

SPPEXA and beyond: Perspectives of Exascale Computing in the Context of AI

H.-J. Bungartz, W. E. Nagel, S. Reiz



computational algorithms
system software
application software
data management and exploration
programming
software tools

3rd France-Japan-Germany trilateral workshop:
Convergence of HPC & Data Science for Future Extreme-scale Intelligent Applications
Tokyo, November 6–8, 2019


SPPEXA – A strategic Priority Program by DFG (Germany) in co-operation with ANR (France) and JST (Japan)

Contents

*SPPEXA - The story
behind*



HPC – Enabler for Computational S&E: Simulation, Optimization, Analytics



Mathematical model


$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} + \frac{1}{\rho} \nabla p - \nu \Delta \mathbf{u} = 0$$

$$\nabla \cdot \mathbf{u} = 0$$

Discretization & solver

$$M \mathbf{u}_h = 0$$

$$A \dot{\mathbf{u}}_h + D \mathbf{u}_h + C(\mathbf{u}_h) \mathbf{u}_h - M^T p_h / \rho = 0$$



- #1: no cycle/pipeline any more, interplay of all phases
- #2: less space for single-field experts
- #3: outer loops (multi-physics, control, UQ, multi-fidelity...)
- #4: in-situ data integration, rather than pre-/post-processing



Parallel code, HPC



Software



Validation



Exploration

Insight,
Design

A Brief History of “Computational”

1st generation: qualitative (forward) simulation

- *HPC: one large run*
- *DA: only through post-processing*

2nd generation: more complexity by optimization, inverse problems

- *HPC: arrival of outer loops*
- *DA: parameter identification & reduction*

3rd generation: getting quantitative through data integration

- *HPC: data-driven / data-enriched simulations*
- *HPC: more outer loops (inverse problems, statistics, UQ, multi-fidelity, ...)*
- *DA: in-situ integration of simulation & experimental data, analytics, ML (i.e. AI)*

Ultimate goal: “predictive” science

For that: overcome the gap between theory & experiment – HPC helps, and AI helps



Where HPC Meets Big Data / Learning / AI

Data analytics uses computational algorithms, needs HPC

- Statistics, stochastics, numerics – esp. high-dimensional
- Examples: eigenvalues, SVD, low-rank approximation, kernel-based methods
- Increasing share of analytics applications on HPC systems (cf. Gordon Bell)
- Re-appearing requests for “more maths, please!” in CS/DS education 😊

Modern HPC applications are data-related or data-driven

- Parameter identification: search in high-dim parameter spaces
- Model-order reduction: reduce computational demand by smart model design
- Multi-fidelity: combine models/simulations/data at different scales/resolution
- Statistics, uncertainty quantification: increase reliability of simulation results

Machine Learning for an HPC context

- Learn from experiments & simulations for more predictive & efficient runs
- Learn from application behaviour for a better dynamic job scheduling
- Same for energy management etc.

Performance & efficiency are crucial for both



Challenges: Multi-X Increases Complexity

From data / images / numbers ...
... to information / insight

From qualitative descriptions ...
... to quantitative prediction

From counting operations ...
... to energy awareness

From simulation ...
... to optimisation

From sequential algorithm design ...
... to massive parallelism

From parameter assumptions ...
... to identification
& estimation

From one-way batch jobs ...
... to user interaction

From forward problems ...
... to inverse problems

From simple tools & codes ...
... to 2x complex ones

From deterministic models ...
... to random & uncertainty

From heroic PhD codes ...
... to large teams / SW

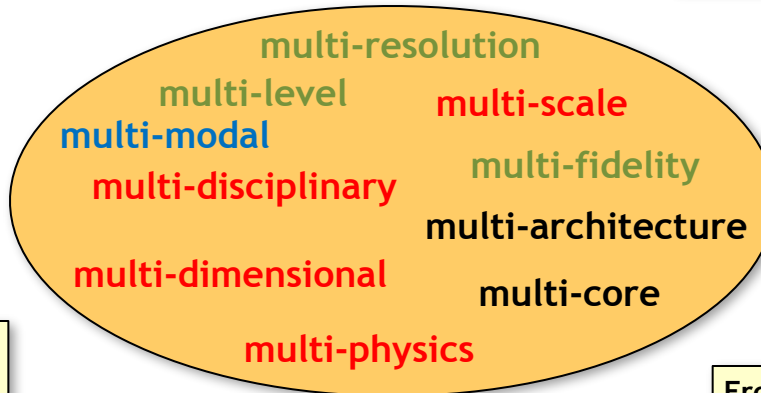
From one (spatial/temporal) scale ...
... to cascades of scales

From hacker's delight ...
... to complex workflows

From single-physics problems ...
... to coupled scenarios

From island fun ...
... to embedding & integration

From flat algorithms & data ...
... to hierarchy



... also for Analytics-Driven Computations

From data / images / numbers ...
... to information / insight

From qualitative descriptions ...
... to quantitative prediction

From counting operations ...
... to energy awareness

From simulation ...
... to optimisation

From sequential algorithm design ...
... to massive parallelism

From parameter assumptions ...
... to identification
& estimation

From one-way batch jobs ...
... to user interaction

From forward problems ...
... to inverse problems

From simple tools & codes ...
... to 2x complex ones

From deterministic models ...
... to random & uncertainty

From heroic PhD codes ...
... to large teams / SW

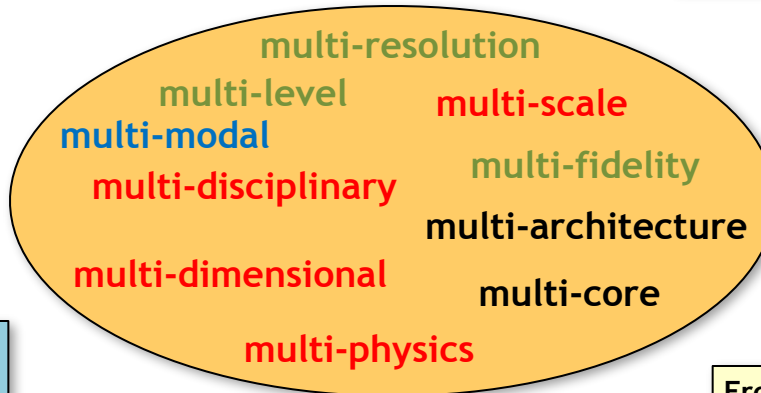
From one (spatial/temporal) scale ...
... to cascades of scales

From hacker's delight ...
... to complex workflows

From single-physics problems ...
... to coupled scenarios

From island fun ...
... to embedding & integration

From flat algorithms & data ...
... to hierarchy



For Performance, Algorithms are Crucial

Broad range of computational algorithm functionalities:

- Discretization schemes, numerical solvers
- Parallelization, communication, load balancing
- Compression, dimension reduction
- Combinatorial, graph-based
- Statistic, stochastic
- Analytical (machine learning, ...)

Algorithm life cycle: design – analysis – implementation – tuning

- **Algorithm Engineering:** design – analysis – implementation – tuning
- **Performance Engineering:** design – analysis – implementation – tuning
- **Efficiency comprises both: implement the “best” algorithms in the best way!**

Notions of efficiency – what is “best”?

- It's been a long way from the classical Landau *order-of-N-to-the-something* to today's multi-faceted performance landscape ...



Notions of Efficiency

- **counting operations, $O(N^k)$ – the classical complexity-related one**
- **#ops – #flops – % of peak performance – the speed-oriented one**
 - both at node level and at system level – weak / strong scalability
- **#bytes (or #bytes/flop) – the communication-related one**
 - communication-avoiding / memory-efficient / compute-bound as good guidelines
- **#Watt (or #Watt/flops or #flops/Watt) – the energy-related one**
 - “cool” solvers – are they really different from classical “fast” ones? May slower be cooler?
- **time-to-solution – the pragmatic one**
 - reasonable at first glance – how long do I have to wait?
 - Q: what to include into consideration, i.e. “time”? Is runtime-only sufficient?
- **digits per flop – the accuracy-related one**
 - brings the model back to stage: a better model → less computation → more digits / flop
- **science (results? pubs?) per flop – the insight-related one**



Contents

SPPEXA – Facts



SPPEXA – It Has Been a Long Way...

2006 – First discussions within DFG’s Commission on IT Infrastructure

- HPC SW runs into problems – lack of funding mechanisms; cf. international situation

2007/2008 – Memorandum initiated by the geosciences

- Title *Scientific Software in the PetaFlop Era*, Roundtable discussion in Tutzing, April ‘08

2010 – Suggestion by German participants in the exascale initiatives

- Against the background of (1) massive investments in high-end systems world-wide and (2) massive investments in HPC software in the USA (DoE-SciDAC-1/2, NSF-OCI), e.g.

2010 – KfR takes responsibility

- Strategic paper and a discussion with DFG’s president, M. Kleiner (November 2010)
- Outcome: suggestion of a flexible, strategically initiated PP, financed via Strategy Funds

2011 – Increase of speed & DFG decision

- Roundtable expert meeting to shape the baby
- DFG-internal discussions about process
- Submission in August; international reviewing in September; “Go!” in October
- Call in November



SPPEXA – It Has Been a Long Way...

2012 – Review of proposals

- 68 sketches, 24 consortia pre-selected and invited to submit full proposals
- Review workshop in Bonn to select the successful consortia

2013 – Launch of SPPEXA: 13 projects

2014 – Call for proposals SPPEXA-2, incl. F/J partner institutions

2016 – SPPEXA-2: 17 projects, > 50 institutions

- 12 continuations of phase-1 projects
- 4 additional new projects
- 7 of 16 with a bi-national (J&D) or tri-national (J&F&D) appearance
- later: a 17th project joins SPPEXA

2019 – October: International symposium in Dresden

2020 – April: end of SPPEXA



SPPEXA Characteristics

- **A *Priority Program* in DFG's set of funding formats**
- **Strategic initiative of DFG to fund HPC SW in Germany**
 - Fundamental research
 - Establish collaborative, interdisciplinary co-design of HPC applications and HPC methods through several research consortia
- **SPPEXA research is ...**
 - ... driven by domain sciences / CSE applications
 - ... powered by scientific computing & informatics / CSE methodology
 - ... in parts smooth/evolutionary, in parts radical/revolutionary
- **Untypical for a Priority Program: intense coordination**
 - Central events, establish and foster international collaboration, focus on education: doctoral retreats & coding weeks
 - Support project-specific as well as cross-project activities
 - Gender-related measures



SPPEXA's 6 Research Directions

- **Computational algorithms**
 - Aiming at large-scale machines
 - Efficient w.r.t. “modern” complexity measures
- **System software and runtime libraries**
 - Process scheduling
 - System health monitoring
 - Resilience handling
- **Software tools**
 - Compiling, running, verifying, testing, optimizing
- **Application software**
 - Key driver for exascale
 - Hardware-software co-design necessary
- **Programming**
 - Make traditional approaches exascale-ready
 - New programming models
- **Data Management**
 - Process large data sets
 - Archive, make data available



SPPEXA Facts

- **13/17 research consortia funded in Phase 1/2**
 - Interdisciplinary, international research consortia
 - Involving 2-5 groups each
 - Addressing at least 2 out of the 6 SPPEXA topics
 - About 60 PIs and 60 PhD candidates/postdocs per first/second phase
 - Overall budget of 3.8m € per year
- **Two three-year funding phases**
- **Launch of first phase in January 2013**
- **Launch of second phase in January 2016**
 - Strong internationalization component: joint call with France and Japan
 - More than 7 international project consortia (7 w/ F/J)

EXA-DUNE

ExaFSA

Terra-Neo

EXASTEEL

GROMEX

ExaStencils

Smart-DASH

EXAHD

EXAMARK

ExtraPeak

EXAMAG

FFMK

ESSEX

ExaSolvers

ADA-FS

AIMES

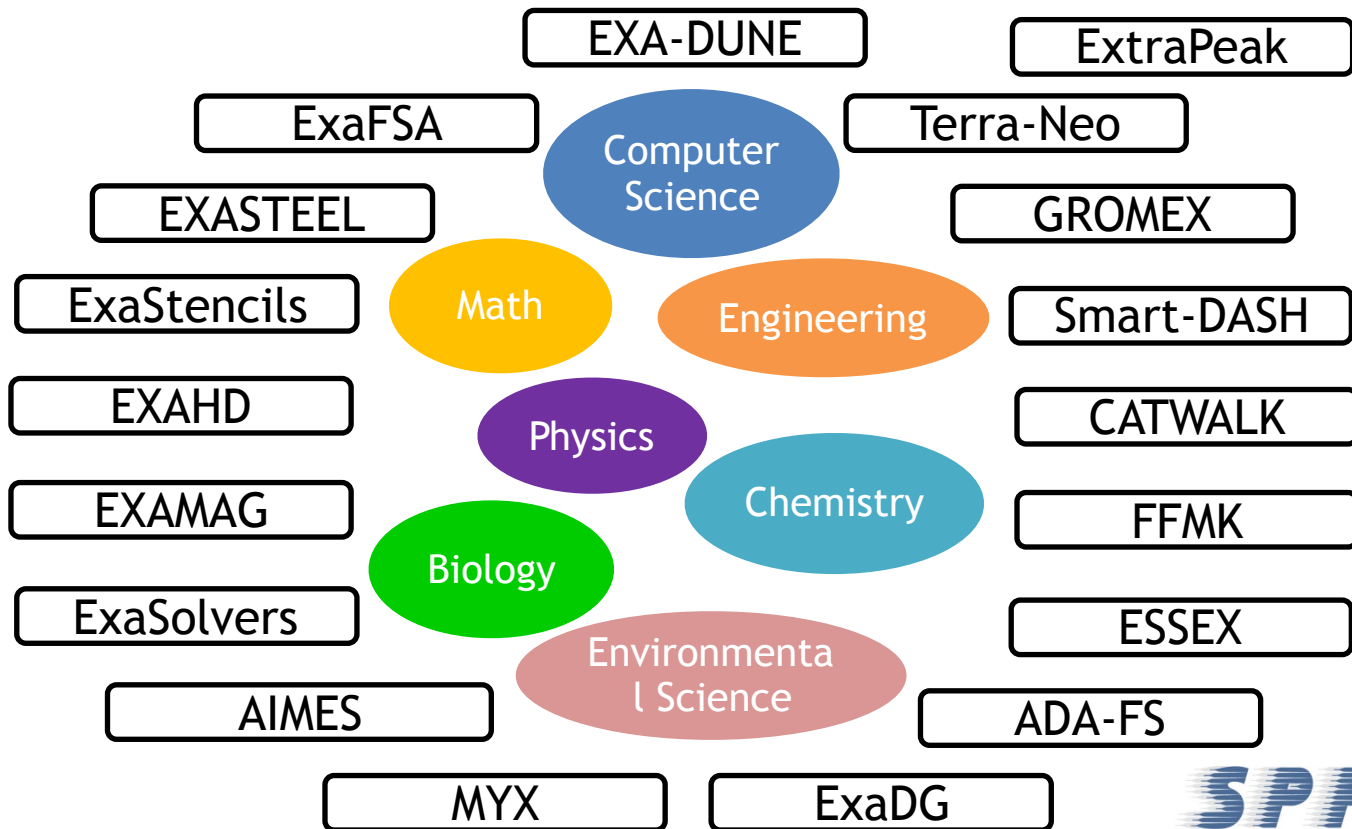
ExaDG

MYX



A Really Interdisciplinary Endeavor

- **Highly interdisciplinary projects and project consortia**
 - Requires close collaboration within and among SPPEXA consortia
 - The central coordination fosters synergistic effects within SPPEXA



A Really International Endeavor: Partner Institutions from...



- **Japan:**

- RIKEN
- Tokyo Tech
- University of Tsukuba
- University of Tokyo
- Tohoku University
- Tokyo University of Science
- Toyo University

- **France:**

- Université de Versailles
- Université de Strasbourg
- Maison de la Simulation, Saclay

- **other Countries:**

- TU Delft, Netherlands
- USI Lugano, Switzerland
- Royal Institute of Technology, Sweden
- UCLA, USA
- ANU, Australia
- Hebrew University Jerusalem, Israel

Projects with F/F&J connection:

- **ExaFSA:** Exascale Simulation of Fluid-Structure-Acoustics Interactions
- **ExaStencils:** Advanced Stencil-Code Engineering
- **EXAMAG:** Exascale Simulations of the Magnetic Universe
- **ESSEX:** Equipping Sparse Solvers for Exascale
- **EXASOLVERS:** Extreme-Scale Solvers for Coupled Problems
- **AIMES:** Advanced Computation and I/O Methods for Earth-System Simulations
- **MYX:** MUST Correctness Checking for YML and XMP Programs



Steering Committee & Project Management

- **Steering Committee: 2 coordinators and 4 PIs**

- **Coordinators:**

- Hans-Joachim Bungartz (Munich)
 - Wolfgang Nagel (Dresden)



- **6 Project PIs, (re-)elected for 3 years:**

- Sabine Roller (Siegen)
 - Christian Lengauer (Passau)
 - Hans-Peter Bunge (Munich)
 - Dörte Sternel (Darmstadt/Berlin)
 - Takayuki Aoki (Tokyo)
 - Nahid Emad (Saclay/Versailles)




- **Scientific Program Manager**

- Severin Reiz (Munich)
 - Project management on day-to-day basis



Scientific Advisory Board

- **George Biros**
Oden Institute, UT Austin
- **Rupak Biswas**
Head, NASA Advanced Supercomputing (NAS) Division
- **Klaus Becker**
Airbus
- **Rob Schreiber**
Assistant Director, Exascale Computing Lab @HP Labs
- **Craig Stewart**
Executive Director, Pervasive Technology Institute, Indiana U



Central Coordination: Portfolio

- **Strategic design of joint/central SPPEXA activities**
 - Visibility at conferences (SC, ISC, EuroPar and many more)
 - Representation at national and international level
- **Funding of project-specific and cross-project workshops**
 - So far > 45 events supported
- **Organization of annual assemblies**
 - Foster scientific exchange between projects
 - Report on scientific and strategic achievements
- **Educational activities**
 - Doctoral retreats and coding weeks
 - Foster exchange between young researchers – international research stays
 - Prizes for outstanding master's or PhD theses
- **Public relations**
 - Sharing/Publishing SPPEXA news
 - Presenting the projects to the wider community
- **Gender activities and childcare funding**



Some Numbers – Phase 2 so far

- **25 workshops supported**
 - project-specific or cross-project
 - stand-alone or at conferences (SC, ISC, Computational-X, SIAM, DATE, ...)
- **3 J-F-D workshops in Tokyo**
 - 2017/2018/2019
- **~ 20 extended research stays of doctoral candidates**
- **3 annual doctoral retreats (one in France)**
- **6 gender-related trainings or workshops (“women in ...”)**
- **Theses directly related to SPPEXA:**
 - > 50 PhD theses
 - > 75 bachelor’s theses
 - > 50 master’s theses

(source: infos from projects – and we all know how diligent people are in responding to surveys ... 😊)



Contents

And Now?




HPC & AI

- **HPC4AI: HPC as enabler (esp. for computationally more demanding “scientific ML” cases)**
 - Project *Numerically based and HPC-empowered approaches for next-generation ML* (with George Biros: kernel-based methods, or class./regr./clust. in discretized high-dim feature spaces)
 - More examples to come in this workshop
- **AI4HPC: less obvious**
 - Dynamic, malleable resource allocation or load balancing for outer-loop scenarios (OPT, PAR IDENT, UQ, multi-physics & multi-fidelity, Sparse Grids, ...)
 - Example: ExaHD project – Sparse Grid combination technique for 5-d plasma simulations
 - ML for the load balancing of the coarse grid solutions
 - ML for circumventing detecting & repairing (on-the-fly!) silent data corruption



Indirect access via combination technique (a “multi-grid” view):

- extrapolation-type approach - variants: classical, dimension-adaptive, opti-com, ...
- superposition (combination) of several, but much smaller, full grids
- analogous combination of solutions: use standard codes on standard (full) grids









- ... minimizes communication and synchronization, allows for efficient load balancing
- ... provides an additional, asynchronous level of parallelism: more accurate/efficient than Monte Carlo, more decoupled than Domain Decomposition

Advantages:

- #1: scalability (communication-light level)
- #2: resilience (error compensation without C/R *PLUS* error and outlier detection)

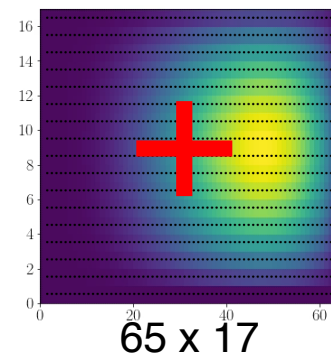




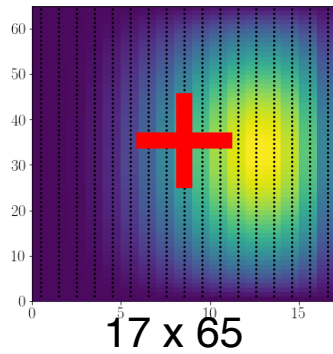
Combination Technique



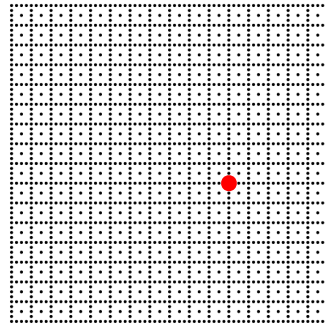

$O(d \log(N)^{d-1})$
grids of size $O(N)$

5d plasma physics:
 10^8 vs. 10^{15} grid pts.

ABFT: Tackle Silent Data Corruption








Algorithmic Solution: Detect Outlier Fields



1. Identify grid point with largest difference
2. Perform robust regression to fit values

→ Find outliers

→ No assumptions about SDC



HPC Beyond SPPEXA – A German View (1)*

* ... from an individual perspective!

German Research Foundation – DFG

- HPC often inside collaborative initiatives (SFB – Collaborative Research Centers, GK – Research Training Groups, Priority Programs), but rarely in the branding / lead – here, SPPEXA is still somewhat unique
- Smaller specific calls on HPC-related topics
 - Performance Engineering (*to increase “use hi-end systems decently” skills*)
 - Sustainability of research software (*not only, but also CSE-/HPC-related; 2nd call*)
- Individual projects

Federal Ministry of Education & Research – BMBF

- 5 successive calls on Software for HPC with varying focus
- Successor initiative not in sight – complete shift towards AI

Excellence Strategy 2019+

- 57 “Excellence Clusters” selected out of about 200 proposals
- Déjà vu: HPC often inside Excellence Clusters, but rarely in the lead
- Even not in the only clear CSE-cluster “SimTech” in Stuttgart



→ A “hidden technology” – crucial for many, but relevance/science itself?



HPC Beyond SPPEXA – A German View (2)*

* ... from an individual perspective!

NHR – National HPC as a research infrastructure

- **2015:** recommendations of the German Science Council
 - Consider HPC (i.e. access to cycles & support) as a research infrastructure
 - Highest strategic relevance for science in Germany
 - Organize it in a more sustainable, holistic way
- **2016–2018:** discussions of Federal & State Governments (GWK)
- **2018:** NHR implemented in November, central office in Berlin 
 - Tier-2 only – a number of centers Germany-wide, more sustainability
 - Call(s), competition-based process, evaluation by DFG; established for several years, prolongation after positive evaluation
 - No more procurement-driven application/funding, but center-based
 - Complementary profiles, a Germany-wide system, jointly run
 - Governance: Strategy Board (science & politics)
 - Step towards total-cost-of-ownership: pure system  $cost \times \text{some factor}$
- Support included, but no HPC-related research



HPC Beyond SPPEXA – A German View (3)*

* ... from an individual perspective!

Europe

• FP 8

- Hi-end research infrastructure: PRACE – Partnership foR Advanced Computing in Europe (2 levels of members; cycles collected, cycles granted – “juste retour”)
- HPC technology (partnership with industry): ETP4HPC
- Application expertise: Centers of Excellence in Computing Applications – communities & centers
- Projects / project consortia: FET-HPC (probably closest to SPPEXA), ...
- ECI, EDI, EXDCI/EXDCI-2
- No direct investment in machines / machine development

• FP 9

- **EuroHPC Joint Undertaking** as overarching legal & funding entity: *“pool EU and participating countries’ resources to build in Europe a world-class supercomputing and data infrastructure and a competitive innovation ecosystem in relevant technologies and applications”* or *“to develop top-of-the-range exascale systems”*
- Procurement/placement: pre-exascale and exascale systems
- Partnership with industry
- Risk: focus on a few / one center(s) in each country
- More holistic view: HPC (PRACE/EuroHPC) & Networks (GEANT) & EOSC
- Not that clear yet: research related to HPC software / applications



HPC Beyond SPPEXA – A German View (4)*

* ... from an individual perspective!

What about international consortia (beyond Europe)?

- G-8 countries
 - G-8 Research Council's Initiative on Multi-lateral Funding
 - First call in 2010: *Application Software towards Exascale Computing for Global-scale Issues*
- SPPEXA phase 2
 - Possibility of D-J or D-F-J consortia
 - Joint call after intense preparations: DFG & ANR & JST
 - Still rather unique (in contrast to university-university things)
- Tri-lateral initiative in AI at ministry level (workshop in Tokyo, Nov '18)
 - Space for HPC in AI programs declared as possible
 - DFG-call in summer purely AI-related
- To some extent: disappointing outcome of the 2018 workshop – but there are plenty of ideas, and there is convergence happening



Thank You for Your Attention!

German Priority Programme 1648 Software for Exascale Computing



computational algorithms
system software
application software
data management and exploration
programming
software tools

EXAMAG - Exascale Simulations of the Magnetic Universe
U Heidelberg +++ U Würzburg +++ U Tokyo +++
U Strasbourg

Smart-DASH - Smart Data Structures and Algorithms with Support for Hierarchical Locality
LMU München +++ U Stuttgart +++ HLRS Stuttgart +++
TU Dresden +++ KIT Karlsruhe

EXASTEEL - From Micro to Macro Properties
U Köln +++ TU Bergakademie Freiberg +++ U Essen +++
TU Dresden +++ U Lugano +++ FAU Erlangen-Nürnberg

Terra-Neo - Integrated Co-Design of an Exascale Earth Mantle Modeling Framework
LMU München +++ FAU Erlangen-Nürnberg +++
TU München

AIMES - Advanced Computation and I/O Methods for Earth-System Simulations *new!*
U Hamburg +++ U Versailles +++ RIKEN +++ Tokyo Tech

ExaStencils - Advanced Stencil-Code Engineering
U Passau +++ FAU Erlangen-Nürnberg +++ U Kassel +++
U Tokyo

EXAHD - An Exa-Scalable Two-Level Sparse Grid Approach for Higher-Dimensional Problems in Plasma Physics
U Stuttgart +++ TU München +++ U Bonn +++
ANU Canberra +++ MPG Garching +++ UC Los Angeles

GROMEX - Unified Long-Range Electrostatics and Dynamic Protonation for Realistic Biomolecular Simulations on the Exascale
MPI BPC Göttingen +++ JSC Jülich +++ Stockholm U

About SPPEXA The Priority Programme *Software for Exascale Computing (SPPEXA)* of the German Research Foundation (DFG) addresses fundamental research on the various aspects of HPC software. SPPEXA runs 2013-2019, and it is implemented in two three-year phases, consisting of 13 (phase 1) and 16 (phase 2) project consortia and more than 50 institutions involved. With SPPEXA's second-phase projects funded by DFG as well as the French National Research Agency (ANR) and the Japan Science and Technology Agency (JST), SPPEXA strives for basic and trilateral research to pave the road towards exascale computing.

EXA-DUNE - Flexible PDE Solvers, Numerical Methods, and Applications
U Heidelberg +++ U Münster +++ U Stuttgart +++
TU Kaiserslautern +++ TU Clausthal +++ TU Dortmund

CATWALK - A Quick Development Path for Performance Models
ETH Zürich +++ RWTH Aachen +++ JSC Jülich +++
TU Darmstadt +++ GU Frankfurt

ESSEX - Equipping Sparse Solvers for Exascale
FAU Erlangen-Nürnberg +++ DLR Köln +++ U Greifswald +++
U Wuppertal +++ U Tsukuba +++ U Tokyo

ExaSolvers - Extreme Scale Solvers for Coupled Problems
RWTH Aachen +++ Tokyo U of Science +++ U Lugano +++
HLRS Stuttgart +++ U Trier +++ GU Frankfurt +++ Toyo U

ADA-FS - Advanced Data Placement via Ad-hoc File Systems at Extreme Scales *new!*
TU Dresden +++ JGU Mainz +++ KIT

ExaFSA - Exascale Simulation of Fluid-Structure-Acoustics Interactions
U Stuttgart +++ TU Delft +++ U Siegen +++
TU Darmstadt +++ Tohoku U

ExaDG - High-Order Discontinuous Galerkin for the Exa-Scale *new!*
U Heidelberg +++ TU München

FFMK - A Fast and Fault Tolerant Microkernel-Based System for Exascale Computing
TU Dresden +++ ZIB Berlin +++ Hebrew U Jerusalem

MYX - MUST Correctness Checking for YML and XMP Programs *new!*
RWTH Aachen +++ MDLS Saclay +++ U Tsukuba +++ RIKEN

Coordinators

Hans-Joachim Bungartz

bungartz@in.tum.de

Wolfgang E. Nagel

wolfgang.nagel@tu-dresden.de

Scientific Program Manager

Severin Reiz

s.reiz@tum.de

or visit www.sppexa.de!

