RÉPUBLIQUE FRANÇAISE
Liberté
Égalité
Fraternité

FRANCE 2030
PROGRAMME DE RECHERCHE
NUMÉRIQUE POUR L'EXASCALE

Exa-MA: Methods & Algorithms for Exascale

Scientific Advisory Board Meeting – Year 2 Review

Christophe Prud'homme (UNISTRA) & Hélène Barucq (Inria)
on behalf of the Exa-MA PC1 Team

January 20, 2026

Exa-MA Scientific Advisory Board

Meeting Agenda

Meeting Schedule: 17:30–18:30

17:30	–	Welcome & Meeting Format
17:35	–	Year 2 Progress & Key Results
17:45	–	**Scientific Highlights** (WP1–WP6)
18:10	–	Software Stack & Strategic Challenges
18:20	–	SAB Feedback & Next Steps

Key Topics for Discussion

- Year 2 achievements & measurable outcomes
- Scientific highlights (WP1–WP6) with live Q&A
- Strategic challenges & SAB feedback

Format

Discussion-oriented

Interactive format

Deep-dive on demand

Materials Available

Slides: `slides-sab.pdf`

D7.1 v2.0.0 Executive Summary

Publication highlights

Software dashboard screenshots

Presentation serves as support for interactive discussion

Exa-MA Vision: Exascale-Ready Methods & Algorithms

Grand Challenges

- Reduce carbon footprint (GHG)
- Design advanced materials
- Climate prediction & assessment
- Personalized health & drug design
- Fusion energy & ITER

Objectives

- **Methods & algorithms** for exascale
- **Software libraries** for exascale stack
- **AI algorithms** at scale
- **Demonstrators**: mini/proxy-apps

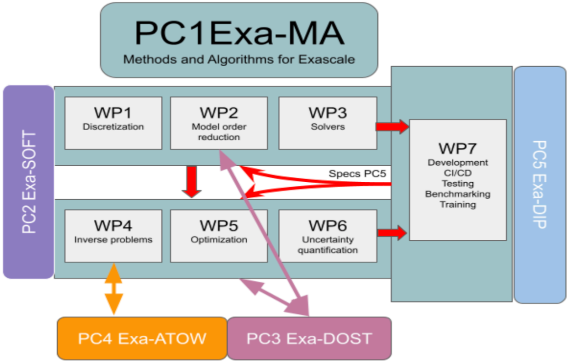

Year 2 Project Metrics (December 2025)

38
publications
97% OA

Publications
34% with DOI

25
recruited
personnel

Team Growth
12 PhD, 5 postdoc, 8 eng.

16
software
packages

Active Software
18 frameworks impacted

55
person
years

Funded Effort
55 funded PY
committed

D7.1 v2.0.0 Year 2 research report (WP1–WP7) | **D7.3 v1.0.1** training materials

Headline Results 2024–2025

50×
speedup
GenEO

Cobra cavity
45M DoFs, 2916 cores

8000
GPUs
LUMI

Bayesian Opt.
Fractal decomposition

80–200%
CPU savings
FNO-Newton

Neural operators
Accelerated
Newton-Krylov

1024
MPI
scaling

Thermoelastic
Near-perfect scalability

Ktirio Urban Meshing: GIS → watertight city-scale meshes (Ktirio/CGAL 6.0)

RÉPUBLIQUE FRANÇAISE
Liberté
Égalité
Fraternité

FRANCE 2030

PROGRAMME DE RECHERCHE
NUMÉRIQUE POUR L'EXASCALE

cea cnrs inria

Scientific Highlights

Feature-Preserving Alpha Wrapping

- Robust watertight mesh from **defect-laden** CAD
- Handles triangle/polygon soups and point clouds
- **Visibility-based** salient Steiner point insertion
- Integrated in **CGAL 6.0** (open source)

Ktirio-geom: City-Scale Reconstruction

- **CGAL-based** urban geometry pipeline
- Reconstruct & remesh **entire cities** (e.g., Paris 150K buildings)
- Deployed on **EuroHPC** (Discoverer, Karolina)

High-Order Space-Time Coupling

- Dynamic time step adaptation (CWIPI)
- High-order explicit & implicit coupling
- Demonstrated on conjugate heat transfer

Benchmarked Applications

- **Thermal bridges** – building energy (Feel++)
- **FDA nozzle** – biomedical flow (Feel++)
- **ParMMG** – anisotropic adaptivity (FreeFEM++)

Frameworks

CGAL, Feel++, FreeFEM++, MMG/ParMmg, CWIPI

WP1: Exascale Thermoelastic Coupling – Strong Scalability

Partitioned Coupling Strategy

- Thermomechanical formulation on **initial configuration**
- Piola-Kirchhoff effective conductivity
- **Block Gauss-Seidel** fixed-point coupling
- Ported to **MFEM** (HPC library)

Scalability Results

- **Near-perfect scalability** to 1024 MPI processes
- **Hundreds of millions** of degrees of freedom
- Tested on CEA Topaze (AMD EPYC Milan)
- Preserves mesh quality at scale

Key Innovation

Large-scale HPC thermomechanical simulations while **preserving mesh quality** by working on initial (undeformed) geometry

Observation: Mechanical problem-solving phase dominates total computation time

Impact: Easy implementation of Gauss-Seidel fixed-point strategy enables flexibility & efficient solver reuse

P. Dubois, PhD thesis (2024–2027); P. Dubois *et al.*, proc. CSMA 2026

WP2: Scientific Machine Learning

Domain-Decomposed PINNs

- FBPINNs/MFBPINNs with PML for Helmholtz (Scimba)
- Energy Natural Gradient optimization; semi-Lagrangian time stepping

Neural Operators for Acceleration

- FNO initial guesses for Newton-Krylov solvers
- **80–200% CPU savings** on benchmark PDEs

Nonlinear Compressive Reduced Basis

- Mitigate Kolmogorov n-width barrier for transport problems
- Linear encoder + nonlinear decoder for parametric PDEs

Applications

- **Helmholtz waves** – seismic & acoustic
- **Tokamak MHD** – fusion energy
- **Vlasov** & advection-diffusion benchmarks

Key Innovation

Nonlinear compressive RB breaks the **Kolmogorov n-width** barrier while enabling fast online evaluation

WP3: Scalable Solvers – 50× Speedup with GenEO

Domain Decomposition Beyond SPD

- GenEO extended to saddle-point systems (HPDDM/PETSc)
- RAS with inexact local solves for nonsymmetric problems
- GPU offloading and mixed precision for GenEO

Mixed Precision & Autotuning

- Mixed-precision GMRES with adaptive precision
- PROMISE: auto-tuning arbitrary precision formats
- Fault detection for Krylov solvers

Cobra Cavity Benchmark

Cores	1-level	GenEO
32	51	**8**
256	359	**14**
2048	3711	**60**
2916	3398	**74**

Iteration counts – 45M DoFs
Near-constant iteration count at scale!

WP4: Inverse Problems & Data Assimilation

4D-VAR Variational Assimilation

- Sequential-in-time optimization of initial conditions
- Applied to shallow-water and large-scale models

Limited Memory Preconditioner (LMP)

- Truncated EVD reduces condition number
- Multifidelity stochastic LMP via randomized low-rank
- Enables parallel matrix–vector products

Observation Strategies

- Bayesian optimal sensor placement with gradient surrogates

Highlights

- 4D-VAR with efficient preconditioning
- Sequential-to-parallel strategies for large-scale DA
- Multifidelity schedules for exascale inversion

Key Innovation

LMP + randomized low-rank approximations for scalable variational DA

WP5: Optimization at Exascale

Fractal Decomposition Algorithms

- Exascale-ready decomposition with MPI + Kokkos (Zellij)
- Strong scaling to **8000 GPUs** on LUMI

Bayesian & Surrogate Optimization

- Fractal Bayesian optimization with multi-surrogate scoring

Shape & AutoML Optimization

- PDE-constrained shape optimization (GeSONN)
- LLM pruning and fine-tuning via multi-objective BO

Scale

- 84% strong-scaling efficiency at 1000 LUMI-G nodes
- Superlinear speedup with distributed surrogates

Software

Zellij, PyParadisEO, GeSONN

WP6: Uncertainty Quantification

Physics-Informed Gaussian Processes

- Constrained field prediction with GP kernels
- High-dimensional CFD fields with linear constraints

Graph Kernels for Mesh Inputs

- Sliced Wasserstein Weisfeiler-Lehman kernel
- Handles large graphs (10^5 nodes) for UQ

Uranie at Scale

- Parallel TRUST ensembles via ICoCo
- URANIE v4.9–v4.10 releases

Applications

- Turbine blade UQ (SAFRAN)
- TRUST multiphysics ensembles
- Urban-scale UQ for decision support

Platforms

Uranie, OpenTURNS, TRUST

RÉPUBLIQUE FRANÇAISE
Liberté
Égalité
Fraternité

FRANCE 2030

PROGRAMME DE RECHERCHE
NUMÉRIQUE POUR L'EXASCALE

cea cnrs Inria

Software & Infrastructure

Software & Application Framework

Application Taxonomy (D7.1)

Mini-app	Single kernel isolation
Extended	+ I/O scripts & data
Demonstrator	Multi-WP workflow
Proxy-app	Full-stack (≥3 WPs)

Framework Ecosystem

- Arcane, CGAL, Feel++, FreeFEM++, TRUST Platform
- Composyx, Hawen, MaHyCo, Samurai, Scimba, Uranie
- +13 additional frameworks across the consortium

18 Frameworks: C++ (12), Python (5), Fortran (3) | MPI (13), GPU (8)

Packaging & Benchmarking

- **Packaging:** Spack/Guix, Docker, Apptainer
- **CI/CD:** feelpp.benchmarking + ReFrame + GitHub Actions
- **FAIR:** DOIs, Zenodo archival
- **Targets:** LUMI, Discoverer, Karolina, Vega, ...

Solver Packages

MUMPS, Composyx, PETSc, HPDDM, Fabulous

RÉPUBLIQUE FRANÇAISE
Liberté
Égalité
Fraternité

FRANCE 2030

PROGRAMME DE RECHERCHE
NUMÉRIQUE POUR L'EXASCALE

cea cnrs inria

Perspectives

Next Steps & Perspectives

Scientific Priorities (Year 3)

- **GPU portability** (eg. Kokkos)
- **Mixed-precision** solvers at scale
- GenEO for **multi-physics**
- Neural operators for **real-time**
- Expand **Data Assimilation** (WP4)

Application Domains

- Fusion energy (ITER)
- City energy & climate
- Subsurface modeling
- Space & aeronautics

Ecosystem Integration

- PC1–PC5 **common milestones**
- **Upstream:** PETSc, MUMPS, HPDDM, CGAL
- **Apps**: Specification → Benchmarking
- EuroHPC **CoE collaboration**
- 2 NumPEx projects 2026 (SAGE-HPC, DAIMOS)

Deliverables 2026

- D7.1 v3.0 (Year 3 report)
- D7.3 training expansion
- Proxy-app demonstrators

RÉPUBLIQUE FRANÇAISE
Liberté Égalité Fraternité

FRANCE 2030

PROGRAMME DE RECHERCHE
NUMÉRIQUE POUR L'EXASCALE

Challenges & Discussion

Open Challenges & Strategic Questions for Discussion

Strategic Challenges

1. **Energy efficiency** at exascale (20 MW target)
2. **Reproducibility** across heterogeneous architectures
3. **Continuation** strategy post-PEPR
4. **Resilience** against soft & hard errors

Bottlenecks Addressed

- (B7) Exascale algorithm redesign
- (B8) Ensemble-based discovery algorithms
- (B9) Resilience, robustness, accuracy
- (B10) Scientific productivity tools
- (B11) Reproducibility & replicability

Questions for SAB Discussion

- How should Exa-MA position for **post-PEPR** continuation?
- Which **emerging domains** would benefit most?
- Recommendations for **international collaborations**?
- How to better integrate with **EuroHPC CoEs** and other calls?
- **Quantum-HPC** hybrid perspectives?

Risk Assessment

- Personnel retention in competitive market
- GPU architecture fragmentation
- Sustainability of open-source ecosystem

Conclusion

Summary – Key Takeaways

Why Exa-MA Matters

1. **Sovereign, open, reproducible** software
2. **Breakthrough results** with measurable KPIs
3. **Ecosystem integration** – FR/EU/US

Key Achievements

- 50$\times$ solver speedup (GenEO)
- 8000 GPU scaling (LUMI)
- 80–200% CPU savings (neural operators)
- Watertight meshing (Ktirio, CGAL 6.0)
- Automated CI/CD to EuroHPC

Resources

D7.1 v2.0.0:
doi:10.5281/zenodo.15188286

Website:
numpex.github.io/exa-ma

Source:
github.com/numpex/exa-ma-d7.1

Contact:
`helene.barucq@inria.fr`
`christophe.prudhomme@cemosis.fr`

Thank You!

Questions & Discussion