I. INTRODUCTION

A. Why Invest Now

This is a particularly appropriate time for the NSF to launch a program in computational physics. In virtually every sub-field of physics, one finds important problems of such complexity that traditional analytic approaches are difficult, if not impossible, to apply. Numerical computations offer the best possibility of making rapid progress on such problems. At the workshop, we heard fascinating examples from astrophysics, atomic and molecular physics, biophysics, chemistry, condensed matter physics, general relativity, high energy and nuclear physics, and plasma physics. (The organization of the workshop is described in Appendix 1, and a full list of talks is given in Appendix 2). The talks addressed some of the most fundamental questions in physics and problems of great societal importance. Entities being studied ranged in scale from the fundamental building blocks of matter to the universe as a whole. Although the physical phenomena described in the talks were highly diverse, several unifying features were evident. These included the challenges of dealing with systems involving large numbers of strongly coupled degrees of freedom and multiple scales, of developing software for ever-changing computer architectures, and of managing widely distributed data sets of enormous size.

The NSF is responding to the growing opportunities in computational science by dramatically increasing the high end computing resources available to academic scientists through its PACI Program. The Terascale Computing System that recently came on line at the Pittsburgh Supercomputer Center has a peak speed of 6 Tflops, nearly tripling the total computing power of the program. The TeraGrid project, which has recently been funded, will provide another major increase in computing power, and will open exciting new opportunities in grid-based computing and distributed data management. At the same time, the desktop machines and commodity clusters that the NSF funds for individuals and groups are growing rapidly in capability. In order to fully capitalize on these investments in hardware, the Physics Division urgently needs a mechanism for training the young scientists who will be the leaders in using it, and for supporting the development of the applications software that will run on it.

Traditionally, advances in algorithms and computational techniques have played an equal role with increases in hardware performance in advancing computational science. This continues to be the case. For example, the development of density functional theory has opened the way to scalable methods in a number of important areas of chemistry and condensed matter physics, and advances in sparse matrix methods have had a major impact in a wide variety of fields. Without such developments, even the large increases in computing power cited above would be insufficient to support progress in many areas.

The Committee believes that the confluence of the three factors cited above—the emergence of a wide variety of important problems that require large-scale computation, the rapid growth in computing power available to academic scientists, and the development of new algorithms and computational methods—provides a major opportunity for scientific advances, which can best be capitalized on through the creation of a program in computational physics.

B. The Role of Computation in Research in Physics

At one level, the change that has occurred in physics and the other natural sciences as a result of modern computation is simple to understand. Computational simulation has taken its place as the method for doing science that is a bridge between experiment and theory. Often we understand thoroughly and precisely the basic laws governing a system (such as Newton's law or Schrödinger's equation), and these have been well tested in simple situations. Over the past two decades, the dramatically increased power of both computing hardware and numerical algorithms have made possible the treatment of complex systems, and the application of these simple laws has become a task requiring new kinds of expertise as well as new computational hardware resources.

A close examination of computational physics in the areas covered by the workshop reveals that the role of simulation in bridging theory and experiment raises new and deep questions. The problems of how to connect multiple length and time scales, and of predictability itself, arise in many areas of physics, as do other fundamental issues associated with complex nonlinear systems. These intellectual challenges emerged as recurring themes of the workshop that were pursued in many of the presentations. They form a basis for commonality of perspectives and approaches among efforts in computational science in almost all of the subdisciplines of physics, as well as in the other mathematical and physical sciences.

Climate modeling provides an example of these challenges. The earth's climate system consists of many interacting components: the ocean, atmosphere, cryosphere, biosphere, etc. To build a successful computational climate model, we must break the system into manageable elements that can be quantified and whose interactions can be treated numerically. A reasonable division might be volumes 10 km of the earth and ocean, but even at this fine resolution, there are still subgridscale phenomena such as thunderstorms that must be parametrized or understood phenomenologically.

The opportunities for new discovery and understanding in physics that are presented by the increased role of computation are easy to recognize, and only a partial catalog of them is presented in this report. No workshop of practical length could hope to be exhaustive in that respect, but even a partial list of those opportunities is compelling by any standard, as we hope the remainder of this report will demonstrate.

But it was not the opportunities alone that led the steering committee to its central conclusion that the NSF should establish a new program in computational physics. We now clearly recognize shared intellectual challenges and common barriers to progress in computational investigations that make it clear that an NSF investment in computational physics, broadly construed, can affect the course of the intellectual development of physics in this century in ways that were simply not possible as recently as twenty-five years ago. The impact would be felt over the entire spectrum of the physical sciences from subatomic physics to cosmology, and extends to the biological sciences as well.

Hence, the panel discussions during the workshop and the deliberations of the steering committee focused on issues, both intellectual and practical, that must be addressed by such a program at NSF. The history of funding of experimental investigations has resulted in a clear recognition of what is necessary for an individual principal investigator or team of investigators to be able to mount an experiment requiring years or even decades to execute. However, there has not yet emerged a complementary recognition of the requirements of individual principal investigators or teams who undertake computational simulations that require five or more years of effort to execute. The problems facing physicists analyzing petabytes of experimental data, which could not have even been undertaken two decades ago, form another class of challenges in computational physics that span subdisciplines.

Simply put, we must move from a mode where we view computational science as an applied branch of theory to a mode where its true resource needs as a distinct research mode are recognized. Concretely, this means providing support for building the infrastructure (software) of computational science at levels commensurate with their true costs, just as we support construction and operation of experimental facilities.

Hence, questions of support for the development, maintenance, and sharing of major software packages and resources formed part of our discussions. Questions of fundamental intellectual challenges and outstanding problems dominated much of the remainder. In Section II, the results of those discussions are distilled into short statements of our recommendations. The brevity of those recommendations indicates the consensus that was easily reached and that was supported by the workshop participants.

Another theme that emerged in the workshop presentations was the array of outstanding problems in applied mathematics that are immediately relevant to the challenges of computational physics and the opportunity that the NSF has to invest in them. A survey of the presentations at the workshop shows a remarkable spectrum of numerical and theoretical methods that are shared among many disciplines. Examples include sparse linear algebra algorithms, scalable linear solvers, adaptive mesh methods, finite elements, spectral methods, fast multipole methods, fast transforms, and variants of density functional theory. There are frequent similarities between computational approaches and challenges that connect research on problems for which the physics is entirely different. The deeper issues raised by complex systems and multiple scales of space and time also offer major opportunities for progress on fundamental mathematical fronts. For this reason, a subsection of Section III is devoted to these questions.

Finally, the committee recognizes an aspect of computational physics that must urgently be addressed by a new program at the NSF: the pressing need for investment and reevaluation of training and education in computational methods of research. The situation has improved only marginally from that described by one of the Steering Committee members, Professor Steve Koonin, in the introduction to a textbook on computational physics sixteen years ago:

> *"Computation is an integral part of modern science and the ability to exploit effectively the power offered by computers is therefore essential to a working physicist. The proper application of a computer to modeling physical systems is far more than blind "number crunching," and the successful computational physicist draws on a balanced mix of analytically soluble examples, physical intuition, and numerical work to solve problems which are otherwise intractable.*
>
> *"Unfortunately, the ability 'to compute' is seldom cultivated by the standard university-level physics curriculum, as it requires an integration of three disciplines (physics, numerical analysis, and computer programming) covered in disjoint courses. Few physics students finish their undergraduate education knowing how to compute; those that do usually learn a*

limited set of techniques in the course of independent work, such as a research project, or a senior thesis."

In Section IV of this report, we will return to this question and its importance to the pursuit of research in physics and the mathematical and physical sciences in general.