

SPPEXA

Software for Exascale Computing

DFG's Priority Program 1648

Overview

Hans-Joachim Bungartz, Wolfgang E. Nagel, Philipp Neumann



BDEC Meeting, Frankfurt, July 17th, 2015

Analysis

Imbalance between “racks” and “brains” hampers the efficient use of systems

Software in the applications is huge, complex – and badly prepared

The (non-)availability of highly performing simulation software gets more and more the enabler (or disabler) for leading science

No simple transfer of existing solutions – need for novel approaches

We are at a new era’s eve – in the applications, in algorithm design, in software complexity, and, in particular, in system architectures (many-core)

Approaches in Germany so far address ...

- ... mostly mono- or bi-disciplinary domains (IN, MA, PH, ENG ...) – important, but not sufficient
- ... single steps of the simulation pipeline (modelling, algorithms, implementation, software, data exploration), but do not take into account their interplay
- ... more applied than fundamental issues (HPC Software program of BMBF – 4th call in decision)


Huge potential for HPC applications in general (capability computing and capacity computing; in science, research, and industry)

CSE

→ Core Enabler for Science & Industry

HPC

→ Core Enabler for CSE




Mathematical model

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

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

Discretization & solver

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



Parallel implementation, HPC





Software



Impact of each step on all other steps!
 Hence #1: no pipeline any more, but a loop, a complete graph
 Hence #2: less space for single-field experts

Validation

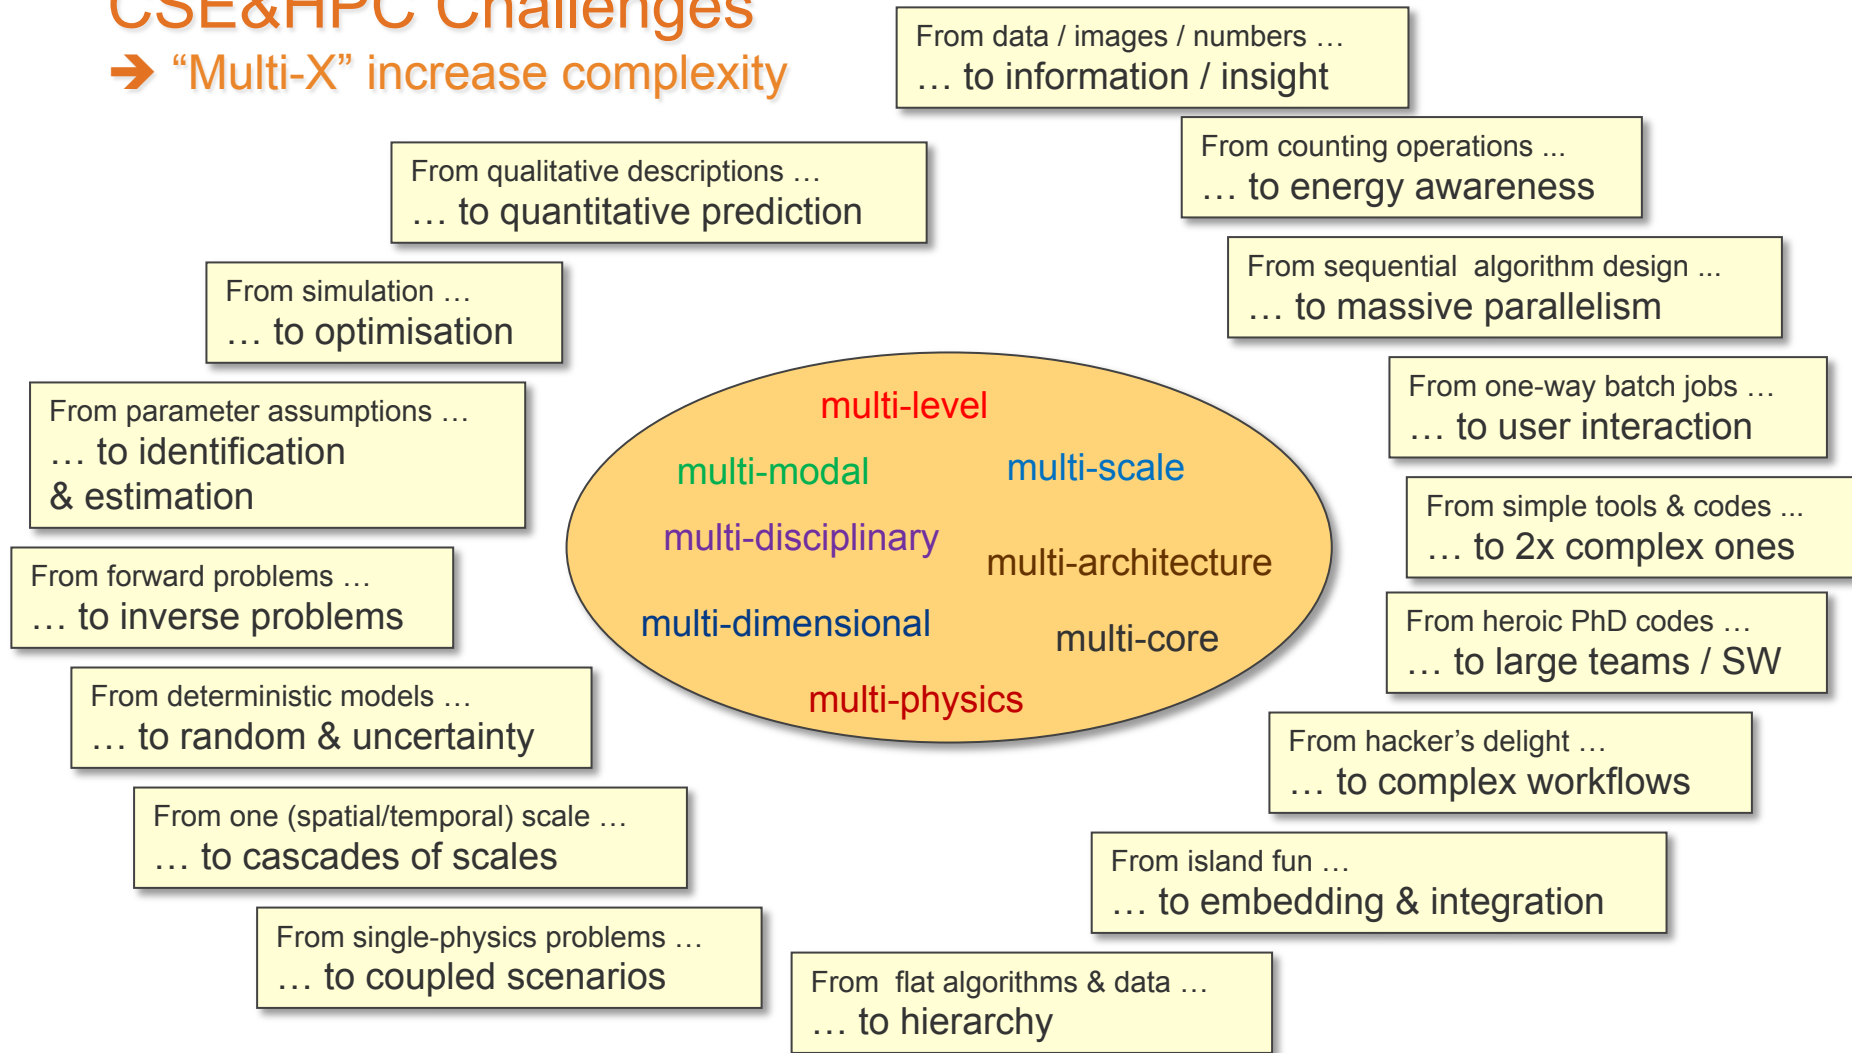


Exploration

Insight, Design

CSE&HPC Challenges

→ “Multi-X” increase complexity



SPPEXA – a Long, but Successful Way...

2006 – first discussions within DFG’s Commission on IT Infrastructure (KfR)

- 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 2008

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, NSF), e.g.

2010 – KfR takes responsibility

- Another strategic paper and a discussion with DFG’s president (Nov. 2010)
- Outcome: suggestion of a flexible, strategically initiated SPP, financed via Strategy Funds

2011 – increase of speed

- Roundtable expert meeting in MAY; DFG-internal discussions (MAY–JUL)
- Submission in AUG; international reviewing in SEP; decision in OCT; call in NOV


2012 – review of proposals

- 68 sketches in JAN, first selection in MAR leading to 24 consortia invited for full proposals
- Submission of full proposals in MAY, review workshop in JUL


2013 – launch of SPPEXA

2014 – Call for proposals SPPEXA-2, incl. international partners (France, Japan)

Yes, We Will Get Exascale!




Yes, We Need Exascale! Cf. Protein Structure




Standard Software & Hardware '11:

- Small scenarios ($< 10^6$ interaction sites)
- Short simulated times ($< 10^{-6}$ seconds)
- Moderately scaling codes ($< 10^3$ processes)
- Simulation times of weeks




Shaw et al., SC '14

Yes, We Need Exascale! Cf. Protein Structure




Gordon Bell Prize 2014:

- Small scenario (10 000 atoms)
- Medium simulated times ($< 10^{-3}$ seconds)
- 64 nodes, specific hardware (Anton-2; ASIC)
- Simulation time of two weeks




Shaw et al., SC '14

Yes, We Need Exascale! Cf. Protein Structure



Gordon Bell Prize 2014:

- Small scenario (10 000 atoms)
- Medium simulated times ($< 10^{-3}$ seconds)
- 64 nodes, specific hardware (Anton-2; ASIC)
- Simulation time of two weeks

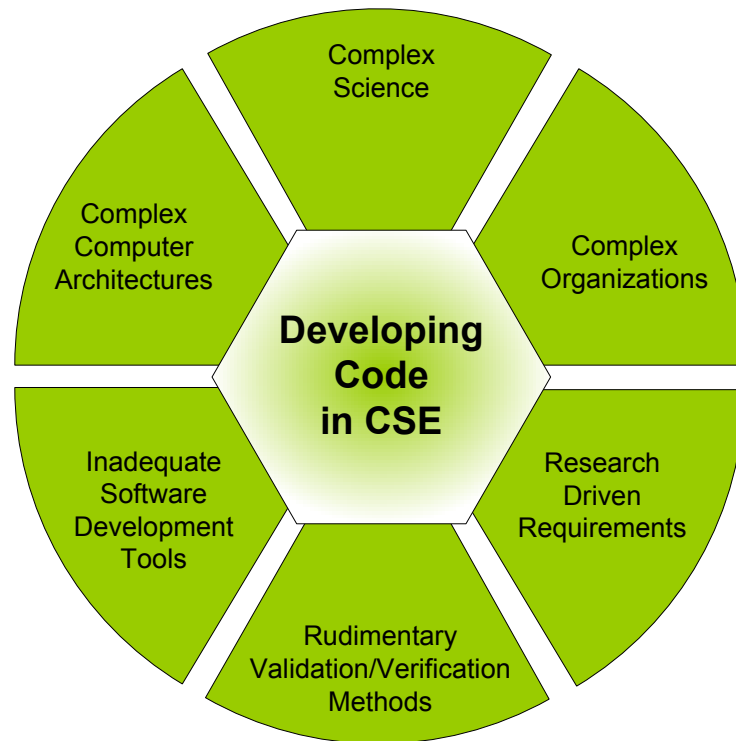


Target:

- Large scenarios ($> 10^{10}$ interaction sites)
- Longer simulated times ($> 10^{-3}$ seconds)
- Massively parallel codes ($> 10^6$ processes)
- Response times of minutes
- Quantitative insights and predictions

Shaw et al., SC '14

Software Is a Problem!



“Software development is the principal bottleneck in CSE” (R. Kendall)

Roadblocks:

- Most developers are domain scientists and engineers, not computer scientists
- **Typical priority: science >> code performance >> software quality**
- **Intellectual level assigned: models >> algorithms >> programs**
- No “software engineering mainstreaming”: design, process models, workflow models, ...
- No “team understanding”: co-operative work, trans-disciplinary, project management, ...
- **Instead, still the lonely heroes with their heroic codes (and sometimes accumulated heroism)**
- **No systematic testing culture**
- **No formal support – verification**
- No best practices
- Many groups working on PSE – in general with limited success

Energy Is a Problem!

State of the art

(Haswell, STREAM/DGEMM)



1 FLOP: ca. 0.42 nJ →

Moving 1 byte: ca. 2.6 nJ →

Projection to exascale


1 ExaFLOP/s: ca. 420 MW

1 ExaByte/s: ca. 2.6 GW



(multi power station ☹)

Keeping Moore Alive Is a Problem!



Main message:

We're no longer getting faster any more ...

... we're getting more instead ...

Source: Herb Sutter
www.gotw.ca/publications/concurrency-ddj.htm


And Parallel Is a Problem!

... would you prefer 4 large jet engines ...



Think about engines for an A 380 ...

... or rather 100,000 hair dryers??




The Topics

1. Computational Algorithms
2. System software
3. Application software
4. Data management and exploration
5. Programming
6. Software tools

SPPEXA – Software for Exascale Computing

- Phase 1: 2013–2015
 - 13 interdisciplinary research consortia
 - about 60 groups
 - International partners
- 6 research fields
 - Computational algorithms
 - Application software
 - System software
 - Programming
 - Software tools
 - Data management
- SPPEXA research is
 - Driven by domain science / CSE applications
 - Powered by SciComp / CSE methodologies
 - Smooth/evolutionary & radical/revolutionary



SPPEXA – Software for Exascale Computing

- Phase 2: 2016–2018
- As before: SPPEXA projects are
 - small consortia (2–5 nat'l partners)
 - inherently interdisciplinary
 - addressing several SPPEXA topics
- New: Focus on internationalization
 - Joint call with France (ANR) & Japan (JST)
 - German, Japanese-German, or French-Japanese-German proposals
 - Joint evaluation
 - It's an endeavor new for us ... and for the agencies
- Response to call
 - 13 + 1 renewal proposals
 - 8 new proposals

German Priority Programme 1648
“Software for Exascale Computing”



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, which is particularly urgent against the background that we have entered the era of ubiquitous massive parallelism. SPPEXA started in 2013 and is implemented in two three-year funding phases, with 13 project consortia and more than 40 institutions involved. Thus, SPPEXA has joined the other national initiatives to pave the road towards exascale computing.

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

DASH - Hierarchical Arrays for Efficient and Productive Data-Intensive Exascale Computing
 CEODE China +++ HLR5 Stuttgart +++ KIT Karlsruhe +++ LMU Munich +++ TU Dresden

TERRA-NEO - Integrated Co-Design of an Exa-Scale Earth Mantle Modeling Framework
 FAU Erlangen-Nürnberg +++ LMU Munich +++ TU München

EXASTEEL - Bridging Scales for Multiphase Steels
 FAU Erlangen-Nürnberg +++ TU Bergakademie Freiberg +++ U Cologne +++ U Duisburg-Essen

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

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

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

EXAHD - An Exa-Scalable Two-Level Sparse Grid Approach for Higher-Dimensional Problems in Plasma Physics and Beyond
 ANU Canberra +++ IPP Garching +++ TU München +++ U Bonn +++ U Stuttgart

EXAMAG - Exascale Simulations of the Evolution of the Universe Including Magnetic Fields
 U Heidelberg +++ U Würzburg

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

ESSEX - Equipping Sparse Solvers for Exascale
 FAU Erlangen-Nürnberg +++ German Aerospace Center +++ U Greifswald +++ U Wuppertal

EXASOLVERS - Extreme Scale Solvers for Coupled Problems
 HLR5 Stuttgart +++ MPI MIS Leipzig +++ RWTH Aachen +++ USI Lugano +++ U Frankfurt +++ U Trier

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



... much more than N projects

Overall – a Thrilling Research Experience

- Significant research impact and output
- Strong involvement of the HPC-related communities in Computer Science, Mathematics, Science, and Engineering
- Re-involvement of some core CS methodology
- Significant cross-disciplinary networking
- High international visibility – cf. feedback on SPPEXA's budget
- Strong further internationalization step fostering SPPEXA's research and pioneering inter-agency activities (it's fun to be the spearhead ...)



... the final message

Important Issues ... from 2012 – implemented & still valid!

Multi-disciplinarity	[within each consortium: progress in methodology + proof-of-concept in lead application]
Project consortia	[critical mass, > 1 nat'l group, > 1 topic, 1 position/group]
Role of the application	[not primary focus, but important for the proven relevance]
Semantics of “exa”	[not literal – advance large-scale parallel computing]
Semantics of “software”	[almost everything above hardware – 6 topics]
“Bring together ...”	[HPC establishment and ideas from outside]
Strategy balance	[evolutionary & revolutionary approaches]
Topical balance	[all 6 topics should be represented]
Application balance	[broad presence; science & engineering]



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

Thank you for your attention!

AGENCE NATIONALE DE LA RECHERCHE
ANR



SPPEXA – a trilateral agency collaboration

Mark Asch / French National Research Agency (ANR)

Yoshi-aki Shimada / Japan Science and Technology Agency (JST)

Marcus Wilms / German Research Foundation (DFG)

G8 and DFG activities concerning BDEC

▶ G8

- ▶ The G8 initiative / First call BDEC projects
- ▶ Follow-up call (almost)



▶ DFG

- ▶ Priority program targeting the Exascale
- ▶ Multilateralization (trilateral call)



G8 Initiative


Finished projects of the first call

- ▶ *G8 Heads of Research Councils* meeting in Kyoto/Japan in May 2008: Initiative for multilateral research funding
- ▶ Call in 2010: *Interdisciplinary Program on Application Software towards Exascale Computing for Global Scale Issues*
- ▶ Six projects on climate, earth system, seismic wave simulations, nuclear fusion simulations, dynamics of large biomolecular systems
- ▶ Final projects review, June 12-13, 2014 at Princeton University, Princeton, NJ, USA



Planned G8 Follow-Up Call

Why, how, and when?



Marcus Wilms (DFG)

4 BDEC, Frankfurt, 2015-07-16

G8 Initiative

Multilateral funding schemes are painful, but are worth the effort!

- ▶ Unique opportunity to enable collaborations between ≥ 3 different countries (4 projects with 5 or 6 countries)
- ▶ Beyond bilateral or regional international collaborations
- ▶ Potential to speed-up research by enabling top experts to collaborate; access to top HPC infrastructure
- ▶ Great opportunity for interdisciplinary and -national student training
- ▶ Heavy administrative challenge for the national funding agencies (test-bed)
- ▶ Need to find ways to foster multi-lateral collaborations



G8 and DFG activities concerning BDEC

▶ G8

- ▶ The G8 initiative / active projects of the first call
- ▶ Planned follow-up call
 - Why, how, and when?



▶ DFG

- ▶ Priority program targeting the Exascale
- ▶ Multilateralization (trilateral call)



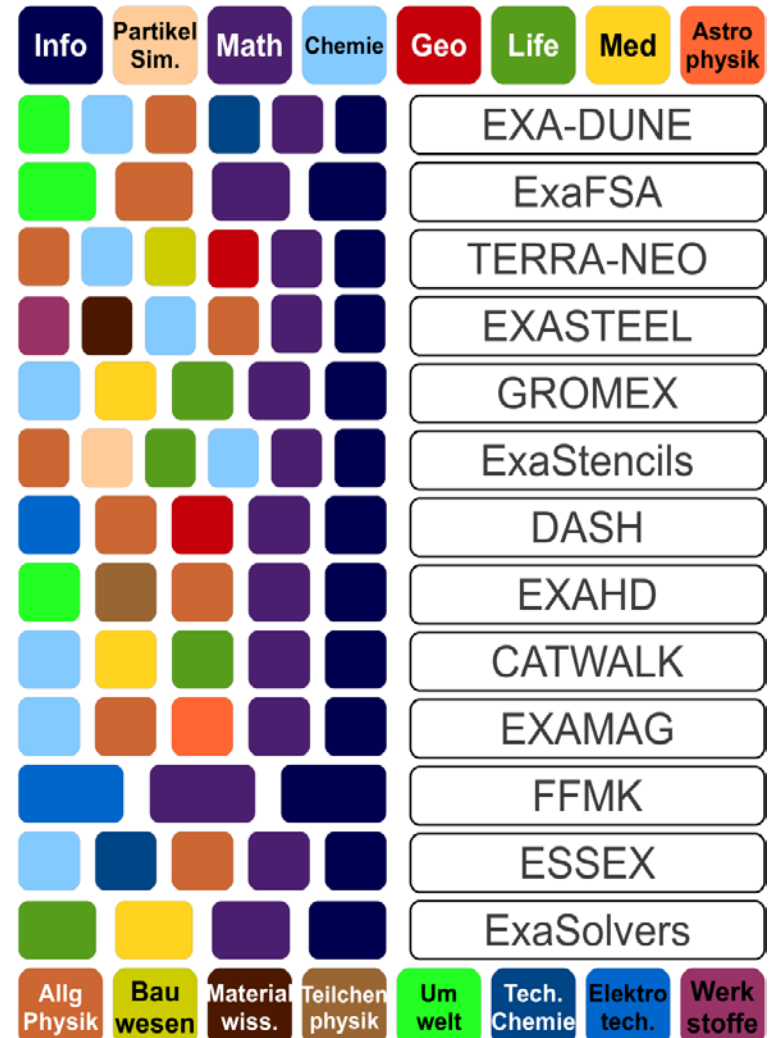
DFG Priority Program

Interdisciplinary Research Consortia



Software for Exascale Computing “SPPEXA”

- ▶ two three-year funding phases
- ▶ phase 1 launched in January 2013
- ▶ phase 1 budget of 3.7m € per year
- ▶ cover > 15 disciplines
- ▶ involve 2-5 groups each
- ▶ address ≥ 2 out of 6 SPPEXA topics
- ▶ close collaboration within and among SPPEXA consortia



Marcus Wilms (DFG)



- ▶ momentum from the G8 activity & visibility of SPPEXA-1
- ▶ scientists approached several funding agencies regarding further funding opportunities for multilateral research in BDEC
- ▶ following negotiations among interested agencies, ANR (France), DFG (Germany), and JST (Japan) agreed on a pragmatic way to open DFG's SPPEXA-2 and JST's CREST programs to connect existing coordinated national funding activities

Multilateralization of SPPEXA



- ▶ fixed agency budgets
- ▶ support collaborative projects of bi- or trilateral research teams, bringing together researchers from France (ANR), Germany (DFG), and Japan (JST)
- ▶ tri-lateral call for proposals was published in October 2014 and invited proposals for SPPEXA's second three-year funding period 2016 to 2018




- ▶ 21 (8 new) project proposals were submitted comprising 11 unilateral (German), 5 bi-lateral (Japanese-German), and 5 tri-lateral (French-Japanese-German) proposals
- ▶ two-day joint review meeting, held in Bonn on 5-6 May 2015
 - plenary applicant colloquium with oral and poster presentations (day 1)
 - review panel meeting (day 2)

Multilateralization of SPPEXA

Decision making pending (October '15)



- ▶ 16 (4 new) proposals recommended for funding
- ▶ 50% of the recommended proposals are to be multilaterally funded:
 - 4 bi-lateral (Japanese-German) proposals
 - 3 tri-lateral (French-Japanese-German) proposals
 - 2 proposals with Swiss co-funding (funding commitment pending)
 - DFG funding is further recommended for two project partners in Israel and the Netherlands.



Multilateralization of SPPEXA

Lessons learned



- ▶ use administrative layouts of the G8 Call (for setting up a streamlined Lol and Implementation plan)
- ▶ IESP/BDEC and other meetings were/are useful breeding grounds
- ▶ connect scientific and agency coordinators of national programs
- ▶ joint review builds up trust
- ▶ separate funding decisions but commitment to follow the scientific review panel recommendations
- ▶ pragmatic approach

Multilateralization of SPPEXA

Lessons learned



- ▶ connecting JST CREST and DFG coordinated programs is now generally established for further Japanese-German collaborations
- ▶ importance of multilateral funding activities
- ▶ trilateral initiative might serve as an effective stimulus in opening up funding opportunities for multilateral research as a more general practice



DFG

Thank you

More information

- ▶ DFG: <http://www.dfg.de>
- ▶ funded research projects: <http://www.dfg.de/gepris>
- ▶ search for more than 17,000 German research institutions: <http://www.dfg.de/rex>