

Approaching EXASCALE



Dr. Robert Rosner, director of Argonne National Laboratory, discusses exascale computing, partnerships with industry, energy efficiency in HPC, and getting people excited about science.

SciDAC Review: *In your past three years as director of Argonne National Laboratory, you have initiated a series of five new strategic initiatives, including one in exascale computing technology and another in sustainable energy production. Why did you identify these as especially important for the laboratory? Do you see the two as separate or as complementary efforts?*

Dr. Robert Rosner: The initiatives are really part of a larger picture aimed at transforming Argonne, and ensuring that its missions align with the missions of our principal sponsor, namely the Department of Energy, and in particular, the DOE Office of Science. In the two areas you mention explicitly, it turns out that computing at the exascale will play a transforming role in the underlying science and engineering disciplines at the heart of energy production—from things nuclear to the fundamental chemistry of catalysts.

Scientists currently are using petaflops-scale computers—the Blue Gene/P at Argonne and the Cray at Oak Ridge. Exascale is the next big step. What will exascale computing enable researchers to do that petascale does not?

Every performance index for computational capability—from the megaflops of the 1970s to the petaflops of today—can be associated with certain key computational capabilities that are enabled by reaching that performance index, and the transition to exaflops within the next decade will not be different. Some examples: exaflop simulations will enable direct numerical simulation (DNS) of the Navier–Stokes equations for realistic laboratory fluids into the fully turbulent regime—without any resort to subgrid modeling; it will allow us to carry out

magnetohydrodynamic simulations for (magnetic) Prandtl numbers characteristic of realistic liquid metals. In both cases, this will allow us to go beyond today’s models for what happens at currently unresolved physical scales to explore the basic physics connecting the dissipative length scales to the energy-containing length scales in realistic (magneto)fluids—a key step in placing much of, for example, astrophysics and engineering science currently reliant on models on a firmer basis established by true physical understanding. There are many similar examples one can draw upon from chemistry, material science, biology, and—perhaps even more exciting—the social sciences, but these two examples come from the science I know best from my own work.

In your experience as an astrophysicist, you have seen computers play an ever-increasing role. Would you please describe how high-performance computing has advanced your own research?

During much of my own career, the theoretical work I’ve done has tried to bridge what can be done by analytical (that is, pencil-and-paper) work to models that apply to real (astro)physical systems. This bridging is difficult, to say the least, and has in the past relied heavily on physical intuitions gained from laboratory experiments—but these are limited in what can be practically explored (and measured!), and are expensive to boot. What modern simulations have enabled is the ability to go beyond what can be done in the laboratory, to broaden and strengthen our abilities to construct models for astrophysical systems, and indeed to broaden the range of our physical intuition into regimes that we have not had (experimental) access to before. Of course, experiments remain key—they are the essential tools for making sure that our simulation results make sense: after all, computer simulations will always provide results, and the trick is to make sure that they are correct results—and that is a key role played by experimentation as part of code validation.

Argonne has announced that a new Theory and Computing Sciences Building will be built at the Laboratory. What role do you envision this building will play in enabling breakthroughs in science and engineering? What is encompassed by the term “computing sciences?” How do you see theory and computation working together? And how does experiment fit in?

The Theory and Computing Sciences (TCS) Building really serves several independent purposes. First, it (obviously) will

INTERVIEW: Dr. Robert Rosner

Biography in Brief: Dr. Robert Rosner

Dr. Robert Rosner became director of Argonne National Laboratory on April 18, 2005. Prior to this position, he served as chief scientist at Argonne for three years and was the architect of Argonne's 20-year strategic plan for science and technology.

An internationally recognized leader in science and engineering, Dr. Rosner is frequently interviewed on subjects ranging from energy research and accelerator science to high-performance computing and nanotechnology. His interviews have appeared on CBS, National Public Radio, and E&E News, and he has been quoted in *Crain's Chicago Business*, *Inside Energy*, and the *Chicago Tribune*.

Background and Accomplishments

Dr. Rosner earned a Ph.D. in physics from Harvard University in 1976. His scientific work has centered on plasma physics, including the development of numerical tools for simulating astrophysical phenomena. He was chairman of astronomy and astrophysics at the University of Chicago from 1991 to 1997, and in 1998 he was named the university's William E. Wrather Distinguished Service



Figure 1. Dr. Robert Rosner.

Professor. Dr. Rosner led a collaboration of Argonne and University of Chicago scientists in establishing the Center for Astrophysical Thermonuclear Flashes (FLASH Center), which he directed from its founding in 1997 until 2002. Funded by the Advanced Simulation and Computing program of the National Nuclear Security Administration, the FLASH

Center is addressing the 50-year-old problem of exploding stars ("Computing the Detonation of a White Dwarf Star," *SciDAC Review*, Summer 2007, p10).

Awards and Honors

Dr. Rosner was elected to the American Academy of Arts and Sciences in 2001 and is a Fellow of the American Physical Society. In 2004, he was the Rothschild visiting professor at the Newton Institute for Mathematical Sciences at Cambridge University. He has received honorary doctoral degrees from the Illinois Institute of Technology (2006) and from Northern Illinois University (2007).

Dr. Rosner serves on a number of advisory committees, including the NAS/NRC Committee on Evaluation of Quantification of Margins and Uncertainty (QMU) Methodology, the External Advisory Committee for the National Ignition Facility at Lawrence Livermore National Laboratory, the Scientific Advisory Committees for the Max Planck Institute for Solar System Research (Lindau, Germany) and the Astrophysical Institute Potsdam (Germany), and the Senate of the Helmholtz Association (Germany).

be the home of Argonne's program in Leadership Computing, and as such, will provide the infrastructure necessary to maintain this program into the next decade. Second, it will provide the venue for theorists, experimentalists, computer scientists, and (applied) mathematicians to mix, to interact, to get to know each other at the level of peers—all crucial ingredients in our effort to strengthen the critical interaction between the (physical) science community and the folks that provide the wherewithal to simulate. The essential element of this interaction is for these peers to recognize and appreciate the intellectual challenges—their difficulties, and their intrinsic attractiveness—recognition and appreciation being the key ingredients in getting first-rate researchers coming from very disparate disciplines to work together effectively. And experiments are a natural partner in all this, as I've already alluded to: experiments are ultimately the source of all of our questions, and the tool by which we can test our ability to "get the right answer" via simulations.

Prior to the establishment of the Argonne Leadership Computing Facility, Argonne and Lawrence Livermore National Laboratory began working closely with IBM to develop a series of computing systems based on IBM's Blue Gene platform. Argonne currently is partnering with SiCortex in exploring another innovative computer design. How

do you view such partnerships with industry? What do you think the national labs can contribute?

Many (if not most) scientists and engineers today view the leading-edge computers as a given, that is, as machines that have been handed to them without much input on their part in the performance, design, etcetera of these machines. In the beginning of the computer era, that is, in the 1940s and early 1950s, this was not so: the science and engineering communities were very tightly coupled to the industry building forefront computers—indeed, the early "super-computers" were designed and built with the direct input from the ultimate users. But since then, we've gone through a long era during which such tight coupling was more notable by its absence than by its presence. I think that from the perspective of computer architecture (and languages), we are entering an era where this close coupling will again be essential—as viewed from the perspectives of both the vendors and the users. One important reason is that the architectures are becoming sufficiently complex (and the software needed for effective utilization of this hardware sufficiently arcane) that isolated scientists will not be in the position to write effective code for these leadership computers on their own. That is, the "marriage" of computer scientists, applied mathematicians, and discipline (physical) scientists

and engineers has turned out to be essential in order to write modern codes that take effective advantage of the power of modern massively parallel machines—and this marriage requires in my opinion an inclusion of the folks who are leading the design of next-generation leadership machines. As for the role of national labs in these “marriages,” I think it is clear from recent history that the labs have proven to be an effective venue for bringing all of these marriage partners together—I have seen this process in person, and close-up, both at my own lab, Argonne, and at Livermore, where I am a regular visitor.

The SiCortex and the Blue Gene systems at Argonne are said to be “energy efficient.” Can you explain what this means? In what other ways besides high-performance computing is Argonne promoting energy efficiency?

A key question regarding the power consumption of large computer installations is the power expended per productive (and unproductive) computational step. In other words, I distinguish between the power consumed by a computer when it is simply sitting there idly, without doing any calculations, and the power consumed when an (application) program is being run. The scale of computers has now

It turns out that computing at the exascale will play a transforming role in the underlying science and engineering disciplines at the heart of energy production—from things nuclear to the fundamental chemistry of catalysts.

reached the point where the power consumed by a machine not doing any computational work is comparable to what is consumed when a calculation is being performed; and furthermore, that power level is daunting, that is, the electricity costs are becoming a significant—and even the dominant—part of the operating costs of a major computing facility. Hence, there is substantial reason for considering how one might design machines that, first, consume as little power as possible when they are idle, and second, consume as little power (per computational step) as possible when they are computing. The SiCortex and the Blue Gene systems represent two avenues towards reaching these two goals. As for what else we are doing in this context: the most obvious additional step is to make sure that the environment in which these computers sit is itself energy-efficient—and that is indeed the aim of the architectural design of the TCS building!

National security and the stability of financial communities have been uppermost in our minds recently. What role do you see that high-performance computing can play in ensuing national security, in understanding the dynamics of organizations, or in predicting behavior of complex systems—whether natural (biological or physical) or social?

I already alluded to the fact that as one approaches exascale computing, it will be possible to transform the role of simulations within the social science disciplines. Just to give one example, the power of exascale computing will allow us to begin the evolution of the present general equilibrium models characteristic of much of current econometric modeling in the direction of (validated) dynamics. This evolution is essential if one is to address inherently time-dependent questions—for example, questions of equilibrium and stability (whether of financial systems or social organizations). Right now, we are not yet in the position to carry out simulations of sufficient fidelity that most social scientists will take seriously; but within the next decade or so, I predict that this will change dramatically.

One of the challenges facing scientists today is convincing the public of the importance of research. As Dr. Orbach emphatically stated: We must “actively and publicly make the case for science.” What suggestions do you have for accomplishing this? Do you think the INCITE program, in focusing on a limited number of high-impact science advances, will help?

This question is deserving of a book-length answer, not what is effectively a “sound bite”—but at least let me try to answer with what I regard as a key missing element in the discussion to date. We as scientists have tried very hard to “popularize” science, that is, to educate the public in the excitement, the intellectual richness, and the plain fun that is science. The strong growth nationally in science museums and planetaria is testament to these efforts; and federal agencies have played a key role, both as funders of these efforts, and as the agents carrying out these efforts themselves. What we have not done so well is to explain the deep connection between basic research that is carried out today and the industrial capabilities that flow from the results of that research, not today, but possibly two or more decades from now. There are numerous examples of this lag between basic and applied research—think of the transistor, the laser, computers, the Internet—but in an age where instant gratification and quarterly results are the coin of the realm, it is difficult to convincingly argue that basic research is not just yet another mouth to be fed at the public funding trough, but is really an activity essential for the economic and social well-being (and competitiveness) of our nation. In this context, the INCITE program has the potential for bringing home this lesson—the deep connection between basic research and what is ultimately wealth creation within our country—in a very concrete fashion. The immediate example that comes to mind is of course the work done by Pratt & Whitney within the INCITE program, and that company’s latest development of highly efficient jet turbine engines for the aircraft industry.

Thank you for taking the time to answer our questions. ●