RÉPUBLIQUE FRANÇAISE
Liberté
Égalité
Fraternité

Secrétariat général pour l'investissement

https://www.numpex.org

FRANCE 2030

PROGRAMME DE RECHERCHE
NUMÉRIQUE POUR L'EXASCALE

Preparing CSE Applications for Exascale
The Exa-DI NumPEx project

J.P. Vilotte (CNRS)

Exa-MA General Assembly, Strasbourg, 14-15 January 2025

https://gitlab.inria.fr/numpex-pc5

anr agence nationale de la recherche

Who should care about this?

- **Productive Exascale systems:** integrate applications and software technologies with an expanded exascale software stack easily deployable on facilities and instantiable into application environments.
- **CSE applications and computing architectures are gaining in complexity and heterogeneity:** emerging hybrid workloads (HPC/HPDA/ML-AI): high costs to develop, build and deploy on Exascale facilities, impairing portability and reproducibility
- **Maintaining and adapting existing scientific application codes to rapidly evolving exascale architectures:** needs a lot of work facing limited resources and expertise.
- **Productive Exascale applications:** require improved scientific software development methodologies leveraging high-quality compatible software components (libraries, frameworks, abstraction layers), programming and execution models, data and compute workflow tools.
- **NumPEx project Exa-DI:**
  - Implement an agile software co-development process across NumPEx in partnership with CSE Applications: deliver Motif-based proxy apps and collections of compatible software components (SDKs) addressing cross-cutting computing and communication workload patterns, including data analytics and AI/ML, covering a wide range of strategic domains
  - Promote the use of meta build technologies and containers (e.g., Spack, Guix-HPC), software community guidelines and certification, CI and performance evaluation methodologies

Challenge in developing exascale scientific applications

Better Scientific understanding and discovery
Software productivity

More complex and hybrid (HPC/HPDA/AI) software, data and compute workflows to map on more complex and specialised hybrid accelerated architectures

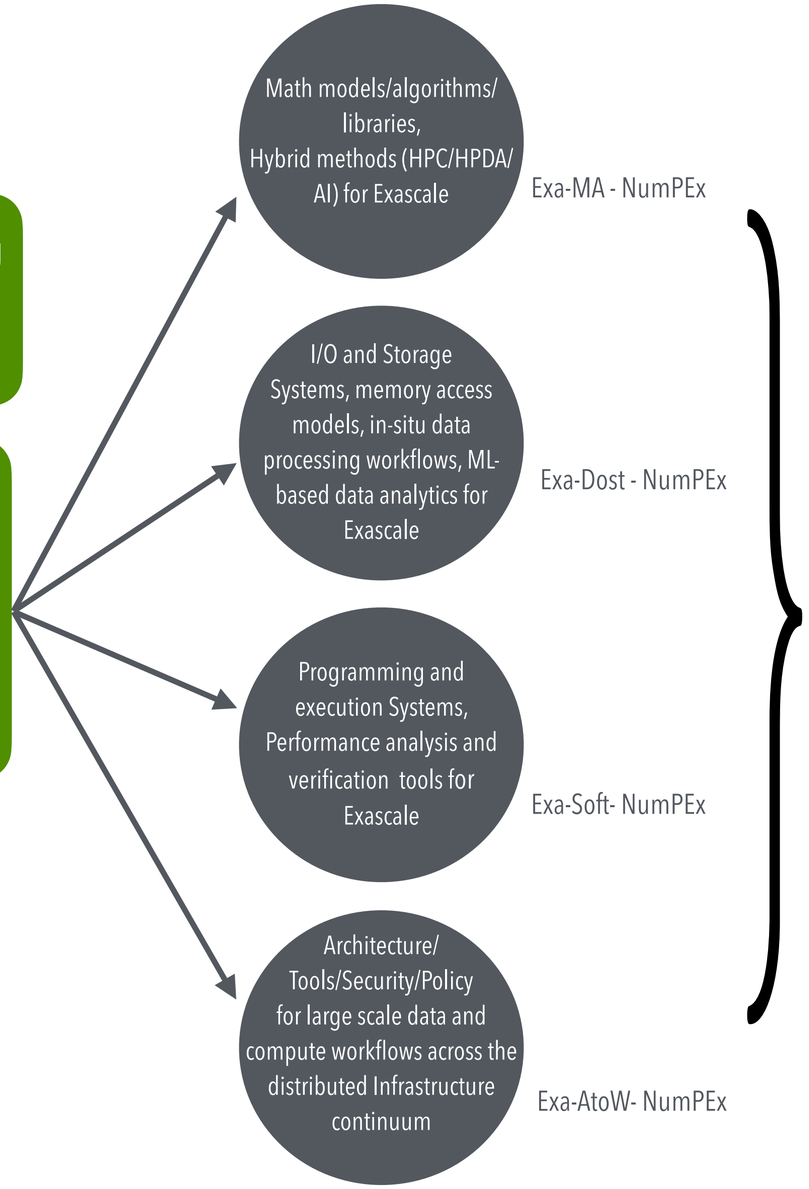

Scientific results can no more be obtained with heroic programming

- ***agile scientific software engineering methodologies:*** deliver and leverage high-quality cross-cutting compatible and reusable software components and metrics with impact to real applications

➡ ***Agile teams in partnership with application teams:*** gather different expertise, roles and responsibilities

Technical challenges

- Model and software system can be under R&D
- Requirements change during the lifecycle as knowledge grows
- Verification complicated by floating point representation
- Increasing architectural complexity, heterogeneity and diversity

Sociological challenge

- *Competing priorities and incentives*:
  - Sponsors often care more about scientific publications than software per se
  - Balancing development and maintenance with research
- *Limited resources by far*

Exa-DI CSE Application Demonstrators

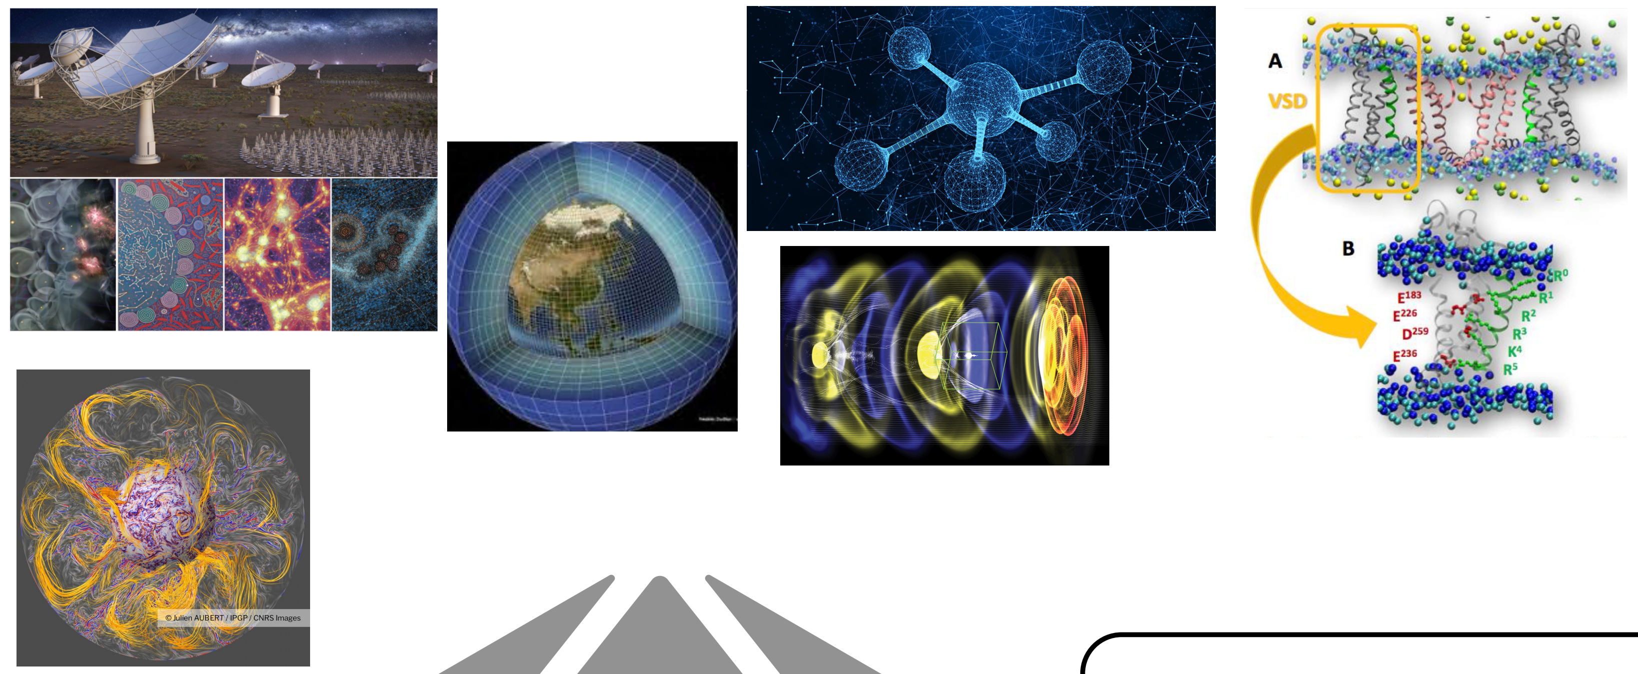

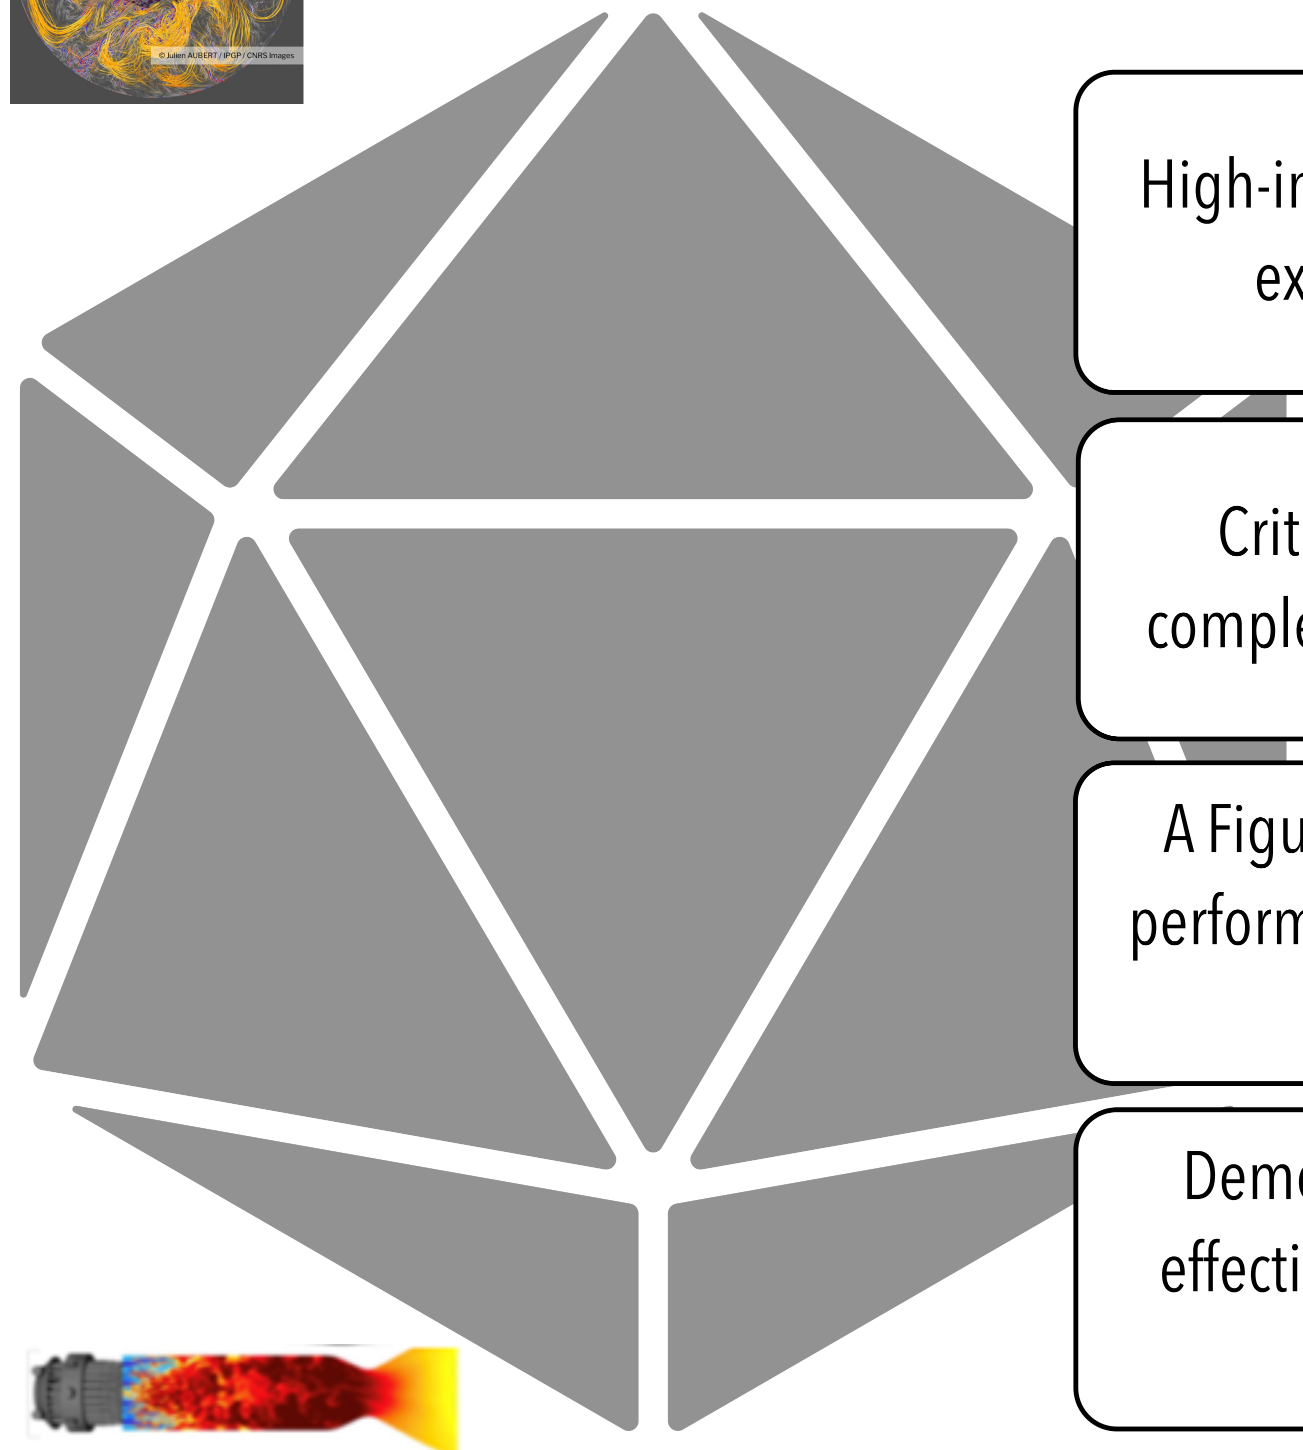

High-impact science and engineering exascale challenge problem

Criteria for assessing successful completion of the challenge problem

A Figure of merit (FOM) quantifying performance enhancement and rate of science

Demonstration and assessment of effective Software Development Kits (SDKs) 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 accelerator-based exascale architectures
- New multi-physics and multi-scale capability
- Hybrid HPC/HPDA/AI coupled workloads
- On-line AI-based data analysis/reduction
- New mathematical approaches, including AI, algorithms and model improvements
- Leveraging robust and accurate logical collections of reusable and composable Motif-based software components
- Improving performance portability through abstraction layers
- Programming and execution models for heterogeneous accelerated architectures
- Build fondations for a sustainable scientific software stack

First Exa-DI ADs

AD name	Category	Short description	Related projects	Teams
AEVOL	Computational Biology & Digital Health	Evolutionary processes, genetic trajectories, insilico experimental evolution simulator, AI)based data analytics	EvoAct, Evoluthon, ANR NeGA	CNRS, INRIA, Michigan State University, UGA, INSA Lyon
BigDFT	Quantum Chemistry and Materials	Ab-initio atomistic simulation, quantum mechanics, AI-based data analytics, molecular structures, Bio-systems	CoE MAX	CEA, RIKEN, U. Bristol
ChEESE-EarthDynamoWave	Earth System & Environment	Earth Dynamo, MHD, data assimilation, Physics-based AI/ML	CoE ChEESE-2P	CNRS, IPGP, UGA
ChEESE-WaveSimulation/ FWI	Earth System & Environment	Wave propagation, inversion, data analysis and control, urgent computing, AI based data analytics	CoE ChEESE-2P	CNRS, BRGM, UGA, U. Nantes
CROCO-Mediation	Earth System & Environment	Costal-regional ocean modelling, in-situ data analytics and control, data assimilation, Physics-based AI/ML, in-situ data analytics	MEDIATION, PEPR TRACCS	CNRS, CEA
DYABLO	Astronomy & Astrophysics	MHD and plasma physics, radiation	PEPR ORIGINS SKA-France	CEA, CNRS, INRIA
GEOSX-MAKUTU	Earth System & Environment	Multi-physics simulator platform, wave propagation, inversion, CO_2 sequestration and monitoring	GEOSX	Total-Energy, INRIA, LLNL, Stanford
GEOXIM-IFPEN	Earth System & Environment	Multi-physics simulator platform, Hydrogen & CO_2 storage, geothermal energy	PEPR sous-sol	IFPEN, CEA, UP Saclay
Gysela-X	Magnetically Confined Plasma Physics, Energy Production	Edge-core confined electromagnetic plasma turbulence, in-situ data analytics	IR* ITER, TSVV, EoCoE	CEA, MdlS, INRIA, CNRS
KATY	Computational Biology & Digital Health	Biological molecular pathways, cell typing, AI-based data analytics	PEPR Digital Health, KATY & CANVAS EU projects	CEA, INSERM, Cosnet lab, U. Zaragoza
LHcb	High-energy Particle Physics	Data analytics and event detection, reconstruction, wide-area workflows	LHCb, IR* LHC, CERN	CNRS, CEA, CERN

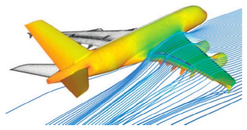

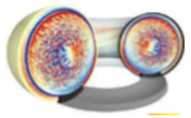

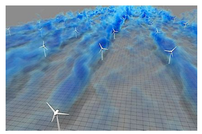

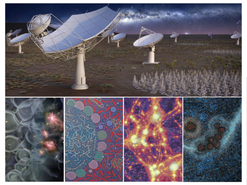

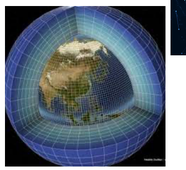

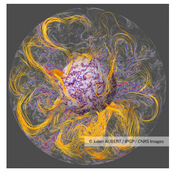

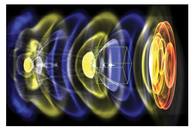

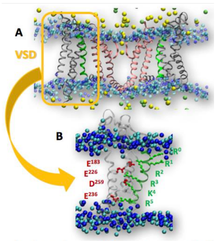

AD name	Category	Short description	Related projects	Teams
MANTA	Industrial installations Safety	Structure mechanics, multi-component and reactive fluid dynamics, fluid-structure interactions, nuclear plants safety, accidental scenarios		CEA, CNRS, INRIA, U. Lorraine
SAFRAN	Sustainable Transport & Mobility, Energy production	Turbulent reactive flows in complex geometries, fluid-structure interactions, aerodynamics, Combustion		INRIA, SAFRAN
SKA SDP	Astronomy & Astrophysics	Streaming data analytics, imaging, calibration, TOAs	SKA-France, SKAO, LOFAR/ NenuFAR,	CNRS, ICEA, NRIA, OP, OCA
SKA-SRC	Astronomy & Astrophysics	Multi-modal data analysis, wave scattering statistics, High-dimensional sampling, Simulation Based Inference	SKA-France, SKAO, LOFAR/ NenuFAR ERIC	CNRS, CEA, INRIA, OP, OCA
SMILEI	Plasma physics	Plasmas under extreme condition, laser particle beams, electron laser wakefield acceleration	CoE PlaXma, AIDAS Joint Virtual Lab	CNRS, CEA, UP Saclay, MdlS, Polytechnique
SONICS	Sustainable Transport & Mobility, Energy production	Turbulent flows in complex geometries, fluid-structure interactions, aerodynamics, combustion, engineering design	EuroHPC NextSim, Clean Sky 2 Joint European Technology Initiative	ONERA, INRIA, CERFACS
TINKER-HP	Molecular Dynamics, Digital Health	High resolution molecular dynamics simulation, drugs discovery, digital health, in-situ data analytics	PEPR EPIQ, EMC2 ERC SYG, ITN PhyMol, ANR MAPPLE	Sorbonne Université, INRIA, CNRS, Qubit Pharmaceuticals, CNAM, U. Paris Nord, UP Saclay, ENS Lyon
Urban-Building-Air-Quality	Sustainable Urban System Planning	Urban wind-field simulation, air quality, energy consumption & comfort at building and urban scales, heat transfer and CFD	CoE Hidalgo-2	UNISTRA, CNRS
YALES2	Sustainable Transport & Mobility, Energy Production	Turbulent 2-phase and reactive flows, combustion, aerodynamics , engineering design	CoE CoEC	CNRS, UGA, IU. Lorraine, NRIA, Sorbonne Université, SAFRAN Tech, ONERA, Airbus

Exa-DI initiatives and on-going activities

- **Co-analysis of initial ADs and identification of initial cross-cutting computation and communication motifs**
  - Efficient discretisation for PDE @exascale; and Block-structured AMR @ exascale
  - supported by 13 ADS
- **Workshop on transactional software packaging and deployment technologies (**April 3-4, 2023, INRIA Paris):
  - Software packaging experts (Spack, Guix-hpc, Nix), regional and national (TGCC, IDRIS, CINES) centres, GENCI, Atos
- **Co-development workshop on Efficient Discretisation for PDE@Exascale** (November 7-8, 2023, CEA Saclay)
  - members of NumPeX, Exa-DI leaders, AD codes developers, and external subject matter experts (30 participants)

  ➡ **Motif-based proxy apps WGs: high-order FEM (since February 2024, Henri Calandra); unstructured meshes (since September 2024, J. Vanharen)**
- **Co-development workshop on block-structured AMR @ exascale** (February 6-7, 2024, IPGP, Paris)
  - members of NumPeX, Exa-DI leaders, AD codes developers, and external subject matter experts (35 participants)

  ➡ **Motif-based proxy apps WG (Starting early November, 2024)**
- **Guix-HPC technical meeting** (July 1-5, 2024, Bordeaux)
  - R. Wurmus (MDC Berlin), C. Baines (London), A. Enge (INRIA), F. Avats, R. Grabage, L. Courtes (Exa-DI)
- **Co-development workshop on AI4HPC @ scale** (October 2-3, 2024, Paris)
  - members of NumPeX, Exa-DI leaders, AD codes developers and external subject matter experts (40 participants)
  - **AI-coupled HPC/HPDA workflows:**
    - Image analysis @ exascale;
    - Data analysis and robust simulation based inference @ exascale

First cross-cutting algorithmic Motifs

Supported by 13 ADs

Computational & communication motif	Application Demonstrators	Science & Engineering Domain
Efficient PDEs discretisation @ exascale	CHEESE-EarthDynamo, CHEESE-WaveSimulation/FWI, GEOSX-Makutu, GEOXIM-IFPEN	Earth System & Environment, Geo-resources
	CROCO-Mediation	Climate System Models, Ocean
	MANTA	Industrial Installations Safety
	SAFRAN, SONICS, Yales2	Sustainable Transport & Mobility, aerodynamics, energy production
	Urban-Building-Air-Quality	Sustainable Urban Planning
Block-structured Adaptive Mesh Refinement @ exascale	CHEESE-WaveSimulation/FWI, GEOSX-Makutu, GEOXIM-IFPEN	Earth System & Environment, Geo-resources
	CROCO-Mediation	Climate System Models, Ocean
	DYABLO	Astronomy & Astrophysics
	SAMURAI	AMR framework for PDEs
	SAFRAN, SONICS, Yales2	Sustainable Transport & Mobility, Aerodynamics, energy production
	Thinker-HP	High-resolution Molecular Dynamics simulation
	MANTA	Industrial Installations Safety

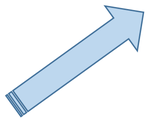

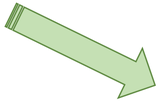

Algorithmic sub-motifs

Structured/unstructured PDEs discretisation
Mesh & discretisation libraries
Adaptive meh refinement & dynamic remeshing
Structured and unstructured mesh representation & abstraction
Efficient operator format & algebraic construction
Sparse & distributed linear algebra
Scalable matrix-free solvers
Batched dense tensor contration
Tree & Graph based matrices
General interpolation
Time integrator schemes
I/O & memory access
Portable programming models & abstraction layers
Task-based execution models

Block-structured mesh representation
Data abstraction and containers (mesh, particles, embedded boundaries)
Operation stubs and operator formats
AMR time stepping (sub-cycling)
Implicit-explicit time integrator schemes
Linear multi-grid solvers
I/O optimisation & memory management
Communications (particle-particle, particle-mesh)
Portable programming models & abstraction layers
Task-based and dynamic execution models
Dynamic load-balancing

Other potential cross-cutting Motifs will be further collectively defined (e.g., birds-of-a-feather sessions, including related national PEPR):

- *Particle-based methods @ exascale*; *On-line data analysis @ exascale*; multi-modal *image analysis @ exascale, AI-coupled HPC/HPDA workflows @ exascale*

The Exa-DI iterative agile process

Develop efficient agile scientific software co-development methodologies

- ***Co-identify cross-cutting computation and communication Motifs*** with impact across the ADs in partnership with application teams and NumPeX projects
- ***Co-develop well-curated Motifs-based proxy apps*** for software components performance and portability evaluation
- ***Stream-aligned co-development of Motif-based software components (libraries, frameworks, tools)*** across NumPEx
- ***Integrate and deliver logical collections of Motif-based software components (SDKs)*** to accelerate Exascale application codes development and respond to trend of continuous refactoring with improved software design
- ***Foster the use of software transactional software management technologies (Spack, Guix-hpc) and CI methodologies*** in synergy with the computing facilities and the adoption of NumPeX Software Community Guidelines
- ***Improve scientific software methodologies and productivity*** by developing training materials and beacon of good practices directed to the CSE community at large

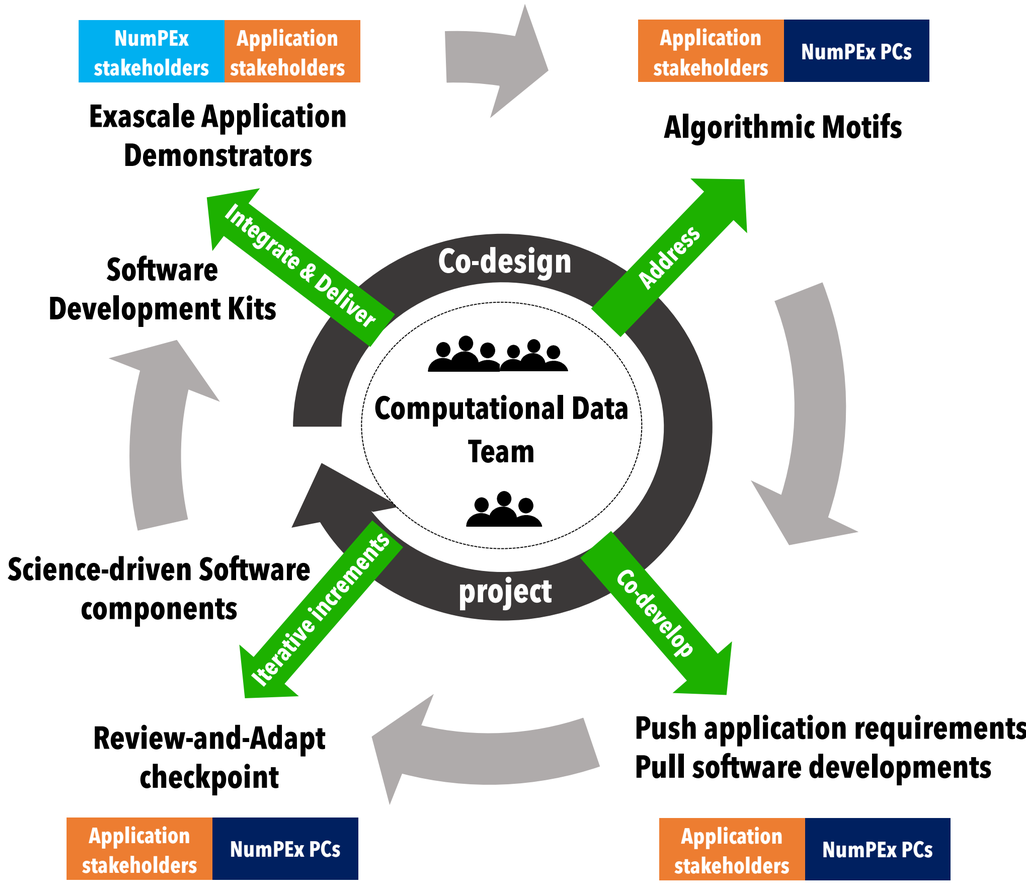

➡ Interdisciplinary CDT at the interface between the application and the R&D teams in NumPeX

➡ Important aggregation layer: steering coordination between initially loosely coupled software R&D in NumPeX, while fostering longer term research.

Exa-DI multi-layer organisational elements

- **The Application Demonstrator teams:** focused on targeted development addressing a significant domain-specific exascale challenge problem. Application demonstrator codes are typically a general software capability of the domain supported by specific software Application Development teams. that are not directly funded by Exa-DI and are responsible of defining, programming and executing their activities and and managing their internal processes. The Exa-DI software co-design and co-development process provides a framework for building and coordinating partnership with the ADs, co-identifying cross-cutting computation and communication Motifs needs and planning to address the co-development and the integration across NumPEx of logical collections of reusable Motif-based software components.

- **The NumPEx Research and Software Product teams:** working in the different projects (PC1-PC4) and that are the fundamental organisational building block of the NumPEx software development efforts. The teams are responsible for defining and executing their activities and managing their own internal processes. However, within NumPEx, their work is coordinated at a higher-level by the Exa-DI software co-design and co-development process supported by the CDT, and driven by the needs of the Application Demonstrator teams. This provides the framework for coordinating, tracking, and assessing collective development progress of logical collections of compatible and complementary Motif-based software components (SDKs) that can be re-used and assessed in well-specified and well-curated proxy apps.

- **The Exa-DI SDKs and proxy apps co-design and co-development teams:** addressing Motif-based Software Development Kits (SDKs) together with well-specified proxy apps. Each team involves CDT staff that co-develop, in partnership with the Application Development teams, logical collections of compatible and complementary Motif-based software components that are integrated and assessed in proxy apps. Teams depend on contributing research and software product teams in the NumPEx projects. Motif-based SDKs and proxy apps facilitate interaction and coordination among the software products teams, application demonstrators and CDT in a variety of activities such as requirements, design space exploration, training, and evolution of scientific software development practices and tools. Coordination of versioning, design space exploration, and software delivery at the SDK level helps amortise costs across related products and reduces complexity at the top software stack level. This level also spurred a cooperative dynamics where teams collaborate on some activities and learn from each other.

- **The Exa-DI software integration and delivery enabling team:** provides co-design and co-development teams with access to skills and capabilities to build, test, deliver, and deploy the co-developed Motif-based SDKs together with proxy apps, while managing versioning across software products for reproducibility, correctness and security patching. The team hardens and supports the use of meta-build technologies such as Spack and Guix-hpc, that can be configured to create custom builds targeting different environments to provide the Motif-based SDKs and proxy apps in containerised environments and with binary caches of previously built products to enable rapid rebuilds.

Motif-based proxy apps

Co-develop well-documented community-driven Motif-based proxy apps with a programable high-level abstraction interface for algorithmic specifications and problem parametrisation that:

- **Capture cross-cutting complex operations or entire algorithms in workload phases** of interest *among* Applications
- **Address increasingly complex AI-coupled HPC/HPDA workloads**
- **Handle different operation sequences, data structures, layouts and logistics** for the same problem;
- **Focus on consistency, verification, validation, explainability, reproducibility and uncertainty quantification** with a solid determination of generalisation errors
- **Enable integration of logical collections of Motifs-based software components** (libraries, frameworks, workflow tools) with improved cross-layer optimisation
- **Harden and extend the use of abstraction layers, programming and parallel execution models** to improve performance portability
- **Leverage transactional software management and deployment technologies**, together with **profilers and analysing tools, simulators, continuous integration**
- **Improve scientific software methodologies**

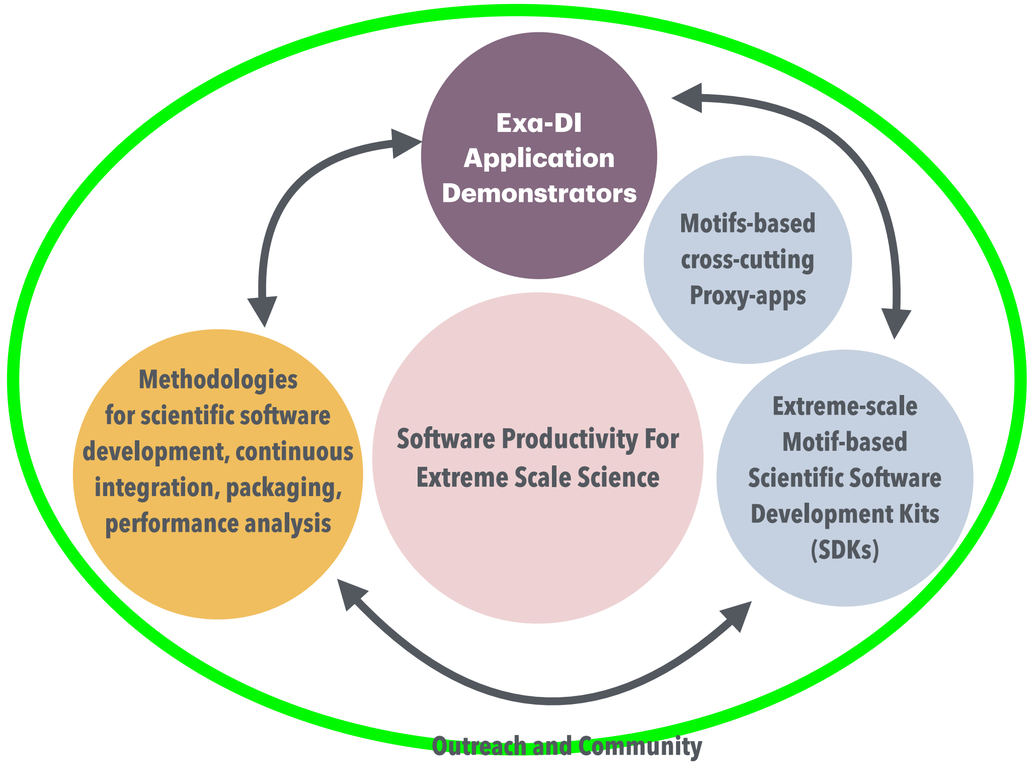

Partnership with application teams: delivery of crosscutting methodologies and metrics with impact on real applications.

Multi-expertis Exa-DI CDT

Contributing Research & Software production teams of the NumPEx projects (PC1-PC4)

Motif-based Software Development Kits (SDKs)

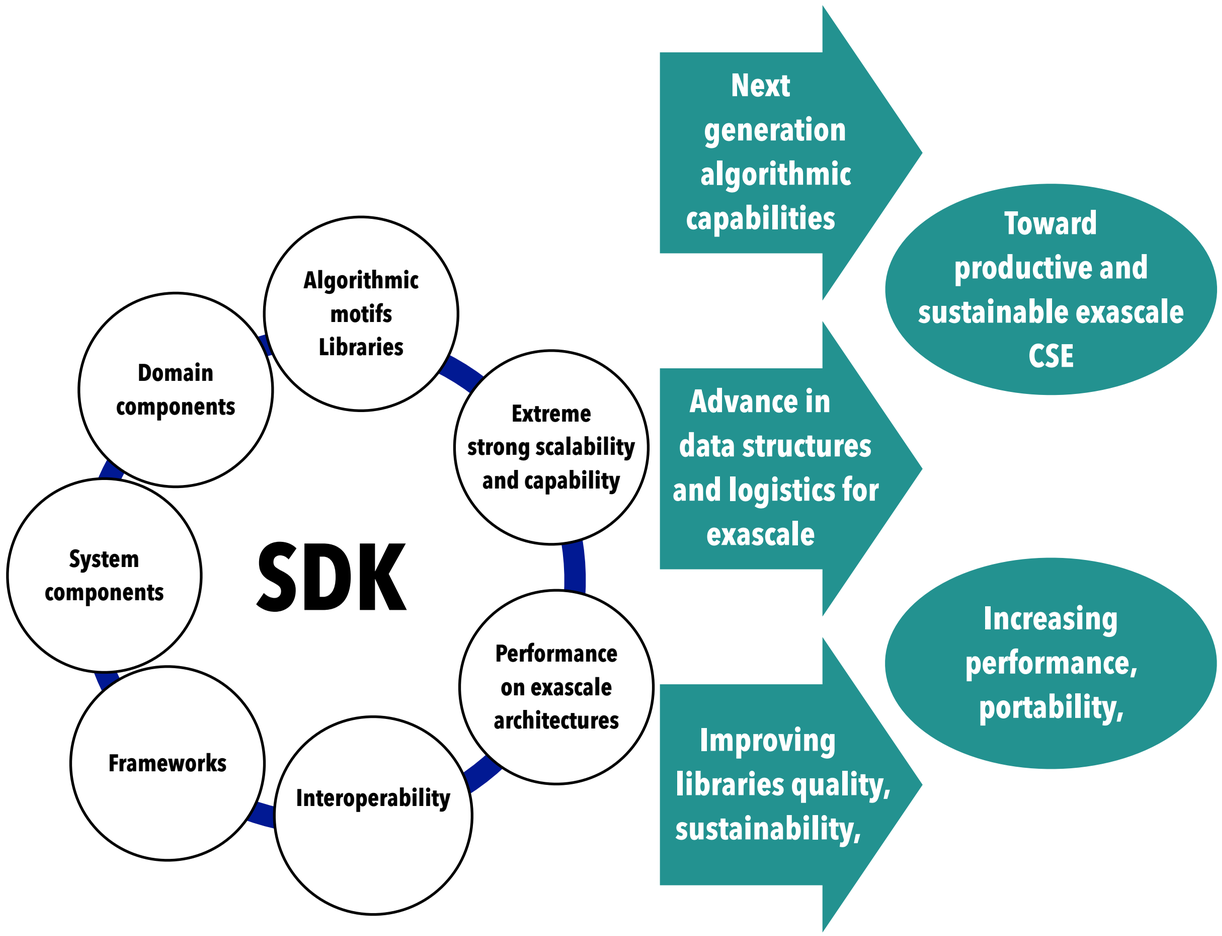

Timeline: SDK — SDK — . . . . . — SDK

- SDK provide the foundation for a modern extreme-scale scientific software ecosystem, where application development is accomplished by composition of high-quality, reusable and composable motif-based software components rather than by tangential use of libraries.
- Proxy apps teams produce a small domain components code that expresses the particular purpose of the Motif-based software and gain the bulk of functionality by parameterised use of SDK components.
- SDK frameworks for documentation, testing, and code quality, as well as established software policies and best practices, can be adapted and adopted as appropriate by the application developers to provide compatible, high-quality, and sustainable software.
- As we move toward this new ecosystem, application development times from first concept to scalable production code should be greatly improved.
- Success hinges on the quality, interoperability, usability, and diversity of SDKs and ability to deliver SDKs with meta build technologies (e.g. Spack, Guix-HPC) to scientific application developers.

Exa-DI contacts

Exa-DI gitlab: https://gitlab.inria.fr/numpex-pc5

Exa-DI CDT Program Manager: Jérôme Charousset (jerome.CHAROUSSET@cea.fr)

Exa-DI CDT enabling team: Bruno Raffin (bruno.Raffin@inria.fr) & Benoit Martin (bmartin@cea.fr)

Exa-DI PIs: Jean-Pierre Vilotte (jean-pierre.vilotte@cnrs-dir.fr) & Valérie Brenner (valerie.brenner@cea.fr)