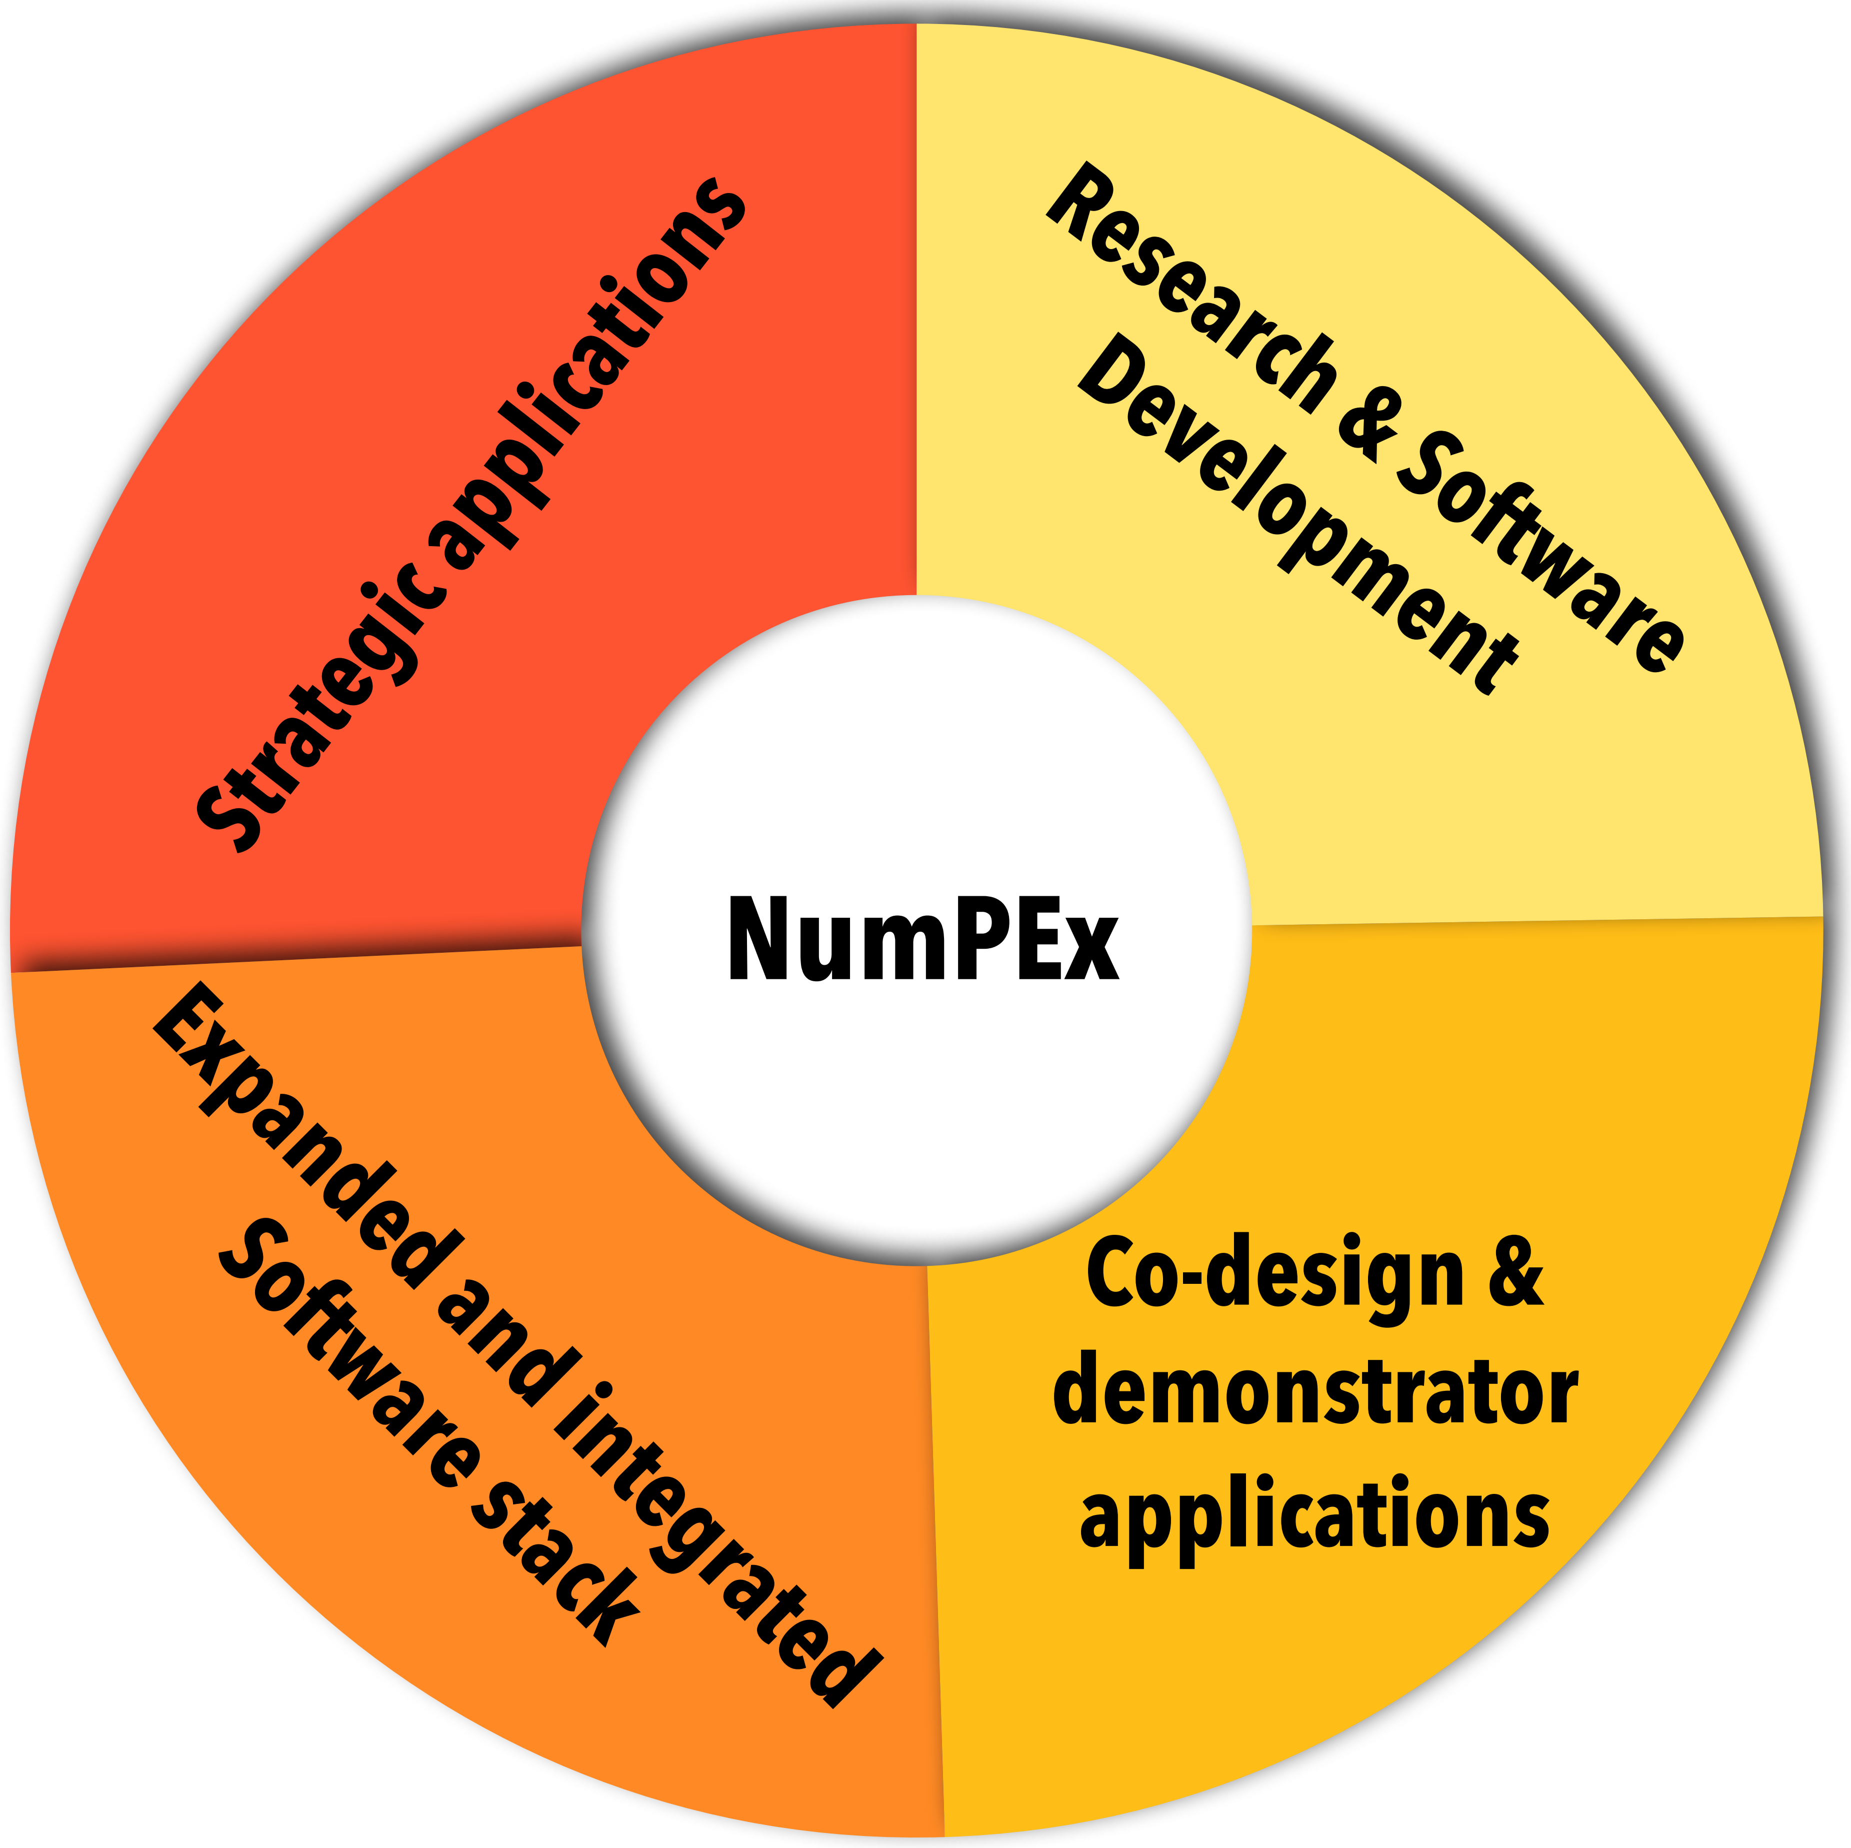

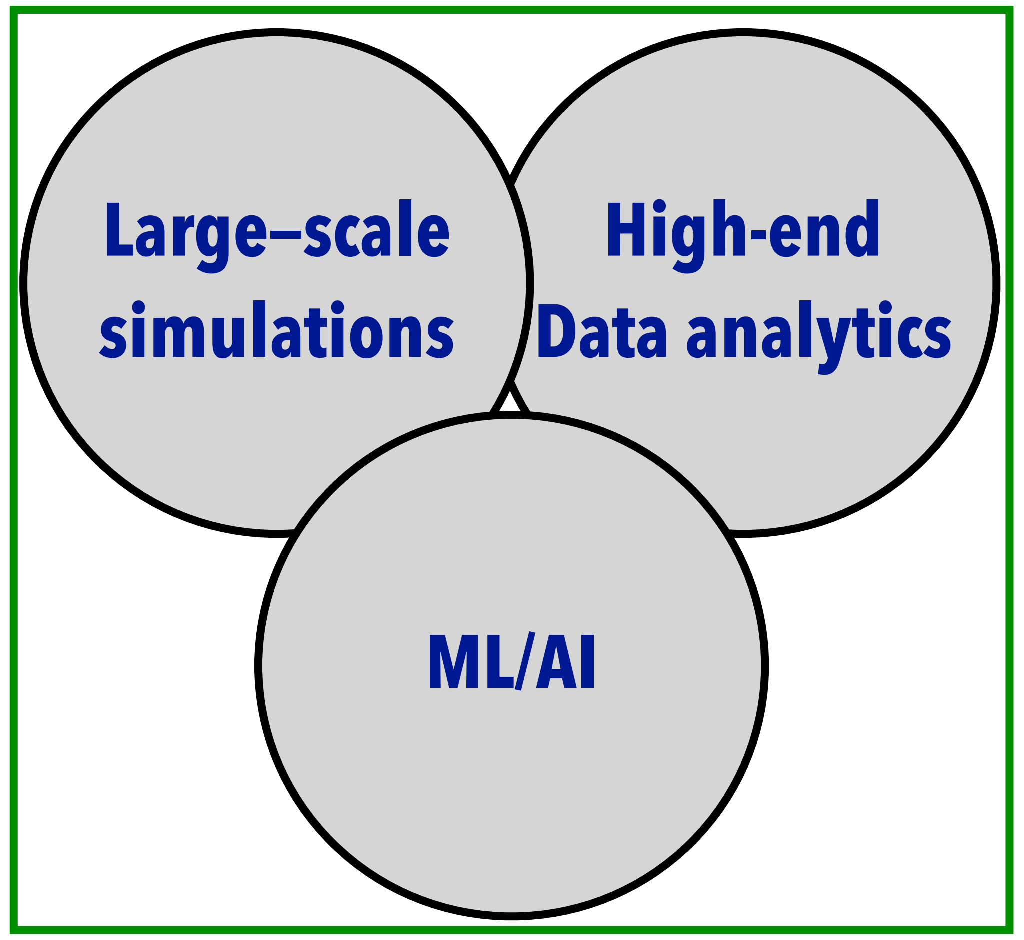

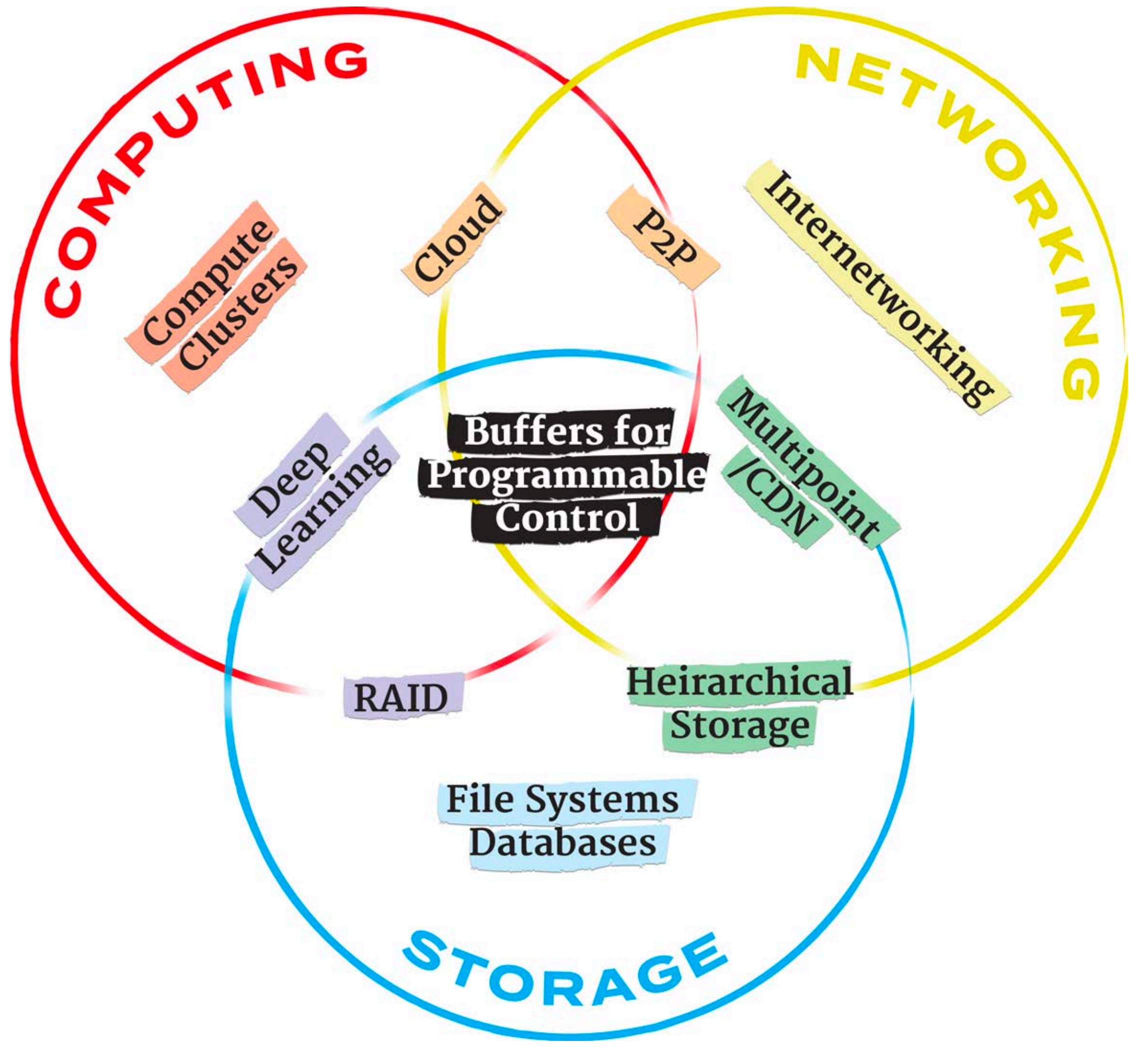

NumPEx
Exascale computing

NumPEx work plan

CSE Applications

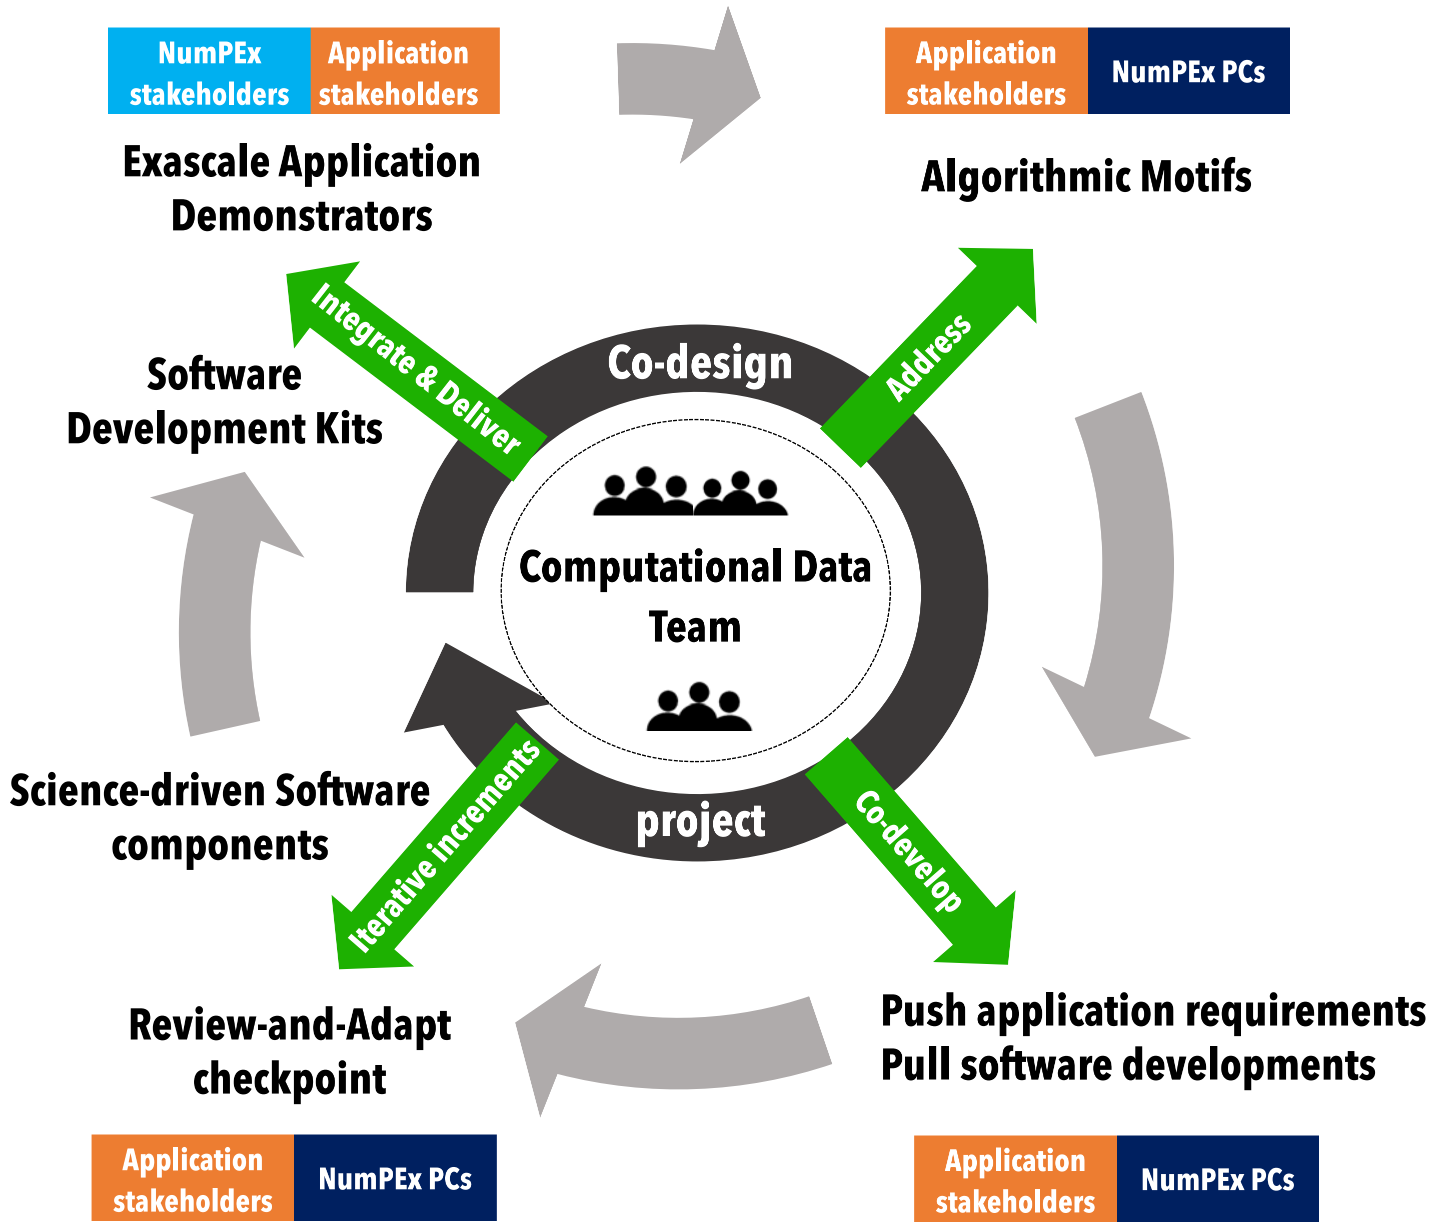

ExaDIP
software co-design & co-development, integration and delivery
J.-P. Vilotte (CNRS)
V. Brenner (CEA)

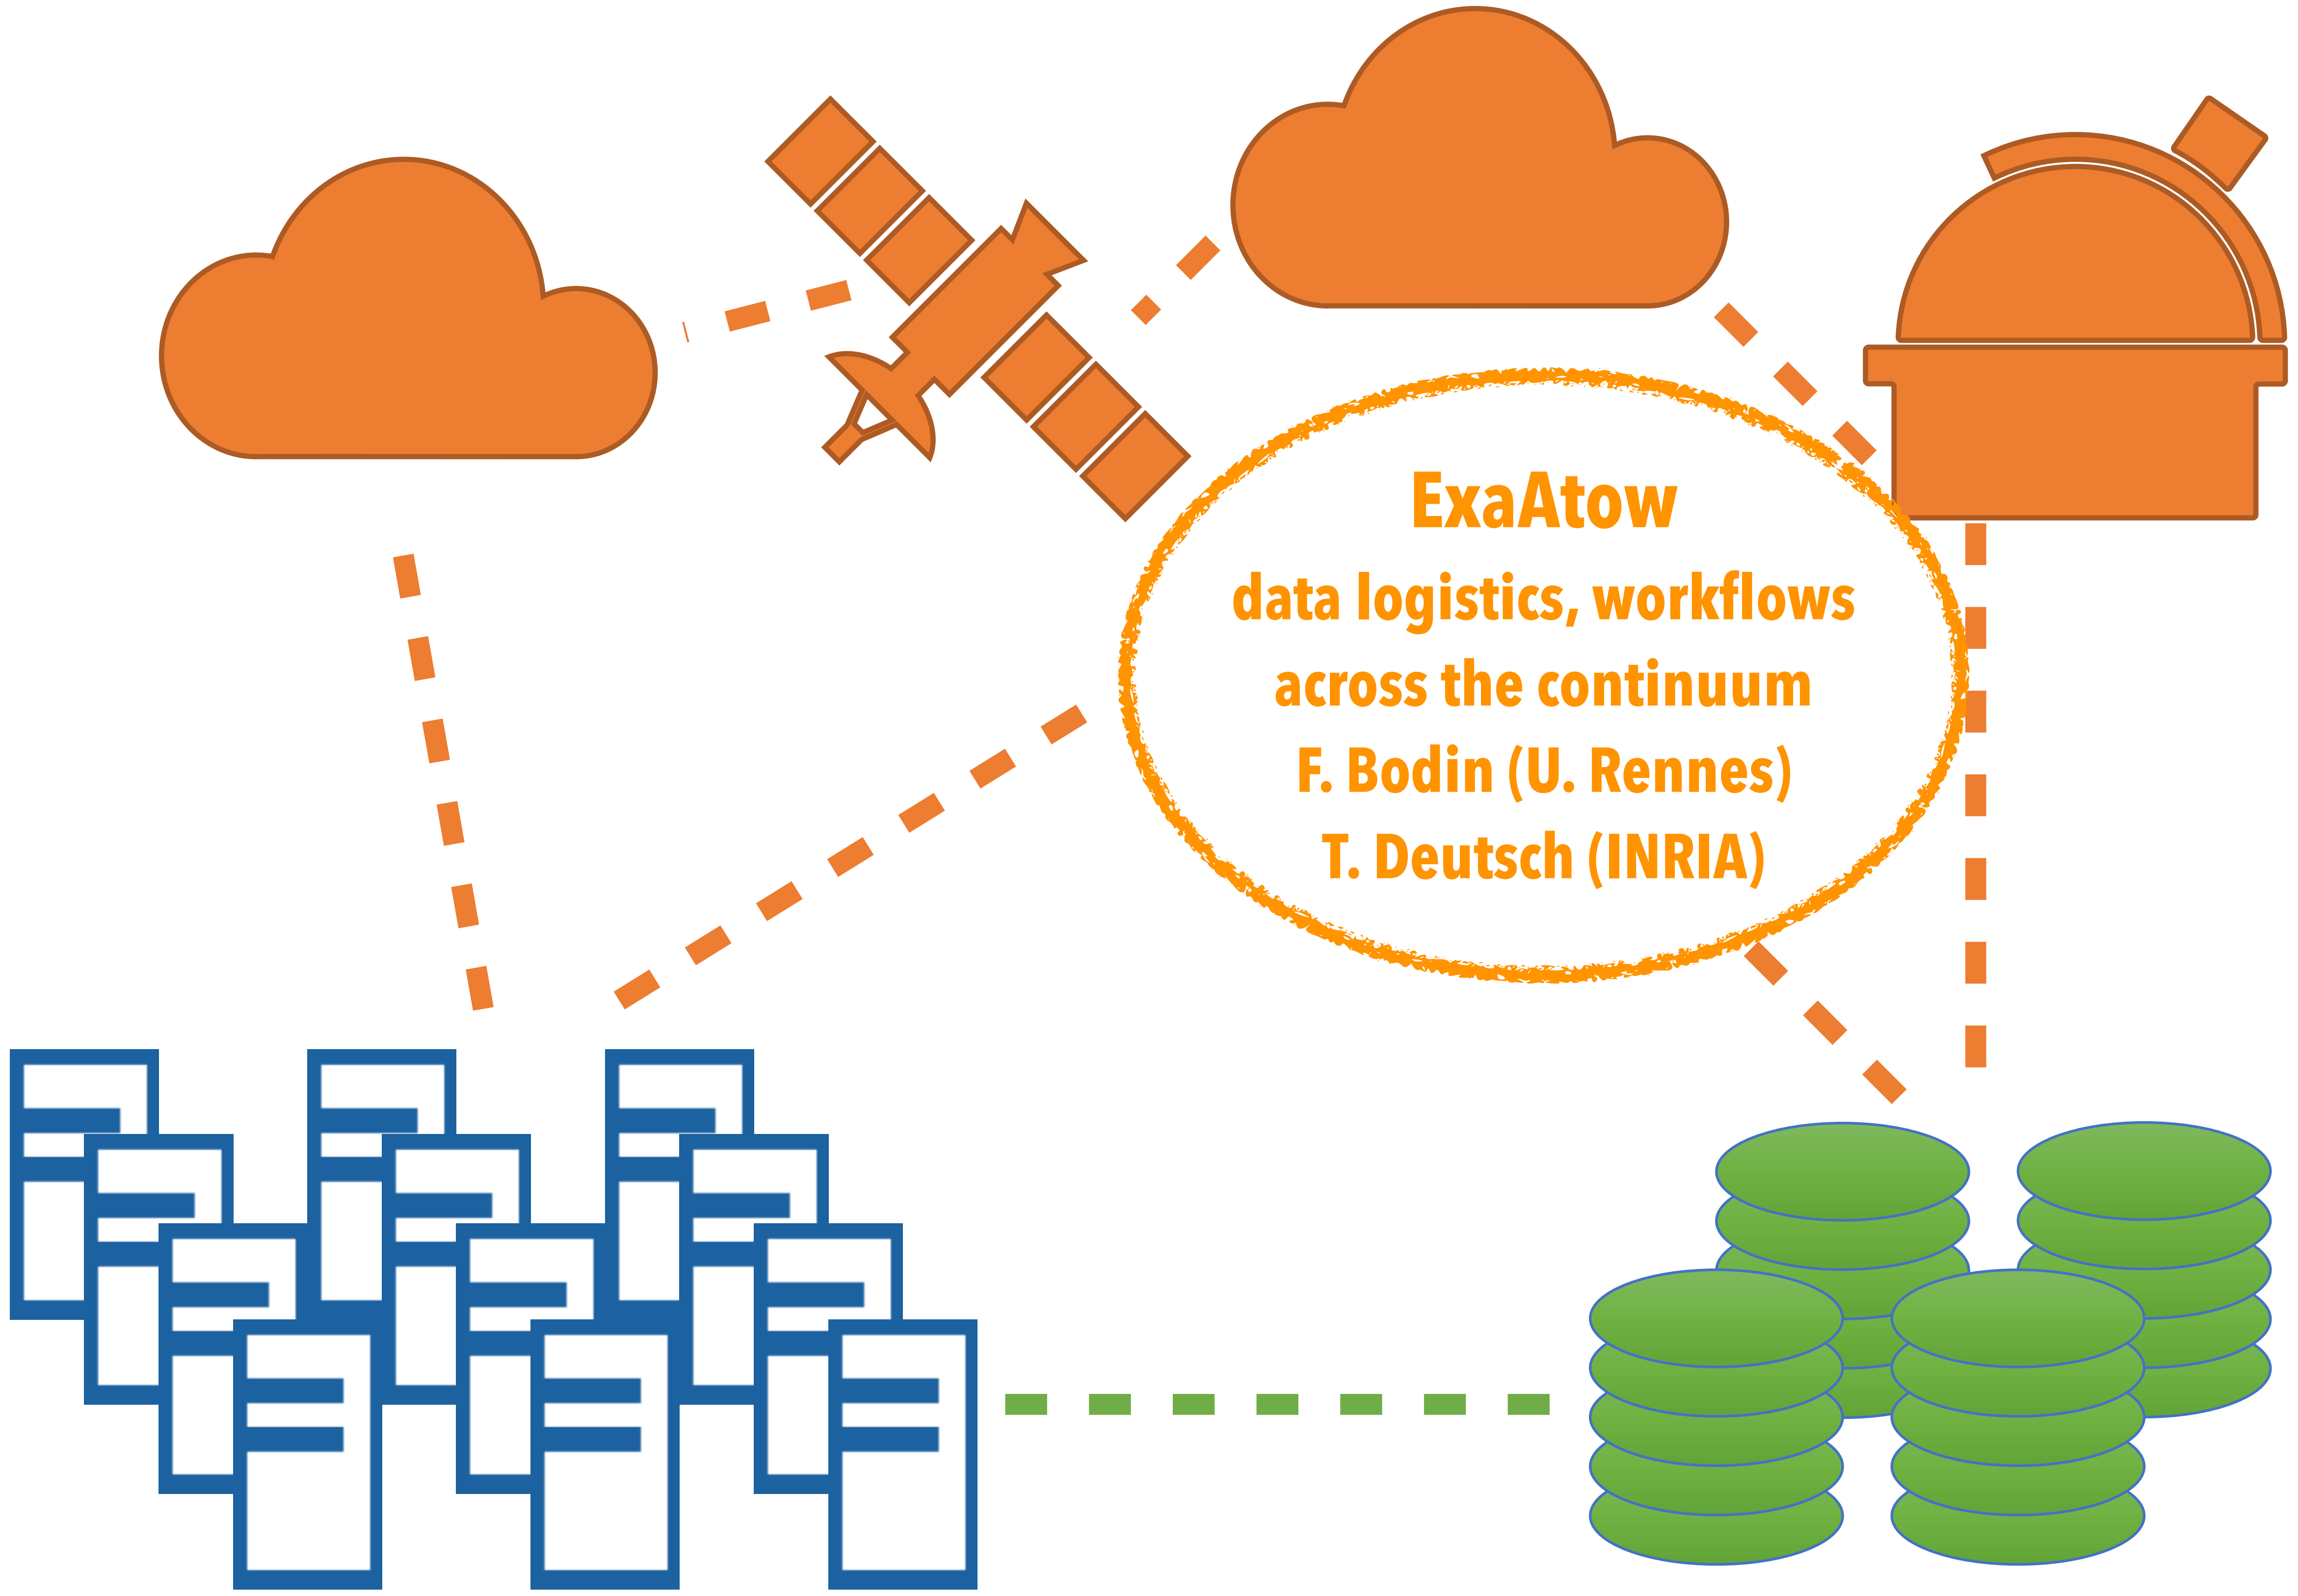

ExaAtow
data logistics, workflows across the continuum
F. Bodin (U. Rennes)
T. Deutsch (INRIA)

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)

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 Software Development Kits (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 Computational Data Team to enable agile co-design projects***
- *Broaden awareness/use of robust software components in exascale application development methodologies*

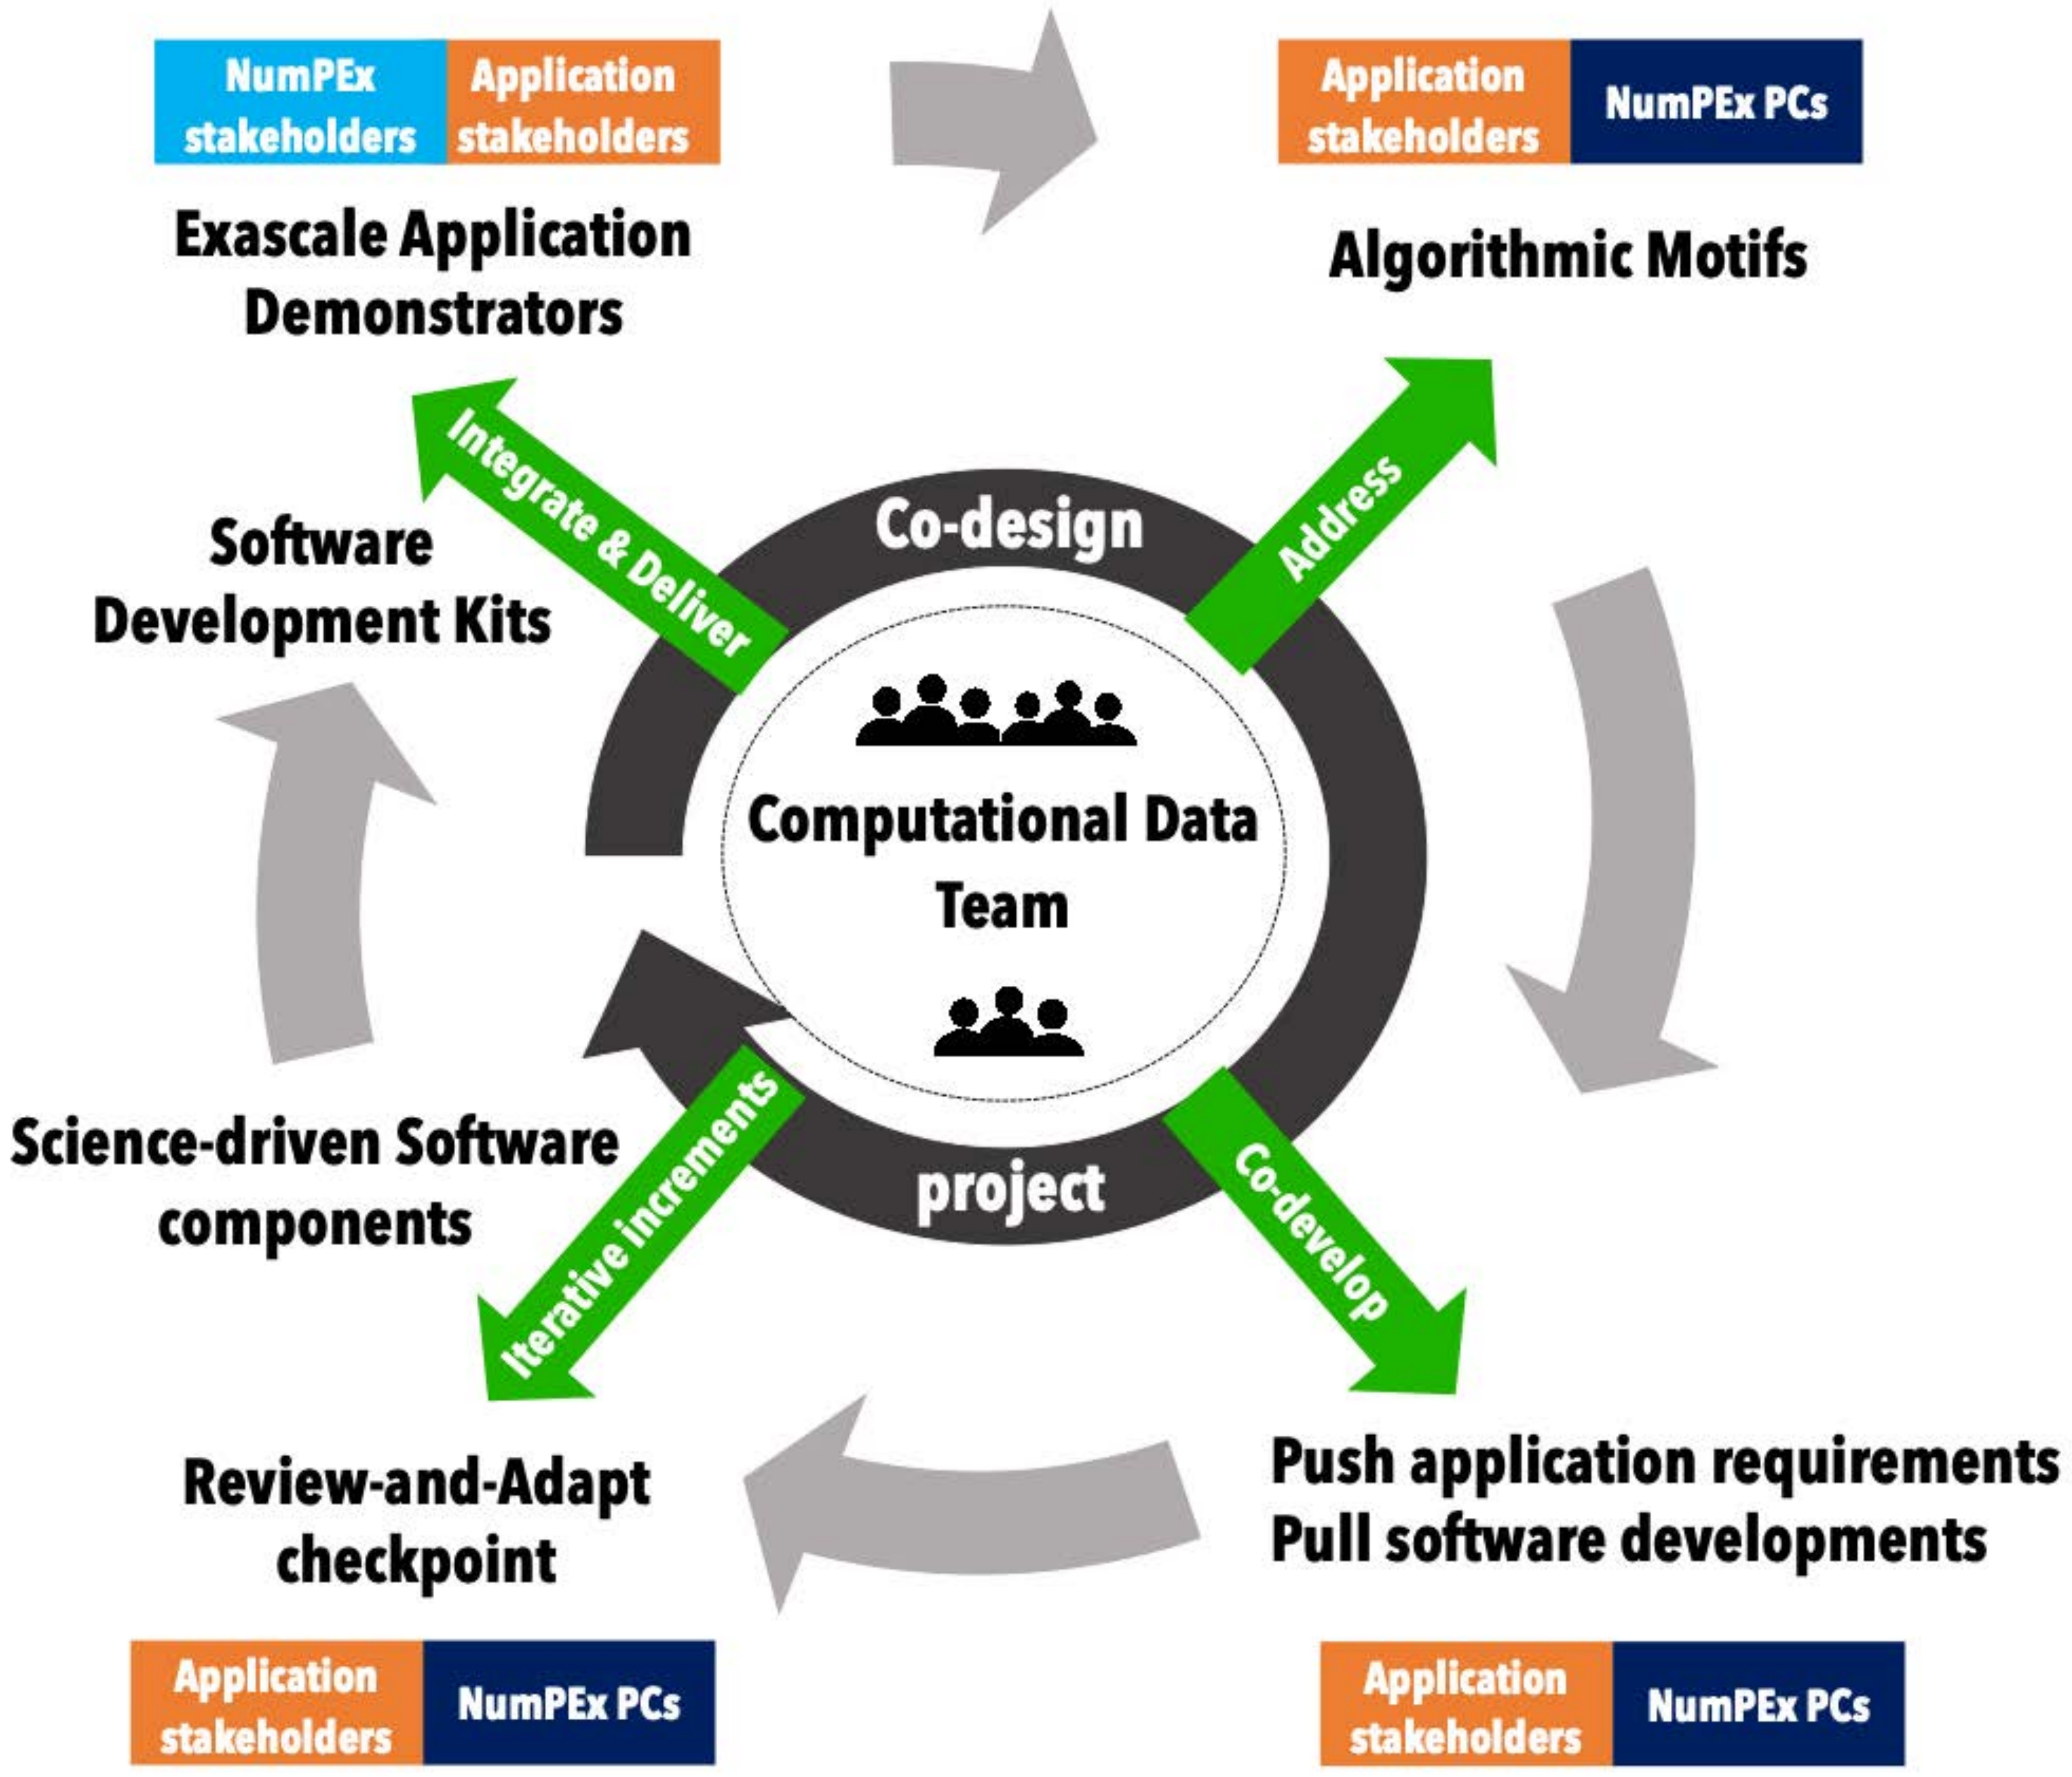

Team: CNRS, CEA, INRIA

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

NumPEx

Exascale computing

CSE application demonstrators


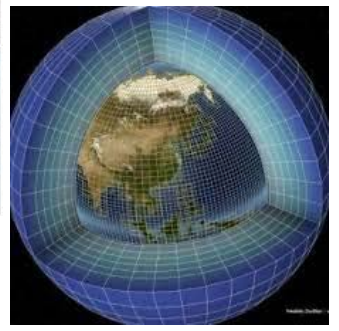

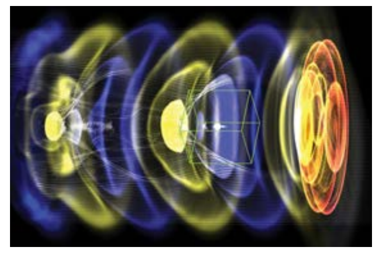
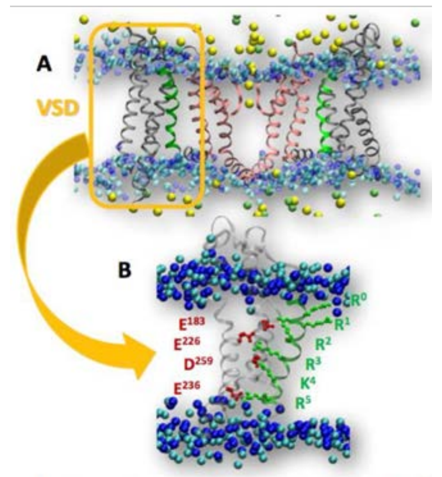
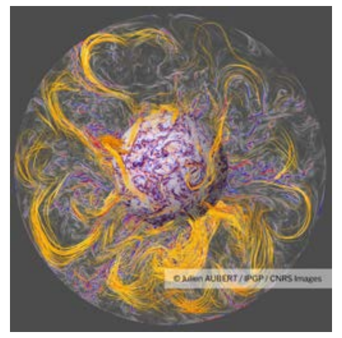

High-impact science and engineering exascale challenge problem

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


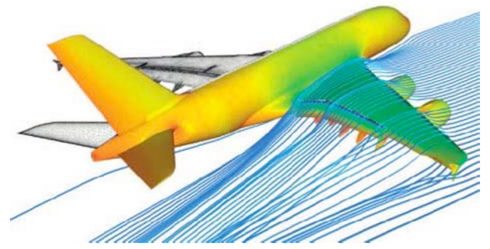

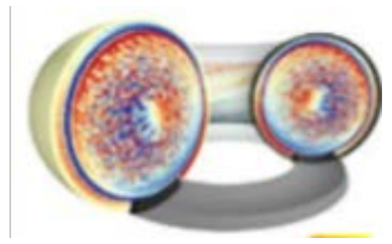
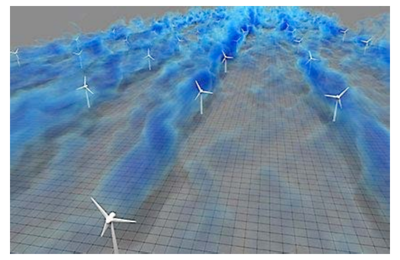

Astronomy & Astrophysics
Earth System Models & environment
Environmental extreme events
Computational biology & Life science
Laboratory laser-plasma physics
High-energy particle physics
Quantum chemistry and materials
Digital health
Environmental & societal risks
Urban systems planing
Magnetically confined 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

Identified cross-cutting algorithmic motifs

Efficient discretisation for PDEs @ Exascale

- *PDE-based multi-physics multi-scale simulations (FV, FEML, SEM, HDG), unstructured hexahedral and tetrahedral meshes, isotrope & anisotrope 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*

Goals

- proceed and integrate logical suites of software components and tools addressing cross-cutting computation and communication patterns with emerging hardware technologies
- Deliver them as software development kits easily deployable and instantiable within CSE application environments for testing, accelerating exascale development and feedback.
- Improve scientific software development methodologies balancing trade-off between fine-grained control and modularity based on integrated/interoperable libraries

Flexibility/Interoperability

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

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
- support collaborative and coordinated continuous integration/benchmarking experiments with standardised performance tools to guide optimisations

NumPEx
Exascale computing

ExaDI Software Development Kits (SDKs)

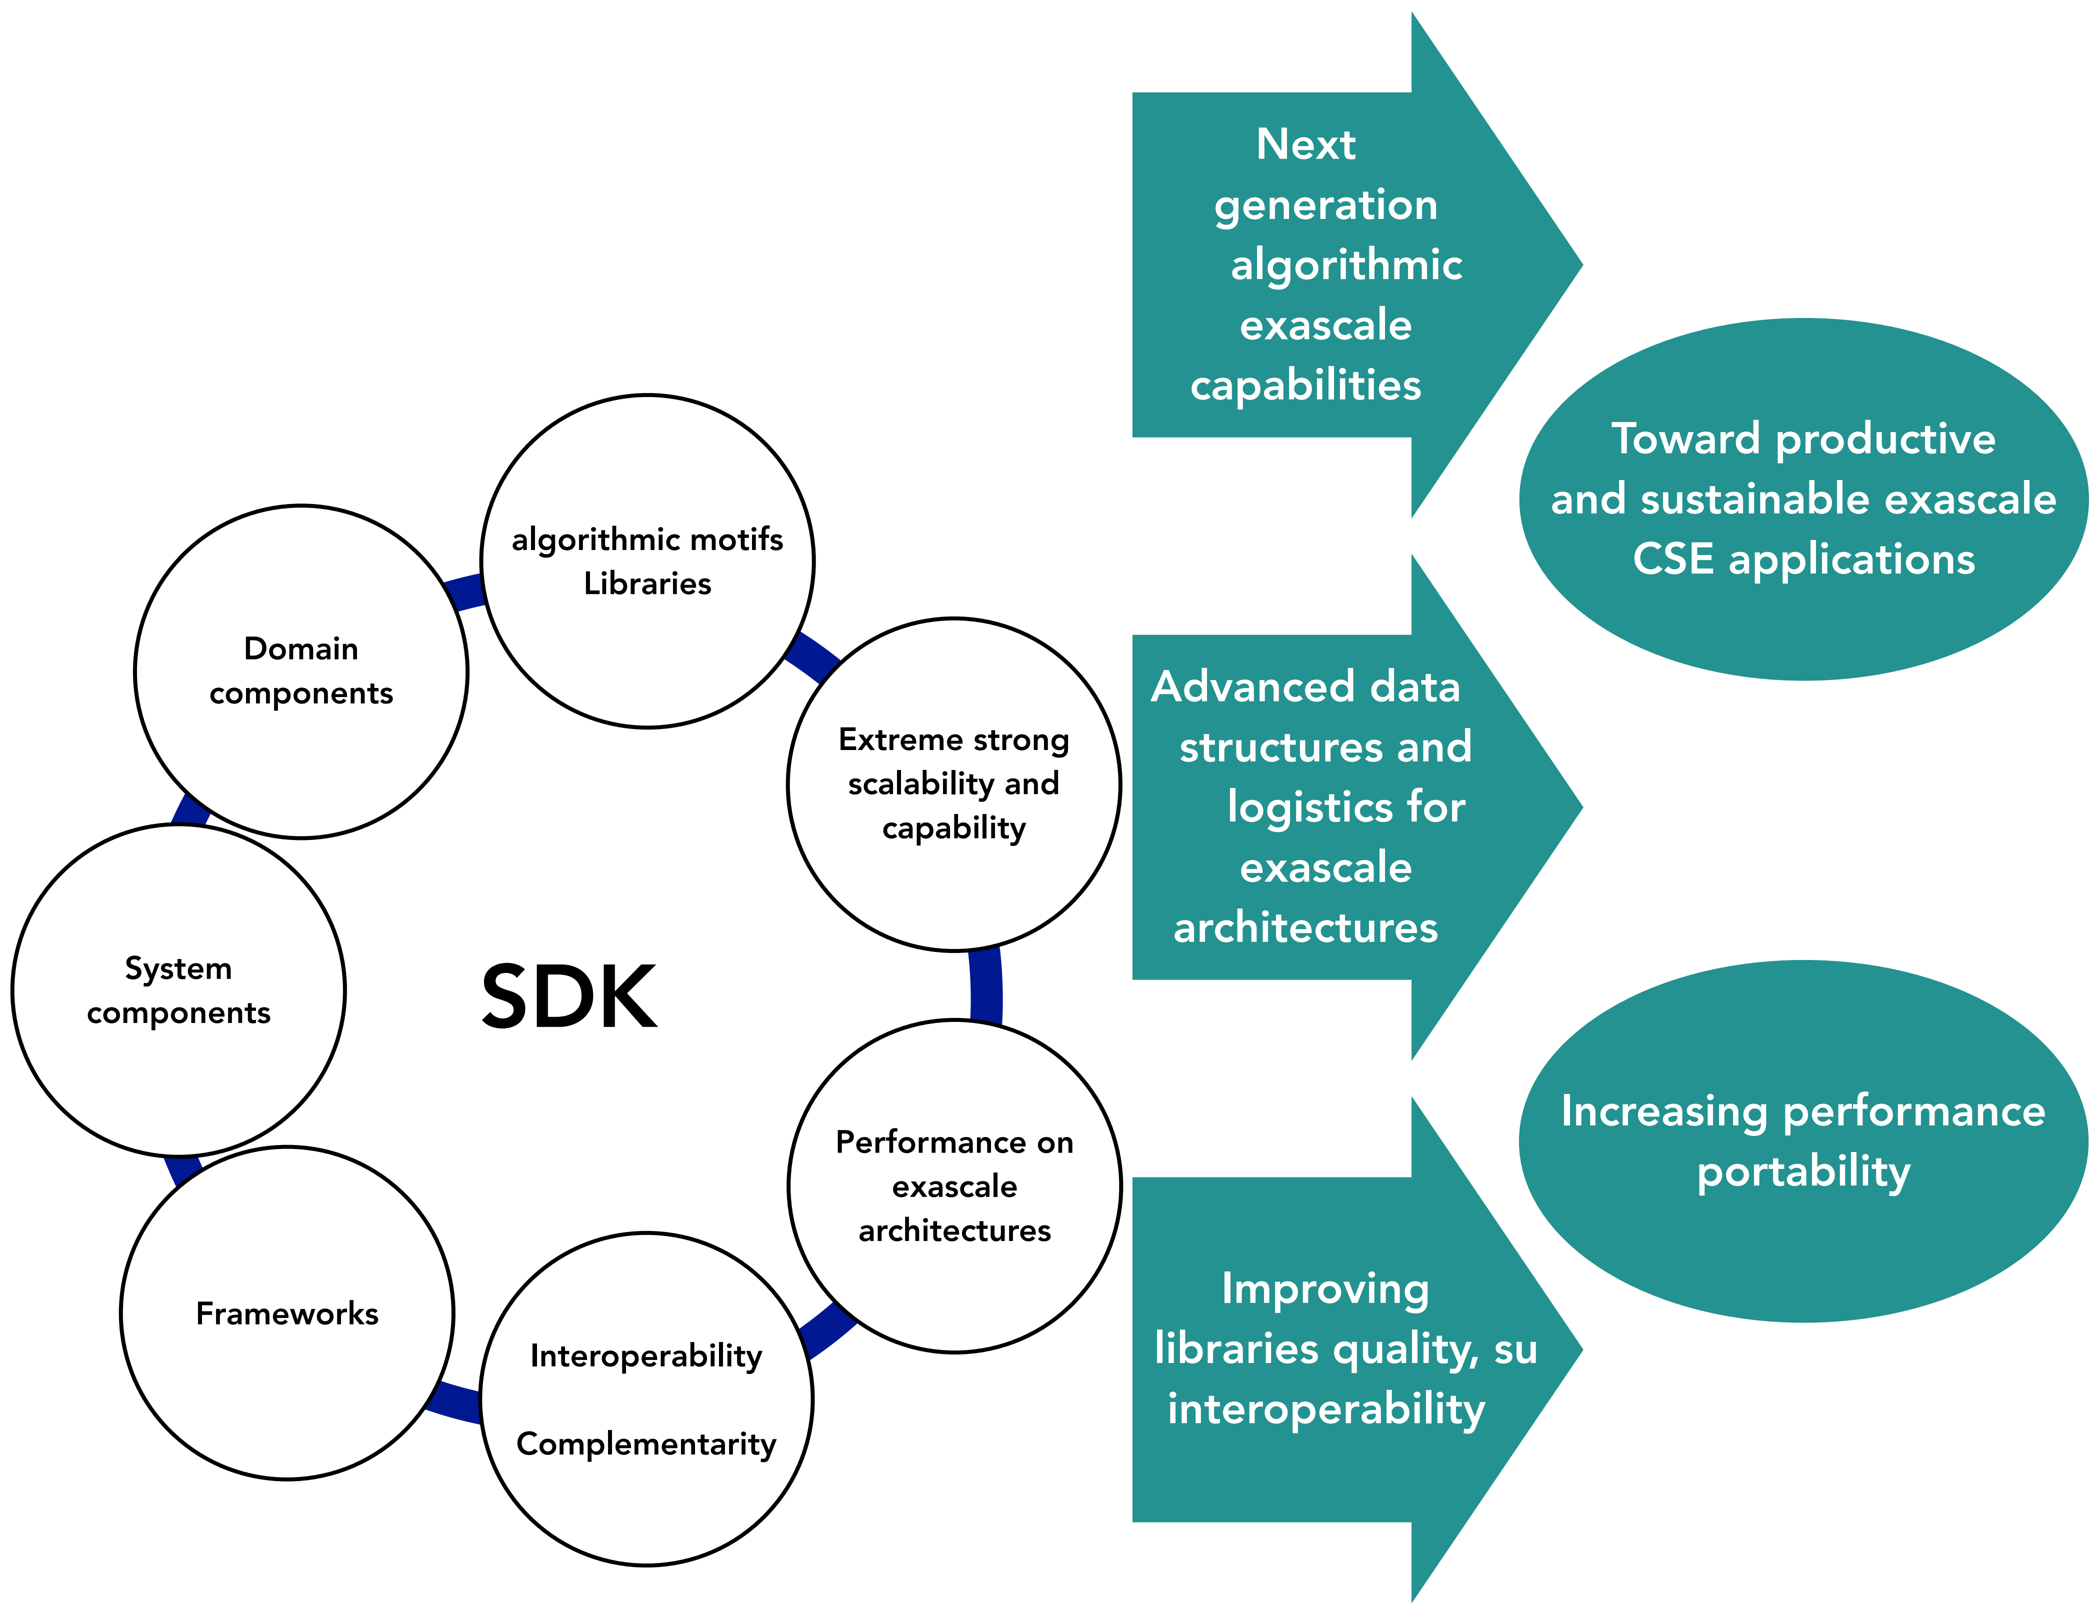

Timeline: SDK release 1 → SDK release 2 → .... → SDK release n

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

NumPEx
Exascale computing

Block-structured AMR @ Exascale

Exa-DI NumPEx

- ***Multi-scale and particle-based multi-physics and multi-scale simulations***
- ***Separating data structures and basic operations from algorithms that use them***

ECP AMReX SDK https://amrex-codes.github.io/amrex/

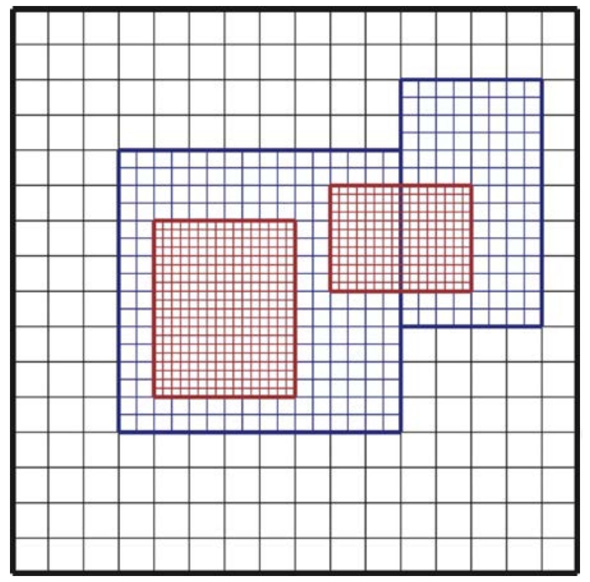

Distributed data structure (MultiFab)
User-specified subtask communication strategy

Grid-levels based AMR

Algorithms

- Block-structured mesh representation
- Data abstractions and containers (mesh, particles, embedded boundaries)
- Operation stubs and operator formats
- AMR time stepping (sub-cycling) approaches
- Linear Multigrid solvers
- ODE solvers
- Memory management and I/O optimisation
- Communications (particle-particle, particle-mesh)
- Task-based parallelism and dynamic execution model
- Load balancing
- abstraction layers (portable programming model)

Application Demonstrators

- ***Cosmology/astrophysics***
  - Dyablo-GINEA, Dyablo-WholeSUN, PHARE
- ***Ocean Modelling***
  - Croco/NEMO
- **Magnetically confined fusion plasma**
  - Gisela-X
- ***Earth & environment***
  - GEOXIM-IFPEN, GEOSX-Total
- ***Aeronautics/CFD Combustion***
  - Sonics/Onera, Safran, Yales2
- ***Industrial Risk & Safety evaluation***
  - *Manta*

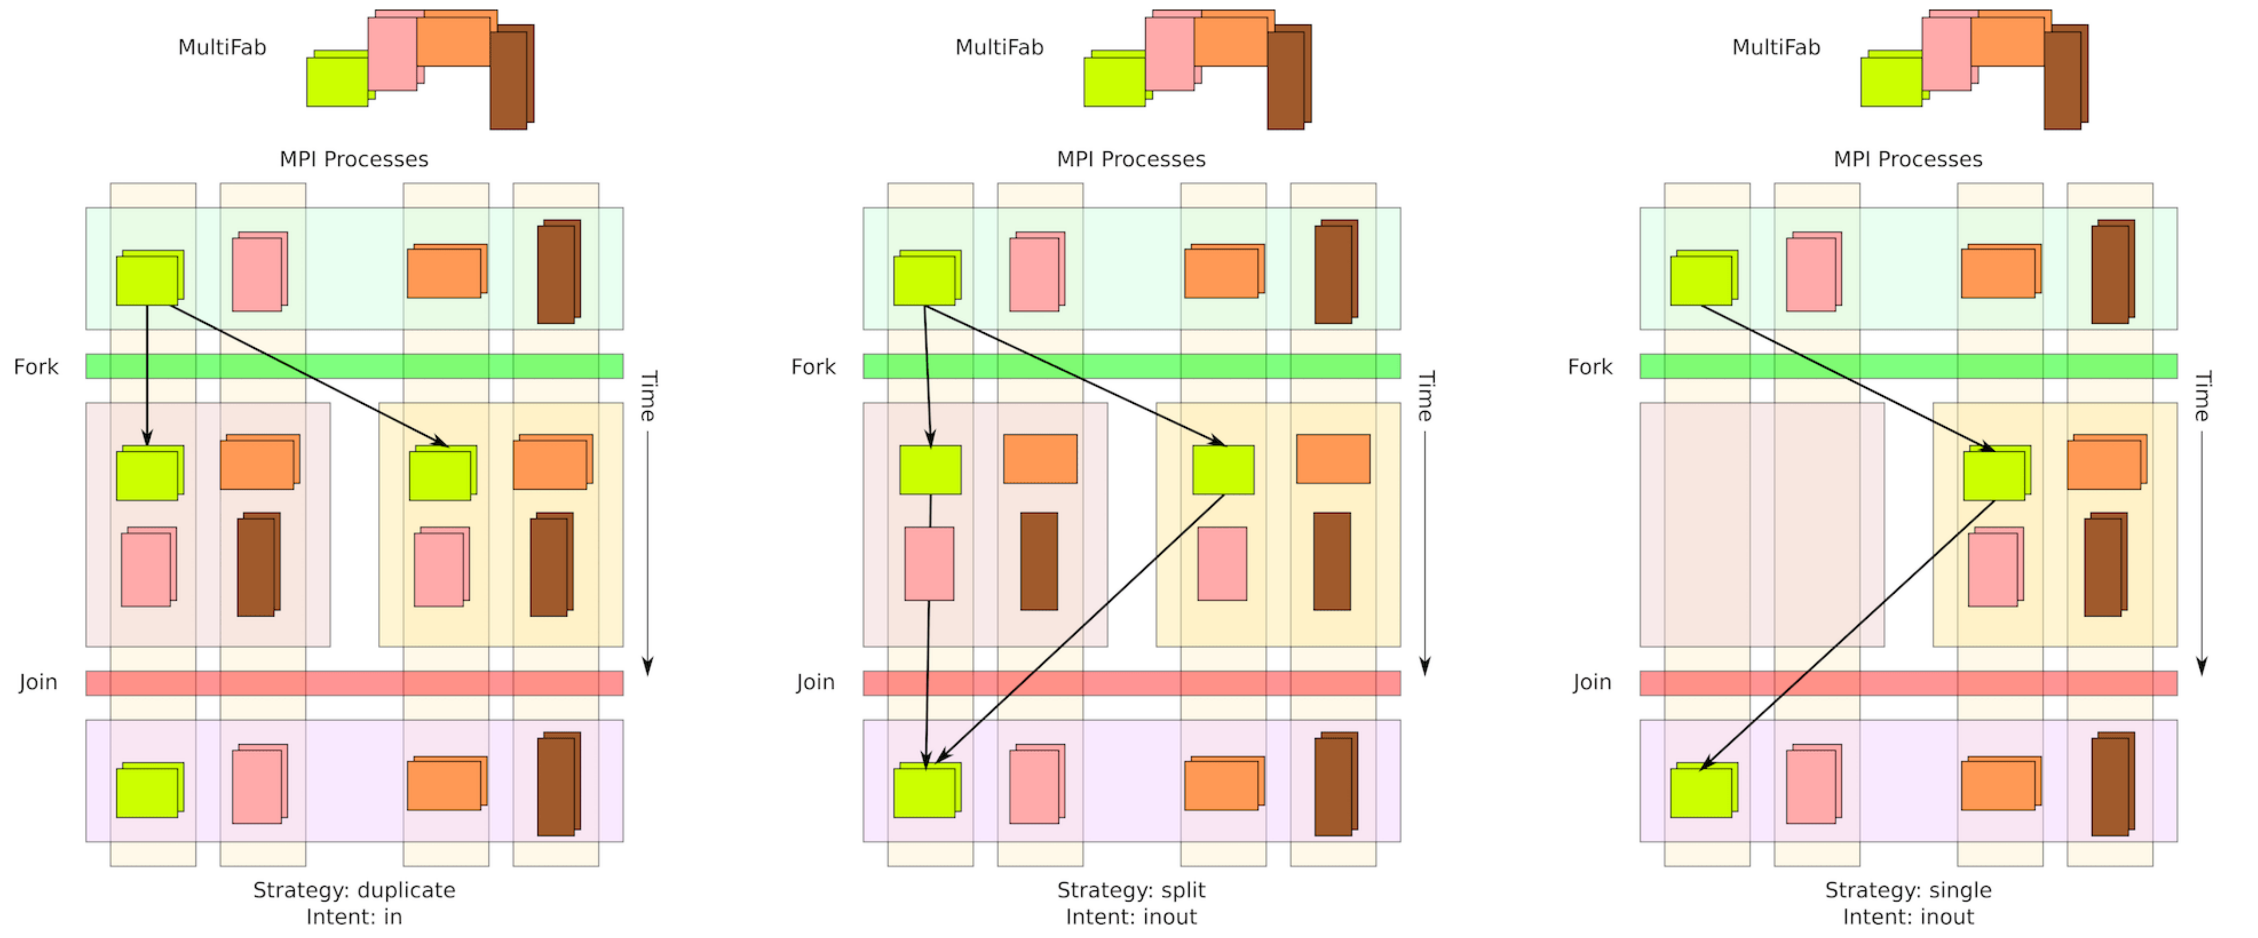

https://xsdk-project.github.io/MathPackagesTraining2022/lessons/amrex/

Software components, standards

- **ECP AMRex:** AMRex, ALPINE/ZFP, PETsc, Hypre,SUNDIALS, CodeGen MxUI/MUI, AMRVis, SENSEI, VisIt, ParaView, Exa IO/HDF5, Kokkos, RAJA
- **NumPEx:** Dyablo, SAMURAI, ArcaneFramework,GEOS, MGRIT/Parareal, PyMGRIT, MCGSolver, MUMPS, Scotch, LvArray, DDC, AGIOS, Damaris, METIS, HDF5

NumPEx
Exascale computing

Exa-DI flow of product integration and delivery

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 logical software components portfolios
- Co-develop well-specified proxy-apps / many-apps suites
- Software integration, testing & profiling tools, benchmarking specifications

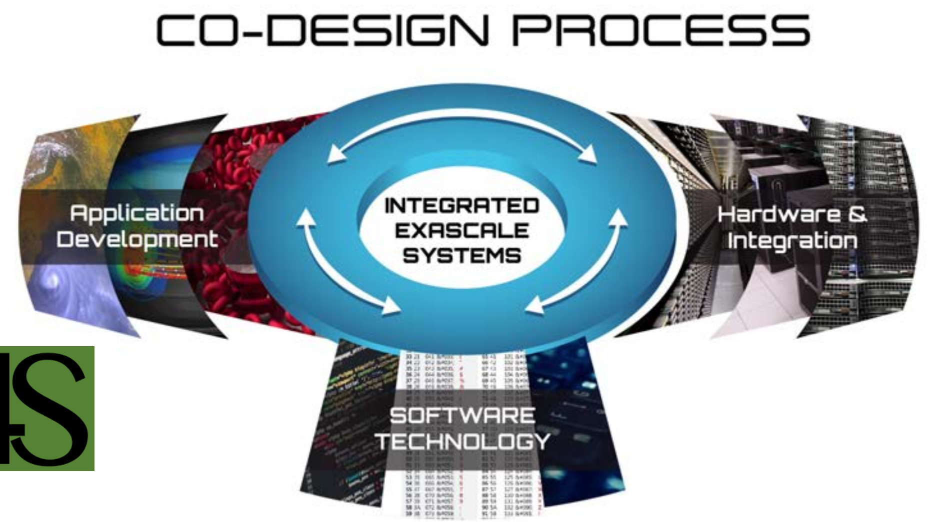

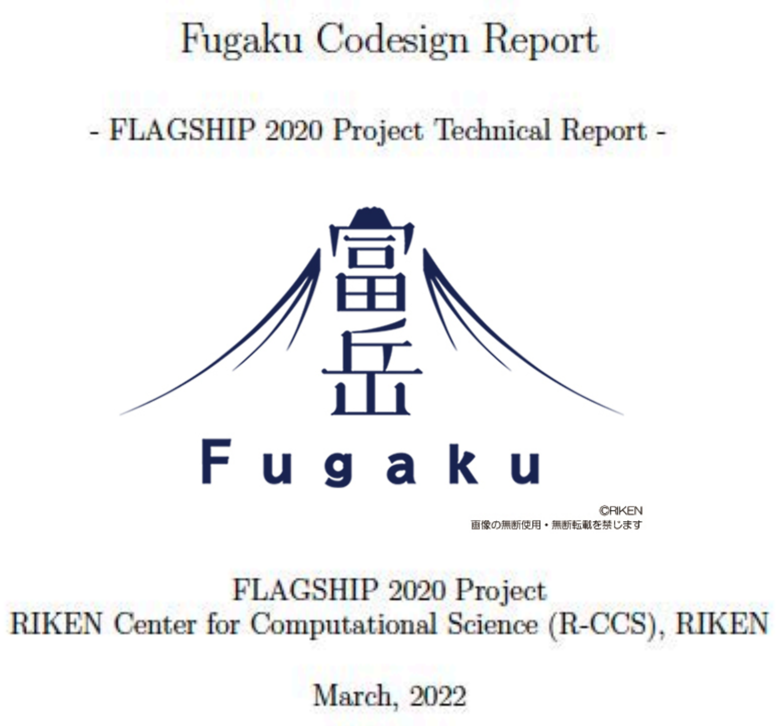
Fugaku Codesign Report

- FLAGSHIP 2020 Project Technical Report -

Fugaku

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

March, 2022

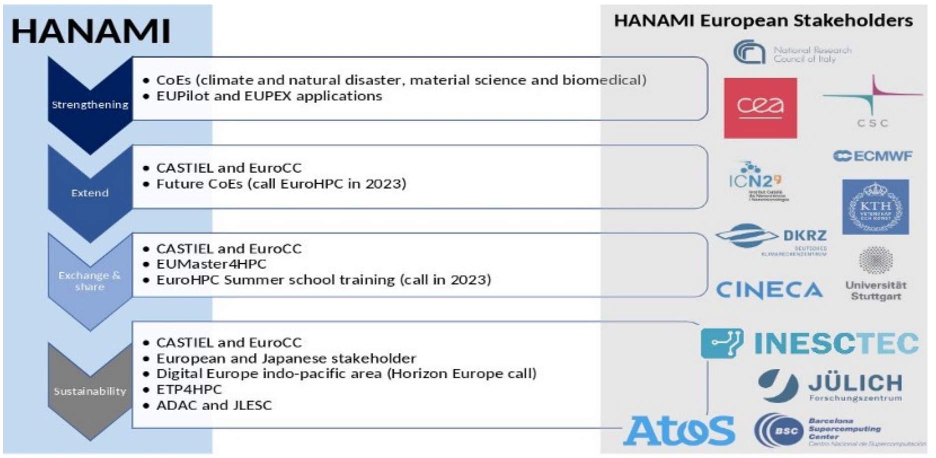

Workshop objectives and return

Develop and share a common understanding of:

- Exa-DI agile co-design and co-development process driving collaboration between software packages development and integration with overall NumPEx technologies, streamline developper and user workflows, maintaining testing and benchmarking, and coordinate SDK releases
- role and responsibility 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 block-structured AMR @ exascale

Co-identify across several ADs:

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

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

- what software components can be leveraged,
- what gaps and missing functionalities need to be addressed
- what new and/or alternative components need to be co-developed by agile teams (objectives, deliverables)
- what mini-apps and proxy-apps can 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 a follow-up with a mini-apps / proxy-apps GT