


Inhalt

Editorial	2
Iterationsschleife	5
New Report on Research and Education in Computational Science and Engineering	7
The Coupling Library preCICE	11
Summer of Simulation 2016	21
Student Cluster Competition 2016: - Team TUM	28
- Team FAU	30
The Supercomputing Conference at Salt Lake City	36
BGCE Absolventenfeier 2016	33
SPPEXA News - Workshops 2017	38
Gastprofessor Marco Donatelli von der Universita Dell'Insubria in Como	39
High-Dimensional Approximation-what?	41
PARNUM 2017 - Call for Contributions	43



Quartl* - Impressum

Herausgeber:

Prof. Dr. A. Bode, Prof. Dr. H.-J. Bungartz, Prof. Dr. U. Rüde

Redaktion:

S. Herrmann, Dr. S. Zimmer

Technische Universität München, Fakultät für Informatik
Boltzmannstr. 3, 85748 Garching b. München

Tel./Fax: +49-89-289 18636/18607

e-mail: hermasa@in.tum.de, **www:** <http://www5.in.tum.de/quartl>

Redaktionsschluss für die nächste Ausgabe: **01.03.2017**

* Quartl: früheres bayerisches Flüssigkeitsmaß,

→ das Quartl: 1/4 Kanne = 0.27 l

(Brockhaus Enzyklopädie 1972)

Das Quartl erhalten Sie online unter <http://www5.in.tum.de/quartl/>



Das Quartl ist das offizielle Mitteilungsblatt des Kompetenznetzwerks für
Technisch-Wissenschaftliches Hoch- und Höchstleistungsrechnen in Bayern
(KONWIHR) und der Bavarian Graduate School of Computational Engineer-
ing (BGCE)

Editorial

Erfreulicherweise bekomme ich aus der Quartl-Leserschaft ja des öfteren einen Hinweis, ein Dokument oder einen Link zugespielt – offensichtlich verbunden mit der freudigen Erwartung, dass das dann auch im Quartl zum Zuge kommen möge. Sehr oft klappt das, denn naturgemäß erfüllen die mir so zugespielten Dinge zumeist alle meine Anforderungen an prachtvollle Sottisen. Diesmal fand ein solches Juwel gleich auf vier verschiedenen Wegen zu mir – fraglos ein Zeichen, eine höhere Fügung, der man sich bekanntermaßen tunlichst nicht in den Weg stellen sollte.

Es geht um eine Email-Nachricht aus Österreich, präziser vom dortigen Business Circle. Falls jemand jetzt schon lachen sollte – halt, die Pointe kommt doch erst noch! Also schnell gegoogelt, und schon landet man auf einer bunten Webseite mit der Begrüßung „*Willkommen bei Business Circle! Im Kreis der Spitzenvertreter aus Wirtschaft, Wissenschaft und Politik. Die Nr. 1 bei Konferenzen in Österreich. Seit 1994 Ihr Partner für Ihre Pole Position.*“ Das reicht mir schon, um mich hinreichend bedeutend zu fühlen (das erinnert ja fast an die unzähligen Emails „*Dear esteemed professor ... your renowned institution ...*“ aus Indien et al.; aber es reicht auch, um diesen geschäftigen Kreis richtig einordnen zu können – es riecht förmlich nach hohem Sottisen-Potenzial. Also die Webseite schnell beiseite gelegt und die Email mit dem interessantesten Titel „*Vertrags-Know-how für Nicht-Juristen*“ angeschaut.

Da heißt es: „*Sehr geehrter Herr Professor Bungartz, Sie sind kein Jurist, haben aber laufend mit Verträgen zu tun? Wir machen Sie sattelfest: In unserem Intensiv-Seminar lernen Sie von der aktuellen Rechtsprechung über heikle Inhalte bis hin zu den 5 Todsünden im Vertragsmanagement alles, was Sie für eine optimale Vertragsgestaltung wissen müssen.*“

Zunächst einmal liegt die Absenderin, eine Frau Schwägerl, mit ihrer Einstiegsfrage richtig: Ja, ich bin kein Jurist (ich weiß nicht, ob ich dafür dem Herrgott, meinen Eltern, meiner Schule oder mir selbst danken muss – meine Dankbarkeit in alle Richtungen ist jedenfalls gigantisch), und ja, ich habe zumindest öfter mit Verträgen zu tun, als mir lieb ist. Lässt das auf einen Volltreffer schließen? Mitnichten!

Erstens würde man auch einen Einstieg der Art „*Sie sind kein Kfz-Mechaniker, sitzen aber laufend im Auto?*“, oder „*Sie sind kein Lungenarzt, atmen aber laufend durch Ihre Lungen?*“ verwenden können, um jemandem ein Seminar zu Autoreparatur oder einen Steilkurs „Lungentransplantation für jedermann“ zu verkaufen. Zweitens liegt meine Frustration bei Verträgen im Allgemeinen weniger in den Verträgen an sich als vielmehr in den durch Vertragsentwürfe magisch angezogenen Juristinnen und Juristen begründet. Also alles andere als ein Volltreffer.

Daher in aller Deutlichkeit: Schlimm genug, dass ich mich mit Verträgen und den um diese herum scharwenzelnden Juristen herumärgern muss – die Vorstellung, da sattelfest zu werden, beunruhigt mich extrem. Am Ende könnte ich mich gar nicht mehr über Juristen ärgern, sondern würde so etwas wie Verständnis für sie empfinden?

Überhaupt Juristen – warum bloß sind sie so, wie sie sind? Und warum sind manche von denen, mit denen ich es zu tun habe, nicht im Ansatz so wie die aus diversen Filmen bekannten? Im Kino machen sie Dinge möglich, die niemand für möglich gehalten hätte; im echten Leben verhindern sie Dinge, die jeder und jede für ganz einfach gehalten hätten. Im Kino nutzen sie jeden vorhandenen (oder auch nicht vorhandenen) Spielraum zur Interpretation des Gesetzestexts, manchmal über die Grenzen des Sauberen hinweg – schließlich sollen die eigenen Interessen durchgeboxt werden; im echten Leben sind sie erschreckend phantasielos, kommentieren ein „XYZ ist nicht möglich“ in irgendeinem Gesetz oder irgendeiner Verordnung doch glatt mit dem nun wirklich mehr als einfallslosen „XYZ geht nicht“.

Aber irgendwie komme ich vom Thema ab. Der Business Circle mit seinen tollen Angeboten steht doch im Fokus, nicht irgendwelche Juristen. Also schauen wir uns dessen Portfolio noch etwas genauer an. Ein Seminarartikel lautet „*Persönliche Haftung als Geschäftsführer und Vorstand vermeiden*.“ Ja, ja – uns etwas von unternehmerischen Risiken und gigantischer Verantwortung erzählen, um einen satten Bonus einzustreichen; aber eine Risikovermeidungsstrategie fahren, die jedem Beamten zur Ehre gereichen würde.

Sie sehen schon: Der Business Circle wird ohne mich als Kunden, und ich werde ohne seine Unterstützung bei meiner Pole Position auskommen müssen. Mit und ohne Juristen!

Doch genug der einleitenden Worte – die gesamte Quartl- Redaktion wünscht Ihnen allen Frohe Weihnachten, ein paar geruhsamere Tag zwischen den Jahren und dann einen vollendeten Start in ein rundum erfolgreiches Jahr 2017! Und zunächst wünschen wir vor allem viel Spaß mit dieser neuen Ausgabe Ihres Quartls!

Hans-Joachim Bungartz

Important dates:

- January 13, 2017
- January 30, 2017
- February 6, 2017
- March 13, 2017

Invited speakers:

- Iain Duff
- Katarina Gustavsson
- Selime Girol
- Lois Curfman McInnes
- Linda Stals

Program chairmen:

- Ulrich Rüdte
- Marian Vajtersic

Program committee:

- Hans-Joachim Bungartz
- Dietmar Fey
- Gundolf Haase
- Axel Klawonn
- Günter Leugering
- Miriam Mehl
- Gabriel Oksa
- Roman Trobec
- Pavel Tvrdik
- Roman Wyrzykowski

Organizing committee:

- Dominik Bartuschat
- Julia Deserno

For further information, registration, and abstract submission, please see our workshop website www.parnum2017.fau.de. In case of any further questions, please do not hesitate to email us at conference-parnum2017-orga@fau.de.

We are looking forward to welcoming you at the PARNUM 2017.

Dominik Bartuschat

Iterationsschleife

N=21

25. November 2016

Zoppo Trump ist der Gegenspieler von Kalle Wirsch. Durch List und Betrug ist es ihm zunächst gelungen, Kalle Wirsch zu besiegen bzw. kalt zu stellen. Unterstützt wird er dabei von einer Ratte und einer Spinne. Kalle Wirsch gelingt es aber, sich mit Hilfe zweier Kinder zu befreien und schlusslich Zoppo Trump im fairen Wettbewerb zu bezwingen. So beschreibt Tilde Michels 1969 die Geschichte, und die Augsburgsburger Puppenkiste hat das ganze im ARD im November 1970 dargestellt. 46 Jahre ist es also her, dass ein Wirsch dem Treiben des bösen Trump ein Ende machte.

2016 sieht die Welt ein wenig anders aus. Ein Trump ist amerikanischer Präsident. Der Wahlkampf der vorausging wird übereinstimmend als Tiefpunkt demokratischer Auseinandersetzungen in den USA angesehen, und der zukünftige Präsident machte sich einen Namen durch interessante Aussagen über politische Themen, deren Zitierung das Erscheinen dieses Textes womöglich problematisch machen könnte. Fast scheint es, als sei hier ein Märchen zu einem bösen Ende gekommen. Gewonnen – so scheint es – hat derjenige Kandidat, der durch Untergriffe aller Art – auch solche aus früheren Jahren sowie sein eigenes Sprechen darüber – der Gegnerin letztlich die Schneid abgekauft hat. Zurück bleibt ein Land in dem Männer mit Begeisterung davon erzählen, dass sie diesmal Trump gewählt haben und dass ihre eigenen Frauen darüber geweint und/oder getobt haben.

Betrachtet man das Ganze etwas genauer, so ist die Sache ein wenig komplizierter und gleichzeitig ein wenig einfacher. Zunächst zum Einfachen. Tatsächlich ist Trump nicht mehr als er behauptet zu sein: ein erfolgreicher Unternehmer der einmal politisch auf den Putz hauen wollte. Vom Format eines Berlusconi – dessen Verflechtungen aus Firmen, Politik, Geschäften und Sex legendär sind – ist er weit entfernt. Aber auch vom Österreicher H.C. Strache, der ein klares sehr rechtes Profil präsentiert, ist Trump weit weg. Trump ist also – wie das Erdmännchen Zoppo Trump – einfach Trump - offen vulgär sozusagen als politisches Credo.

Das Komplizierte dabei ist: das Volk hat offenbar einfach nur noch das Gegenteil von dem gewollt was es bislang von Washington bekommen hat. Aufzudröseln aber, was denn das Volk durch diese Wahl möglicherweise noch wollte – was also die „positive“ Botschaft sein sollte – ist erheblich schwieriger bzw. unmöglich.

Diskussionen zeigen, dass das alles Mögliche sein kann. Und im Grunde ergibt sich, dass Trump nicht gewählt wurde weil er das Gegenteil von dem tut oder sagt was Washington tut oder sagt sondern weil er den Eindruck erweckt, er würde das Gegenteil dann schon tun wenn es denn so weit wäre. Gewählt wurde also die Hoffnugn auf das Drinschlagen und noch nicht das Dreinschlagen selber - der Schein von politischem Weltuntergang aber nicht der politische Weltuntergang selber. Womit wir beim Kern des Problems sind. Ganz offensichtlich braut sich für die demokratischen Regierungen dieser Welt etwas zusammen, was diese – und ihre politikwissenschaftlichen Berater - in ihre pauschalen und normierten Konzepte von Politik nicht mehr integrieren und damit nicht mehr fassen können. Bleibt man dabei in Deutschland, so wechselt der Wähler gar nicht von einer zur anderen Partei (wie die Logik der Politikwissenschaft erwartet) sondern er verschwindet zunächst in einem großen Topf der Nichtwähler. Dieser wurde lange ignoriert, denn Nichtwähler können auch nichts „Falsches“ wählen. Genau aus diesem Topf aber kommen nun die Wähler, die „das Gegenteil“ von allem was in X passiert, haben wollen. Und so bilden sich im politischen System Parteiblasen aus, die – je nach Klugheit des politischen beteiligten Personals – manchmal schneller (Piraten) und manchmal langsamer (AfD?) platzen oder vielleicht auch gar nicht platzen. Wesentliches Merkmal ist die Unklarheit der Forderungen und die Heterogenität des Personals. Beides sind lebenswichtige Faktoren für die politische Blase. Die Unklarheit erst erlaubt es dem Wähler, seine eigenen Hoffnungen, Ängste oder auch Vorurteile, auf die Blase zu projizieren und sich damit in dieser zu spiegel und mit dieser zu identifizieren. Und je schillernder und unfassbarer das Personal der Blase bleibt, umso höher die Chance, dass diese Projektionsfähigkeit der Blase nicht durch die Realität zerstört wird.

Einziges Problem für den designierten Präsidenten der USA: er muss nun 4 Jahre konkret arbeiten. Mit jeder Entscheidung aber, wird seine Blase konkret als solche erkennbar. Die Gegenmittel sind die weitere Ausdehnung der Blase oder das Platzen. Beides aber wünscht man sich nicht. So geht ein Hilferuf an die Physiker dieser Welt: wie hält man eine Blase 4 Jahre stabil so dass sie trotzdem wirkt als würde sie wachsen? Am Ende muss es doch wieder die Naturwissenschaft richten.

M. Resch

PARNUM 2017 - Call for Contributions

PARNUM 2017, the 11th workshop in the series of "Parallel Numerics" workshops, will take place in Waischenfeld, Germany from Wednesday, April 19 until Friday, April 21, 2017. The PARNUM series started in 1994 and from that date, ten workshops were held in Austria, Slovakia, Slovenia, and Poland.

The objective is the exchange of research results in the area of parallel scientific computing, parallel algorithms, and high performance computing.

Specific topics of interest include, but are not limited to:

- Large-scale parallel applications
- Scalable parallel algorithms
- Efficient numerical methods
- Parallel computer architectures
- Performance analysis, tuning, and debugging
- Fault-tolerance
- Scientific workflows

Abstracts (of at most one page in length) for contributed talks should be submitted before Friday, January 13, 2017 via www.parnum2017.fau.de.

The estimated number of PARNUM 2017 participants is 30 to 40.

For participants registering before Monday, February 6, 2017 a reduced early bird registration rate applies. This workshop will include an excursion to Bamberg, whose medieval old town is a UNESCO World Heritage Site, and which is also famous for its varied beer brewing tradition.



Figure 1: Markus Hegland received the TUM Ambassador certificate.

During the award presentation to Markus there was a funny mishap: the name of the research topic of Markus – high-dimensional approximation – just refused to roll of the tongue of Madam Vice President. We can only speculate about whether she was thinking about “approbation” or “appropriation” at the moment.

On Monday Markus took the flight back to summerly Australia to give a talk at the CTAC conference just a day later. But now, officially, as a TUM Researcher Alumni representative.

Valeriy Khakhutskyy

New Report on Research and Education in Computational Science and Engineering

Over the past two decades, computational science and engineering (CSE) has become an increasingly important part of research in academia, industry, and laboratories. Mathematics-based advanced computing is now a prevalent mean of discovery and innovation in essentially all areas of science, engineering, technology, and society, and the CSE community is at the core of this transformation. SIAM has been a driving force in this development by hosting the activity group on Computational Science and Engineering (SIAG/CSE), organizing the biennial flagship SIAM CSE conference, and publishing the SIAM Journal on Scientific Computing – one of the top journals in the field. A 2001 report, Graduate Education in Computational Science and Engineering¹ by the SIAM Working Group on CSE Education, has helped to define the role and scope of CSE during the past two decades. However, a combination of disruptive developments - including the architectural complexity of extreme-scale computing, the planet’s ongoing data revolution, and the penetration of mathematically-based CSE methodology into more and more fields - is currently redefining CSE’s reach.

While CSE is rooted in the mathematical and statistical sciences, computer science, the physical sciences, and engineering, today it increasingly pursues its own unique research agenda. The field is now widely recognized as an essential cornerstone that drives scientific and technological progress in conjunction with theory and experiment. Scientific experimentation and theory, the classical paradigms of the scientific method, both strive to describe the physical world. However, high-fidelity predictive capabilities can often be achieved only by numerical computation. CSE’s overarching goal of achieving truly predictive scientific capabilities is its distinguishing factor. It accomplishes this through advances that combine modeling, numerical analysis, algorithms, simulation, big data analytics, high performance computing, and scientific software. The development of predictive capabilities lies at the core of CSE as a new discipline in its own right and has already impacted

¹<http://epubs.siam.org/doi/pdf/10.1137/S0036144500379745>



Figure 1: CSE at the intersection of mathematics and statistics, computer science, and core disciplines from the sciences and engineering. This combination gives rise to a new field whose character is different from its original constituents.

a number of disciplines, including but not limited to simulation-based design in the automotive industry, simulation-based decisions in computational medicine, and simulation-based predictions of global climate. It is also set to catalyze fundamental changes in many more areas of technical, economic, societal, and political decision processes.

A new report, titled Research and Education in Computational Science and Engineering,² analyzes the current status of CSE and the aforementioned new developments. The report, available in preprint form and from the SIAG/CSE wiki page,³ summarizes the status of CSE as an emerging discipline and presents the field's trends and challenges in research and education for the next decade. The report is based on the outcomes of a 2014 workshop sponsored by SIAM and the European Exascale Software Initiative, a minisymposium, and a panel discussion held during the 2015 SIAM CSE conference, as well as feedback from the CSE community collected over the

²<https://arxiv.org/abs/1610.02608>

³http://wiki.siam.org/siag-cse/index.php/Main_Page

High-Dimensional Approximation-what?

“Excellence attracts excellence, and flaschen attract flaschen” as recently stated by TUM President Prof. Wolfgang A. Herrmann. And so it is no coincidence that when the committee considered the most outstanding candidates to represent the TUM Researcher Alumni in the international arena, they inevitably decided to bestow the prestigious title of a “TUM Ambassador” to our friend and IAS alumnus Prof. Markus Hegland from the Australian National University. His TUM Excellency Prof. Hegland re-visited Munich in the week from November 21st to 27, to receive this honorable title.

Markus Hegland is our long-time collaborator. In 2010 he became a Hans Fisher Senior Fellow at the TUM Institute for Advanced Study (IAS) in the Focus Group “High-Performance Computing” hosted by Prof. Hans-Joachim Bungartz. As Hans Fisher Senior Fellow, he supervised two PhD students, Christoph Kowitz and me. And it is not for nothing that it says “once a Fellow – always a fellow” as Markus continues to visit the chair for scientific collaboration and to receive guests from Munich (cf. Quartl 75), although his IAS fellowship already ended in 2014.

And also this time Markus’s week-long stay in Munich was filled not only with *glühwein* and Christmas markets. During the week he examined the PhD defense of the author and participated in a panel discussion of TUM Junge Akademie on the opportunities for internationalization of young and aspiring scientists. The event program reached the apogee on Sunday in the Munich philharmonic hall where, as a part of the Vivat TUM concert, Markus received the TUM Ambassador certificate from President Herrmann and the Senior Vice President for International Alliances Dr. Hannelore Keidel.

Die Verbindung TUM-Como hat schon zu einigen wechselseitigen Besuchen und Blockkursen geführt. So war vor zwei Jahren Professor Stefano Serra aus Como bereits Gast in Garching. Seine ehemalige Doktorandin Mariarosa Mazza hat inzwischen eine Postdoc-Stelle am IPP Garching inne; dabei geht es um die Analyse und effiziente Lösung von linearen Gleichungssystemen, die bei der Verwendung von isogeometrischen Elementen entstehen.

Als kulturelles Beiprogramm bot sich in der Adventszeit ein gemeinsames Plätzchenbacken und ein Besuch des Weihnachtsoratoriums im Herkulesaal an.

Thomas Huckle




Figure 2: CSE pipeline, from physical problem to model and algorithms to efficient implementation in simulation software with verification and validation driven by data. The pipeline is actually a loop that requires multiple feedbacks.

past two years. Despite CSE’s fundamental importance, the report finds that many current institutional structures do not adequately reflect the needs of the discipline. Examples of barriers preventing CSE advancement include a dearth of appropriate interdisciplinary structures at universities and funding institutions, lack of recognition for the important role of scientific software, and institutional challenges in creating suitable educational programs. The new report elaborates on these arguments in detail and reveals the following central findings: 1. CSE has matured to a discipline in its own right. 2. Computational algorithms lie at the core of CSE progress, and scientific software, which codifies and organizes algorithmic models of reality, is the primary means of encapsulating CSE research to enable advances in scientific and engineering understanding. 3. CSE methods and techniques are essential to capitalize on the rapidly-growing ubiquitous availability of scientific and technological data. The report also highlights a number of specific CSE “success stories” – application examples in which CSE research is significantly impacting the real world.

These accounts emphasize both the long-term payoff of investment in fundamental CSE research and the criticality of sustaining that investment to leverage current and future opportunities -as articulated in the report's recommendations - for CSE research and education over the next decade.

This text is being republished from the December 2016 issue of SIAM News, see <https://snews.siam.org/TabId/919/ArtMid/2285/ArticleID/1772/New-Report-on-Research-and-Education-in-Computational-Science-and-Engineering.aspx>

Ulrich Rüde

Ulrich Ruede, Karen Willcox, Lois Curfman McInnes, Hans De Sterck
2013-2014 Officers of the SIAM Activity Group on Computational Science and Engineering (SIAG/CSE)

Gastprofessor Marco Donatelli von der Universita Dell’Insubria in Como

Professor Marco Donatelli war vom 28.11. bis 9.12. als Gast am Lehrstuhl für Wissenschaftliches Rechnen der TUM-Informatik-Fakultät. Finanziert wurde der Aufenthalt dankenswerter Weise vom Bayerischen Staatsministerium für Bildung und Kultus, Wissenschaft und Kunst. Während seines Aufenthaltes hielt Herr Donatelli einen Blockkurs ”Structured matrices, Multigrid, and Image Processing”, der sich speziell an BGCE-Studierende richtete, aber darüber hinaus auch 19 Teilnehmer aus unterschiedlichen Studiengängen angelockt hat. Außerdem gab Herr Donatelli einen Vortrag im IPP.




Figure 1: Marco Donatelli

Marco Donatelli hat in Florenz Computer Science studiert. Den PhD in Computational and Applied Mathematics erwarb er 2006 in Mailand. Seit 2014 ist er Associate Professor of Numerical Analysis in Como. Er ist ein sehr aktiver Forscher auf den Gebieten strukturierte Matrizen, Multigrid Verfahren, Bildverarbeitung, insbesondere Inverse Probleme und Regularisierungsmethoden. Neben seinen Kontakten zur TUM arbeitet er zusammen mit Martin Hanke-Bourgois in Mainz, Lothar Reichel an der Kent State University, James Nagy an der Emory University sowie mit Tony und Raymond Chan, beide in Hong Kong. Weiterhin leitet er das Graduiertenprogramm der naturwissenschaftlichen Fakultät seiner Universität.

SPPEXA News – Workshops 2017

Der November 2016 Workshop-Call von SPPEXA traf auf großes Interesse der Projekte. Damit können wir 2017 und somit die zweite Phase von SPPEXA mit einigen spannenden Workshops einläuten. Bei dieser Gelegenheit möchte ich auch an das Annual Plenary Meeting erinnern. Die Anmeldung dazu ist auf www.sppexa.de freigeschaltet.

	Garching	Annual Plenary Meeting
20.-22.03.2017	Heidelberg	Fast High Order DG Methods for Future Architecture
23.-24.03.2017	Hamburg	Understanding I/O Performance Behavior (UIOP)
06.04.2017	Tokyo	Parallel Programming Models – Productivity and Applications
25.-28.05.2017	Göttingen	Exascale Data Generation and Analysis for MD Simulation
26.-28.06.2017	Lugano	Exa-scale Solver for Application-Driven Science (@PASC)
26.-27.09.2017	Hamburg	Exascale I/O for Unstructured Grids

Benjamin Uekermann

The Coupling Library preCICE Past, Present and Future

preCICE is a library that enables black-box coupling of existing simulation software. The standard example is the coupling of a computational fluid dynamics solver and a computational structural mechanics solver to perform a fluid-structure interaction (FSI) simulation. The concepts of preCICE apply, however, also for more general multi-physics applications. The library offers the user three groups of features: equation coupling schemes for acceleration and stabilization of iterations between different solvers, methods for communication between separated executables, and data mapping schemes for the interpolation between non-matching discretizations. In this article, I give an overview on the historical development of preCICE from first ideas to an efficient, parallel and usable software.

1 Introduction

The development of preCICE⁴ is driven by three design goals.

1. The easy reuse of existing simulation software. Therefore, the user should be able to integrate the application programming interface (API) into an existing simulation software in a minimally invasive way, without altering the structure of the coupled code.
2. The coupling complexity should be hidden from the possibly non-experienced user, to allow for a focus on the physical results.
3. The coupling should be efficient enough for the computational effort to hide completely behind the costs of the solvers – also on massively parallel systems.

⁴ preCICE is open-source under the LGPL3 license. For more information, please visit <http://www.precice.org>. The official reference of preCICE is [1].

The first development started with the predecessor of preCICE, FSI*ce, and focused on the first two goals. To achieve highest flexibility for the code integration, a pure black-box coupling was considered. Black-box means that the coupling solely works with nodal values of both solvers at the common coupling interface. In particular, no internal discretization details of the two solvers should be used. This restriction, of course, complicates the numerical coupling, but it increases the flexibility drastically. Markus Brenk started the code development in 2003 as part of the DFG research unit *Fluid-Structure Interaction: Modelling, Simulation, Optimization*⁵ FSI*ce was a server-based coupling tool [2], meaning that a central instance, the server, takes responsibility of all three coupling feature groups – the equation coupling, the communication, and the data mapping – cf. Figure 1 a). Although the layout has obvious efficiency drawbacks, the design allowed to hide both solvers from each other and focus on the, at the time, still underdeveloped equation coupling and data mapping schemes.




Figure 1: Various coupling layouts. a), Server-based layout as used by FSI*ce [2], b), peer-to-peer layout as used by the original preCICE [3], c), fully parallel layout as used by the current preCICE [4].

Bernhard Gatzhammer re-developed the server-based concept of FSI*ce into the peer-to-peer concept of preCICE from 2008 on [3]. The coupling logic needed no central instance any longer. Still, for parallel solvers, a server thread was used per solver as a proxy, cf. Figure 1 b). Furthermore, the API of preCICE operated on a higher, more abstract level as the original *send-receive* API from FSI*ce, allowing for a higher flexibility. In section 2, I briefly introduce the API of preCICE.

⁵<http://fsw.informatik.tu-muenchen.de>




Figure 1: The Salt Lake Temple located on Temple Square in Salt Lake City, Utah.

Bakhtiari, doctoral candidate at the SCCS chair, co-authored with Prof. Hans-Joachim Bungartz, Prof. Miriam Mehl and Prof. George Biros, was accepted at the technical program. Arash Bakhtiari attended the conference at Salt Lake City and presented the results of his PhD work in the HPC community.

Arash Bakhtiari

The Supercomputing Conference at Salt Lake City

International Conference for High Performance Computing (HPC), Networking, Storage and Analysis (SC16) is one of the most distinguished conferences in HPC community. This year, the SC conference was held from 13th till 18th November in Salt Palace convention center in Salt Lake City, Utah.

Salt Lake City has been hosting many scientific computing events like *SIAM Conference on Computational Science & Engineering* at 2015. The city is the capital of the U.S. state of Utah with an estimated population of 190,000. It was founded in 1847 by *Mormon* followers. Since the headquarters of The Church of Jesus Christ of Latter-day Saints (LDS Church) and Salt Lake Temple (Figure 1), are located in this city, it is considered a holy city by members of the LDS church.

The SC conference has a diverse community of participants: researchers, scientists, application developers, agency program managers and journalists from all over the world participate in the conference. The conference consists of several components like technical program, awards, exhibits, student program and SCInet.

The SC conference hosts several society awards like *ACM Gordon Bell Prize* and *ACM/IEEE-CS George Michael Memorial HPC Fellowships*. The Gordon Bell prize was established in 1987 and recognizes outstanding achievement in high-performance computing applications. Since 2011, a cash award of \$10,000, funded by Gordon Bell, accompanies the recognition. The George Michael (one of the founding fathers of the SC Conference series) Memorial Fellowship honor exceptional PhD students throughout the world who has conducted research in HPC area.

The technical program is the heart of SC. Its invited talks, panels, research papers, tutorials, workshops, posters, and Birds of a Feather (BoF) sessions have inspired new and innovative areas of computing. The technical program with an acceptance rate around 20% for papers is very competitive. The paper "A Parallel Arbitrary-Order Accurate AMR Algorithm for the Scalar Advection-Diffusion Equation" submitted to the SC conference by Arash

In a collaborative effort, we rendered the server processes completely unnecessary by introducing a fully parallel concept [4], cf. Figure 1 c), as part of the first phase of the SPPEXA ExaFSA project⁶ from 2013 on. The feature groups are now executed on distributed data directly on the solver ranks. A point-to-point communication scheme allows for high scalability, finally achieving the third design goal of preCICE. With this last development step, the range of coupled simulations could be increased from moderately to massively parallel applications, marking a significant difference in many fields.

Over the last decades, other coupling libraries with a similar focus as preCICE appeared, such as DTK, EMPIRE, MpCCI, OASIS3, or OpenPALM. Nevertheless, the combination of the high-level API, the fully parallel layout, and the sophisticated equation coupling schemes marks a unique selling point of preCICE. In the next section, I briefly introduce the reader to the features and the API of preCICE. Afterwards, Section 3 lists some application examples. I close the discussion by drawing future lines in Section 4.

2 User Perspective of the Library

Figure 2 sketches the current layout of preCICE. In the following, I first give an overview of the features of preCICE, followed by an exemplary study of the library's API.

2.1 Features

Coupling schemes can be serial or parallel, and explicit or implicit. preCICE allows for both choices at runtime. Here, serial refers to a sequential execution order of the solvers, one after the other, parallel to a simultaneous execution. Explicit schemes compute a fixed amount of solver steps in one timestep, while implicit schemes sub-iterate until convergence. Implicit schemes are, in particular, necessary if the scenario features a strong interaction, which needs to be stabilized. For FSI, this corresponds to a high added-mass effect.

⁶<https://ipvs.informatik.uni-stuttgart.de/SGS/EXAFSA>

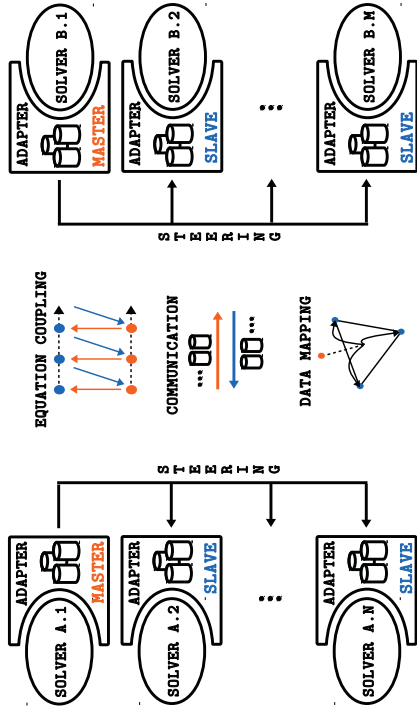


Figure 2: preCICE in a nutshell. Two parallel solvers A and B, on N and M ranks, are coupled via preCICE. The three feature groups, equation coupling, communication, and data mapping, are executed on distributed data directly on the solver ranks. No central instance is necessary.

For such a situation, preCICE offers acceleration schemes, such as a simply (dynamic) underrelaxation, but also sophisticated quasi-Newton schemes as the IQN-ILS scheme (also known as Anderson acceleration) and the IQN-MVJ scheme (also known as generalized Broyden).

The library offers three basic variants for communication: MPI, TCP/IP sockets, implemented with `Boost.Asio`, and communication via files. For MPI, start-up in a common communicator is similarly supported as start-up in separated communicators. preCICE can use every basic variant to build up a global point-to-point communication, meaning that only the ranks of neighboring sub-domains of the coupled solvers communicate with each other.

For the data mapping, projection-based mapping methods and radial basis function (RBF) interpolation are supported. All mapping schemes come in two flavors: a consistent variant, guaranteeing the exact mapping of constant values, and a conservative variant, guaranteeing the conservation of integral



Abbildung 3: Interne Absolventenfeier COME & CSE: Gastredner Prof. Frank Jenko (links) und Verabschiedung von Christoph Riesinger (rechts).

Mehrere großartige studentische Beiträge boten Gelegenheit das zurückliegende Studium Revue passieren zu lassen. So wurde neben Anekdoten aus dem Studentenleben die Gemeinschaft und die Internationalität der Programme positiv herausgestellt. Außerdem bot die interne Absolventenfeier die Möglichkeit den langjährige CSE-Koordinator Christoph Riesinger feierlich zu verabschieden. Prof. Hans-Joachim Bungartz und Alfredo Parra würdigten seine Leistungen und bedankten sich mit den besten Wünschen für die Zukunft und einem kleinem Geschenk.

Beide Absolventenfeiern boten somit einen angenehmen Rahmen für eine eigentlich unangenehme Sache – den Abschied von vielversprechenden Studierenden und netten Kollegen, die zu neuen Ufern aufbrechen – und helfen uns, das lachende mehr als das weinende Auge in Erinnerung zu behalten.

Michael Rippl, Tobias Neckel

Insgesamt stellte die Absolventenfeier eine gelungene Veranstaltung in bewährtem Format dar. Das ersetzt natürlich nicht die eigene, familiärrere Absolventenfeier, die z.B. die Studiengänge “Computational Mechanics” (COME) und “Computational Science and Engineering” (CSE) zusammen abhalten. Hier ist auch etwas mehr Platz für Begleitpersonen der Absolventen und fachlichen Austausch. Dieses Jahr fand diese Feier am 24. November in München mit CSE als Gastgeber statt.



Abbildung 2: Interne Absolventenfeier COME & CSE: Absolventen von COME (links) und Absolventen von CSE (rechts).

Gastredner Prof. Frank Jenko (Department of Physics & Astronomy, University of California) begeisterte die Zuhörer dabei mit einem engagierten und interessanten Vortrag über Kernfusion und stellte neben den aktuellen und interessanten Entwicklungen auch die wichtige Rolle der Simulation dabei heraus. So ist bei der Kontrolle der Kernfusion der Einschluss des Plasmas eine große Herausforderung, wobei vor allem die Turbulenz im Inneren des Plasmas problematisch ist. Diese kann durch geeignete Methoden simuliert werden und daraus kann wiederum die Anordnung von Magnetspulen zum Einschluss des Plasmas errechnet werden. Dabei spielen die Fortschritte im Bereich High-Performance-Computing eine große Rolle, da sich damit immer komplexere Szenarien sowohl für das Plasma als auch für die Magnetfelder simulieren lassen.

values. The projection-based mapping methods comprise a nearest-neighbor mapping and a nearest projection mapping. While nearest-neighbor is a first order method, nearest-projection is second order if projection distance from one mesh to the other is much smaller than the mesh width, which typically holds for practical applications. As RBF basis functions, local as well as global functions are supported, marking a trade-off between accuracy and efficiency.

2.2 Application Programming Interface

Figure 3 lists a code example for an existing standard fluid solver that is prepared for FSI coupling. The example uses the C++ interface of preCICE. Additionally preCICE offers interfaces in C, Fortran90/95, Fortran2003, and Python. The coupling is configured in line 3 by means of an xml file, where all coupling details can be specified at runtime. In line 12, the coupling interface mesh is defined as an array of nodal coordinates. The main steering functions of preCICE are `initialize` and `advance` in lines 23 and 30. These two functions apply all specified features, such as communication of coupling data, equation coupling schemes, and data mapping. preCICE uses no explicit sending and receiving functions, but the high level `advance` allows to decide at runtime whether the coupling should be serial or parallel. Instead, the user can access coupling data through buffers in lines 29 and 31. The code example of Figure 3 currently only allows for explicit coupling. Extending it for implicit coupling, however, is simple, and only skipped due to space restrictions.

Ready-to-use coupling adapters already exist for several codes as listed in Table 1.

3 Applications

preCICE has been used for numerous applications. Some of the highlights are an aortic blood flow, depicted in Figure 4, and a multi-fluid coupling depicted in Figure 5. The scalability has, for example, been tested by means

```

1 turnOnSolver(); //e.g. setup and partition mesh
2 precice::SolverInterface precice("FluidSolver", rank, size);
3 precice.configure("precice-config.xml");
4
5 int dim = precice.getDimension();
6 int meshID = precice.getMeshID("FluidMesh");
7 int vertexSize; // number of vertices at wet surface
8 // determine vertexSize
9 double* coords = new double[vertexSize]; // coords of vertices at wet surface
10 // determine coordinates
11 int* vertexIDs = new int[vertexSize];
12 precice.setMeshVertices(meshID, vertexSize, coords, vertexIDs);
13 delete[] coords;
14
15 int displID = precice.getDataID("Displacements", meshID);
16 int forceID = precice.getDataID("Forces", meshID);
17 double* forces = new double[vertexSize*dim];
18 double* displacements = new double[vertexSize*dim];
19
20 double dt; // solver timestep size
21 double precice_dt; // maximum precice timestep size
22
23 precice_dt = precice.initialize();
24 while (not simulationDone()){ // time loop
25   beginTimeStep(); // e.g. compute adaptive dt
26   dt = min(precice.getMaxDt(), dt);
27   computeTimeStep();
28   computeForces(forces);
29   precice.writeBlockVectorData(forceID, vertexSize, vertexIDs, forces);
30   precice_dt = precice.advance(dt);
31   precice.readBlockVectorData(displID, vertexSize, vertexIDs, displacements);
32   setDisplacements(displacements);
33   endTimeStep(); // e.g. update variables, increment time
34 }
35 precice.finalize();
36 delete[] vertexIDs, forces, displacements;
37 turnOffSolver();

```

Figure 3: Code example for an existing fluid solver, prepared for an FSI coupling via preCICE.

of an artificial coupling of two Euler domains, cf. Figure 6. For details of all three scenarios, I kindly refer the reader to [4].

4 Future Work

After discussing the past of preCICE in Section 1 and the present in Sections 2 and 3, I want to close this article with an outlook on the future of the coupling library. The future development angle of preCICE is twofold. First,

BGCE: (Absolventenfeier)²

Am 15. November 2016 fand die Absolventenfeier des Elitenetzwerks Bayern (ENB) statt, diesmal in der Residenz in München. In feierlichem Rahmen wurden die diesjährigen Absolventinnen und Absolventen des ENB verabschiedet. 27 Elitestudiengänge, 18 Internationale Doktorandenkollegs, Max-Weber-Stipendiaten und internationale Nachwuchsforschergruppen stellen die vier Säulen des ENB mit insgesamt etwas über 3000 Mitgliedern. Staatssekretär Bernd Sibler ließ es sich bei der Veranstaltung nicht nehmen, jeder Absolventin bzw. jedem Absolventen persönlich zu gratulieren.



Abbildung 1: BGCE-Absolventen mit Staatssekretär Bernd Sibler bei der Übergabe der ENB-Zertifikate.

Anschließend war zum Buffet und gemütlichem Beisammensein in den feierlichen Sälen der Residenz mit launiger Jazzmusik geladen. Dort konnten Kontakte mit anderen Absolventen geknüpft und mit den Betreuern über die zurückliegende Zeit im ENB resümiert werden.



Abbildung 2: v.l.n.r. Kai Streitferdt, Benedikt Öhrlich und Jan Laukemann an der Booth während der Competition.




Abbildung 3: Infobildschirm an der Booth.

Wir bedanken uns bei den Sponsoren, allen voran HPE, dem GCS und SPPEXA, unseren Unterstützern vor und während der Competition und unseren Betreuern Alexander Ditter und Johannes Hofmann und hoffen, noch die eine oder andere Erfahrung im Bereich HPC und Cluster Competitions machen zu dürfen!

Team SegFAUit

Solver	Physics	Discr.	Legal	Group
Ateles	A, CF	DG	in-house	U Siegen
Alya	IF, S	FE	in-house	BSC
A*STAR Flow	CF	FV	in-house	A*STAR
Calculix	S	FE	open-source	A*STAR, SimScale
Carat++	S	FE	in-house	TUM STATIK
Code_Aster	S	FE	open-source	SimScale
COMSOL	S	FE	commercial	-
EFD	IF	FD	in-house	TUM SCCS
FASTEST	IF+A	FV	in-house	TU Darmstadt
FEAP	S	FE	in-house	TU Darmstadt
FEM-shell	S	FE	open-source	U Stuttgart SGS
Fluent	IF	FV	commercial	-
OpenFOAM	CF, IF, S	FV	open-source	U Delft, SimScale
Peano1	IF	FE	in-house	TUM SCCS
SU2	CF	FV	open-source	-

Table 1: List of single-physics solvers that are currently coupled to preCICE. Physics: A – acoustics (linearized Euler equations), CF – compressible flow (Euler or Navier-Stokes equations), IF – incompressible Navier-Stokes equation, S – structural mechanics. Discretization (in space): DG – discontinuous Galerkin, FD – finite differences, FE – finite elements, FV – finite volumes.

preCICE should be made ready for exascale as part of the second phase of the ExaFSA project. Second, the user group should be enlarged by supporting industrial applications.

To reach exascale, we have to face several challenges. The initialization of preCICE is currently based on collective operations of the full mesh data. For non-changing coupling interfaces and up to approximately 10,000 cores, this concept is tolerable. To reach beyond this limit, and to allow for dynamic coupling interfaces, however, a multi-level approach needs to be incorporated.




Figure 4: Fluid-structure interaction of an aortic blood flow. Two snapshots during one cycle show velocity vectors besides the structural deformation. The fluid and the structure solver from the finite element Alya System are used, coupled with preCICE.

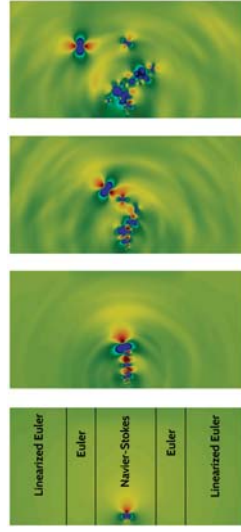


Figure 5: Three field flow coupling with the discontinuities Galerkin solver Ateles, coupled with preCICE. Pressure values of a subsonic jet are shown at various time instances. The leftmost snapshot also depicts the sub- domains.

Collective operations should only be used for higher-level structures, such as bounding boxes. Furthermore, the RBF mapping faces scalability limitations as it currently uses a global polynomial to stabilize the local basis functions. Only a fully local scheme allows for the necessary higher scalability. Finally, interpolation in time is necessary to handle the increasing difference in scales, which we are facing for exascale applications.



Abbildung 1: Fotoshooting zum ersten Treffen des Student Cluster Competition Teams.

Knoten wollten schließlich getestet werden. Am nächsten Morgen begann schließlich die Competition: Neben klassischen Performance-Benchmarks wie Linpack und HPCG hatten wir ebenso Aufgaben wie Password Recovery, die Simulation astrophysischer Daten oder die Reproduktion eines wissenschaftlichen Papers auf unserem Cluster zu bewältigen. Hinzu kam eine Aufgabe, die erst zu Beginn der Competition veröffentlicht wurde – Molekulardynamiksimulation mit gromacs.

Dies alles bedeutete vor allem: Wenig Schlaf. In Schichten wurde mehr als zwei Tage und Nächte rund um die Uhr gecoded, deployed und gescrpted. Letztlich erreichten wir bei unserer ersten Teilnahme mit unseren Leistungen den 9. Platz mit einem von vielen Preisrichtern als “interessantestes System des Wettbewerbs”bezeichneten CPU-only Systems unter vielen GPU-Clustern. Auch abseits und nach der Konferenz entdeckten wir mit anderen Teilnehmern der Competition viele schöne Ecken der Stadt und die Hälfte des Teams blieb noch eine weitere Woche, um den Westen der USA zu erkunden. Die anderen Teammitglieder mussten jedoch nicht weniger zufrieden sein, da ein Upgrade in die Premiumclass den zehnstündigen Heimflug doch sehr verstüßten.

Student Cluster Competition 2016 - Team FAU 12.-17.11.2016

Für das Team SegFAUit der Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU) war es eine Woche mit wenig Schlaf, jeder Menge Softdrinks, Junk-Food und scheinbar nicht enden wollenden Problemen. Klingt erst einmal gar nicht lustig, war es letztendlich jedoch schon, da es ebenfalls eine Woche vieler neuer Bekanntschaften, interessanter Vorträge, frischen Ideen, Sichtweisen und voller Abwechslung war.

Beginnen wir jedoch am Anfang: Nach einer knapp 16-stündigen Reise von Nürnberg nach Salt Lake City, Utah, kamen wir erschöpft im Plaza Hotel direkt neben dem Ausstellungsgelände an. Das Wetter überraschte uns mit sommerlichen Temperaturen und strahlendem Sonnenschein. Im Laufe der Woche wurde uns jedoch klar, dass unser voreiliger Ärger über das Einpacken von langen Unterhosen und Pullovern statt kurzen Shorts und Sonnenbrillen unnötig war, da hier in der Wüste das Wetter innerhalb einer Nacht von 20 Grad Celsius und Sonnenschein auf 5 Grad Celsius und Schneefall umschlagen kann. Der erste Eindruck der Stadt war durchaus faszinierend. In der Innenstadt findet sich ein gesunder Mix aus Neubauten mit großen Glasfassaden und Gebäuden, die zwar ebenso neu wirken, jedoch die Architektur europäischer Barockschlösser aufweisen.

Am zweiten Tag ging es sogleich los mit der ersten Sicherheitsunterweisung und einem Event zum Kennenlernen aller Studenten auf der SC16. Gleichzeitig ging es für uns an die Arbeit, denn es galt ein Clustersystem aufzubauen, ins Netzwerk einzubinden und zu konfigurieren. Währenddessen legten wir ebenso Aufmerksamkeit auf die Verschönerung der Booth, z.B. durch Sponsorenposter, LED-Lichtern oder auch Analyse- und Teampräsentationen, um möglichst viele Besucher zum Anhalten zu bewegen. Schon vor der Competition hatten wir so alle Hände voll zu tun, insbesondere da wir ohne Vorwarnung ein neues Chassis bekommen haben und die Firmware desselbigen zu alt für unsere brandneuen Cartridges war. Aber auch dies ließ sich beheben und am Abend war schließlich der Cluster betriebsbereit und es konnten nochmal Benchmarks und Scaling-runs durchgeführt werden, 37

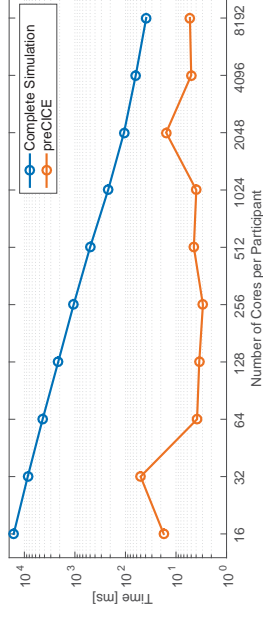


Figure 6: Strong scaling of the work per timestep for a density pulse traveling through an artificial coupling interface in an Euler domain. The time spent for the coupling does not influence the overall scalability.

To allow for more industrial applications, the RBF mapping, in particular, has to be made more robust. Therefore, interface parts, such as the back and front side of a thin flap, have to be detected automatically. Additionally, the amount of tuning parameters has to be reduced and replaced by auto-tuning procedures. Finally, the list of ready-to-use adapters has to further grow with a policy to make adapters as general as possible and also open-source.

Benjamin Uekermann

Acknowledgements preCICE is a joint software project currently developed by Florian Lindner and Klaudius Scheufele in the group of Miriam Mehl in Stuttgart and by myself in the group of Hans-Joachim Bungartz in Munich. The presented results in Section 3 are joint work with Verena Krupp et al. from University Siegen, and Juan-Carlos Cajias et al. from the Barcelona Supercomputing Center. I, furthermore, thankfully acknowledge the financial support of the priority program 1648 – Software for Exascale Computing of the German Research Foundation and of the Institute for Advanced Study of the Technical University of Munich as well as the provided computing time on the SuperMUC and the CoolIMAC at the Leibniz Supercomputing Centre.

References

- [1] Bungartz, H.-J., Lindner, F., Gatzhammer, B., Mehl, M., Scheufele, K., Shukaev, A. and Uekermann, B. preCICE – A fully parallel library for multi-physics surface coupling. *Computers and Fluids*, to appear (2016).
- [2] Brenk, M. *Algorithmische Aspekte der Fluid-Struktur-Wechselwirkung auf kartesischen Gittern*. Dissertation, Department of Informatics, University Stuttgart, 2007.
- [3] Gatzhammer, B. *Efficient and flexible partitioned simulation of fluid-structure interactions*. Dissertation, Department of Informatics, Technical University of Munich, 2015.
- [4] Uekermann, B. *Partitioned fluid-structure interaction on massively parallel systems*. Dissertation, Department of Informatics, Technical University of Munich, 2016.

managed to get back on track and ended up producing outstanding results in many categories after two sleepless days. It also excelled in conference engagement and the committee interviews that were held on multiple talks and workshops that had to be attended in the meanwhile.



Figure 1: The PhiClub: Stefan Haas, David Schneller, Svilen Stefanov, Sharru Møller, Maximilian Horning, Jan Schuchardt (from left to right, top to bottom).



Figure 2: The PhiClub phighting at SCC16.

After the closing ceremony, the team from upper Bavaria did not reach the first place, but succeeded in beating their Frankian “rivals” from Erlangen. Finally, our team was free to enjoy the less stressful side of the conference, including a dinner with interesting people from the HPC world and the closing tech party, before heading on a road trip through the mid-west or back to Munich respectively.

The people met, skills learned and experiences made are sure to be remembered.

Sebastian Rettenberger

Student Cluster Competition - Team TUM 12.-17.11.2016

One week – 14 teams – cutting-edge hardware – no sleep

For the second time, the Technical University of Munich has participated in the Student Cluster Competition under the supervision of Prof. Dr. Bader, Sebastian Rettenberger and Roland Wittmann.

The Student Cluster Competition is part of the annual Super Computing Conference. Since ten years, student teams of six members each compete in executing a variety of scientific applications in a non-stop 48-hours event on a high performance cluster, while staying within a strict power limit, participating in the conference and impressing the judges with their knowledge on high performance computing. Sponsored by the Gauss Computing Center, ExaHype and their vendor partner RSC, our team PhiClub set out to Salt Lake City in the middle of November, together with their supervisors, their 8-node Intel Xeon Phi cluster, prepared posters for their booth and their team-shirts. To be successful, the last six months had been spent in preparation for this highly competitive event. The main focus was on getting ready for the three applications that had been announced beforehand: Scientific visualization with ParaView, cracking password hashes and reproducing scientific paper results on graph connectivity in metagenomics data. Regular meetings were followed by a day-long preparation workshop on the weekend before the trip.

After brief hours of visiting the impressive Mormon temple district and a security briefing, the clusters were brought onto the exhibit floor and the stressful final preparation began. Countless hours were spent on tweaking the LinPack and HPCG settings for the eight-hour-long benchmarking session that always precedes the other applications. Even though our team did not come out on top in this category, it was still able to outscore last year's winner by a factor of more than 2. At 7pm on Monday evening, the hot phase began. Last announcements were made, another yet unannounced application (simulation of molecular dynamics with GROMACS) was revealed and then the teams rushed to their clusters. After a rough start of essentially killing the entire software stack through very optimistic memory allocation, our team

Summer of Simulation 2016 by BioLab@LRZ

Current multi and many core architectures pose a challenge to simulation codes with respect to scalability and efficiency. Particularly life and material science simulations often cannot just increase their system sizes because underlying algorithms scale unfavorably and insights from larger systems are limited. Still, computational demands are high due to the required abundant sampling of phase space or molecular structures and more accurate physical descriptions.

In May 2016 the BioLab of the LRZ applications support initiated the “Summer of Simulation” (SoS16) to help young scientist in these fields to tackle their problems on current supercomputers. Masters and PhD students employing molecular dynamics or quantum chemical simulations were called to submit a short one-page proposal describing their project. The aim was to port these projects on SuperMUC, find an optimally scaling setup and run their applications during the summer semester break.

The SoS16 started with a kickoff meeting in July, where the eight participants from the Ruhr-University Bochum (RUB), the University of Bonn (UB), the Friedrich-Alexander University Erlangen-Nurnberg (FAU), and the Technical University of Munich (TUM), respectively, presented their projects and were assigned one of the four tutors from the LRZ BioLab. The projects spanned a broad range from highly accurate surface chemistry over enzymatic reactions to simulations of nanoporous gold with a few hundred millions of particles.

In the following five weeks the students had to get their code and simulations running on SuperMUC and to optimize the setup. Here, each project had an initial budget of one million CPUh for preparatory simulations and to demonstrate the scalability of their project. With the guidance of their tutors the students prepared follow up proposals to apply for up to nine million additional CPUh. After a speedy review process, a total of 50 MCPUh had been granted and were available until 15th of October for the simulations.

At the closing workshop end of October each student presented the progress and results made over summer and handed in a final report. Currently, most projects are evaluating their data for scientific publications. Furthermore, follow up proposals for CPU time based on their experiences obtained during SoS16 are under way.

Three of the projects involved the density functional theory code CPMD (www.cpmd.org) for first principles molecular dynamics to explore chemical reactions on metal oxide surfaces. For these many electron systems ultra-soft Vanderbilt pseudopotentials are applied, for which comparatively small plane wave basis sets suffice and drastically reduce the overall computational costs.




Figure 1: Improved scaling of CPMD on SuperMUC Phase 2.

On the downside, this limits the number of MPI processes that can be used in the simulation because CPMD is parallelized by distributing planes of the real space mesh that supports the plane wave basis and the electron density. During the SoS16 and with the help of the LRZ BioLab the hybrid OpenMP/MPI parallelization of the Vanderbilt code within CPMD was revised and permitted an optimal setup of 4 MPI processes times 7 OpenMP threads per node of SuperMUC Phase 2. Together with the tuning of the Intel MPI library a notable speed up was achieved as is shown in Figure 1.

Paul Schwarz (group of Prof. B. Meyer, FAU) used this CPMD version to study the reaction pathways of the condensation of methylsilanetriol on aluminum oxide (see Figure 2). Hannah Schlott from the same group developed

size of 6400 cores per run. Finally, the largest systems were investigated by Zhuocheng Xie (group Prof. Bitzek, FAU) exploring the material properties of nanoporous gold. Nanoporous metals are a popular field in materials science due to their potential technological applications in actuation, catalysis and sensing. Based on experimental structures of nanoporous gold, a model system containing 450 million gold atoms was constructed and its material properties were investigated. Simulations with the LAMMPS molecular dynamics program were run on SuperMUC Phase 2 using up to four islands (2048 nodes). Figure 7 shows a snapshot of a simulation compressing the gold cluster in z-direction.

All projects were carried out by curious, industrious and eager students and it was a great pleasure for the tutors to work with them. Moreover, the close contact with the different projects showed hurdles and pitfalls, whose fixing helps to improve the usability of SuperMUC in general.




Figure 7: Compression of the nanoporous gold model system.

Acknowledgements: This article will appear in the spring 2017 edition of inSiDE (Innovatives Supercomputing in Deutschland). Kind support by the SuperMUC steering committee (Prof. Wellein) is gratefully acknowledged.

Gerald Mathias, LRZ




Figure 5: Collective variables used for umbrella sampling. Purple, cyan and brown molecules correspond to different conformational states.



Figure 6: Mutated residues around the active center of CE6 Compression of the nanoporous gold model system.

The challenge of this project was to adapt the python scripted workflow, which steers the different applications, to be used on SuperMUC Phase 1. NAMD could be run efficiently on two nodes, whereas the quantum chemistry application Turbomole did not scale beyond node boundaries. Here, however, the two independent Turbomole calculations per mutation step were distributed on the two nodes. A further massive parallelism was achieved by running 200 Monte Carlo paths simultaneously, which totals to a job

a simulation protocol to generate amorphous surface structures of zinc oxide and other oxides by an annealing procedure. The generated structures will serve to study catalytic reactions on such unordered interfaces.




Figure 2: Methylsilanetriol condensating on an aluminum oxide surface.

The third project by Niclas Siemer (group Prof. D. Marx, RUB) dug even deeper into the CPMD code. For the simulation of gold clusters on titanium dioxide surfaces he had to employ a so called Hubbard U correction to account for the poor description of the localized electrons in the d-shells of the titanium atoms by standard DFT functionals. The initial implementation of the U correction in CPMD, however, was developed for small metal cluster only, and never tested for larger systems. Correspondingly, it showed a rather poor scaling behavior. During the first stage of the project he revised, thread parallelized, and optimized the U correction code at the LRZ. Figure 3 shows that the originally more than 40% overhead for about a hundred U corrected atoms is reduced to well below 10%. Furthermore, the overall time per integration step on 20 nodes was reduced from 55 sec to below 11 sec, yielding a speedup of more than 5. With these improvements computing the dissociation energy of O₂ on the TiO₂/Au surface was feasible during the SoS16 period.

Two other quantum chemistry projects focused on small metal containing clusters. Sascha Thinius (group Prof. T. Bredow, UB) computed the structure

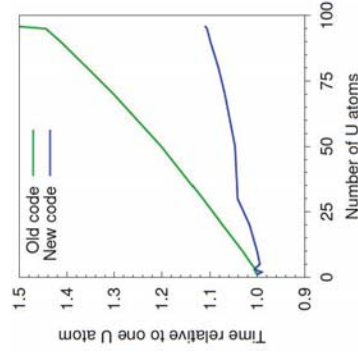


Figure 3: Scaling of the CPMD code with the number of U corrected atoms.

of chalcopyrite nano particles in order to improve the description of leaching processes of such copper ores. The particular focus of the project was selecting the optimal program and setup because even small models of such nanoparticles easily contain a few thousand atoms, for which the computation of a wave function is very demanding. Here, the GPAW program was shown to scale to up to 900 cores and to outperform VASP by up to a factor of twenty. To treat even larger metal oxide systems at high accuracy Martin Paleico (group PD J. Behler, RUB) is developing an interaction potential that is based on a neural network (NNP) for zinc oxide/copper systems. ZnO Cu is used as catalyst in the industrial synthesis of methanol and a better understanding of this reaction by simulation could help to improve its efficiency. To reach the required accuracy of a few meV/atom the neural network has to be trained with energies for a large set (105) of different ZnO/Cu structures covering up to a few hundred atoms each. Generating these data was carried out on SuperMUC with the Vienna Ab-initio Simulation Package (www.vasp.at) using 20 nodes per structure. The resulting accuracy of the NNP is shown in figure 4 for the ZnO zincblende and wurtzite crystal structures.

Two projects originated from the biophysics/biochemistry community. Florian Kandzia (group Prof. Zacharias, TUM) analyzed the domain dynamics of the heat-shock protein Hsp90. This protein is an essential molecular chap-




Figure 4: Comparing NNP and first principles results for crystal structures.

erone that facilitates the folding of many other proteins and plays a vital role in the assembly of functional complexes. In order to bind to a variety of proteins, its homologues adopt very different 3-dimensional conformations ranging from tightly closed to loosely connected dimers.

Based on a pre-equilibrated structure, 560 initial conformations for a 2D umbrella sampling simulation scheme along a center-of-mass distance and center-of-mass torsion collective variable (see Figure 5) were generated by preparatory simulations. Structural exchange between these windows by a replica exchange algorithm enhanced the sampling and convergence during the production run with the molecular dynamics package Amber16 (ambermd.org) on SuperMUC Phase 1. Using 80 cores per replica the resulting 44800 cores for each run cover approximately one third of the supercomputer partition, which is close to the maximally allowed job size during normal operation.

The biochemistry project by Sophie Mader (group Prof. Kaila, TUM) employed a new method for computational enzyme design. Here, amino acids of enzymes are mutated randomly, the properties of the mutant are calculated and a Metropolis Monte Carlo procedure decides if the mutation is accepted or declined based on its catalytic activity. The target system to be enhanced by this QM/MM Monte Carlo method was the computationally designed enzyme “CE6” for the Diels-Alder reaction (see Figure 6). Mutations are performed by the tool VMD, then structures are relaxed using the molecular dynamics package NAMD (www.ks.uiuc.edu/Research/), and properties are calculated with the quantum chemical program Turbomole (www.turbomole.com).