European & Asian Supercomputing

Gary M. Johnson
Office of Advanced Scientific Computing Research
Office of Science

Gary.Johnson@science.doe.gov

Advanced Scientific Computing Advisory Committee Meeting
American Geophysical Union
14-15 August 2007

Caveats

- This presentation contains opinions and impressions, rather than the results of a scientific study
- There are many excellent sources of information on European & Asian supercomputing. To cite just a few:
  - Supercomputing conferences in Europe & Asia
    - International Supercomputing Conference
      http://www.supercomp.de/isc08/index.php5
    - International Conference on High Performance Computing
      http://www.hipc.org/
    - HPC China (2006全国高性能计算学术年会)
      http://www.sccas.cn/hpcchina2006/index.html
  - Special sessions at domestic supercomputing conferences
    - e.g. 3rd Workshop on Chinese High Performance Computing to be held at SC 2007
      http://www.atip.org/node/96
  - Special reports and studies
    - Asian Technology Information Program
      http://www.atip.org/
    - World Technology Evaluation Center
      http://wtec.org/
      (currently conducting a study of simulation-based engineering and science)
  - The Top 500 List
    http://www.top500.org/
- The membership of ASCAC is generally quite well informed on European & Asian supercomputing

International Events have Domestic Impact

The New York Times

April 20, 2002

Japanese Computer Is World's Fastest, as U.S. Falls Back

By JOHN MARKOFF

SAN FRANCISCO, April 19 — A Japanese laboratory has built the world's fastest computer, a machine so powerful that it matches the raw processing power of the 20 fastest American computers combined and far outstrips the previous leader, an I.B.M.-built machine.

The new Japanese supercomputer was financed by the Japanese government and has been installed at the Earth Simulator Research and Development Center in Yokohama, west of Tokyo. The Japanese government spent $350 million to $400 million developing the system over the last five years…

Assembled from 640 specialized nodes that are in turn composed of 5,104 processors made by NEC, the new Japanese supercomputer occupies the space of four tennis courts and has achieved a computing speed of 35.6 trillion mathematical operations a second. The processors are linked in a way that allows extremely efficient operation compared with the previously fastest "massively parallel" computers, which are based on standard parts rather than custom-made chips.

U.S. Department of Energy

Office of Science

Science AAAS

July 18, 2003

Advanced Scientific Computing Research Program

Simulators Face Real Problems

Scientists trying to model everything from the death of a star to the structure of a protein don't have the computing power they need. But help may be on the way

LIVERMORE, CALIFORNIA—Chris Fryer has a modest goal: He wants to recreate a supernova. The astrophysicist at Los Alamos National Laboratory in New Mexico knows that the titanic explosion from the collapse of an aging star doesn't always spew mass evenly in all directions. Indeed, it can fling the newborn neutron star from its birthplace with enough force to send it right out of the galaxy. What he doesn't know, however, is whether the asymmetries that contribute to the ejection also help trigger the explosion.

To find out, Fryer needs to model the convective boil of the star in three dimensions, while at the same time accurately portraying the pressure-cooker physics of neutrinos trapped near the surface of the star's collapsing iron core. But that requires a supercomputer at least 100 times more powerful than what's now available. And Fryer doesn't want to wait for Moore's Law—the historical doubling every 18 months of the number of transistors that fit on a silicon chip—to take care of the problem. "We're overdue for another supernova in the Milky Way," he says. "The goal is to make as many predictions as we can before that happens. But to make predictions, you need to have all the physics."

Fryer is not alone among U.S. scientists in pining for a large increase in computing power. "There are a surprising number of clear milestones that could be achieved with a factor of 100 [improvement]," says Michael Norman, an astrophysicist at the University of California, San Diego, who chaired a session last month at a government-sponsored meeting on high-end computing and its applications. A 100-fold boost in power, for example, would allow scientists to increase the use of quantum mechanics in modeling electron distributions and improve simulations of protein folding, enzyme actions, and nanoscale materials, and address many other problems.

A new climate-modeling machine that the Japanese government brought online last year has brought the issue to a head (*Science*, 1 March 2002, p. 1631). The top processing speed of the Earth Simulator far outstrips anything else around (see table), giving users an unprecedented ability to capture and predict the dynamics of the oceans and atmosphere. Its debut has also energized U.S. science policymakers. In March, the National Academy of Sciences began a study to assess the future of supercomputing, followed by last month's meeting of a task force assembled by an interagency working group under the White House Office of Science and Technology Policy. One week later, a panel convened by the Department of Energy (DOE) spent 2 days coming up with examples of exciting, just-around-the-corner milestones that were stymied by a lack of computing firepower. The goal is to craft a supercomputing initiative for the president's next budget. "The world of supercomputing is churning," says Horst Simon, director of the National Energy Research Scientific Computing Center (NERSC) at DOE's Lawrence Berkeley National Laboratory in California, "and something new will come out."

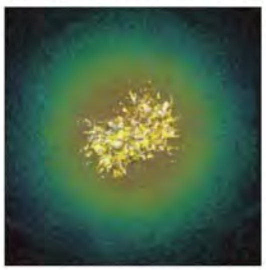

Power play. Chris Fryer needs more computing power to study lopsided explosions in collapsing stars about to become supernovae.

CREDITS (LEFT TO RIGHT) HUNGERFORD AND FRYER, 2003; C. FRYER

Different strokes

Supercomputers appeared soon after the dawn of computing in the 1950s. Defined then as anything faster than what was available to the general public, the original machines had specially designed architectures that were well suited to scientific problems. Consumer demand soon eclipsed scientific demand, however, leading companies to focus on the faster processors sought by offices and game enthusiasts. Supercomputers came to rely on many fast processors working in parallel rather than specialized architectures. A few companies, such as Cray Research, tried to buck the trend by improving connections between processors or the way in which data are sent to them. These machines performed well on scientific tasks, but they were not big sellers.

Now the status quo may be changing again. In the last few years, the idea of simply hooking up more and more processors in parallel has run up against problems in programmability, power consumption, and space. At the same time, it has become clear that supercomputers optimized for one type of operation may not be well suited for tackling other tasks. "It's beginning to look like different strokes for different folks," says Bill Feiereisen, head of computer and computational sciences at Los Alamos. "There is likely to be a plethora of needed architectures."

Supercomputer designers face three major challenges that are not critical to the function of ordinary single-processor PCs. The first is how to increase the amount of information that can be passed between the processor and the memory banks. Explains Bill Pulleyblank, director of exploratory research systems at IBM Research in Yorktown, New York, "As processors get faster and faster, they spend more and more time waiting for the data to get to them from the memory of the computer." The waiting time is idle time. The second challenge is achieving greater bandwidth between individual processors. A deficiency here means a longer wait as data are passed between processors. The third hurdle is lowering the amount of time it takes a processor to initiate contact with another processor. In some supercomputing applications, small amounts of data are passed frequently between neighboring processors. If this interaction, called latency, is not optimized, then these numerous short connections can bog down the system.

Those features are not equally important to every supercomputer user, however. Machines composed of ordinary PCs hooked up in parallel are well suited to tackle problems that can be broken down into independent pieces. Examples include the "SETI@home" project, which parcels up and ships the radio frequencies obtained from one section of the sky to computers around the world searching for telltale signs of extraterrestrial intelligence, and "Folding@home," which sends individual processors the data needed to make one attempt to fold a protein. Because these problems don't require interaction between processors, there is no need for elaborate hardware or software to facilitate communication.

Another type of supercomputer may excel in tackling problems that require frequent communication of small amounts of information between processors. These supercomputers are ideal for materials scientists studying atomic-scale interactions. Simulations of nuclear explosions also require knowledge of such fine-scale interaction, hence the popularity of ASCI machines at the Department of Energy's national laboratories.

Then there are problems such as climate modeling that demand a massive coupling of the interacting parts across the entire system of processors. In addition to broadcasting information widely, computers built for these tasks may also be required to move long strings of data between processors. The Earth Simulator is such a machine; it employs techniques that allow simple calculations to be performed simultaneously on massive amounts of data and puts a premium on bandwidth to handle the large chunks of data being shuffled about. However, it is less well equipped for the neighbor-to-neighbor problems that demand low latency.

Fryer needs all three features in order to understand how supernovas behave. Modeling the hydrodynamics of the convective motion relies on heavy-duty communication among processors, whereas the breakneck speed of radiation transport requires the rapid broadcasting of results in one region across the entire star.

No time to wait

Shortly after the debut of the Earth Simulator, DOE put forward a proposal to build an "ultrascale simulation for science" machine. "People were shouting, 'Computenik, computenik,' " says David Keyes, an applied mathematician at Columbia University. "It was like Sputnik." DOE also asked Keyes to assemble a group of researchers to draw up a wish list of what they'd like to study and what level of computing power would be needed.

The group, which met last month, will share its findings with the interagency High End Computing Revitalization Task Force. The latter group hopes its recommendations, which won't be made public, will have an impact on the 2005 budget that the president sends to Congress next winter. Daniel Reed of the National Center for Supercomputing Applications at the University of Illinois, Urbana-Champaign, who chaired the task force's June meeting, characterizes them as a way "to coordinate increased investment in high-end computing."

If these studies do result in new funding, there will be no shortage of technical problems to spend it on. One problem plaguing U.S. machines is efficiency—the amount of work actually performed compared with the amount of work that would get done if each processor could work continuously without needing to wait for data to arrive. For climate problems, the Earth Simulator operates at 30% efficiency or higher, well above what U.S. machines can achieve. "Inadequate bandwidth and network latency limit us to about 1% of the peak performance of the computer we run on," says San Diego's Norman. "That's typical of the [off-the-shelf component–based] architectures that we've grown so fond of in the United States. We've spent a decade lowering price for peak performance, and there's nothing wrong with that," Norman adds. "But the metric should really be sustained performance."

Another major obstacle is writing code to harness the power of an escalating number of processors. "While it was difficult to program for 16, 32, or 64 processors, we always believed that it would become easier for larger numbers once we learned how to do it," says Juan Meza, head of high-performance computing research at NERSC. "But the further up we went, to 500, 1000, and 10,000 processors, the harder and harder it became."

A related problem is the reliability of these machines. The more processors on board, the more likely it is that one will crash and bring down the entire system. IBM is working with scientists at Lawrence Livermore National Laboratory on a 65,000-processor machine, called Blue Gene, that will overcome processor failures so that data being collected during a month-long run aren't lost. Designed for modeling molecular biochemistry, Blue Gene's 360 teraflops of computer power will dwarf anything now available. "When we looked, for example, at protein simulations, we said, 'You know that the machine we need is approximately 1000 times more powerful than anything that exists now,' " Pulleyblank says. "If we just wait on things doubling every 18 months, it's going to take 10 or 15 years to tackle those problems. But the scientists can't wait."

The appeal of the top-of-the-line supercomputers is their ability to attack problems that are beyond the reach of ordinary machines, says Fryer. "They allow us, every once in a while, to run simulations that we couldn't do any other way," he says. And because innovations often trickle down to the masses, Fryer adds, the next supercomputers keep hope alive among scientists: "They show us what our future will be in computing." —KATIE GREENE

Katie Greene is a science writer in Oakland, California.

World's Fastest Supercomputers

Rank	Manufacturer	R_{max}*	Site	Country	Installed
1	NEC	35,860	Earth Simulator Center	Japan	2002
2	HP	13,880	Los Alamos National Laboratory	U.S.	2002
3	Linux Networx	7,634	Lawrence Livermore National Laboratory	U.S.	2002
4	IBM	7,304	Lawrence Livermore National Laboratory	U.S.	2000
5	IBM	7,304	NERSC/LBNL	U.S.	2002
6	IBM	6,586	Lawrence Livermore National Laboratory	U.S.	2003
7	Fujitsu	5,406	National Aerospace Laboratory of Japan	Japan	2002
8	HP	6,160	Pacific Northwest National Laboratory	U.S.	2003
9	HP	4,463	Pittsburgh Supercomputing Center	U.S.	2001
10	HP	3,980	Commissariat à l'Énergie Atomique	France	2001
11	HPTi	3,337	Forecast Systems Laboratory, NOAA	U.S.	2002
12	IBM	3,241	HPCx	U.K.	2002
13	IBM	3,164	National Center for Atmospheric Research	U.S.	2002
14	IBM	3,160	Naval Oceanographic Office	U.S.	2002
15	IBM	2,560	Euro. Ctr. for Medium-Range Weather Forecasts	U.K.	2002

* Processing speed in gigaflops.

SOURCE: TOP500.ORG

Facilities for the Future of Science

A Twenty Year Outlook

In **November, 2003** Secretary Abraham announced the Office of Science's 20 Year Science Facility plan

The second ranked Near Term Priority was UltraScale Scientific Computing Capability

DOE Office of Science issued to SC laboratories a call for proposals to provide Leadership Class Computing Capability for Science with a funding profile of $25M/year for five years in **February, 2004**.

Federal Plan for High-End Computing

May 10, 2004

- **Research and Development:**
  - A coordinated, sustained research program over 10-15 years to overcome major technology barriers that limit effective use of high-end computer systems.
- **Resources:**
  - Providing high-end computing resources across the full scope of critical Federal missions
- **Procurement:**
  - The HECRTF Plan proposes several pilot projects for improving the efficiency of Federal procurement processes, benefiting both government and industry.

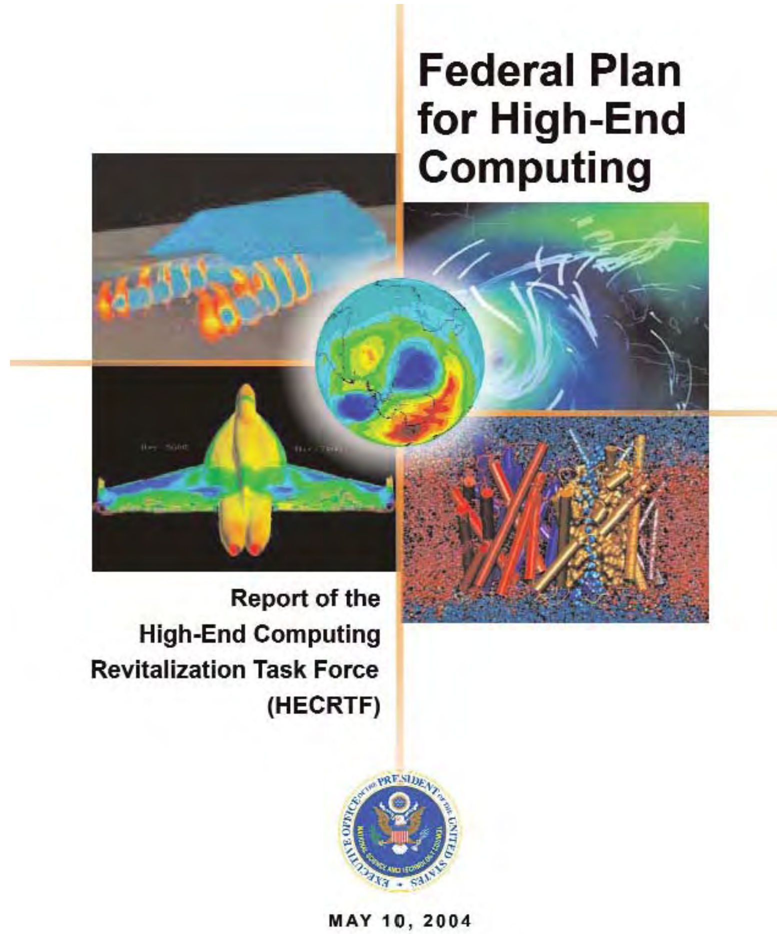

For Immediate Release
May 12, 2004

DOE Leadership-Class Computing Capability for Science will be Developed at Oak Ridge National Laboratory

WASHINGTON, DC – Energy Secretary Spencer Abraham announced today that the U.S. Department of Energy (DOE) will grant Oak Ridge National Laboratory (ORNL) in Oak Ridge, Tenn., and its development partners, Cray Inc., IBM Corp. and Silicon Graphics Inc., $25 million in funding to begin to build a 50 teraflop (50 trillion calculations per second) science research supercomputer. The department selected ORNL from four proposals received from its non-weapon national labs.

"This new facility will enable the Office of Science to deliver world leadership-class computing for science," said Secretary of Energy Spencer Abraham. "It will serve to revitalize the U.S. effort in high-end computing."

The supercomputer will be open to the scientific community for research.

ORNL won the award in a peer-reviewed competition with three other Office of Science national laboratories. In response to a solicitation, the four laboratories submitted proposals designed to improve substantially the national research community's computing capability – or ability to perform the largest, most complex simulations – and thereby enhance prospects for important research advances and scientific breakthroughs in all science disciplines supported by DOE and other federal science agencies.

Federal Agency Responses
one perspective

- DARPA
  – Research and develop high productivity computing systems
- NSF
  – Create cyberinfrastructure to support open computational science and engineering
- NASA
  – Focus on HPC for aeronautics and space applications
- DOE (Office of Science)
  – Lead the general effort to apply HPC to pathfinding applications in science and engineering and to discover new applications

- **Leadership Class Computing**
  - Oak Ridge National Laboratory
    - http://nccs.gov/leadership/index.html
  - Argonne National Laboratory
    - http://www.alcf.anl.gov/

- **INCITE**
  - Innovative and Novel Computational Impact on Theory and Experiment
    - http://www.sc.doe.gov/ascr/incite/index.html

- **SciDAC**
  - Scientific Discovery through Advanced Computing
    - http://www.scidac.gov/

- **E^3SGS Initiative**
  - Simulation and Modeling at the Exascale - for Energy, Ecological Sustainability and Global Security

Hardware

U.S. Department of Energy

Office of Science

Top Dozen Supercomputers

Advanced Scientific Computing Research Program

TOP500 List - June 2007 (1-100)

R_{max} and R_{peak} values are in GFlops. For more details about other fields, check the TOP500 description.

Rank	Site	Computer	Processors	Year	R_{max}	R_{peak}
1	DOE/NNSA/LLNL United States	BlueGene/L - eServer Blue Gene Solution IBM	131072	2005	280600	367000
2	Oak Ridge National Laboratory United States	Jaguar - Cray XT4/XT3 Cray Inc.	23016	2006	101700	119350
3	NNSA/Sandia National Laboratories United States	Red Storm - Sandia/ Cray Red Storm, Opteron 2.4 GHz dual core Cray Inc.	26544	2006	101400	127411
4	IBM Thomas J. Watson Research Center United States	BGW - eServer Blue Gene Solution IBM	40960	2005	91290	114688
5	Stony Brook/BNL, New York Center for Computional Sciences United States	New York Blue - eServer Blue Gene Solution IBM	36864	2007	82161	103219
6	DOE/NNSA/LLNL United States	ASC Purple - eServer pSeries p5 575 1.9 GHz IBM	12208	2006	75760	92781
7	Rensselaer Polytechnic Institute, Computional Center for Nanotechnology Innovations United States	eServer Blue Gene Solution IBM	32768	2007	73032	91750
8	NCSA United States	Abe - PowerEdge 1955, 2.33 GHz, Infiniband Dell	9600	2007	62680	89587.2
☆ 9	Barcelona Supercomputing Center Spain	MareNostrum - BladeCenter JS21 Cluster, PPC 970, 2.3 GHz, Myrinet IBM	10240	2006	62630	94208
☆ 10	Leibniz Rechenzentrum Germany	HLRB-II - Altix 4700 1.6 GHz SGI	9728	2007	56520	62259.2
11	NNSA/Sandia National Laboratories United States	Thunderbird - PowerEdge 1850, 3.6 GHz, Infiniband Dell	9024	2006	53000	64972.8
☆ 12	Commissariat a l'Energie Atomique (CEA) France	Tera-10 - NovaScale 5160, Itanium2 1.6 GHz, Quadrics Bull SA	9968	2006	52840	63795.2

Japan

Background

- Supercomputing Technology has been selected as one of the "Nation's Key Technologies"
- To enhance the competitiveness of Japanese science, technology and industry.
- To maintain capability of development of a supercomputer within the country, and to enable continuous development.

Current MEXT Supercomputer Installations

(MEXT = Ministry of Education, Culture, Sports, Science and Technology)

Advanced Scientific Computing Research Program

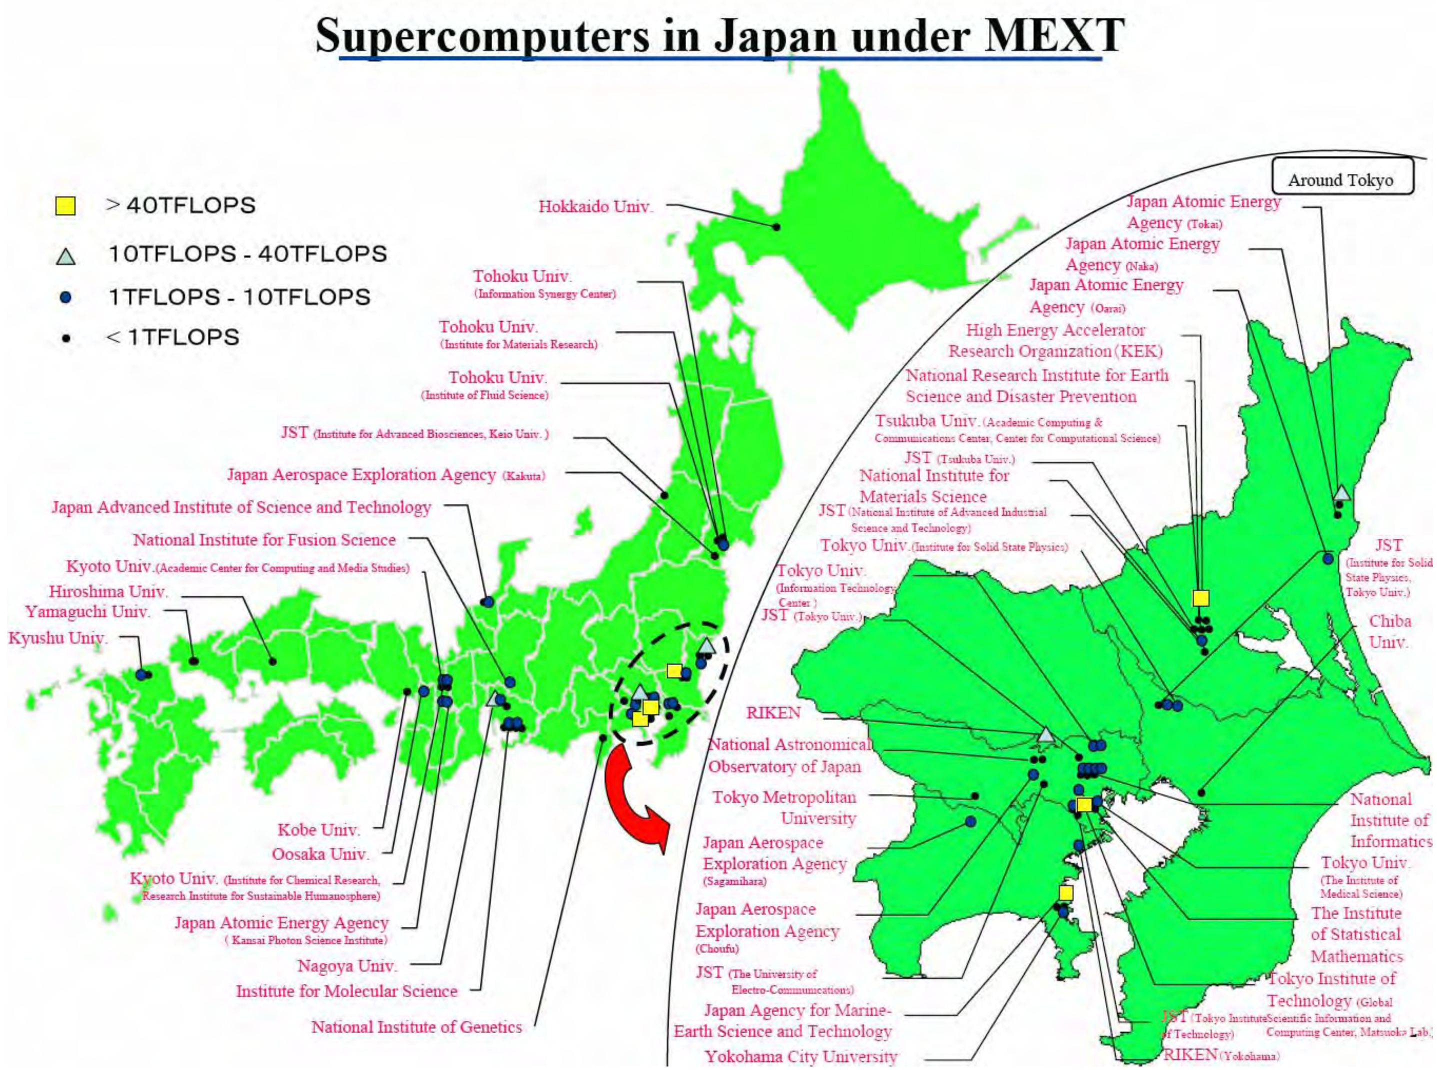

MEXT Vision: Continuous Development of Supercomputers

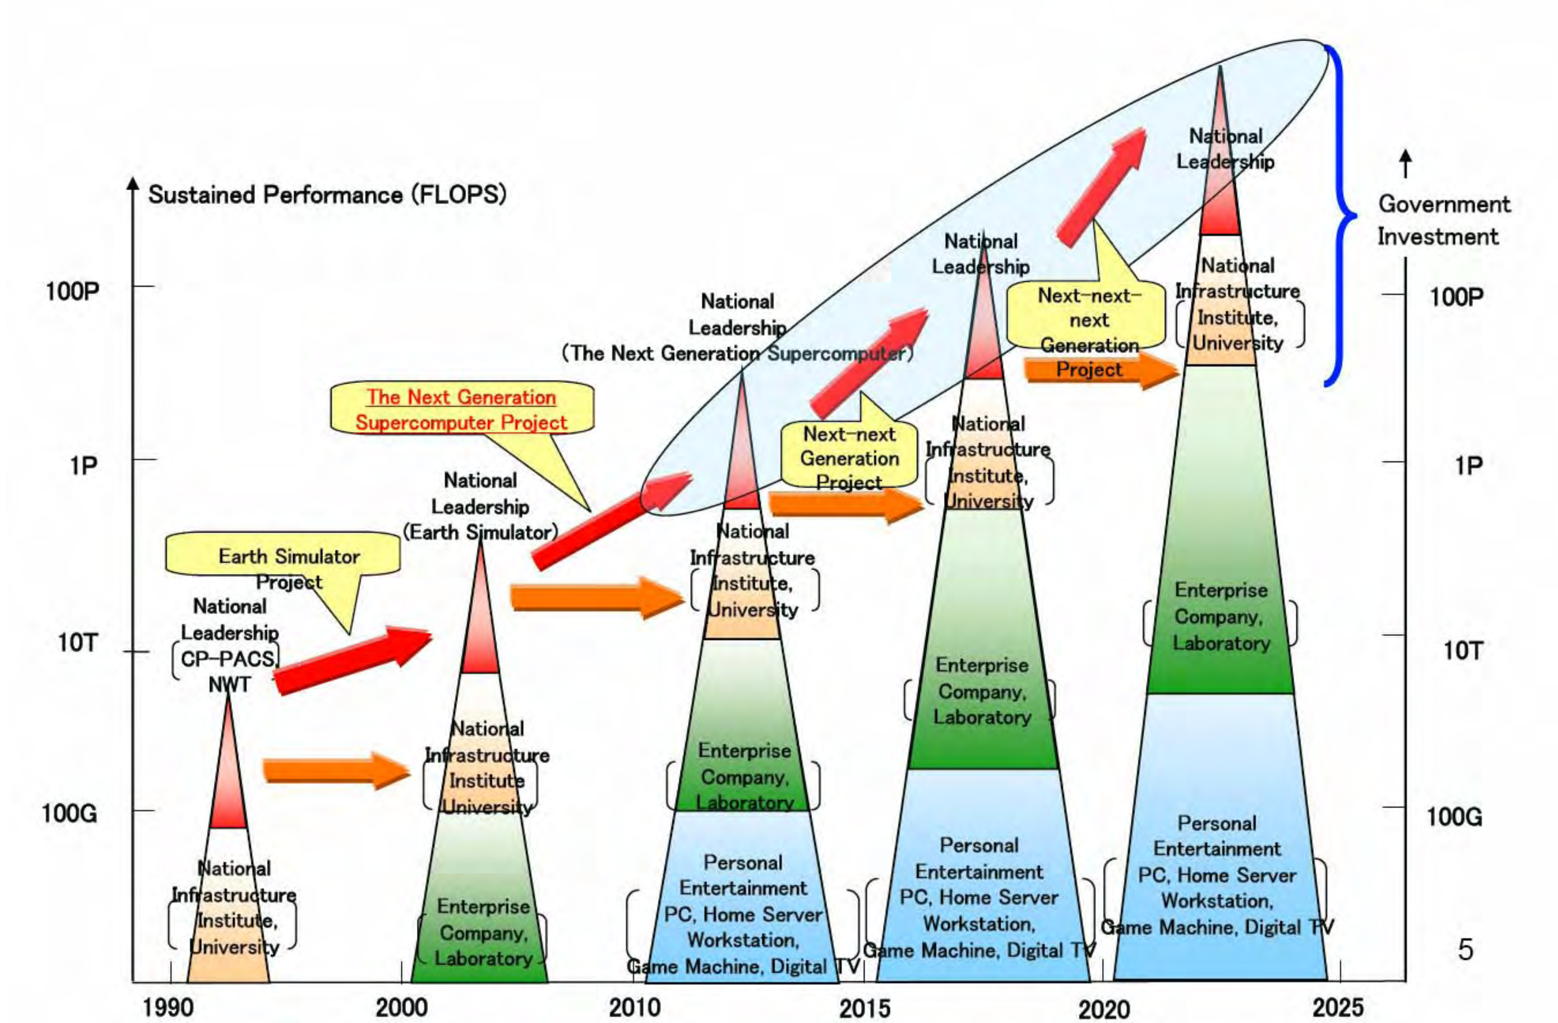

RIKEN Next-Generation Supercomputer Project

- "The Next-Generation Supercomputer R&D Center (NSC) was established by RIKEN on January 1, 2006, to design and build the fastest supercomputer in the world. The new supercomputer will be used in a wide range of fields of scientific research, from life sciences to nanotechnology."

  http://www.nsc.riken.jp/index-eng.html

- "RIKEN is aiming at the development and usage of a 10 PFLOPS-class Next-Generation Supercomputer as the highest-standard generic computer to lead further supercomputer development"

Next-Generation Supercomputer Project Organization

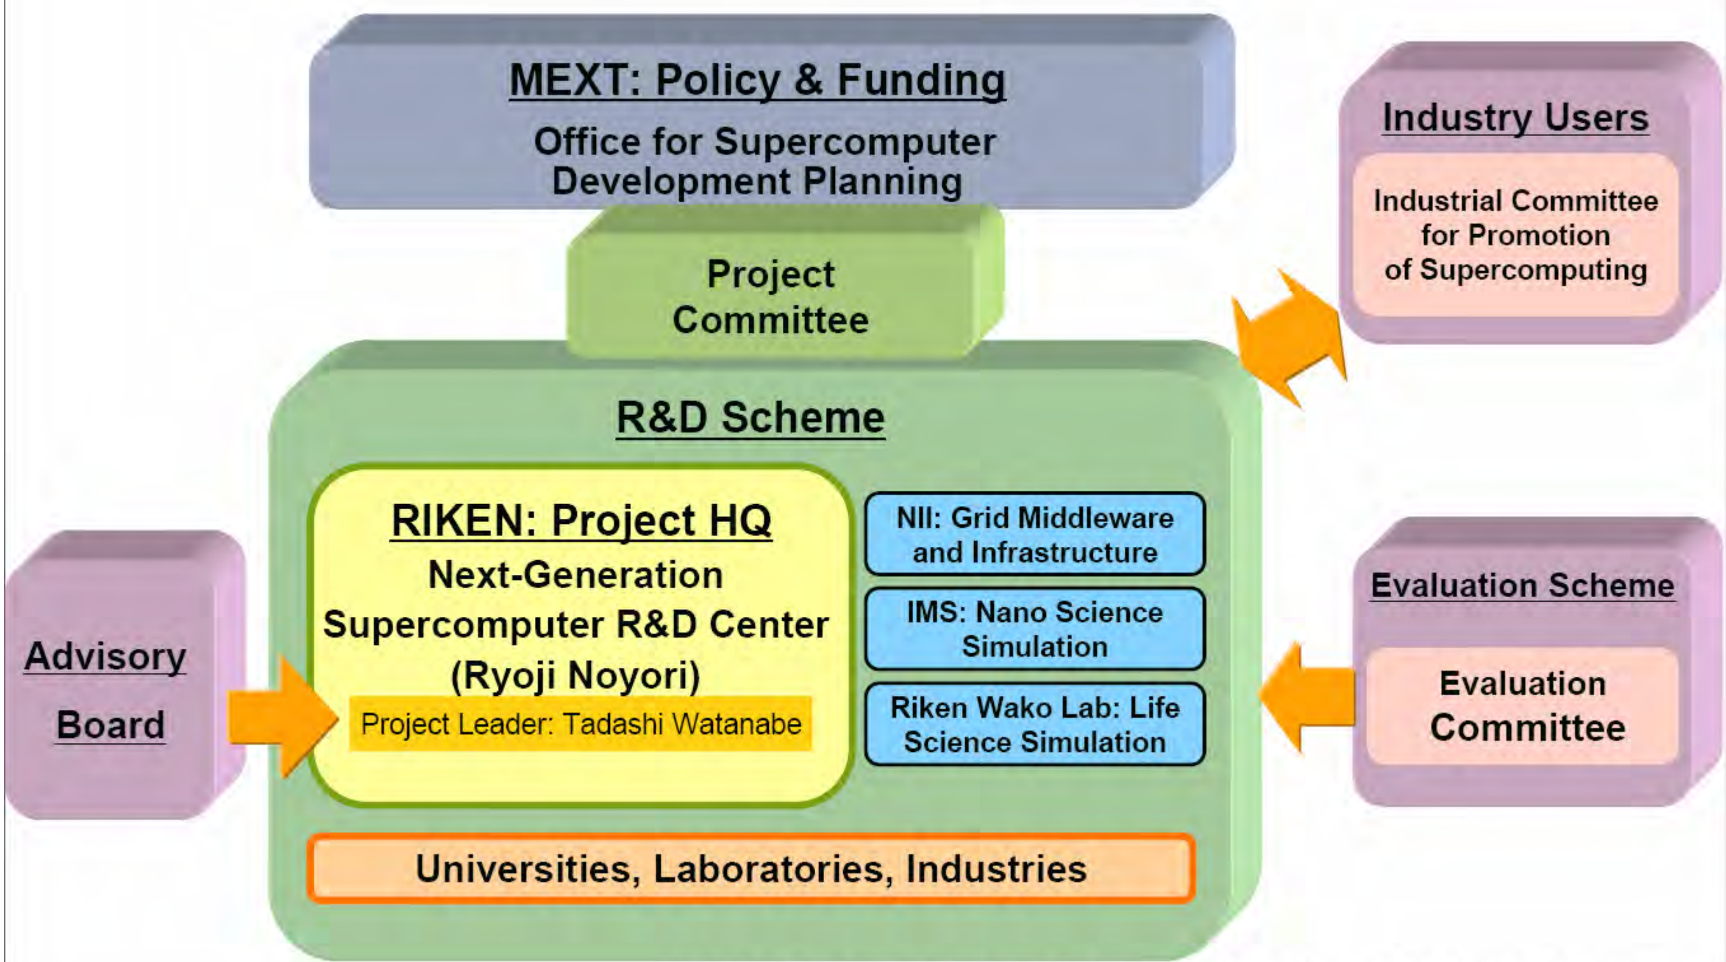

(Note) NII: National Institute of Informatics, IMS: Institute for Molecular Science

21 Target Applications

June 22th 2007
RIKEN Next -Generation
Supercomputer R&D Center,
Strategy Council,
Application Committee

In order to make the best use of resources of the Next-Generation Supercomputer, we have investigated which research fields and what application programs need the Peta-scale computational power.

Our aim is to develop an architecture design of the Next-Generation Supercomputer based on the characteristics of each program.

The RIKEN Application Committee has selected 21 application programs appropriate for execution on the Next-Generation Supercomputer.

In the project, we have developed benchmark-test programs regarding these 21 applications, and examined the proposed architectures based on benchmark-test results.

U.S. Department of Energy

Office of Science

Target Applications List

http://www.nsc.riken.jp/target-application/target-application-eng.html

Advanced Scientific Computing Research Program

Target Applications		
No.	Title	PDF
1	Cavitation	
2	COCO	
3	FrontFlow/Blue Large Eddy Simulation (LES)	
4	FrontSTR	
5	GAMESS FMO	
6	GNISC	
7	LANS	
8	LatticeQCD	
9	MC-Bflow	
10	MLTest	
11	Modylas	
12	NICAM	
13	NINJA/ASURA	
14	Octa	
15	PHASE	
16	ProteinDF	
17	RISM/3D-RISM	
18	RSDFT	
19	Seism3D	
20	sievgene/myPresto	
21	SimFold	

Industrial Applications not possible without Next-Generation Supercomputer

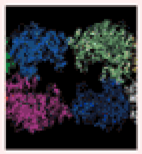

Development of New Pharmaceuticals
Newly viable high polymer simulation enables volume selection of prospective pharmaceutical constituents in a short period of time.

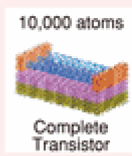

Development of New Semiconductor Materials
Realization of individual atom simulation revolutionizes material development, which currently relies on trial and error methods.

Car Chassis Structural Strength Analysis
A 1 to 2-hour automated simulation can replace the current crafting of car crash simulation settings that usually takes several months, and this leads to safety improvements and increased industrial competitiveness.

Aseismic Analysis of Bridges and Other Complex Structures
This type of analysis enables more precise and realistic structure collapse predictions, contributing to the reduction of earthquake damage.

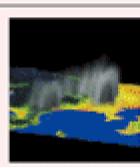

Prediction of Typhoon Paths and Localized Storms
Weather simulations within a 1 km^2 area enable prediction of localized storms and typhoon paths with greatly increased precision.

Next-Generation Integrated Life Simulation Software

Advanced Scientific Computing Research Program

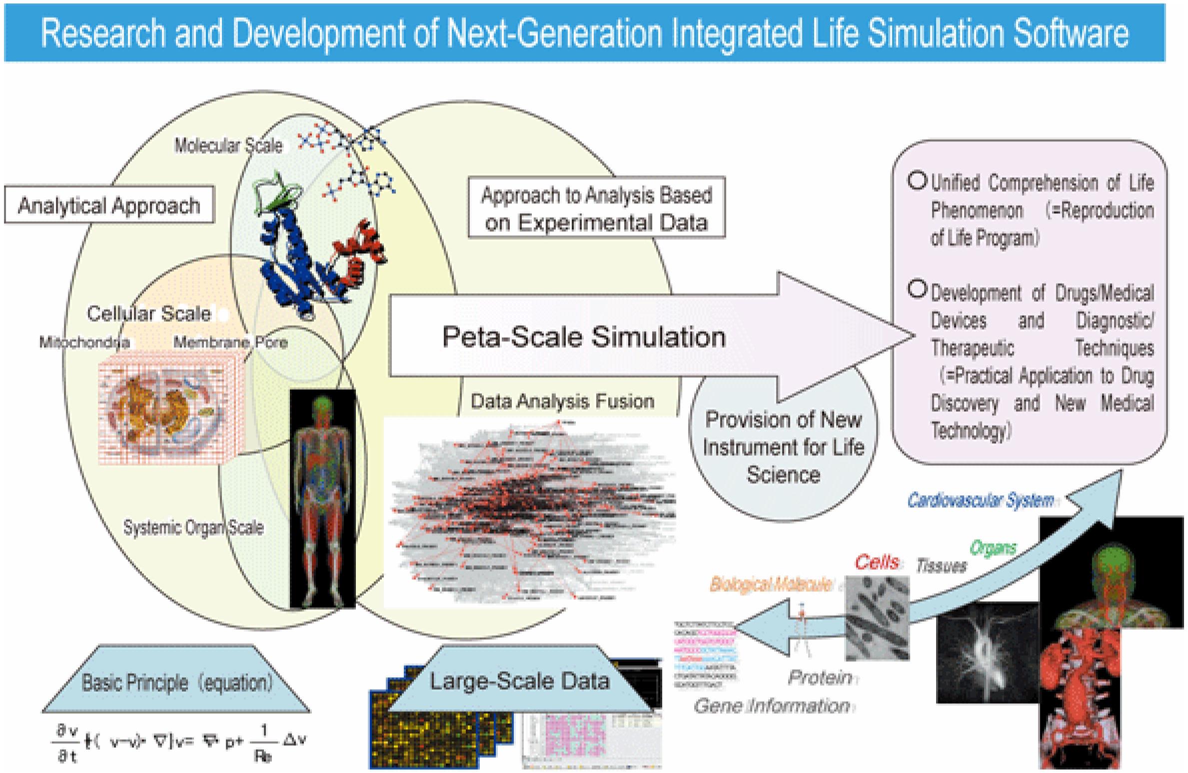

U.S. Department of Energy

Office of Science

Next-Generation Integrated Nanoscience Simulation Software

Advanced Scientific Computing Research Program

The Institute for Molecular Science plays a central role in the development of the "Grand Challange to Next-Generation Integrated Nanoscience" simulation software, which is a major challenge for next-generation nanoscience. The institute is establishing a basis for computational nanoscience, focusing in particular on (1) next-generation nano-informational function and materials, (2) next-generation nano-biomolecules, and (3) next-generation energy. The software is designed to make full use of Next-Generation Supercomputer. A nationwide team from national institutes, universities and industry has been organized for this project.

http://groups.ims.ac.jp/index.html#

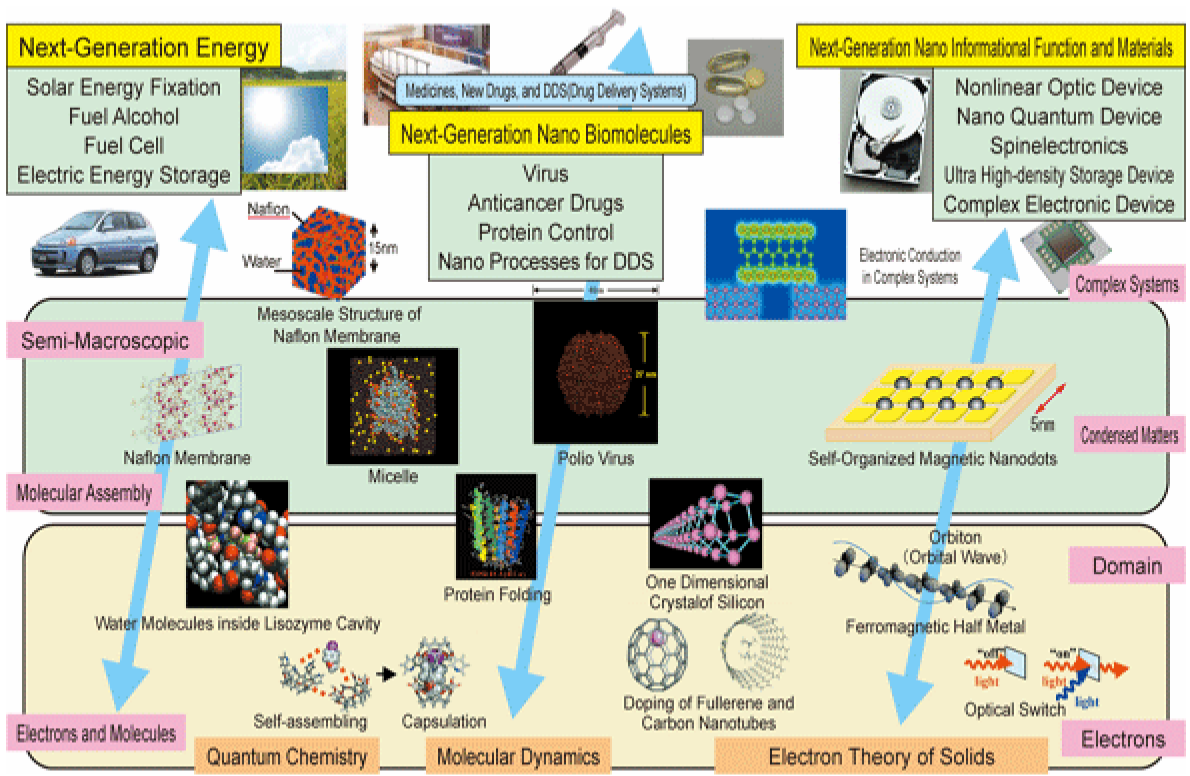

Cyber Science Infrastructure

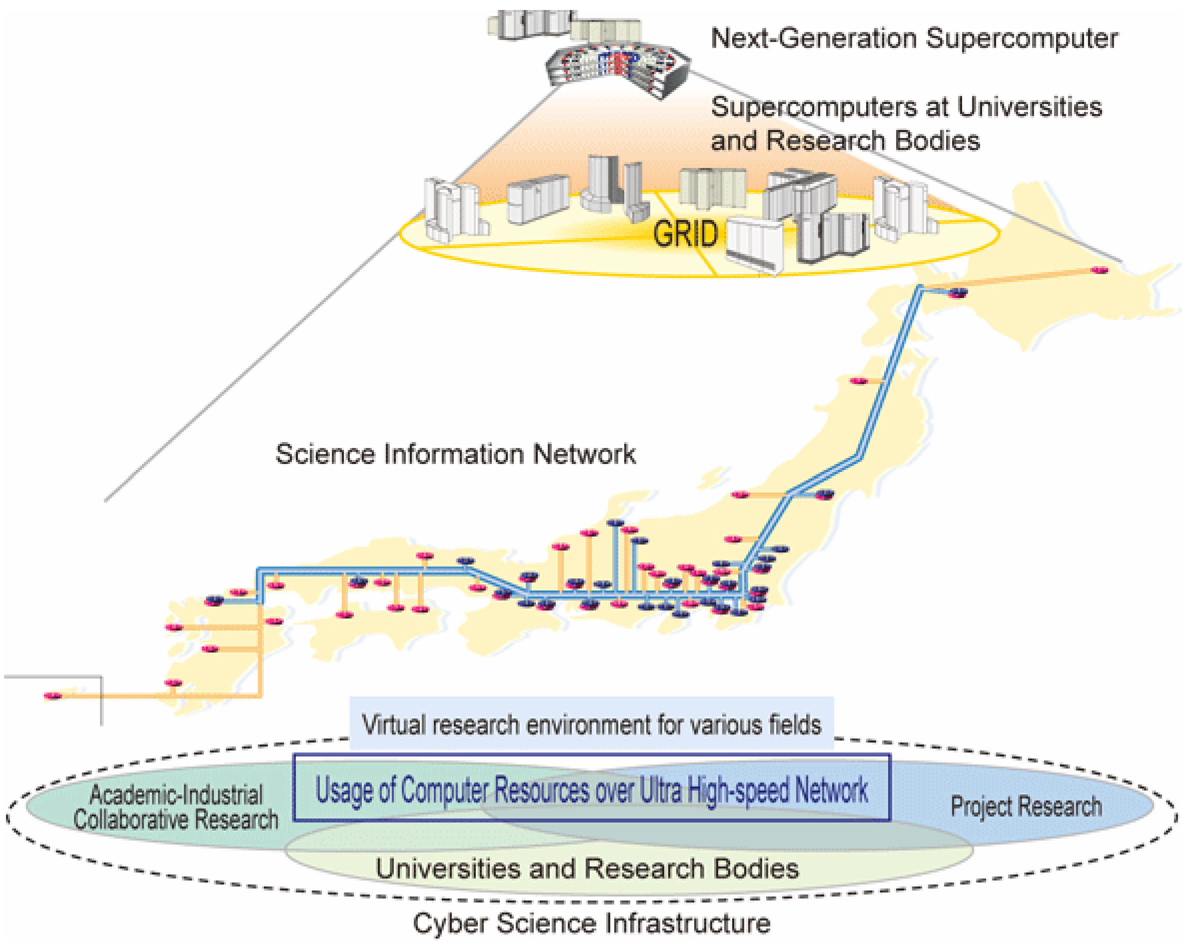

Interlinking with the Earth Simulator

"The Earth Simulator owned by the Japan Agency for Marine-Earth Science and Technology (JAMSTEC) started operations in 2002, and it is still one of the world's fastest supercomputers. In order to harness the expertise gained through the development and operation of the Earth Simulator, RIKEN and JAMSTEC have concluded a partnership agreement. RIKEN and JAMSTEC will jointly develop application software that runs on the Next-Generation Supercomputer and perform a variety of work under this agreement."

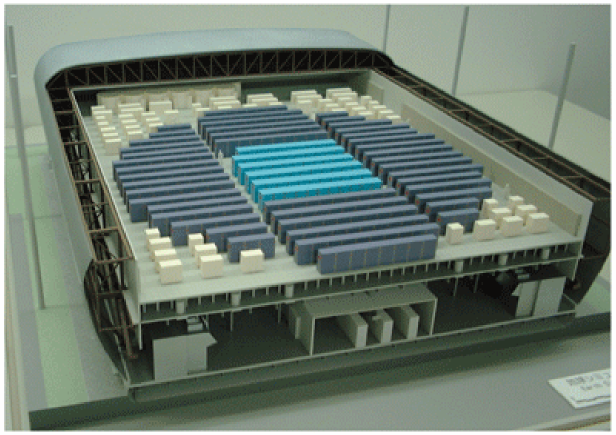

Towards Wider Industrial Usage

"To promote industrial applications, the Industrial Committee for Promotion of Supercomputing was established in December 2005, and more than 150 corporations from a wide range of industries, including chemical, medical, and manufacturing participate in the council.

RIKEN will promote industrial usage of the Next-Generation Supercomputer working together with the Council."

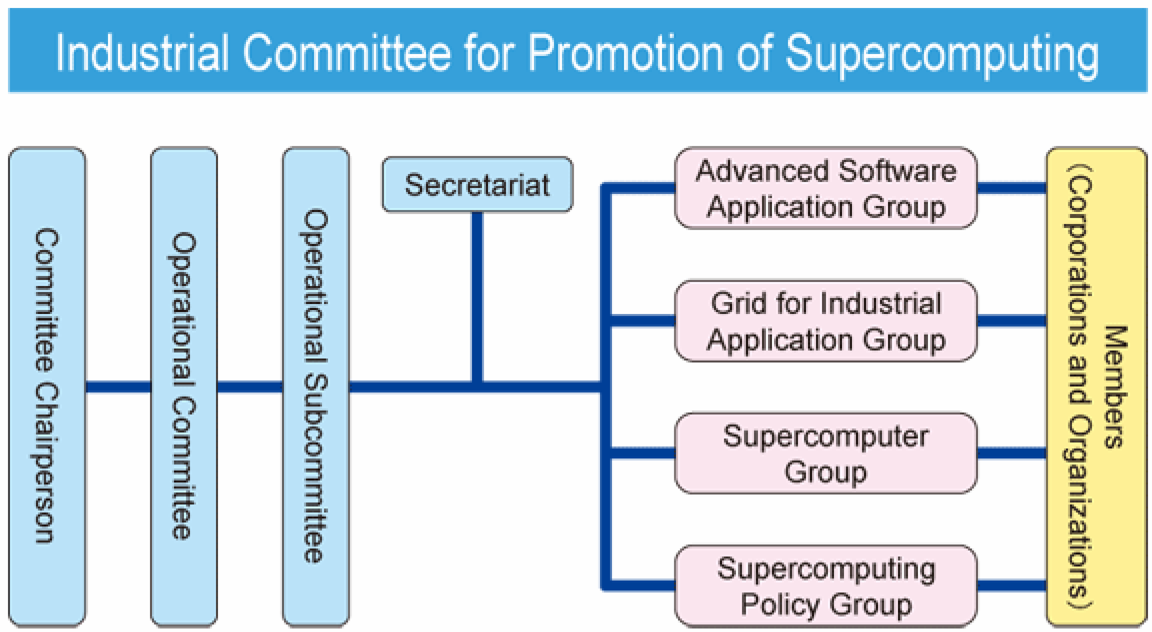

Center of Excellence for Supercomputing

"RIKEN is aiming to establish a Center of Excellence (COE) for Next-Generation Supercomputer research and human resource development. From hardware component technology to application software development, RIKEN will establish a research and development environment and develop human resources, collaborating with researchers and research bodies from countries all around the world."

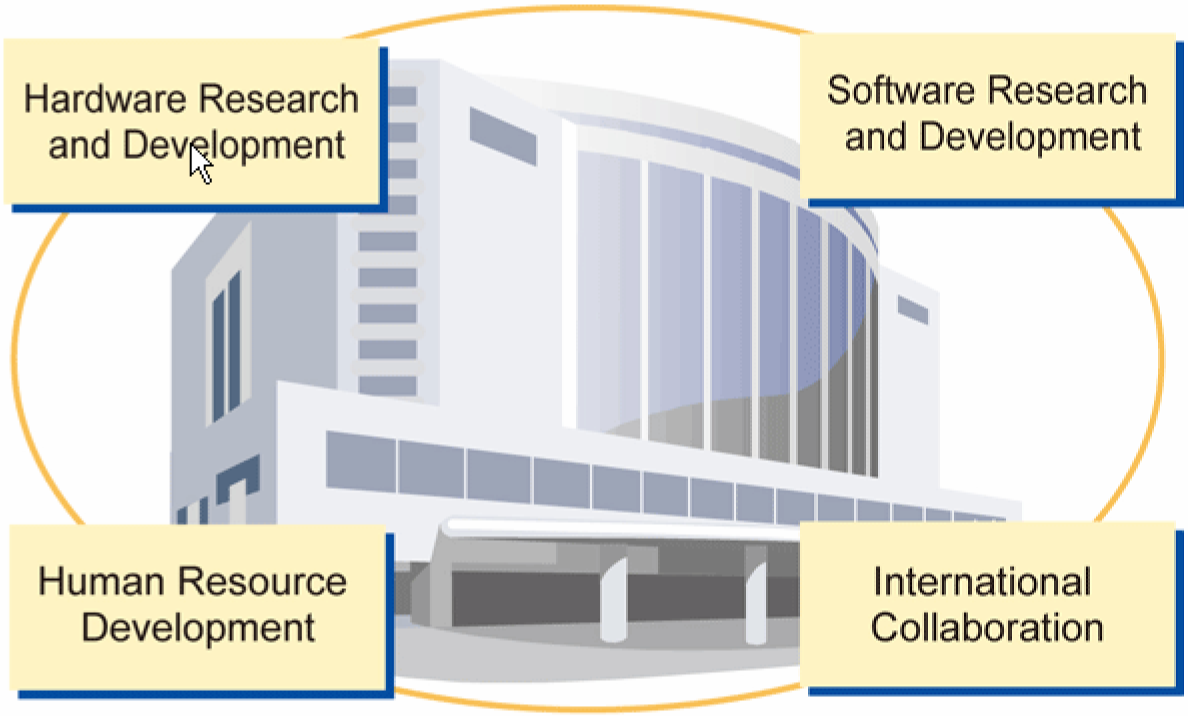

European Union

Partnership for Advanced Computing in Europe

PACE


Partnership for Advanced Computing in Europe MoU signed April, 17th 2007
PACE
Partnership for Advanced Computing in Europe

U.S. Department of Energy

Office of Science

ESFRI – European Roadmap for Research Infrastructures

Advanced Scientific Computing Research Program

The European roadmap is a first exercise at European level and is the result of wide stakeholder consultation.

1000 high-level experts from all fields of research were involved and consulted in the process of preparing the roadmap, of which 200 were involved in the peer-review.

Research Infrastructures are one of the crucial pillars of the European Research Area, in particular for capacity building.

Motivation for HPC Focus

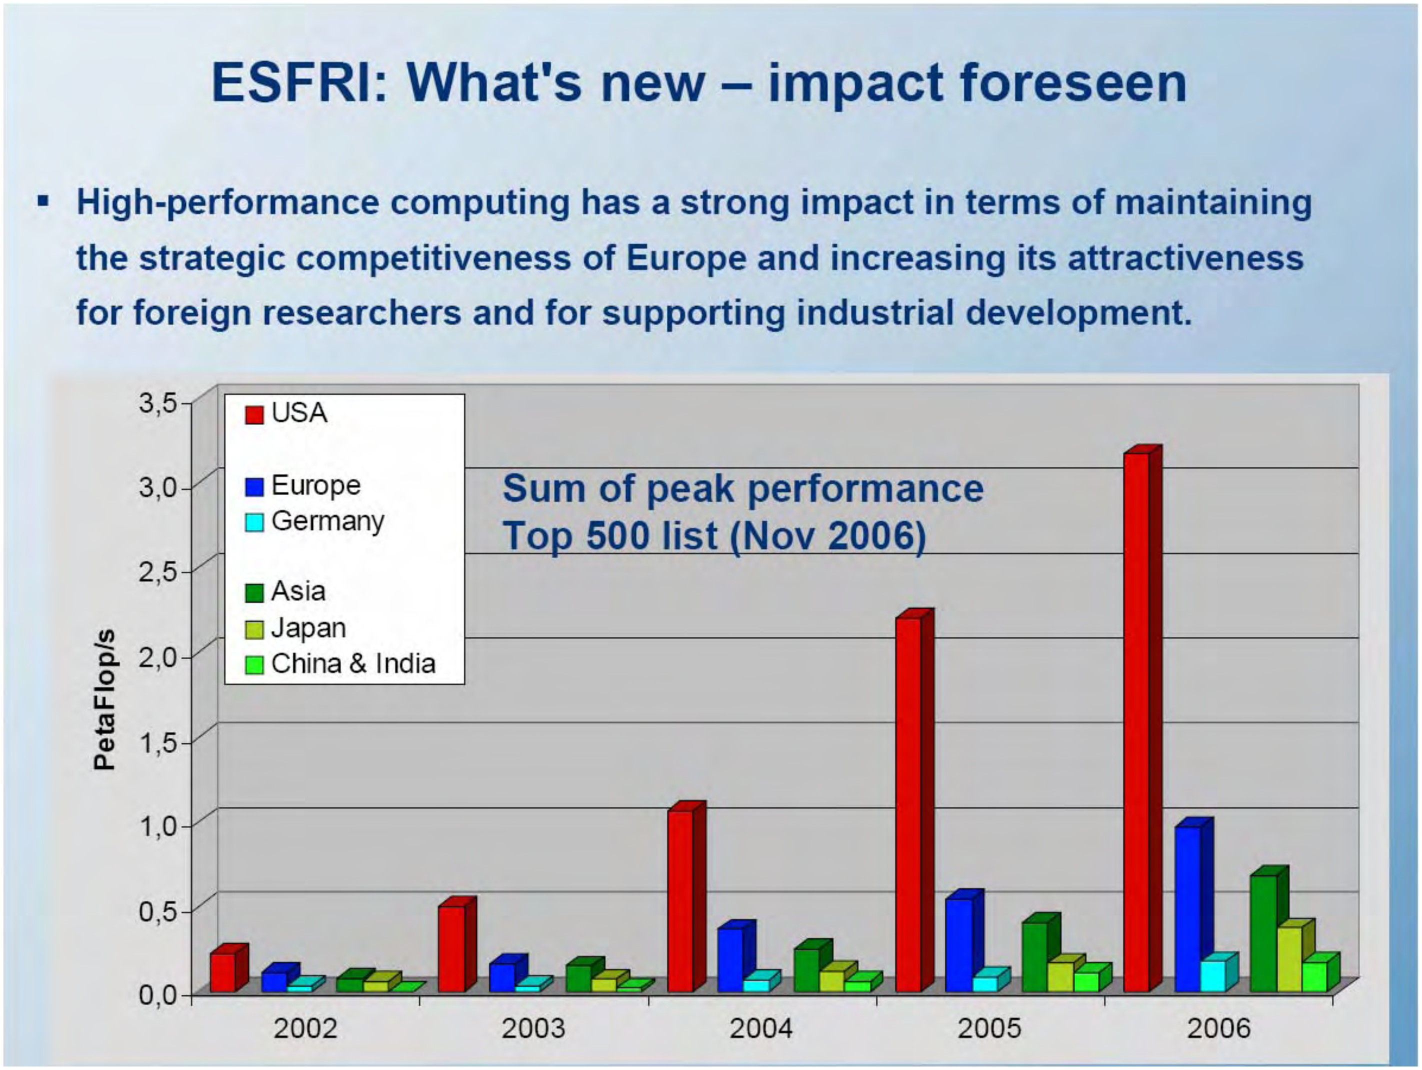

Strategy

- The high-end (capability) resources should be implemented every 2-3 years
- Construction cost 200-400 Mio. €, Running cost 100 – 200 Mio. €, i.e. about 1,6 – 4,0 billions € for a 10 year period.
- with supporting actions in the national/regional centers to maintain the transfer of knowledge and feed projects to the top capability layer

PACE Vision

From cooperative High Performance Computing in Europe

to

leadership class

European HPC-facilities,

The Large Scale European Supercomputing Infrastructure

integrated in a High Performance Computing Service

– the European Ecosystem -

PACE Membership

Advanced Scientific Computing Research Program

PACE Next Steps

The next step: European leadership computing facilities embedded in an European Ecosystem

- European access to a tier 0 capability HPC research infrastructure reaching and exceeding the PetaFlop/s and maintaining world-wide competitiveness
- A sustainable infrastructure based on a sustainable and funded renewal process to follow technological progress
- Complementing the national HPC centres and services

➔EU-support is required to open access to all European countries and to industries and is a chance to promote HPC-competence on a European level for a competitive European IT-Knowledge

Simulation Laboratories

Advanced Scientific Computing Research Program

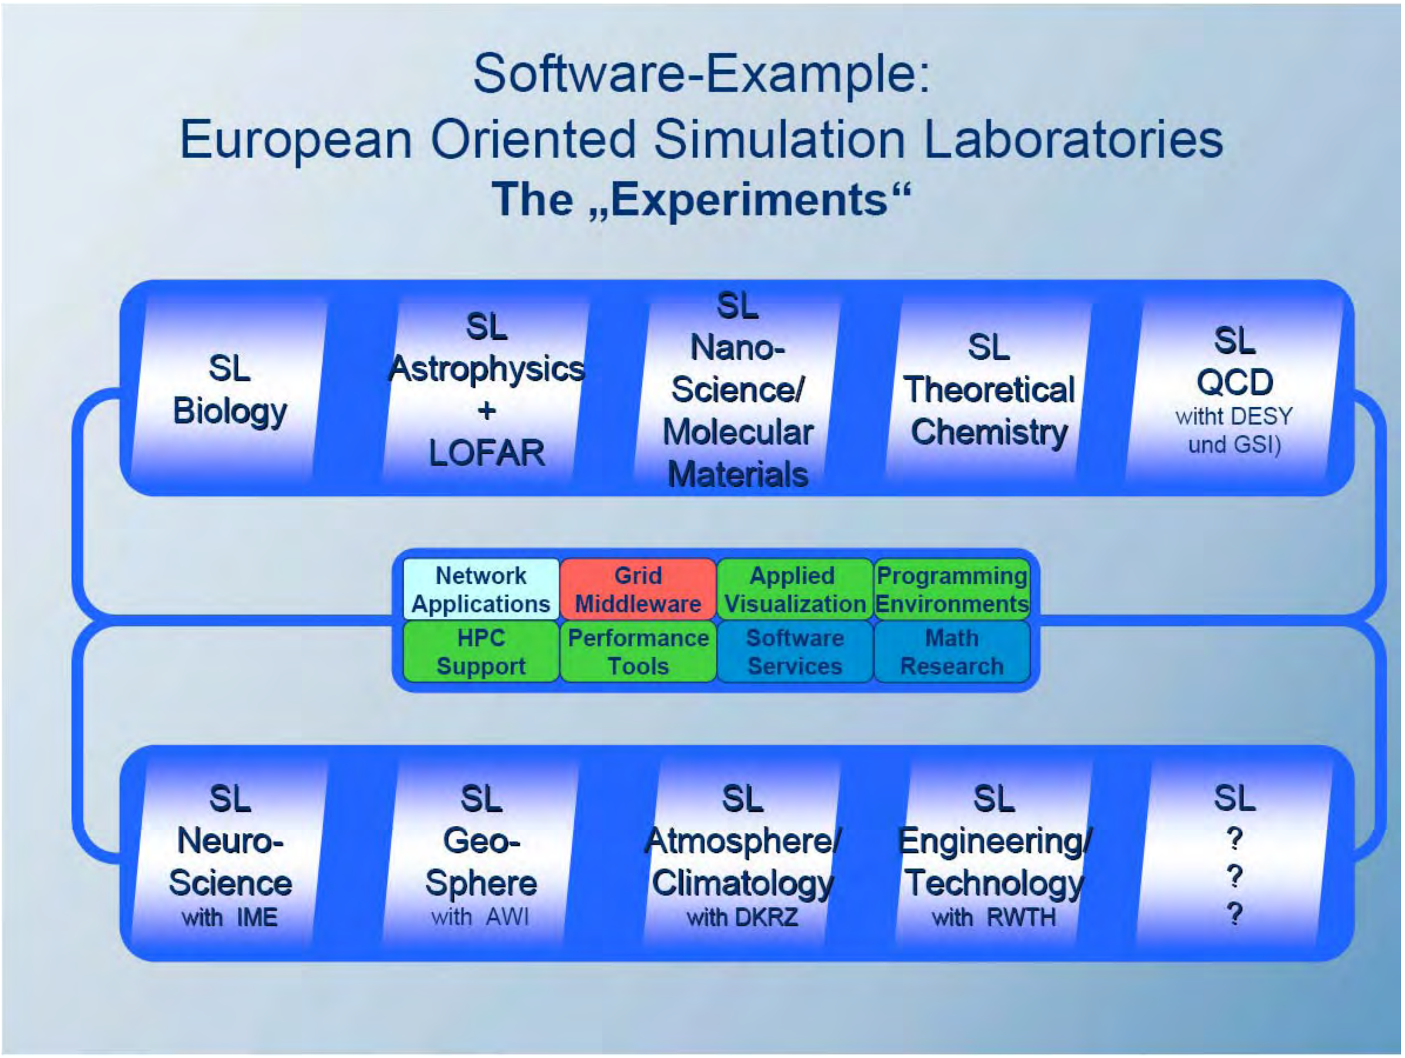

Computational Science Education

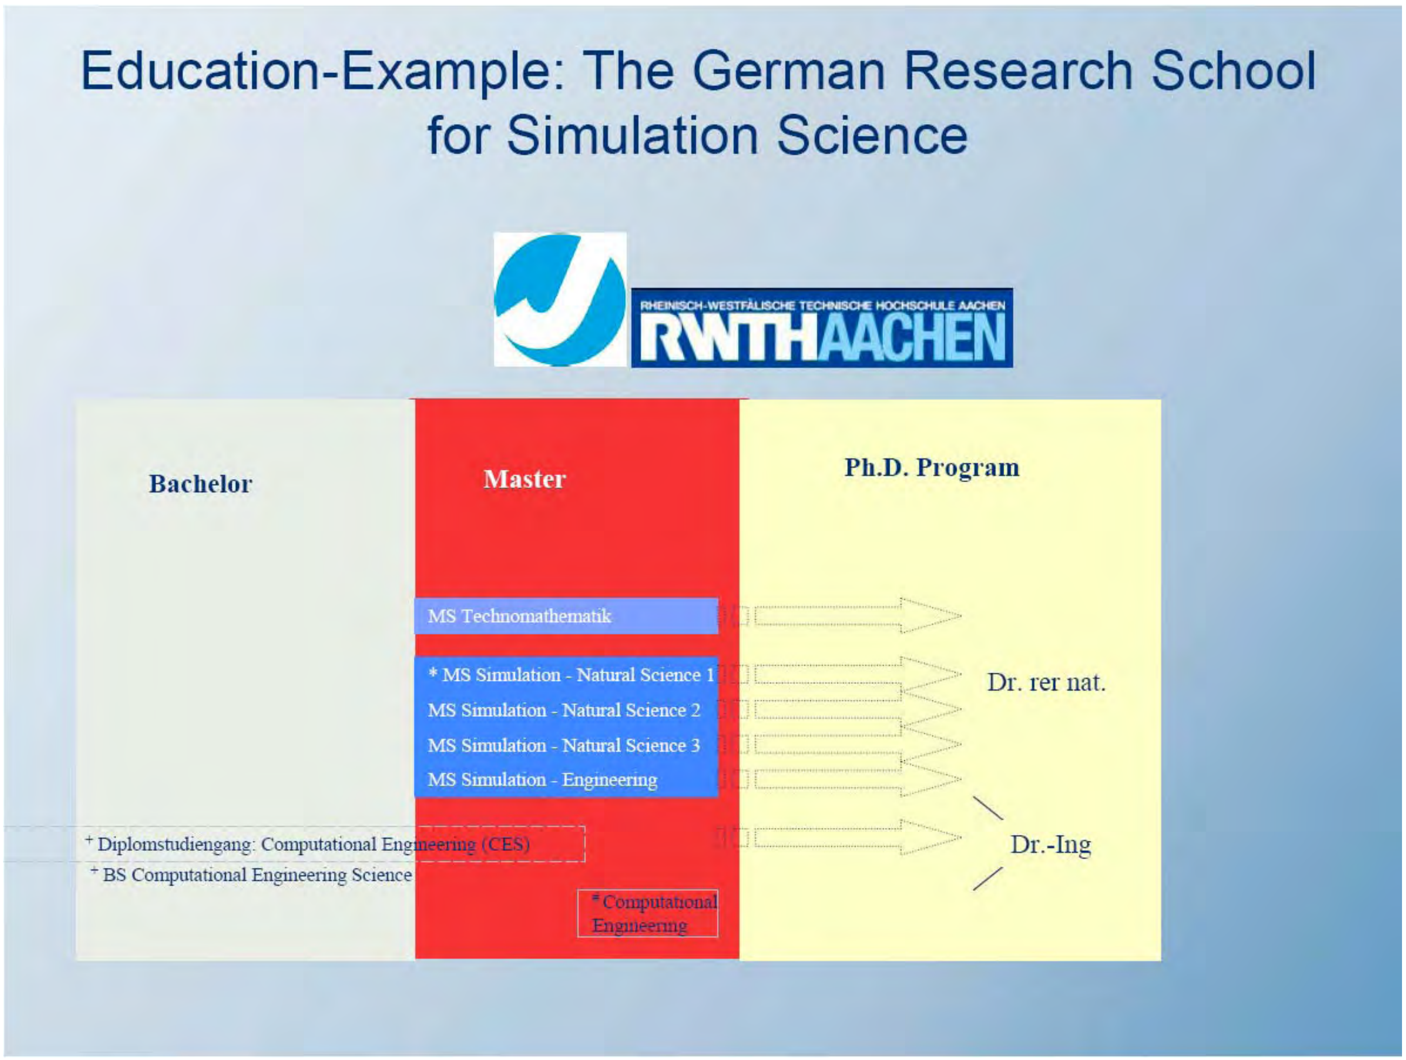

PACE Roadmap

- **April, 27th 2007: Signing MoU**
- **May 2th, 2007: Delivering EU-FP7-Call supporting the preparatory phase**
- **June 2007 – May 2009: Building the European legal entity**
- **2008 – 2009: Building prototypes of Petaflop Computers**
- **2009 – 2010: Starting the production phase with Petaflop computing, servicing the scientific communities**

Conclusions

- The HPC hardware race is on.
- Applications development has become much more central to HPC planning.
- Both Japan and the European Union are building out their HPC infrastructure.
- There is a growing body of opinion in the EU that Europe needs an independent capability to produce HPC hardware.