PRACE: Partnership for Advanced Computing in Europe

FZ-JULICH

Professor Dr. Achim Bachem is Director of the Partnership for Advanced Computing in Europe (PRACE) and Chairman of the Board of Directors of Germany's Jülich Research Centre.

SciDAC Review: *You are the Director of PRACE, a new European supercomputing organization. Since most of our readers are in the U.S., let me first ask you to tell us about PRACE. What is PRACE and what are your goals for high-performance computing (HPC) in Europe?*

Professor Dr. Achim Bachem: The Partnership for Advanced Computing in Europe, PRACE, is an initiative started in 2007 by the major players in HPC from 14 European countries — which has in the meantime grown to 18 countries — with the mission to create a leading persistent high-end HPC infrastructure for European public and private research and development, contributing to the advancement of European competitiveness in industry and research.

PRACE is based on previous activities, in particular HET, the HPC in Europe Taskforce, which has produced a scientific case that clearly has demonstrated the need for a world-class HPC service for European researchers. Through this preparatory work, the European HPC service successfully made its way onto a list of 34 Research Infrastructures proposed to be built in Europe with high priority. This so-called ESFRI list was published by a policy advisory group named European Strategy Forum for Research Infrastructures in 2006.

Our short- to medium-term goals are to set up pan-European HPC services provided by a small number of leading (Tier 0) centers which will commence operation in 2010. The HPC services will be provided under a common peer review system. These services will not comprise cycles alone but will also include user support and training. Our medium- to long-term goal is to foster and revitalize the HPC industry in Europe through collaborations with vendors.

The PRACE project is funded by the European Commission (EC) for a period of two years with the goal to complete the legal, administrative, and technical work to make this happen. I am actually the coordinator of this project; the chair of the PRACE initiative is rotating between Germany, France, The Netherlands, the United Kingdom, and Spain.

At the end of those two years, how will you measure your progress in achieving the goal of changing from being individual centers to being a larger partnership?

A number of the partners in PRACE have actually worked for years, often in competition, to be one of the first centers to get funding from the European Union (EU), as opposed to funding from our own country. We saw each other as competitors, and you know that competitors don't talk to each other about their systems and goals. The question was, who would become the EU-funded center?

Then in 2006, another idea came along — a much better idea. We realized that money would not be falling from heaven, so as a group we decided to start with our national funding and build up PRACE in this way. Suddenly, the whole atmosphere changed. Money was no longer the reason for trying to work together. In contrast, we started looking at what kind of advantage we would have if we built up the PRACE infrastructure using national funding. We put away a lot of our competitiveness and, more or less, became friends.

We still have to finalize the legal framework for PRACE, but in spirit we are approaching this partnership in a European way. We realize there are advantages in pursuing different architectures and in working together to develop a petascale software project. This also enables us to approach industry as a European-wide partnership, rather than on a center-by-center basis, and this offers many, many advantages.

I think it was at this step that we could say PRACE was truly born. Now we are in the end phase and deciding questions such as where to locate our headquarters. But 70% of the questions have been solved.

So, PRACE is already in existence, even without a formal legal partnership. And what is really important for us is that it was hard for any one center in Europe to be a partner with a U.S. program, but with PRACE, we could form stronger partnerships. In fact, we are already in discussions on some joint initiatives.

Looking back in the history of HPC, the European countries — in particular France, Germany, and the UK, and more recently Spain — have been very successful in fielding large computing facilities. Your own center in Jülich is home to the largest system in Europe. So one might ask, what will PRACE add that did not exist before?
As I mentioned, until recently the HPC landscape in Europe was mainly characterized by fragmentation and competition between countries. PRACE adds a strong element of cooperation and coordination which adds value in many aspects.

We expect at least one large system to be deployed each year in Europe within PRACE. This guarantees that European users will always have access to a top-level system. One country alone is not able to fund a continuous provision at top level.

These systems will be open to researchers from all European countries, including those with smaller HPC budgets. The motto "the best systems for the best scientists" will no longer be restricted to some richer European countries.

Through coordination, we can deploy a greater variety of architectures to provide a larger user base with an optimal system — instead of having very similar off-the-shelf systems in several countries. This includes risk-sharing for investments in more advanced technologies.

When acting jointly, Europe can play a more active role and contribute better to international cooperations addressing future challenges, such as the International Exascale Initiative. In the long run, we also expect to be quite influential on vendor roadmaps.

One effect of PRACE that already yields fruit is that the political awareness for HPC has increased in many European countries, which now strengthen and structure their HPC activities. An example is Germany with the formation of the Gauss Centre for Supercomputing (GCS) and the first GCS petaflop/s system that has been deployed in Jülich.

You are very familiar with the international scene in HPC and computational science. On the international playing field, in which areas do you see Europe having intellectual strengths?
To cite EU Commissioner Reding, responsible for information technology (IT), Europe is excellent at turning money into ideas, but needs to become better at turning ideas into money. This means that Europe has strengths, in particular in areas which are mainly driven by academia. When you look at the value chain of HPC and IT in general, this covers material science and nanotechnology at one end, and software at the other end. Here, European researchers are making important contributions in many areas, from operating systems, middleware, and tools, to numerical mathematics and algorithms.

Biography in Brief

Professor Dr. Achim Bachem has served as Chairman of the board of directors of Research Centre Jülich since October 2006. He is a member of the board of directors of the Gauss Centre for Supercomputing in Germany and project coordinator of the European PRACE project (Partnership for Advanced Computing in Europe).

Research Centre Jülich pursues cutting-edge interdisciplinary research on solving the grand challenges facing society in the fields of health, energy, and the environment, and also information technologies. With a staff of about 4,400, Jülich – a member of the Helmholtz Association – is one of the largest research centers in Europe. In May the first petaflop/s computer in Europe at GAUSS@JÜLICH was inaugurated in Jülich. At the same time Jülich celebrated two other events of great significance for Europe: the launch of the first European computer for fusion science and the JUROPA supercomputer, which will be used for a wide range of applications in European research.

That is an interesting comment about the challenge of turning ideas into money. Why do you think this has been the case?
In Europe, we have a very different culture of entrepreneurship. If you have a good idea and are thinking about starting a company, there is a lot of fear: you don't know where to get the money or how to go about it. One survey I read said that 60% of Europeans fear that if they fail at something, it will be the end of their career. In the United States if you fail, you just start anew. Interestingly, in German we call startup funding "risk capital," but in the United States it's called "venture capital." This accurately reflects the different perspectives.

Except for the fairly recent entry of Bull, there is no domestic European supercomputer manufacturer. Do you see this as a potential weakness for HPC in Europe, compared to the United States and Japan, who both have a fairly large and established number of vendors?
What we are aiming for in Europe is an independent access to HPC, which is considered one of the key technologies of the 21st century. In a globalized world, this does not mean that Europe has to manufacture all system components, but it needs sufficient capability to set its own agenda, such as putting a focus on green IT. To support this, PRACE has formed STRATOS, a group including PRACE partners and industry consortia. Its goal is to stimulate joint technology projects and developments to revitalize the European HPC industry in the long run. Therefore, I highly welcome the success of Bull to establish itself as an HPC manufacturer on the market. This is not a small achievement in times that are characterized by consolidation of vendors. By the way, Bull has been active and successful in the French HPC market for quite some time.

And you also have a new system built by Bull and Sun Microsystems at Jülich.
For me, supercomputing is a key technology and I want the competence of supercomputing to be in Europe. I don't care whether

the machine is a U.S. or Japanese system, but it is important that these manufacturers have research labs in Europe. If we can convince IBM to double its research staff here, that's a real value for the EU in terms of investment in research and production. But with the AMD plant in Dresden, on the other hand, I don't know if it makes sense to build another chip factory — chips are only 20–30% of the value added in supercomputing.

I think we should focus more on software — this is the biggest challenge for the EU. I think we have a chance to contribute to HPC software, which would enhance European competence in this area. We need more firms like SAP.

When you presented the PRACE Award at the 2009 International Supercomputing Conference (figure 1), you made a point of saying that PRACE supports not only systems, but also supports people. Can you elaborate on this idea?

We think the systems are really the easiest part and that we will probably continue to follow Moore's Law — the vendors will continue to do their job of delivering high-performance components. But it is much, much more difficult to configure these machines so that they can be used for science, which is becoming much more complex in terms of models, algorithms, and codes. For PRACE, we see a bigger task to support thinking in terms of new models and new algorithms.

> This is a new way for us and we would like to adopt some of the collaborative ideas of SciDAC.

We are now at a paradigm for solving new problems, and for these bigger problems we need a new way of modeling them. For example, in aeronautics, all problems have boiled down to Navier–Stokes equations. But is there any chance to create solvers for Navier–Stokes equations that will scale to exaflop/s? We need to think of new physical models that take into account, from the very beginning, that they shall be simulated on massively parallel computers. This will only work if physicists, chemists, and biologists closely work together with mathematicians and computer scientists. We see this as another challenge for PRACE.

Which application areas will be the focus of PRACE? How have they been selected? Do you see any difference from the United States with respect to applications?

Part of PRACE's mission is to provide a horizontal research infrastructure that serves the scientific community at large. Therefore, we will not set the priorities ourselves. There will be a scientific steering committee involving scientific communities, and also funding agencies will have a word. Early in the project we conducted a survey of resource usage on more than 20 of the largest HPC systems in Europe. We found that the largest consumers today are particle physics and computational chemistry, followed by condensed matter physics, computational fluid dynamics (CFD), Earth and climate research, astronomy and cosmology, and life sciences, in that order. We expect to see the same areas also on the future PRACE systems.

With its resources, PRACE shall also contribute to the solution of the major European societal challenges — such as energy, health, environment, transportation — but will be open to any kind of research. Unlike in the United States, defense applications are not expected to play a role in PRACE.

What about the allocations process for PRACE? Horst Simon at LBNL has talked about the "democratic socialism" approach in which users receive uniform allocations — but it is rarely as much as they would like.

We are going to reserve PRACE allocations for the very, very big problems. Horst's "socialism" will still be applied to our Tier 1 centers, but with the Tier 0 centers, we foresee a different system. We will ask the top scientists and they will decide who gets the cycles. We won't take into account which country a researcher is from, just the quality of science. The goal is to support the best science.

One thing we would like to change as we build up the infrastructure is to put the user — and not the provider — in the driver's seat. But as you can imagine, this is not an easy thing to change. At Jülich, we have little power over the allocations — 70% of the cycles are determined outside of our center. However, if the users of these allocations are successful, it reflects well on our center.

We also had early discussions about whether it would be better for users to pay for their time, but we decided this is a very bad solution. It puts pressure on the user to get the funding, but has no provision for having experts judge the quality of the science.

You are quite familiar with the SciDAC program in the United States. From a European perspective, what do you see as the strengths of SciDAC? Will there be a similar program in the EU?

As a key strength of the SciDAC program, in particular in comparison to funding programs in Europe, I consider the integrated and interdisciplinary approach bringing together computational and computer scientists to attack the full spectrum of challenges in advancing application sciences using HPC, from algorithms, libraries, and tools to applications. Sustained funding over an extended period is also an important success factor. Unfortunately, there has been no similar action by the EC so far. Funding was either directed to disciplinary research or to certain aspects of computer science, leaving a lot of potential synergies unused. Meanwhile it seems that the EC has taken notice of this deficit and will act. A multi-million Euro HPC program is envisioned which will also address collaboration of application scientists and computer scientists. PRACE is prepared to play its role in this interdisciplinary endeavor.

Also, SciDAC has made a concerted effort to reach out to universities. But in Europe, many of our centers are affiliated with universities, so maybe we need to work more closely with other research centers. Our scientific research system is also different — the Helmholtz Society provides a scientific research infrastructure, such as supercomputers, synchrotrons, and polar research ships. But we are looking for ways to break down barriers and work together.

ISC'09

Figure 1. *Prof. Dr. Achim Bachem (right) presents the 2009 PRACE Award to J. C. Mouriño of Spain's Supercomputing Center of Galicia on behalf of the winning group. The award, which recognizes the best scientific paper by a European student or young scientist, was presented during the opening session of the 2009 International Supercomputing Conference in Hamburg, Germany.*

For example, the Jülich Center and University of Aachen have formed a "merger by competencies" to create JARA — the Jülich Aachen Research Alliance. Our two campuses are 15 miles apart and with this almost-merger we hope to raise the visibility of both institutions. And the Karlsruhe Institute of Technology (KIT), has been formed between a national lab and a university.

The driver for this in Germany has been the Excellence Initiative, started five years ago. The universities compete for funding by earning a "stamp of excellence." At first, there were just five universities, but now there are 12, including Aachen and Karlsruhe. This is a new way for us and we would like to adopt some of the collaborative ideas of SciDAC.

International competition or collaboration in HPC — how do you see this choice evolving in the near future?

It has always been a mixture of both and I expect and hope that it will continue to be so. Most HPC centers of international relevance have long-standing bi- or multi-lateral cooperations on specific topics. On the other hand, the competition between these centers, even on a national level, is a major driving force for progress. The increasing importance of HPC for science and the economy is of course pushing towards stronger international competition. On the other hand, the unprecedented challenges of exascale computing ask for worldwide coordination and cooperation. This might not be a dilemma: for example, in the automotive sector, one can find competitors combining their development efforts in fuel cell or battery developments without compromising competition.

Do you see any funding advantages in closer collaboration?

Funding will get tighter, but so far the government is putting more money into research despite the economic crisis. If at some point we don't get more money, it will be easier to survive in a collaborative mode than in a competitive environment. This is true not only on a national scale, but also at the European level.

We also have strong ties to Oak Ridge National Laboratory, with 20 people from Jülich working at the Spallation Neutron Source (SNS). This is helping to open the door to the SNS for the European research community.

But still, the competition should be there in science. Otherwise, we will not be pushing ourselves to achieve more.

Even with the EU, the Euro, and the opening up of borders, there are still rivalries among European nations — just look at the emotions surrounding last year's European championship series in soccer. And this competition even shows up in the TOP500 list. How is PRACE overcoming — or perhaps leveraging — these differences?

In PRACE, we are fair sportsmen. We congratulated our colleagues from the Barcelona Supercomputing Centre for Spain's victory in the European soccer championship in 2008, just as they had congratulated us for getting a number two system in the TOP500 list in 2007.

But seriously, what we experience in many European projects — and in PRACE in particular, where more than 270 researchers from 14 countries work together closely — is that there is no rivalry at the working level, but a spirit of learning and benefiting from each other. On the political level, the expected added value for science and economy in each country and in Europe at large is the main driver for PRACE. Besides that, also the participating future Tier 0 centers benefit, by gaining visibility as national representatives in a European success story.

Do you see any parallels between the PRACE approach and the partnerships at CERN?

If we could reach the same stage as CERN, that would be both a good target and a good benchmark. But CERN has a well-defined community. Supercomputing is not really a community for all scientists — we have no real profile. In the field of multi-particle physics, I think it is clear the community would come together to fight for funding. But I'm not sure all scientists would band together to fight for money to support supercomputing.

What do you think are the biggest challenges for computational science in the next decade? What approaches to overcoming these challenges do you think are appropriate?

The currently foreseeable technologies to achieve multi-petaflop/s to exaflop/s peak performance — excessive parallelism and heterogeneous processing units — are somewhat disruptive and will be hard, if not impossible, to leverage with today's applications, tools, and even programming paradigms. There was a similar situation with the transition from vector to MPP and large SMP systems in the early 1990s that we can learn from: it needs time and experimentation before new paradigms of hardware and programming models are mature enough for standardization and general use. A main difference to the vector MPP–SMP transition is, however, that now there is only little possibility to experiment on a small scale. Exascale software needs to be available at the same time as the hardware in order to leverage the huge hardware investments. This will be possible in a global coordinated effort and with the involvement of "heroic" scientific and industrial user communities.

Thank you for taking the time to answer our questions.