

Software for Exascale Computing: Some Remarks on the Priority Program SPPEXA



Hans-Joachim Bungartz, Wolfgang E. Nagel, Philipp Neumann, Severin Reiz, and Benjamin Uekermann

Abstract SPPEXA, the Priority Program 1648 “Software for Exa-scale Computing” of the German Research Foundation (DFG), was established in 2012. SPPEXA was DFG’s first strategic Priority Program—strategic in the sense that it had been the initiative of DFG’s board to suggest a larger and trans-disciplinary funding scheme to support the development of software at all levels that would be able to benefit from future exa-scale systems. A proposal had been formulated by a team of scientists representing domains across the STEM fields, evaluated in the standard format for Priority Programs, and financed via special funds. Operations started in January 2013, and after two 3-year funding phases and a cost-neutral extension, SPPEXA’s activities will come to an end by end of April, 2020. A final international symposium took place on October 21–23, 2019, in Dresden, and this volume of Springer’s Lecture Notes in Computational Science and Engineering—the second SPPEXA-related one after the corresponding report of Phase 1 (see Appendix 3 in [1])—contains reports of 16 out of 17 SPPEXA projects (the project ExaSolvers will deliver its report as a special issue of Springer’s journal *Computing and Visualization in Science*) and is, thus, a comprehensive overview of research within SPPEXA.

While each single project report emphasizes the respective project’s individual research outcomes and, thus, provides one perspective of research in SPPEXA, this contribution, co-authored by the two scientific coordinators—Hans-Joachim

H.-J. Bungartz · S. Reiz (✉)
Technical University of Munich, Garching, Germany
e-mail: bungartz@in.tum.de; reiz@in.tum.de

W. E. Nagel
Technical University of Dresden, Dresden, Germany

P. Neumann
Helmut-Schmidt-Universität Hamburg, Hamburg, Germany
e-mail: philipp.neumann@hsu-hh.de

B. Uekermann
Eindhoven University of Technology, Eindhoven, Netherlands

Bungartz and Wolfgang E. Nagel—and by three of the four researchers that have served as program coordinator over the years—Philipp Neumann, Benjamin Uekermann, and Severin Reiz—emphasizes the program SPPEXA itself. It provides an overview of the design and implementation of SPPEXA, it highlights its accompanying and supporting activities (internationalization, in particular with France and Japan; workshops; doctoral retreats; diversity-related measures), and it provides some statistics. It, thus, complements the papers from SPPEXA’s research consortia collected in this volume.

1 Preparation

While *supercomputers* were recognized early as an important research infrastructure for German science and have been since then on the agenda (recommendations of the German Science Council (Wissenschaftsrat), introduction of the performance pyramid, Gauss Centre for Supercomputing, Gauss Alliance, NHR—Nationales Hochleistungsrechnen), the situation for *supercomputing* has always been quite different. First, the funds for HPC systems are typically limited to investments, i.e. the machinery; the current NHR initiative takes a more comprehensive view. Second, software development is frequently not considered as “science”, which entails that neither typical projects in informatics or mathematics nor their counterparts in fields of application cover more than prototype development. Recently, BMBF’s HPC software program and DFG’s sustainable scientific software initiative, fortunately, have acknowledged the crucial role of software for HPC and support software development explicitly. Third, HPC software development has happened in Collaborative Research Centers or similar formats before, but mostly in an isolated way: an informatics initiative contained an HPC software project as an application, or a physics initiative contained a simulation- or HPC-oriented project. But all this hardly ever looked at more than one peculiar aspect at a time, and it was at most an interdisciplinary endeavor of two fields.

However, when Moore’s law at least gets exhausted a bit and performance gains are more and more achievable through a more and more massive parallelism only, it is obvious that software and its performance and scalability play an increasingly crucial part. Therefore, the challenges at the eve of the exa-scale era required more—and that’s actually what happened elsewhere, for example in the U.S. or in Japan: a significant, concerted initiative, bringing together informatics, mathematics, and several domains of application, comprising all relevant aspects of HPC software. That’s where SPPEXA entered the stage.

2 Design Principles

SPPEXA was designed to provide a holistic approach to HPC software, comprising the aspects most relevant for ensuring the efficient use of current and upcoming high-end supercomputers, and to do this via exploring both evolutionary and

disruptive research threads. Six research directions were identified as crucial ones: (1) Computational Algorithms, (2) Application Software, (3) System Software and Runtime Libraries, (4) Programming, (5) Software Tools, and (6) Data Management. Computational algorithms, such as fast linear solvers or eigensolvers, are a core numerical component of many large-scale application codes—both classical simulation-driven and recent data analytics-oriented ones. If scalability cannot be ensured here, the battle is already almost lost. Application software is the “user” of HPC systems, typically appearing as legacy codes that have been developed over many years. Increasing their performance via a co-design that addresses both the “systems—algorithms” and the “algorithms—applications/models” interfaces and combines algorithm and performance engineering is vital. Performance engineering can’t succeed without progress in compilers, monitoring, code optimization, verification support, and parallelization support (such as auto-tuning)—which underlines the importance of system software and runtime libraries as well as of tools. Programming, including programming models, is probably the topic where the need for a balance of evolutionary research (improve and extend existing programming models, e.g.) and revolutionary approaches (explore new programming models, new language concepts such as Domain-Specific Languages) gets most obvious. Data management, finally, has always been HPC-relevant in terms of I/O or post-processing and visualization, and it is of ever-increasing importance since more and more HPC applications are on the data side.

To ensure the impact of this holistic idea, it was clear that having a set of projects in our Priority Program where some address this issue and others that one, and where they may collaborate or not, would not suffice. Therefore, SPPEXA’s concept was to have a set of larger projects, or project consortia (research units—Forschergruppen), that would all have to address at least two of the six big topics with their research agenda; and that would all have to combine a relevant large-scale application with HPC-methodical advancements. This means that neither a merely domain-driven research (“improve my code, and this is a contribution to HPC in itself”), as we see it frequently in domain-driven research initiatives (Collaborative Research Centers in physics, life sciences, or engineering, e.g.), nor a generic purely algorithmic research (“if I improve my solver, this will help everyone”), as we see it frequently in mathematics- or informatics-driven research initiatives, would be allowed to find their place in SPPEXA. This was somewhat challenging, since we had to communicate this concept clearly and to convince potential applicants and reviewers that everyone should really comply with this agenda.

Furthermore, there is one property better known from Collaborative Research Centers than from Priority Programs: program-wide joint activities. For example, we wanted to have a vivid collaboration framework of cross-project workshops; networking with the big international programs; a focus on education also through fostering novel teaching formats or coding weeks and doctoral retreats for the doctoral candidates; gender-related activities to understand, evaluate and work towards a more gender-balanced research community; etc. This allowed for sharing mutual best practices in HPC for the mathematics- or informatics- or application-driven areas. Therefore, there was more coordination than we see in typical Priority Programs.

3 Funded Projects and Internal Structure

In the first funding phase, the following thirteen projects or project consortia were funded:¹

CATWALK—A Quick Development Path for Performance Models. Felix Wolf (Darmstadt), Christian Bischof (Darmstadt), Torsten Hoefler (Zürich), Bernd Mohr (Jülich), and Gabriel Wittum (Frankfurt)
ESSEX—Equipping Sparse Solvers for Exa-scale. Gerhard Wellein (Erlangen), Achim Basermann (Köln), Holger Fehske (Greifswald), Georg Hager (Erlangen), and Bruno Lang (Wuppertal)
Exa-Dune—Flexible PDE Solvers, Numerical Methods, and Applications. Peter Bastian (Heidelberg), Olaf Ippisch (Clausthal), Mario Ohlberger (Münster), Christian Engwer (Münster), Stefan Turek (Dortmund), Dominik Göddeke (Stuttgart), and Oleg Iliev (Kaiserslautern)
ExaFSA—Exa-scale Simulation of Fluid-Structure-Acoustics Interactions. Miriam Mehl (Stuttgart), Hester Bijl (Delft), Sabine Roller (Siegen), Dörte Sternel (Darmstadt), and Thomas Ertl (Stuttgart)
EXAHD—An Exa-Scalable 2-Level Sparse Grid Approach for Higher-Dimensional Problems in Plasma Physics and Beyond. Dirk Pflüger (Stuttgart), Hans-Joachim Bungartz (München), Michael Griebel (Bonn), Markus Hegland (Canberra), Frank Jenko (Garching), and Hermann Lederer (Garching)
EXAMAG—Exa-scale Simulations of the Evolution of the Universe Including Magnetic Fields. Volker Springel (Heidelberg) and Christian Klingenberg (Würzburg)
ExaSolvers—Extreme-scale Solvers for Coupled Problems. Lars Grasedyck (Aachen), Wolfgang Hackbusch (Leipzig), Rolf Krause (Lugano), Michael Resch (Stuttgart), Volker Schulz (Trier), and Gabriel Wittum (Frankfurt)
EXASTEEL—Bridging Scales for Multiphase Steels. Daniel Balzani (Bochum), Axel Klawonn (Köln), Oliver Rheinbach (Freiberg), Jörg Schröder (Duisburg-Essen), and Gerhard Wellein (Erlangen)
ExaStencils—Advanced Stencil-Code Engineering. Christian Lengauer (Passau), Armin Größlinger (Passau), Ulrich Rüde (Erlangen), Harald Köstler (Erlangen), Sven Apel (Saarbrücken), Jürgen Teich (Erlangen), Frank Hannig (Erlangen), and Matthias Bolten (Wuppertal)
FFMK—A Fast and Fault-tolerant Microkernel-Based System for Exa-scale Computing. Hermann Härtig (Dresden), Alexander Reinefeld (Berlin), Amnon Barak (Jerusalem), and Wolfgang E. Nagel (Dresden)
GROMEX—Unified Long-range Electrostatics and Dynamic Protonation for Realistic Biomolecular Simulations on the Exa-scale. Helmut Grubmüller (Göttingen), Holger Dachselt (Jülich), and Berk Hess (Stockholm)
DASH—Smart Data Structures and Algorithms with Support for Hierarchical Locality. Karl Furlinger (München), Colin W. Glass (Stuttgart), José Gracia (Stuttgart), and Andreas Knüpfer (Dresden)
Terra-Neo—Integrated Co-Design of an Exa-scale Earth Mantle Modeling Framework. Hans-Peter Bunge (München), Ulrich Rüde (Erlangen), Gerhard Wellein (Erlangen), and Barbara Wohlmuth (München)

¹Some Principal Investigators have changed affiliation during the SPPEXA program. We specified the most recent main affiliation here.

After 3 years, twelve of those got a prolongation for the second funding phase, some with an “international extension” (bi-national with Japanese partners or tri-national with French and Japanese partners):

ESSEX-2—Equipping Sparse Solvers for Exa-scale. Gerhard Wellein (Erlangen), Achim Basermann (Köln), Holger Fehske (Greifswald), Georg Hager (Erlangen), Bruno Lang (Wuppertal), Tetsuya Sakurai (Tsukuba; Japanese partner), and Kengo Nakajima (Tokyo; Japanese partner)

Exa-Dune—Flexible PDE Solvers, Numerical Methods, and Applications. Peter Bastian (Heidelberg), Olaf Ippisch (Clausthal), Mario Ohlberger (Münster), Christian Engwer (Münster), Stefan Turek (Dortmund), Dominik Göddeke (Stuttgart), and Oleg Iliev (Kaiserslautern)

ExaFSA—Exa-scale Simulation of Fluid-Structure-Acoustics Interactions. Miriam Mehl (Stuttgart), Alexander van Zuijlen (Delft), Thomas Ertl (Stuttgart), Sabine Roller (Siegen), Dörte Sternel (Darmstadt), and Hiroyuki Takizawa (Tohoku; Japanese partner)

EXAHD—An Exa-Scalable 2-Level Sparse Grid Approach for Higher-Dimensional Problems in Plasma Physics and Beyond. Dirk Pflüger (Stuttgart), Hans-Joachim Bungartz (München), Michael Griebel (Bonn), Markus Hegland (Canberra), Frank Jenko (Garching), and Tilman Dannert (Garching)

EXAMAG—Exa-scale Simulations of the Magnetic Universe. Volker Springel (Heidelberg), Christian Klingenberg (Würzburg), Naoki Yoshida (Tokyo; Japanese partner), and Philippe Helluy (Strasbourg; French partner)

ExaSolvers—Extreme-scale Solvers for Coupled Problems. Lars Grasedyck (Aachen), Rolf Krause (Lugano), Michael Resch (Stuttgart), Volker Schulz (Trier), Gabriel Wittum (Frankfurt), Arne Nägel (Frankfurt), Hiroshi Kawai (Tokyo; Japanese partner), and Ryuji Shioya (Toyo; Japanese partner)

EXASTEEL-2—Dual Phase Steels—From Micro to Macro Properties. Daniel Balzani (Bochum), Axel Klawonn (Köln), Oliver Rheinbach (Freiberg), Jörg Schröder (Duisburg-Essen), Olaf Schenk (Lugano), and Gerhard Wellein (Erlangen)

ExaStencils—Advanced Stencil-Code Engineering. Christian Lengauer (Passau), Ulrich Rüde (Erlangen), Harald Köstler (Erlangen), Sven Apel (Saarbrücken), Jürgen Teich (Erlangen), Frank Hannig (Erlangen), Matthias Bolten (Wuppertal), and Shigeru Chiba (Tokyo; Japanese partner)

FFMK—A Fast and Fault-tolerant Microkernel-Based System for Exa-scale Computing. Hermann Härtig (Dresden), Alexander Reinefeld (Berlin), Amnon Barak (Jerusalem), and Wolfgang E. Nagel (Dresden)

GROMEX—Unified Long-range Electrostatics and Dynamic Protonation for Realistic Biomolecular Simulations on the Exa-scale. Helmut Grubmüller (Göttingen), Holger Dachselt (Jülich), and Berk Hess (Stockholm)

DASH—Smart Data Structures and Algorithms with Support for Hierarchical Locality. Karl Furlinger (München), Colin W. Glass (Stuttgart), José Gracia (Stuttgart), and Andreas Knüpfer (Dresden)

Terra-Neo—Integrated Co-Design of an Exa-scale Earth Mantle Modeling Framework. Hans-Peter Bunge (München), Ulrich Rüde (Erlangen), and Barbara Wohlmuth (München)

Furthermore, four new project consortia joined SPPEXA:

ADA-FS—Advanced Data Placement via Ad-hoc File Systems at Extreme Scales. Wolfgang E. Nagel (Dresden), André Brinkmann (Mainz), and Achim Streit (Karlsruhe)

AIMES—Advanced Computation and I/O Methods for Earth-System Simulations. Thomas Ludwig (Hamburg), Thomas Dubos (Versailles; French partner), Naoya Maruyama (RIKEN; Japanese partner), and Takayuki Aoki (Tokyo; Japanese partner)

ExaDG—High-order Discontinuous Galerkin for the Exa-scale. Guido Kanschat (Heidelberg), Katharina Kormann (München), Martin Kronbichler (München), and Wolfgang A. Wall (München)

MYX—MUST Correctness Checking for YML and XMP Programs. Matthias S. Müller (Aachen), Serge Petitot (Lille; French partner), Nahid Emad (Versailles; French partner), Taisuke Boku (Tsukuba; Japanese partner), and Hitoshi Murai (RIKEN; Japanese partner)

Finally, 1 year later, a seventeenth project joined SPPEXA as associated project:

ExtraPeak—Automatic Performance Modeling of HPC Applications. Felix Wolf (Darmstadt) and Torsten Hoefler (Zürich)

Hence, overall, there have been four Japanese-German and three French-Japanese-German consortia within SPPEXA. On the German side, an overall sum of 57 principal investigators from 39 institutions have been involved, representing informatics (25), mathematics (19), engineering (8), natural sciences (4), and life sciences (1).

Concerning governance, SPPEXA was headed by its two Spokespersons Hans-Joachim Bungartz (Technical University of Munich—TUM) and Wolfgang E. Nagel (Technical University of Dresden). For the everyday organization, a Program Coordinator (in chronological order: Benjamin Peherstorfer, now professor at New York University; Philipp Neumann, now professor at Helmut-Schmidt-University Hamburg; Benjamin Uekermann, now with Eindhoven University of Technology; and Severin Reiz, TUM) as well as an Office were established (both at TUM). Strategic decisions in SPPEXA were taken by the Steering Committee, consisting of H.-J. Bungartz, W. E. Nagel, as well as Sabine Roller (Siegen), Christian Lengauer (Passau), Hans-Peter Bunge (München), Dörte Sternel (Darmstadt), and—in the second funding phase—Nahid Emad (France) and Takayuki Aoki (Japan). Finally, a Scientific Advisory Board supported our activities and planning: George Biros (University of Texas at Austin), Rupak Biswas (NASA), Klaus Becker (Airbus), Rob Schreiber (at that time HP Labs), and Craig Stewart (University of Indiana at Bloomington).

4 SPPEXA Goes International

Extreme-scale HPC has always been an international endeavor. In 2010, as the first call in the framework of the G8 Research Councils' Initiative on Multilateral Research Funding, the topic Application Software towards Exa-scale Computing for Global Scale Issues had been selected. In the sequel of that initiative, the idea arose to give SPPEXA in its second funding phase a more international flavor, beyond the individual international partners present in some of the consortia. DFG's head office contacted several of their partner institutions in other countries. While it turned out to be complicated to synchronize activities with the National Science Foundation (NSF) in the U.S., the discussions with the French Agence Nationale de la Recherche (ANR) and the Japan Science and Technology Agency (JST) became very concrete. Finally, for the first time, a funding phase of a complete DFG Priority Program was linked to funding formats from two other countries, and the three agencies combined their forces in a joint call run by DFG. Due to formal restrictions, two new types of SPPEXA consortia were open for application: bi-national Japanese-German or tri-national French-Japanese-German ones.

Overall, the following French institutions participated in SPPEXA projects: Université de Versailles, Université de Strasbourg, and Maison de la Simulation, Saclay. From the Japanese side, the involved partner institutions involved were RIKEN, Tokyo University of Technology, University of Tsukuba, University of Tokyo, Tohoku University, Tokyo University of Science, and Toyo University. Beyond research in the single consortia, one SPPEXA doctoral retreat was held in France, and SPPEXA co-organized three French-Japanese-German workshops—the first one 2017 in the French embassy in Tokyo, the second one in 2018 in the German embassy in Tokyo, and the third one in 2019, again in the French embassy. The first two focused on exa-scale computing, while the third one did a move towards artificial intelligence (AI) and, in particular, addressed the convergence of AI and HPC.

Further internationalization measures were the SPPEXA guest program, the research stays for doctoral candidates (up to 3 months; overall 25 taken in funding phase 2), and our PR activities at the big international meetings. For example, SPPEXA organized panels or sessions at the Supercomputing Conference (SC) and the International Supercomputing Conference (ISC HPC) and participated in the session and poster exhibition on DFG-funded collaborative research at DATE 2019.

5 Joint Coordinated Activities

As mentioned above, SPPEXA featured a rich program of joint cross-consortium activities (the following numbers refer to funding phase 2, 2016–2019):

Guests Overall, more than 85 guest researchers visited one or more SPPEXA projects.

Workshops Workshops were a particular format to foster exchange and collaboration across project consortia. Central funds had been established for that, and each SPPEXA PI could hand in proposals (two calls per year). The proposal had to depict how the cross-consortium effect was to be ensured (more than one organizing consortium, etc.). Overall, 41 SPPEXA workshops, held at conferences or stand-alone, were supported via this channel.

Doctoral Retreats The SPPEXA Doctoral Retreat had two main goals—first, to offer an additional educational component to our doctoral candidates; second, to overcome the sometimes narrow borders of research by connecting with international researchers on a doctoral level (guest lectures, own contributions, hands-on sessions, . . .). Overall, three doctoral retreats were organized: Strasbourg (2016), Dresden (2017), and Wuppertal (2018).

Doctoral Research Stays Following the successful model of TUM Graduate School, where each doctoral candidate university-wide can get funds for an international research stay of up to 3 months, we encouraged our doctoral candidates SPPEXA-wide to enrich their PhD phase with such an international component. Overall, 25 such research interns were funded, examples for destinations being ETH Zurich, NORCE Bergen, or University of Tennessee.

Gender Activities Looking at the gender situation in HPC, it is obvious that the presence of women is even worse than in general in informatics. To improve that situation and to provide a more open atmosphere, a couple of measures were taken. At every Annual Plenary Meeting (2016, 2017, 2018, and 2019), we organized gender trainings by external coaches to raise awareness of gender biases in academia, each with 25 participants. Additionally, SPPEXA members organized workshop-like events such as student MINT mentoring days (2016–2018) and women’s networking events in 2019. Moreover, we connected to industry (Bosch and IBM) via gender bias discussion days called “Equality at Exascale”. Exceptional at this event was that not only women participated, but we had an ideal gender-parity in participants.

Impact on Education As a side effect, HPC education also got a boost by SPPEXA. Numerous lectures and lab courses were updated, and a lot of student theses had topics directly related to SPPEXA projects.

Prizes During the second phase of SPPEXA, every year, the best student and doctoral theses SPPEXA-wide were awarded a prize. Over the years, the winners were:

2016: Klaudius Scheufele (Stuttgart, master’s thesis) and Benjamin Uekermann (Munich, PhD thesis);

2017: Sebastian Schweikl (Passau, bachelor’s thesis), Simon Schwitanski (Aachen, master’s thesis), and Moritz Kreutzer (Erlangen, PhD thesis);

2018/2019: Piet Jarmatz (Munich/Hamburg, master’s thesis) and Sebastian Kuckuk and Christian Schmitt (Erlangen, PhD thesis).

Support of Young Researchers For sustainability in academia, supporting young aspiring researchers is indispensable. We took measures by funding research stays for doctoral candidates and awarding prizes for exceptional theses. Additionally, we also supported bachelor and master students for the student cluster competition at the (international) supercomputing conferences SC and ISC HPC 2016–2019.

Public Relations Dissemination of research becomes more and more important. Continuing efforts from the first phase, SPPEXA featured articles in the InSiDE magazine, published by the GAUSS Center for Supercomputing, twice per year in 2016, 2017, and 2018 introducing one project each time. Furthermore, starting 2018, SPPEXA contributed five articles to the online platform Science Node.² Last, in 2018, SPPEXA also featured an article in the EU Research magazine.

Internationalisation See previous Sect. 4.

6 HPC Goes Data

The computational revolution goes on! Computers and sophisticated computational methods have shaped the “third paradigm”, the third path to insight in science, complementing the classical approaches, theory and experiment, but also building a bridge and providing the missing link between those two. An early incarnation of “computational” were numerical simulations, later expanded by so-called “outer-loop scenarios”, in which repeated simulations allow for enhanced results: optimization, parameter identification, stochastics, or uncertainty quantification. All of this, basically, was model-driven, following a deductive regime of model hypotheses and derivations from them. The latest appearance of “computational” can be characterized by the focus on data: data-enhanced simulation, data analytics, machine learning, or artificial intelligence. Instead of being based on models, this approach is much more data-driven, following an inductive regime of collecting data and drawing conclusions from them. In simplified words, the “data from models” turned into, or was complemented by, a “models from data”. Despite that shift of focus, the basic underlying principle did not change: state-of-the-art computer systems and state-of-the-art computational methods are combined and used to advance the frontier of science. Something new is maybe the fact that the club of scientific domains that benefit from the “third paradigm” has become bigger: While numerical simulation was, more or less, driven by natural, engineering, and life sciences, the data-centered approach comprises all domains, including social sciences and humanities.

Of course, this development has a huge impact on HPC. In particular, new fields and new types of applications popped up, as well as new lines of architectures and systems. For example, in 2018, the majority of finalists for the Gordon Bell

²<https://scienode.org/feature/the-race-to-exascale.php>.

Award, the most renowned prize in HPC, already had a significant amount of machine learning in their papers. World-wide, HPC centers observe an increasing share of data-driven jobs on their machines. This is not surprising: as science and science methodology evolve, the kind of studies done in that context also does. Despite all those changes, the role of HPC is astonishingly stable: HPC is a core enabling technology of “computational”. It was and still is an enabler of numerical simulation, and it has become a crucial enabler of data analytics and artificial intelligence. If artificial intelligence, machine learning, or deep learning have become so popular recently, this is much more due to the fact that established methodology can succeed due to HPC, than due to new AI/ML/DL methodology itself.

These developments are also visible at the end of SPPEXA. Several consortia already are on that “data-driven track”, as, for example, our third French-Japanese-German workshop in Tokyo showed.

7 Shaping the Landscape

When SPPEXA started in 2013, the core idea was to significantly improve algorithms, software, and tools, in order to be prepared for the exa-scale age. In the meantime, we are at the eve of exa-scale systems, as the co-design developments in the U.S. and in Japan (Fugaku) or the discussions in the European Union on exa-scale and pre-exa-scale systems show. And research in SPPEXA has definitely contributed to the application landscape in Germany being much closer to “exa-scale-readiness” than before. Several leading application software packages were involved, and significant progress in terms of scalability and parallel efficiency could be achieved. Furthermore, and maybe even more important, the SPPEXA consortia showed the advantages of the multi-disciplinary engagement, brought together a lot of groups and ideas disconnected before, and, thus, justified the concept of larger, cross-institutional, and cross-disciplinary teams instead of single-PI projects.

The visibility SPPEXA got is stunning. SPPEXA was present at the leading international conferences (Euro-Par, Supercomputing, ISC HPC)—through individual presentations and special events, such as minisymposia or panels. But also at “neighboring” events, such as the DATE 2019 (Design, Automation, and Test in Europe), SPPEXA had a presentation slot and a booth. SPPEXA was involved in the activities (workshops, white papers, etc.) of the BDEC Community (Big Data and Extreme-Scale Computing) as well as in the organization of the Long Program “Science at Extreme Scales: Where Big Data Meets Large-scale Computing” at the Institute for Pure and Applied Mathematics (IPAM) in Los Angeles, and it co-organized a French-Japanese-German workshop series in Tokyo (cf. the section on internationalization). Thus, at an international scale, SPPEXA was generally perceived as the “German player” in the HPC software concert.

8 Concluding Remarks

Without any doubt, SPPEXA has written a success story: in terms of its research, concerning the innovative funding format, with its multi-disciplinary approach, its multi-national facets, and—last, but not least—its huge visibility. We are grateful for all the support we got from the German Research Foundation (DFG): the funding, but also for the encouragement during the preparation of SPPEXA and the continued advice during its runtime.

Appendix 1: Qualification

The following achievements have been completed in the SPPEXA program within 1.1.2016 and 30.04.2020:

Projects	Completed PhD theses	Completed habilitations	Calls to professorship
AIMES	0	0	1
ADA-FS	0	0	0
DASH	1	0	1
ESSEX	1	1	0
ExaDG	4	0	1
Exa-Dune	4	0	1
ExaFSA	2	0	0
EXAHD	3	0	0
EXAMAG	9	0	0
ExaSolvers	1	0	1
EXASTEEL	2	0	1
ExaStencils	5	2	4
ExtraPeak	3	0	0
FFMK	1	0	0
GROMEX	2	0	0
MYX	0	0	0
Terra-Neo	3	0	0
Coordination	1	1	2
Overall	43	3	12

The previous table follows the DFG requirements for final reports in priority programs. At least 25 additional PhD candidates are close to being finished; however, due to the lengthy defense procedure they are not counted here.

Also, please take into account that project consortia vary in size (regarding Principal Investigators and PhD candidates) and their start/end date.

Appendix 2: Software from Project Consortia

In the following, a table with links to software that has been developed by the project consortia in SPPEXA Phase-II is given.

Project	Software developed
AIMES	SCIL
	github.com/JulianKunkel/scil
ADA-FS	GekkoFS
	tu-dresden.de/zih/forschung/projekte/ada-fs
DASH	DASH
	www.dash-project.org/
ESSEX	PHIST, GHOST, CRAFT, RACE, ScaMaC
	bitbucket.org/essex/ {PHIST, ..., RACE, matrixcollection}
ExaDG	deal.II
	github.com/dealii/dealii
EXA-Dune	DUNE
	gitlab.dune-project.org/exadune
ExaFSA	preCICE
	github.com/precice/precice
	Ateles
	apes.osdn.io/pages/ateles
EXAHD	SG++
	github.com/SGpp/SGpp
EXAMAG	AREPO
	arepo-code.org/
ExaSolvers	utopia
	bitbucket.org/zulianp/utopia/src/master/
EXASTEEL	FE2TI
	www.numerik.uni-koeln.de/14079.html
ExaStencils	LFA Lab
	hrittich.github.io/lfa-lab/
	ExaSlang
	i10git.cs.fau.de/exastencils/release
ExtraPeak	Extra-P
	www.scalasca.org/scalasca/software/extra-p/
FFMK	FFMK
	ffmk.tudos.org/
GROMEX	GROMACS
	www.gromacs.org

MYX	MUST
	doc.itc.rwth-aachen.de/display/CCP/Project+MUST
Terra-Neo	HyTeG
	i10git.cs.fau.de/hyteg/hyteg
	waLBerla
	www.walberla.net/
	TerraNeo
	terraneo.fau.de/

Appendix 3: Project Consortia Key Publications

This volume represents a continuation of the corresponding report in SPPEXA Phase-I, which is referenced several times in the text above:

1. Bungartz, H.-J., Neumann, P., Nagel, W.E.: *Software for Exascale Computing-SPPEXA 2013–2015*, vol. 113. Springer, Berlin (2016)

SPPEXA Phase-II showed visibility in the research community with numerous publications. In the following we provide a list of two key publications for each project consortium:³

AIMES

1. Jum’ah, N., Kunkel, J.: Performance portability of earth system models with user-controlled GGDML code translation. In: International Conference on High Performance Computing, pp. 693–710. Springer, Berlin (2018)
2. Kunkel, J., Novikova, A., Betke, E., Schaare, A.: Toward decoupling the selection of compression algorithms from quality constraints. In: International Conference on High Performance Computing, pp. 3–14. Springer, Berlin (2017)

ADA-FS

1. Vef, M.A., Moti, N., Süß, T., Tocci, T., Nou, R., Miranda, A., Cortes, T., Brinkmann, A.: GekkoFS—a temporary distributed file system for HPC applications. In: 2018 IEEE International Conference on Cluster Computing (CLUSTER), pp. 319–324. IEEE, Piscataway (2018)
2. Soysal, M., Berghoff, M., Klusáček, D., Streit, A.: On the quality of wall time estimates for resource allocation prediction. In: Proceedings of the 48th International Conference on Parallel Processing: Workshops, pp. 1–8. ACM, New York (2019)

DASH

1. Kowalewski, R., Jungblut, P., Furlinger, K.: Engineering a distributed histogram sort. In: 2019 IEEE International Conference on Cluster Computing (CLUSTER), pp. 1–11. IEEE, Piscataway (2019)

³Following the DFG requirements for final reports in priority programs.

2. Furlinger, K., Glass, C., Gracia, J., Knüpfer, A., Tao, J., HHünichnich, D., Idrees, K., Maiterth, M., Mhedheb, Y., Zhou, H.: DASH: data structures and algorithms with support for hierarchical locality. In: European Conference on Parallel Processing, pp. 542–552. Springer, Berlin (2014)

ESSEX

1. Pieper, A., Kreutzer, M., Alvermann, A., Galgon, M., Fehske, H., Hager, G., Lang, B., Wellein, G.: High-performance implementation of Chebyshev filter diagonalization for interior eigenvalue computations. *J. Comput. Phys.* **325**, 226–243 (2016)
2. Röhrig-Zöllner, M., Thies, J., Kreutzer, M., Alvermann, A., Pieper, A., Basermann, A., Hager, G., Wellein, G., Fehske, H.: Increasing the performance of the Jacobi–Davidson method by blocking. *SIAM J. Sci. Comput.* **37**(6), C697–C722 (2015)

ExaDG

1. Kronbichler, M., Kormann, K.: Fast matrix-free evaluation of discontinuous Galerkin finite element operators. *ACM Trans. Math. Softw.* **45**(3), 1–40 (2019)
2. Fehn, N., Wall, W.A., Kronbichler, M.: Efficiency of high-performance discontinuous Galerkin spectral element methods for under-resolved turbulent incompressible flows. *Int. J. Numer. Methods Fluids* **88**(1), 32–54 (2018)

EXA-Dune

1. Bastian, P., Engwer, C., Göttsche, D., Iliev, O., Ippisch, O., Ohlberger, M., Turek, S., Fahlke, J., Kaulmann, S., Steffen Mütting, S., et al.: EXA-DUNE: flexible PDE solvers, numerical methods and applications. In: European Conference on Parallel Processing, pp. 530–541. Springer, Berlin (2014)
2. Engwer, C., Altenbernd, M., Dreier, N.A., Göttsche, D.: A high-level C++ approach to manage local errors, asynchrony and faults in an MPI application. In: 2018 26th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP), pp. 714–721. IEEE, Piscataway (2018)

ExaFSA

1. Mehl, M., Uekermann, B., Bijl, H., Blom, D., Gatzhammer, B., Van Zuijlen, A.: Parallel coupling numerics for partitioned fluid–structure interaction simulations. *Comput. Math. Appl.* **71**(4), 869–891 (2016)
2. Totounferoush, A., Pour, N.E., Schröder, J., Roller, S., Mehl, M.: A new load balancing approach for coupled multi-physics simulations. In: 2019 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), pp. 676–682. IEEE, Piscataway (2019)

EXAHD

1. Obersteiner, M., Hinojosa, A.P., Heene, M., Bungartz, H.J., Pflüger, D.: A highly scalable, algorithm-based fault-tolerant solver for gyrokinetic plasma simulations. In: Proceedings of the 8th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems, pp. 1–8 (2017)

2. Hupp, P., Heene, M., Jacob, R., Pflüger, D.: Global communication schemes for the numerical solution of high-dimensional PDEs. *Parallel Comput.* **52**, 78–105 (2016)

ExaSolvers

1. Benedusi, P., Garoni, C., Krause, R., Li, X., Serra-Capizzano, S.: Space-time FE-DG Discretization of the anisotropic diffusion equation in any dimension: the spectral symbol. *SIAM J. Matrix Anal. Appl.* **39**(3), 1383–1420 (2018)
2. Kreienbuehl, A., Benedusi, P., Ruprecht, D., Krause, R.: Time-parallel gravitational collapse simulation. *Commun. Appl. Math. Comput. Sci.* **12**(1), 109–128 (2015)

ExaStencils

1. Köstler, H., Schmitt, C., Kuckuk, S., Kronawitter, S., Hannig, F., Teich, J., Rüde, U., Lengauer, C.: A scala prototype to generate multigrid solver implementations for different problems and target multi-core platforms. *Int. J. Comput. Sci. Eng.* **14**(2), 150–163 (2017). <https://doi.org/10.1504/IJCSE.2017.082879>
2. Schmitt, C., Kronawitter, S., Hannig, F., Teich, J., Lengauer, C.: Automating the development of high-performance multigrid solvers. *Proc. IEEE* **106**(11), 1969–1984 (2018)

ExtraPeak

1. Shudler, S., Calotoiu, A., Hoefler, T., Wolf, F.: Isoefficiency in practice: configuring and understanding the performance of task-based applications. In: *ACM SIGPLAN Notices*, vol. 52, pp. 131–143. ACM, New York (2017)
2. Calotoiu, A., Hoefler, T., Poke, M., Wolf, F.: Using automated performance modeling to find scalability bugs in complex codes. In: *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, p. 45. IEEE, Piscataway (2013)

FFMK

1. Weinhold, C., Lackorzynski, A., Härtig, H.: FFMK: an HPC OS based on the L4Re microkernel. In: *Operating Systems for Supercomputers and High Performance Computing*, pp. 335–357. Springer, Berlin (2019)
2. Gholami, M., Schintke, F.: Multilevel checkpoint/restart for large computational jobs on distributed computing resources. In: *IEEE 38th Symposium on Reliable Distributed System (SRDS)* (2019)

GROMEX

1. Beckmann, A., Kabadshow, I.: Portable node-level performance optimization for the fast multipole method. In: *Recent Trends in Computational Engineering-CE2014*, pp. 29–46. Springer, Berlin (2015)
2. Kutzner, C., Páll, S., Fechner, M., Esztermann, A., de Groot, B.L., Grubmüller, H.: More bang for your buck: Improved use of GPU nodes for GROMACS 2018. *J. comput. chem.* **40**(27), 2418–2431 (2019)

MYX

1. Protze, J., Tsuji, M., Terboven, C., Dufaud, T., Murai, H., Petiton, S., Emad, N., Müller, M., Boku, T.: Myx—runtime correctness analysis for multi-level parallel programming paradigms. In: *Software for Exascale Computing: SPPEXA 2016–2019. Lecture Notes in Computational Science and Engineering*. Springer, Berlin (2020)
2. Protze, J., Schulz, M., Ahn, D.H., Müller, M.S.: Thread-local concurrency: a technique to handle data race detection at programming model abstraction. In: *Proceedings of the 27th International Symposium on High-Performance Parallel and Distributed Computing*, pp. 144–155 (2018)

Terra-Neo

1. Bauer, S., Huber, M., Ghelichkhan, S., Mohr, M., Rüde, U., Wohlmuth, B.: Large-scale simulation of mantle convection based on a new matrix-free approach. *J. Comput. Sci.* **31**, 60–76 (2019)
2. Huber, M., Gmeiner, B., Rüde, U., Wohlmuth, B.: Resilience for massively parallel multigrid solvers. *SIAM J. Sci. Comput.* **38**(5), S217–S239 (2016)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

