Community Climate System Model Strategic Business Plan

(2004-2008)

Prepared by the CCSM Scientific Steering Committee:

Jeffrey Kiehl, NCAR
James Hack, NCAR
Peter Gent, NCAR
William Large, NCAR
Maurice Blackmon, NCAR
Christopher Bretherton, University of Washington
Ping Chang, Texas A&M University
Cecilia Bitz, University of Washington
Scott Doney, Woods Hole Oceanographic Institution
Danny McKenna, NCAR

with contributions from:

Jay Fein, NSF
and
David Bader, DOE
and
CCSM Working Group Co-Chairs

June 2003

www.ccsm.ucar.edu

Table of Contents

Executive Summary

- Over the past ten years, the Community Climate System Model (CCSM) activity has grown to include over 100 active participants. Current attendance at the CCSM annual workshop is close to 300 individuals, while 400 people have downloaded either the CCSM code or data from the control simulation, or both.

- Given the rapid growth in the CCSM program, the current infrastructure, both scientific and support, will require substantial enhancement if the program is to remain robust and stable to maintain the existing program.

- The CCSM Scientific Steering Committee (SSC) plans to transform the current physical CCSM into a fully interactive climate system model representing interactions of the physical system with chemical, biogeochemical, and land cryospheric processes.

- To accomplish the goal of creating a fully interactive systems model, substantial enhancement of scientific and support staff will be needed. Approximately five Ph.D. scientists, eight project scientists, and seven software engineers will be required over the next five years to accomplish the goals of the new CCSM Science Plan.

- The SSC strongly supports the establishment of an educational outreach program to facilitate the education and training of the next generation of climate system modelers. This program should support the exchange of scientists and graduate students between universities, national laboratories, and NCAR.

- Computational demands for CCSM exceed NCAR's current resources. The program will need to find computational resources both within and outside NCAR to ensure that the scientific goals of the CCSM are met over the next five years.

- Managing the multidimensional development of the CCSM code has become challenging, and will require the establishment of a board to review comprehensively all changes in the code. This board will need to involve representatives from all sectors contributing to the code changes.

I. Introduction

The following strategic business plan is addressed primarily to the various funding agency managers who support the Community Climate System Model (CCSM) activity. The purpose of the plan is to give agency managers a coherent description of the current resources used to support the CCSM activity and to describe the human and computational resource needs for further growth of the CCSM for the next five years. This plan identifies specific needs in terms of human resources and funds required. The plan also describes ways to expand the CCSM activity to develop a more comprehensive climate system model. It is hoped that the UCAR and the Climate and Global Dynamics (CGD) division management will also find this plan helpful in meeting the future climate goals of NCAR and CGD.

The present plan will hopefully complement the goals of the CGD Strategic Plan currently under development by members of CGD. In particular, the division plan will identify certain scientific human resource needs that should complement the needs identified herein.

Management of the CCSM program has grown in complexity over the past five years. Originally, the Climate System Model (CSM) activity involved mainly scientists and software engineers within CGD at NCAR. The director of the CGD division and two members of the senior scientific staff directed the CSM project. As the CSM project evolved into the CCSM program, the level of complexity in management increased dramatically.

The CCSM program has grown so that it now involves many more people than the professional staff in the CGD division. A much greater level of coordination and communication among the various CCSM working groups is now required as many of these working groups continue to develop key components of the CCSM. Funding for the CCSM activity from the National Science Foundation (NSF) has increased in recent years to facilitate the growing demands of the program. Outreach to and participation from the university community also has become a priority within the program. With all the coordination demands associated with these changes, the program was reorganized to formally recognize the increased level of management. The CCSM program office was established with a single position to oversee all of the CCSM activities. The chairman of the CCSM Scientific Steering Committee (SSC) is the head of the program office. The office also includes the CCSM coordinator, the CCSM administrator, and the head of the CCSM software engineering group. The head of the CCSM program office supervises all of these positions.

The coordinator is responsible for managing the computational resource allocation for CCSM, ensuring close cooperation among the various working groups, and preparing the five-year science plans for the CCSM program. The CCSM administrator looks after the financial budgets of the program and takes care of all logistical needs of CCSM meetings, including the annual workshop. The head of the software engineering group supervises the software engineers whose task it is to maintain the large code that comprises the CCSM.

The mechanics of identifying scientific priorities, translating these priorities into software engineering priorities, and using available computational resources optimally remains a challenging task. One reason for this is continued growth in the scientific breadth and depth of the CCSM program. The community's enthusiasm for continuing to grow the scientific capabilities of the CCSM program aggressively is a strong endorsement of earlier success, but there is a greater need to manage the growth to ensure future scientific success. The CCSM Strategic Business Plan is an attempt to define the current resources of the program, and more importantly to develop a sound plan for managing these resources to allow growth in the program that meets the needs of the community as identified in the CCSM Science Plan. The plan includes a definition of the current resources for the CCSM program, including human, financial, and computational resources. This plan will be used to determine optimal strategies to build CCSM for the future. This future will hopefully include opportunities to expand the scientific capabilities of the CCSM program, and the plan can be used to determine pathways along the way.

The plan is organized as follows: a brief history of the CCSM program is given; the current resources for the program are defined; the vision for the future and the resource needs for the future are provided; sources of support are investigated; and finally a summary of the plan's findings are presented.

II. Brief History

In 1993, a small group of scientists within the CGD division began meeting to discuss the possibility of building a new, comprehensive, coupled climate model. This group evolved into the Climate System Model Advisory Committee. The basic premises were that the model would be composed of existing component models in use within the division, that these models would be coupled together through a separate module or flux coupler, and that no flux corrections or flux adjustments would be used to alleviate biases within the coupled system. The fully coupled model was developed in collaboration with scientists in the Climate Modeling and the Oceanography Sections, with interest from other scientists in the division to further develop the model to include biogeochemical processes.

After discussions among the director of CGD, the director of NCAR, and the president of UCAR, a proposal was written that outlined the need for such a newly coupled model. The directors offered strong support for the development of this model. They recognized that one or two scientists should lead the day-to-day development of the CSM. Two co-chairs for model development were appointed to lead the project, one from the Climate Modeling Section and the other from the Oceanography Section. Development of the coupled model began in 1994. The development process proceeded with weekly meetings among all interested individuals. In 1996, the first coupled simulation was carried out, and this initial simulation showed no indication of surface climate drift, as all previous coupled simulations had exhibited. This was a significant accomplishment in the science of coupled models.

The first workshop on the CSM was held in Breckenridge, Colorado, in 1996. Results from the simulation of the 300-year control were presented at the meeting. It was recognized that further development of the CSM would require a more organized management structure; thus the CSM SSC was formed shortly after the workshop. At this time various working groups, related to the component models and major scientific thrusts, were formed to provide forums for model development and application activities. Since that time the annual workshop has become one of the most comprehensive climate modeling meetings in the United States. Attendance at the workshop has grown considerably over the years (Figure 1), from 100 to close to 300 attendees.

After the release of the first version of the CSM (CSM1) to the community in 1996, it was recognized that development of the model would require expertise in a wide range of disciplines. Also, members of the greater climate community were interested in contributing to the development of the CSM. An indication of the diverse community for this activity is shown in Figure 2.

Figure 1. Number of participants at the annual CCSM workshop, with NCAR participants (blue) and non-NCAR participants (green).

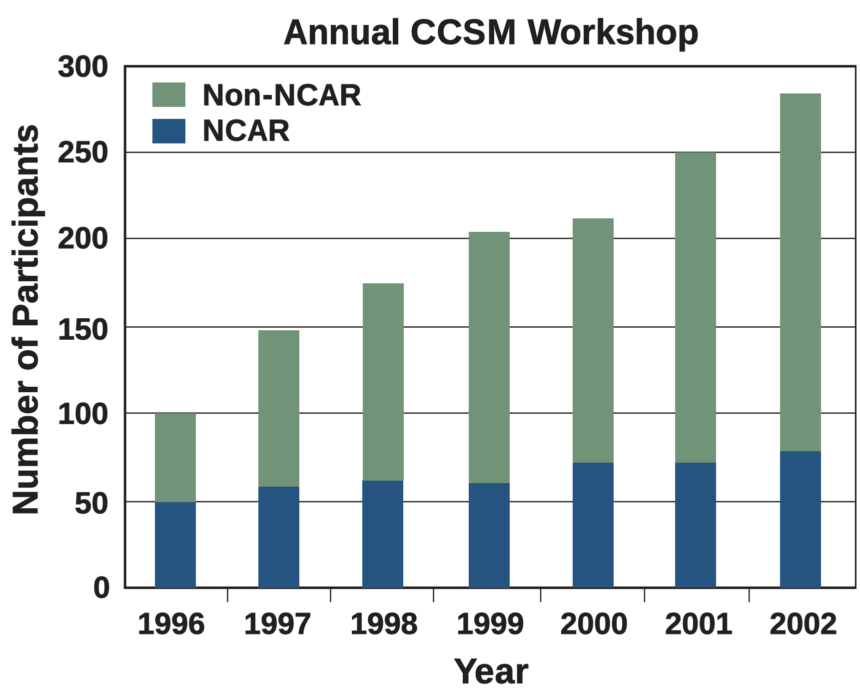

Working groups were formed in areas of the component models (atmosphere, ocean, land, and polar processes) and in areas of diagnosis (climate variability) and applications (paleoclimate, climate change). Development of a biogeochemistry component required the formation of a working group in that area in 1998, and finally the complexity of the software composing the model reached a point where a working group on software engineering was formed. In 2000, the name of the model activity was changed from the CSM to the CCSM to reflect the involvement of the community as a whole. At present these nine working groups compose the working level structure of the CCSM (Figure 3). These working groups report to the CCSM SSC, whose membership is composed of NCAR, university, and national laboratory scientists. Members of the working groups have special privileges with regard to resources. They receive access to prereleased component and coupled models, the Climate Simulation Laboratory (CSL) computational resource pool, and support through the CCSM liaisons.

Figure 2. Attendance at the 2001 CCSM workshop, in percent.

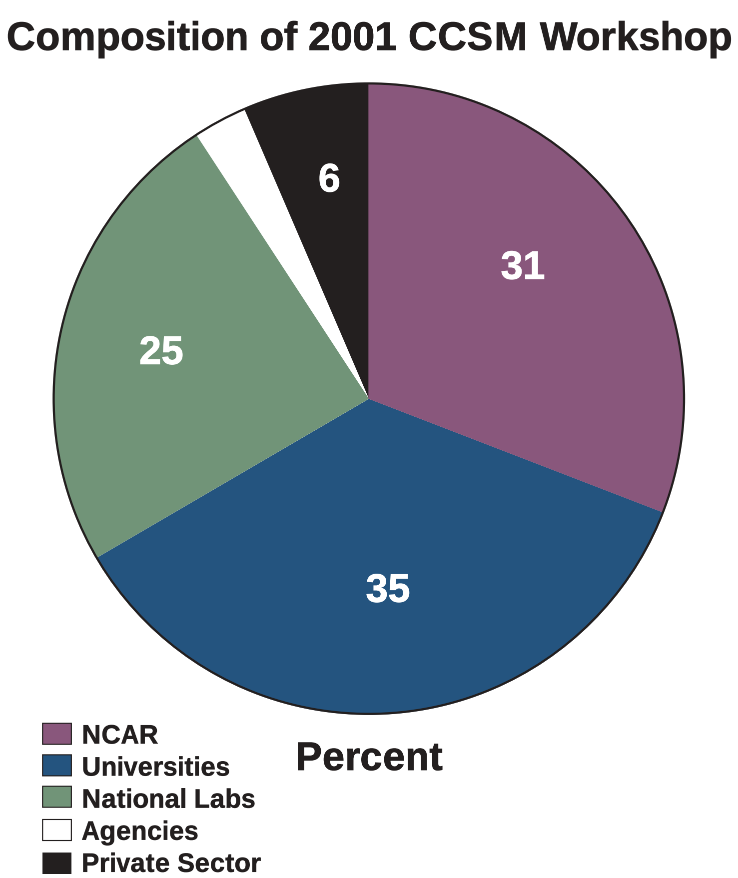

Until 2001, the director of the CGD division was the chairman of the SSC and managed the overall CCSM program. The day-to-day development work was directed by two scientific co-chairs. In June 2001, it was decided that all of these responsibilities would be transferred to the chairman of the SSC. The chairman of the SSC now has responsibility for directing the day-to-day operations of the CCSM, coordinating activities of the SSC, and reporting the status of the CCSM program to the president of UCAR, to the heads of the climate modeling programs at NSF and the Department of Energy (DOE), and to the CCSM Advisory Board (CAB).

The CAB was formed to provide independent assessment and advice about CCSM to management. It reports to the president of UCAR, the director of NCAR, and NSF and DOE management. Advice can be on the direction of model development, the use of the model for scientific problems, and overall management issues. The CAB meets twice a year to deliberate on the status of the whole program. The chairman of the SSC presents a report on the CCSM program to the CAB at these meetings. Once a year the CAB writes a report to UCAR, NCAR, NSF, and DOE management with its observations and advice concerning the program. The CAB also presents these findings to the CCSM SSC, and the chairman of the SSC responds annually in writing to the CAB report. In 2000, to facilitate the CAB's understanding of the SSC process, two members of the CAB were invited to the summer SSC meeting associated with the annual workshop. The CAB has also extended an open invitation to SSC members to attend their biannual meetings. These decisions have led to increased communication between the CAB and the SSC.

In 2000, the CAB raised the issue of how scientists' contributions to the CCSM could be recognized by the community. The SSC decided to establish an annual distinguished achievement award for outstanding contributions to the CCSM effort. The first award was presented to Dr. James Hack of NCAR at the annual CCSM workshop in 2001 and to Drs. Cecilia Bitz of University of Washington and Elizabeth Hunke of Los Alamos National Laboratory in 2002. The recipient of the award receives a plaque recognizing his or her work and an invitation to give a plenary presentation at the workshop. The SSC

Figure 3. Organizational structure of the CCSM working groups.

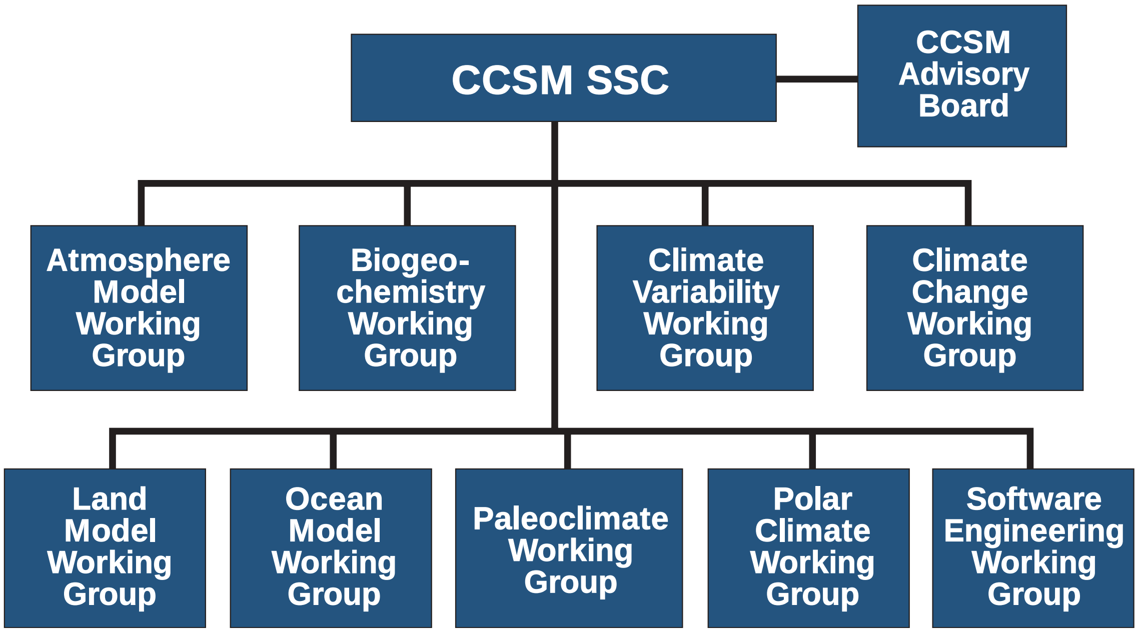

hopes to develop other ways to reward individuals or groups who make significant contributions to the CCSM.

In 2001, it was recognized that the coordination of the various working group activities was beyond the capabilities of a single person, i.e., the chairman of the SSC. A decision was made to create a CCSM coordinator position within the CCSM program office. The tasks of the coordinator are to assure effective communication across the various working groups, manage the CSL computational resource allocation, and provide outreach to the greater climate community.

In 2001, the SSC polled the CCSM community to see if users of the model would like to have a tutorial offered by the support staff to help individuals with running the CCSM and its components. There was significant interest in this idea, and at the end of the 2002 workshop a tutorial on model run scripts was offered to over 60 participants.

The decisions that brought about these changes, and the means for developing the CCSM, occur within biannual meetings of the various working groups, the annual workshop, and weekly meetings of NCAR staff who are integrally involved in the development and application of the CCSM. The NCAR local weekly meetings were found necessary for quick response to questions that arose in the development of this complex model. Attempts are being made to link the greater CCSM community into these NCAR meetings through the use of telecommunications technology and weekly summary reports that are posted on the CCSM Web site for the community at large.

III. Current Resources

a. Human Resources

The development and application of the CCSM involve a diversity of individuals. A substantial number of these individuals are supported through scientific collaborations among NCAR, university, and other national laboratory participants, which are not directly funded. This discussion of available human resources will exclude these contributions to the program since they are not strictly under CCSM management authority. Instead, the discussion will focus on direct-funded human resources assigned to the day-to-day coordination and infrastructure (both scientific and software engineering) needs of the CCSM program.

At present, dedicated CCSM human resources fall into two categories: working group scientific liaisons and software engineers. There are six working group scientific liaisons, one per CCSM model development working group, and one for most of the application-oriented working groups. (The notable exception is the Climate Change Working Group, CCWG.) These liaisons are generally assigned to the NCAR working group co-chairs who set the day-to-day priorities guided by longer-term priorities as determined by the working groups. The main responsibilities for the scientific liaisons are the coordination and dissemination of development activities within the working groups, including the provision of comprehensive diagnostic Web pages, the provision of limited technical assistance to working group members engaged in high-priority development activities, and the management of working group computer allocations.

The software engineering resources are assigned to the CCSM Software Engineering Group, housed in the CCSM program office. Support for this group comes from a variety of sources, including NSF, DOE, and NASA. At the moment, there is one software engineer assigned per component model, one software engineer assigned to the maintenance and development of the model coupler, and two additional full-time employees whose principal responsibilities are the broader coordination and testing of all CCSM software products, including their integrated performance. An important responsibility for the software engineering staff is coordination and testing of software engineering activities both internal and external to NCAR. In some cases, where the number of external participants is small or the component model is relatively simple, this responsibility is minimal. In other cases, such as the atmospheric component model, this responsibility can be more than a full-time job.

A way of measuring the complexity of this coordination responsibility is by the number of developers with write access to the component model code repository. In the case of the atmospheric component, there are more than half a dozen individuals from several institutions regularly updating the development branch of the code repository, with perhaps another dozen or more working on their own branches that will eventually need to be merged with the development branch. A different example is the sea ice component model where only one developer has write access to the code, with a very limited number of working group members working independently on the formulation of the sea ice model.

Setting development priorities in response to the need to run the CCSM or its component models on multiple architectures, or in response to scientific needs, can be an extremely complex process. The Atmosphere Model Working Group (AMWG) has begun experimenting with what is currently called a Change Review Board (CRB) to help establish and guide day-to-day software development priorities. The AMWG CRB consists of the NCAR AMWG co-chair, the head of the NCAR Climate Modeling Section, the CCSM software engineering representative assigned to the atmospheric component, an independent senior NCAR software engineer, and a DOE Scientific Discovery through Advanced Computing (SciDAC) representative (in this case from Oak Ridge National Laboratory). The AMWG CRB meets weekly to review progress and to revise day-to-day priorities as required by immediate software engineering activities (determined to some extent by activities in other model development working groups) and by broader CCSM scientific needs (e.g., new capabilities required by application working groups). Attention to high-priority scientific needs in

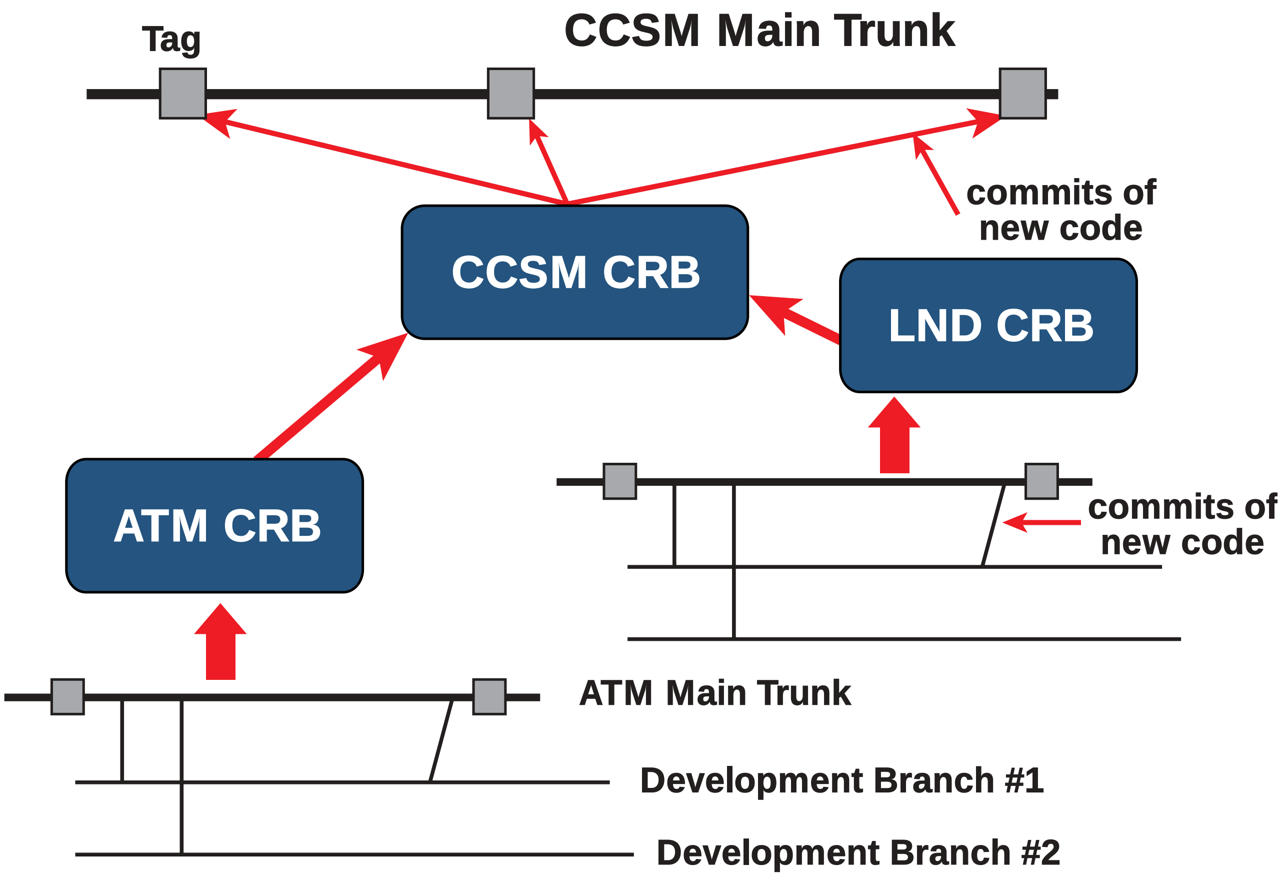

Figure 4. Tree diagram of the CCSM CRB structure. Thin lines denote commits of development versions of the various models to main trunks of model versions. Tags, blue boxes, are model versions released to the community.

the AMWG has been limited by resource issues. Nevertheless, the CRB approach appears to have a great deal of merit in the day-to-day management of software engineering activities, and has helped to identify opportunities for enhancing the level of CCSM software engineering support. The CRB structure is illustrated in Figure 4.

Overall growth in the staff directly funded for CCSM activities is shown in Figure 5.

It is apparent that most of the recent growth has been in software engineering support. To continue to support and develop this program in the future will require substantial enhancement in resources. The current CCSM model is over 500,000 lines of code, and with the addition of many of the features envisioned for the next five years, the code may grow to close to 1 million lines. To manage this type of code and to distribute it to a community of 100 or more users will be a challenge rarely seen in the climate community.

b. Computational Resources

Computational resources must serve a wide variety of needs for a project as complex as CCSM. CCSM requires ample capacity and capability to facilitate rapid turnaround for model development projects, at both the component model and the system level. Developers of the component models should periodically define their current canonical experimental configurations with associated turnaround requirements to ensure that those responsible for ongoing planning for computational resources are regularly apprised of these needs. At the component model level, this type of computing requirement is most often referred to as "capacity

Figure 5. Directed CCSM support positions for fiscal years 2000, 2001, and 2002.

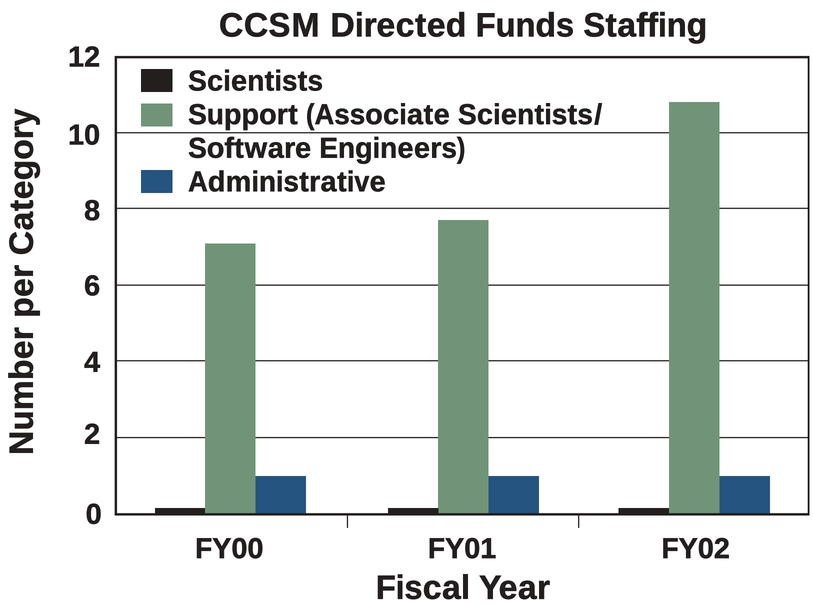

computing," a name that highlights the need for efficiently supporting large numbers of modest high-performance computing applications. CCSM also requires computing resources in support of unique modeling capabilities, such as enabling adequately fast turnaround of the more complex system model so that scientific objectives are tractable. This latter form is often referred to as "capability computing," and it depends on reliable, robust, and genuinely scalable computational platforms. Although there have been some notable enhancements to CCSM computational capabilities at NCAR and elsewhere, these enhancements are still insufficient to meet the needs of the CCSM community. The result is that current computational resources at NCAR, in terms of both capacity and capability, are inadequate to meet the current CCSM Science Plan. This lack has required CCSM researchers to exploit external computational opportunities, some of which have been provided for only short windows of time.

This "reaction-oriented" approach to supporting the scientific program places additional technical demands on a CCSM support infrastructure that is already spread too thin. It has also exposed weaknesses in the overall planning for data storage and data analysis in support of scientific activities.

c. Funding Resources

Although this plan focuses on direct funding for CCSM, it is often asked what the total cost of the program is, where "total" includes an accounting of NCAR staff who are funded from the base NSF budget but whose work focuses on CCSM activities (the above-mentioned non-direct funding). Figure 6 shows the total funding from all agencies that support CCSM activities. Funding predominantly comes from NSF, with important support from DOE, NASA, and NOAA. In 2002, the total funding for the program amounted to $9.69 million, with 59% coming from NSF, 25% from DOE, 9% from NASA, 5% from NOAA, and 2%

Figure 6. Total funding for the past three years of the CCSM program by agency.

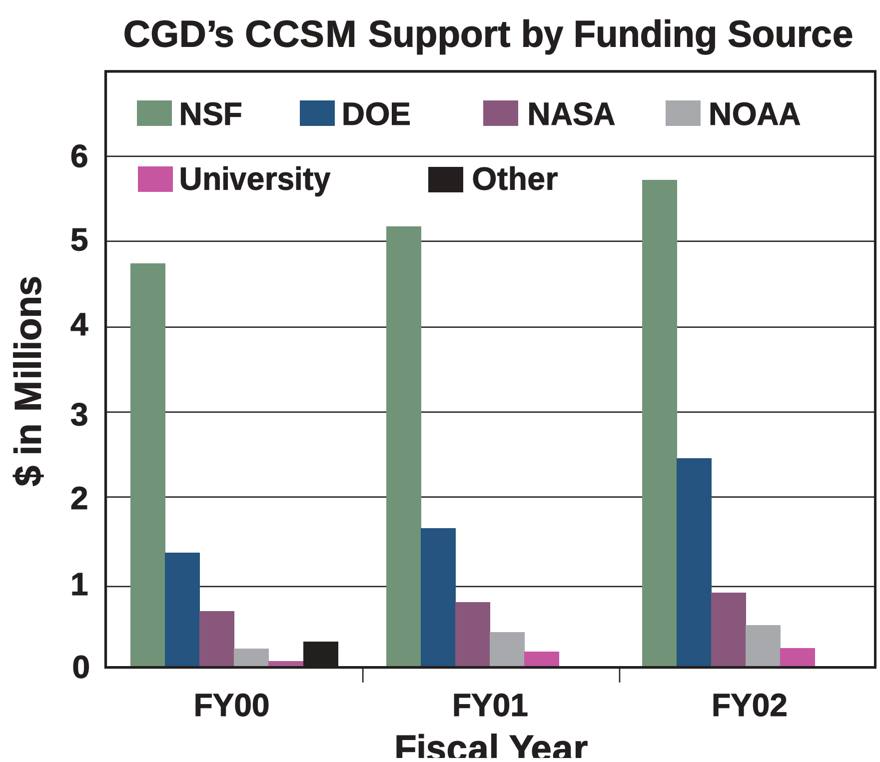

from universities. These funds support 19 scientists; 34 associate scientists, software engineers, and postdoctoral fellows; 3 visitors; and 1 administrator.

The directed NSF funds for the past three years used to support CCSM are shown in Figure 7. This funding information indicates that the majority of program monies have been used to support staff during this time. The annual costs for the program include salaries for staff, which go mainly to support software engineers; funds for the annual workshop; working group meetings; SSC and CAB meetings; and travel for members of these meetings.

For 2002, total directed funds for CCSM amounted to $1.37 million dollars, with 81% going to support staff, 12% for meetings, and 7% for other expenses.

Figure 7. Directed NSF funding for the CCSM program for the past three years. Funding for staff, meetings, and other expenses is shown separately.

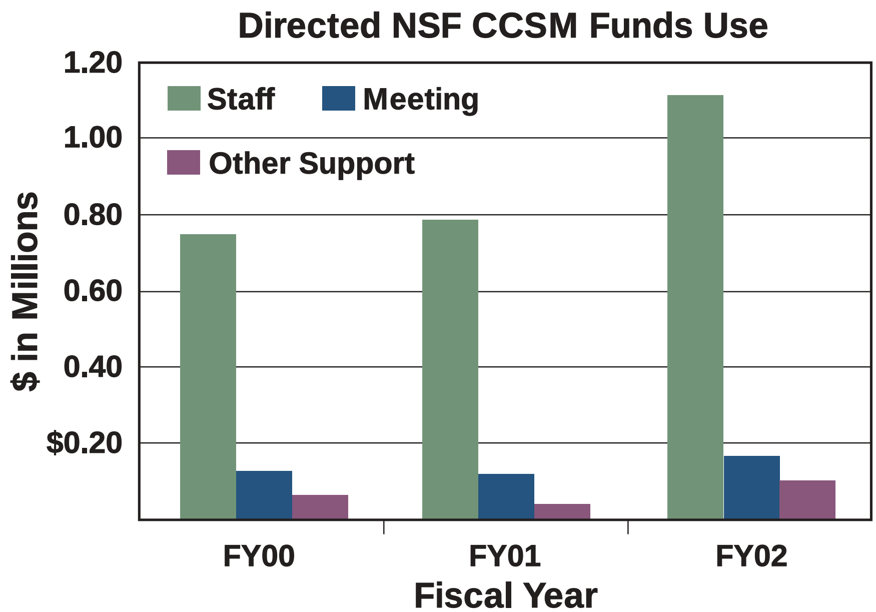

IV. Vision for the Future and Resource Requirements

The strategies described in this document will enable accomplishment of the new scientific goals described in the CCSM Science Plan. A brief summary of the scientific goals is presented here. The goals are defined as either near term (< 1 year) or longer term (> 1 year).

The near-term goals of the program are to address significant biases in the present climate simulation of CCSM2. The mean state biases include a double Intertropical Convergence Zone (ITCZ) structure in the tropical oceans, excessively warm surface temperatures at high latitudes, and associated thin sea ice and a cold tropical tropopause region. The biases in simulated variability include weak tropical inter- and intra-annual variability. There are also deficiencies in the simulated cloud properties, e.g., cloud amount and radiative properties, which may affect the overall climate sensitivity of the CCSM2.

Addressing many of these near-term goals requires significant involvement from members of both the Climate Modeling and Oceanography Sections of CGD. Thus, it is imperative that both of these sections are strongly supported by CGD to ensure these key basic model components are well maintained.

Studies are proposed to investigate the causes of these biases, which will hopefully lead to solutions to the model problems. However, management of the process to accomplish these studies within the near term is a challenge. Much of the scientific expertise needed to address these issues is voluntary, while much of the software engineering activity is focused on software infrastructure issues that had been awaiting the release of the CCSM2. The head of the CCSM software engineering group and the chairman of the SSC can redirect the software engineers to address the support issue, but the voluntary contributions of the scientific staff cannot be directed in such a manner. Despite these limitations, it is hoped that significant progress can be made on addressing existing model biases before the model is used for the Intergovernmental Panel on Climate Change (IPCC) scenario simulations for the Fourth Assessment Report.

a. Science Positions

The long-term goal of the SSC is to transform the CCSM into a more comprehensive climate system model to serve as a primary component of a community earth system model. Such a model will facilitate integrated assessment modeling of climate impacts consequences across the boundaries of natural and societal systems. In this regard, the next generation CCSM will include interactions of the physical climate system with the chemical and biogeochemical systems, as well as land cryospheric processes, such as the growth and decay of glacial ice sheets. The determination of these needs reflect ongoing discussions within the various working groups and views of the SSC.

To accomplish this evolution, the CCSM needs to be implemented to include:

- fully interactive aerosol processes
- interactive atmospheric chemistry
- biogeochemical cycles (C, S, N, stable isotopic tracers)
- dynamic terrestrial vegetation
- glacial ice sheet processes
- and the Whole Atmosphere Community Climate Model (WACCM) and CCSM must be integrated.

To accomplish these science goals will require at minimum an enhancement of the scientific staff in the following areas, which mainly address scientific needs to develop new features of the CCSM:

- Ph.D. scientist in global aerosol microphysical modeling
- Ph.D. scientist in global cloud microphysical modeling
- Ph.D. scientist in global chemical modeling
- Ph.D. scientist in ocean biogeochemical modeling
- Ph.D. scientist in glacial modeling

The support positions required to meet these goals include:

- project scientist in global aerosol modeling
- project scientist in global chemical modeling

These project scientists will work with the Ph.D. scientists in implementing and testing the new model capabilities.

b. Future Support Resources

As the CCSM evolves as the mainstay for a comprehensive system model, there will be a growing demand to reach out to the climate sciences community. Development of many of the components will require close interaction with university and laboratory scientists. Thus, outreach to these communities will need enhancement and additional support. Use of new telecommunication techniques will be explored to aid in outreach and engagement of this community. "Focus workshops" in specific inter- or multidisciplinary areas of research and development will also help increase interactions among those involved in the modeling activities.

The long-term scientific goals for the CCSM include a stabilization of the development and application processes for the physical system. There is a growing need to support the basic science process within CCSM, which requires enhancement of support scientists working in the key areas of atmosphere, ocean, and climate change research. These individuals would help in the implementation and interpretation of physical processes. They would also help in the visualization and interpretation of simulations relevant to climate change research. The additional staff includes:

- support scientist for the AMWG
- support scientist for the Ocean Model Working Group (OMWG)

These positions will augment the existing liaison positions for the AMWG and OMWG and will be supervised by the working group co-chairs.

Managing the scientific development and software engineering dimensions of an earth system modeling framework will be even more demanding than managing the CCSM. The current CRB structure will need to evolve with the complexity of the whole system model. An implementation strategy is required to ensure a stable evolution to this more complex system.

Experience with the present level and distribution of dedicated CCSM resources suggests that much of the current management strategy is workable, but that the level of dedicated support is deficient. It is clear that the scientific liaison positions play a critical role in outreach to and support of the broader scientific community. Additional positions should be immediately funded to support important applications of the CCSM conducted by the CCWG and Climate Variability Working Group (CVWG). The overall level and deployment of scientific liaison support should be reviewed annually to determine whether it is commensurate with the demands placed on each working group by the CCSM scientific program, and by external collaborations with the universities and national laboratories. Thus, to meet the immediate needs of the two application working groups requires:

- liaison for the CCWG
- liaison for the CVWG

The need for a liaison for the CCWG is especially critical given the demands on this group to participate in the IPCC Fourth Assessment Report.

With continued growth in the complexity of CCSM modeling capabilities, software engineering support has become even more critical to sustaining the CCSM effort. Software engineering support will continue to be a pacing factor as needs increase to exploit a variety of evolving computer architectures, both at NCAR and at other high-performance computing facilities around the country. The most immediate and critical problem faced by the software engineering effort is the single-point-of-failure exposure arising from the assignment of one software engineer to each component model. Experience has demonstrated that illness, vacation, or attrition can bring model development, and the associated scientific activities, to a standstill. The CCSM effort is clearly spread too thin for a project of such magnitude and technical complexity and with a need for robust operational availability. There should be at least two software engineering positions associated with each of the principal component models, and each of these engineers should develop additional overlap in technical expertise on related component models (e.g., the coupler) and testing. The development of turnkey testing procedures for the community should be among the higher priority short-term activities for an enhanced software engineering investment. The need for additional support, or redeployment of existing resources, should be reviewed annually to ensure the software engineering effort remains responsive to CCSM scientific objectives. CCSM also needs to formally review community requirements for the support of other tools in the modeling hierarchy (e.g., single-column models, an upper ocean model). There should be an assessment of how best to accommodate these needs with existing software engineering resources, perhaps by implementing such capabilities as special cases of the more complete component model. At the minimum, an additional software engineering position is needed in each of the following areas:

- atmosphere component
- ocean component
- land component
- sea ice component
- coupler

Finally, over the past few years, CCSM scientific leadership has recognized that addressing simulation biases needs to move beyond studying the biases solely within the context of individual component models. An informal interdisciplinary analysis of the coupled simulation properties has evolved through the cooperation of many NCAR scientists and, to a more limited extent, external scientific collaborators. The issues that must be addressed are highly complex, in terms of both the underlying science and the technical difficulties in exploring solutions. Although this informal effort has proven to be highly successful in identifying the more likely underlying causes of simulation biases, it has been severely constrained by an absence of scientific and software engineering support. This constraint represents a significant missed opportunity, particularly with regard to entraining scientists outside NCAR in the effort. This effort needs to be formalized and organized to better leverage other dedicated CCSM resources in pursuit of a better understanding of simulation behavior, and in the incorporation of that understanding in the form of improvements to the modeling framework. A vigorous visitor program needs to be implemented to allow non-NCAR scientists to interact more closely on the diagnosis and development of the fully coupled system.

This activity will require support through

- funds for visitors
- liaison for coupled model experiments

c. Future Computational Resources

The final area of resource growth is computational. As model spatial resolution is increased, and additional physical, chemical, ecological, and biological processes are added to the CCSM, the computational needs will be substantial. The basic requirement of ensemble simulations for most climate change research also places an additional demand on these resources.

CCSM must make a greater effort to better coordinate its computational needs with long-term investments in computational infrastructure made by participating agencies, notably NSF and DOE. CCSM has an obligation to clearly articulate the computational capabilities required to meet the goals of the scientific program, while the computational centers have an obligation to deliver the levels and quality of support required to facilitate such a demanding computational science activity. Some stability in architectural options would be of great benefit to the CCSM program. This can, and should, come in the form of long-term resource commitments to CCSM so as to allow for technical planning related to the implementation of the coupled model, as well as the migration and analysis of simulation data. It is important to acknowledge the many peripheral research activities currently exploring ways to ease the problems of moving code between computer architectures and addressing the challenges of large data set migration and analysis (e.g., Earth System Modeling Framework (ESMF), Earth System Grid). But it is equally important that such research activities not be placed in the critical path of any plan until they have moved well beyond the proof-of-concept stage. This effort has important implications for CCSM plans to manage data and its analysis. For example, the NCAR Scientific Computing Division (SCD) presently plans its data storage needs on the basis of growth in local computational capability. SCD planning needs to be more closely involved in CCSM plans to exploit external computational resources and the implications these plans may have on data storage needs at NCAR, data analysis needs, and demands on external network traffic.

Over the next five years, computational demands are expected to grow rapidly because of the need to improve the spatial resolution of the component models, and because of the additional complexity arising from extensions to the physical climate system (e.g., atmospheric chemistry, ocean and terrestrial biogeochemistry). In estimating future computational needs, we will normalize resource requirements by the current (2003) CCSM CSL allocation.

Over the next year, plans are to double the resolution for the atmospheric component from the current nominal 3° x 3° discretization to 1.5° x 1.5°. Our goal will be to conduct the majority of future development and production work (including IPCC projections) at this higher resolution. This will require at least a five-fold increase in computational resources for uncoupled atmospheric simulation work, and approximately a three-fold increase for work with the coupled model.

By year 2, we anticipate the adoption of a fully interactive atmospheric aerosol package, increasing the cost of atmospheric integrations by a factor of two. Additionally, CCSM working groups are expected to be routinely working on development versions of the CCSM that include biogeochemical capabilities and atmospheric chemistry. During this period, these additional capabilities are expected to increase overall demand by an extra 50%, requiring yet another three-fold increase in computational resources for year 2.

We expect that the CCSM will have incorporated fully interactive atmospheric chemistry capabilities, in addition to a more complete treatment of biogeochemical cycles, by the end of year 3. The biogeochemical modeling is expected to result in a doubling of the computational cost, while a fully interactive chemical package will likely multiply these costs by another factor of three. Assuming that approximately 50% of CCSM work will be conducted with the simpler configuration of the physical climate system, with the remainder conducted using the more complete physical and biogeochemical configuration, we project the need for an additional factor-of-four increase in computer resources by the end of year 3.

Computational needs over the next five years will be driven by the need to conduct multicentury ensemble simulations using the complete system model, at enhanced spatial resolution. Figure 8 shows the projected computational resource needs for CCSM activities over the next five years. To build the fully comprehensive system model with appropriate resolution will require an approximate 144-fold increase in computational resources relative to what is currently available through the CSL allocation.

d. Future Model and Data Distribution Needs

Who is the clientele of CCSM? Current participation in CCSM can be measured in a number of ways. At the most fundamental level are the individuals involved in the actual

Figure 8. Normalized increase in required computational resources relative to FY03 CCSM CSL resources. The total increase is separated into costs related to increases in resolution, the incorporation of fully interactive atmospheric chemistry, and the addition of biogeochemical modeling capabilities. A projected increase in computational capabilities at present expenditure levels (per Moore's Law) is shown for reference.

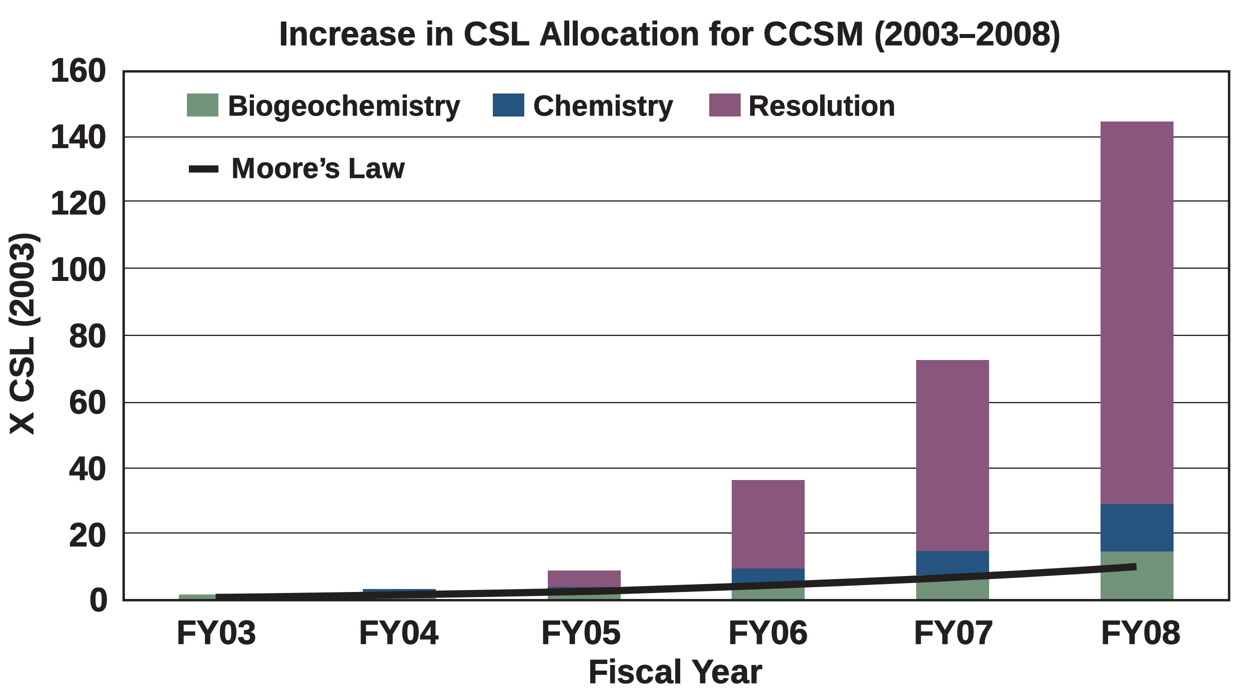

development of the model, who also use the model for basic research. The number of individuals attending the annual workshop (Figure 1) is another measure of community participation. There has been a three-fold increase in those attending the meeting since its inception. Another measure of the clientele is based on who has downloaded the model or output from the CCSM control simulation. Figure 9 shows these statistics for NCAR, university, government laboratory, international, and other users.

This user base represents a wide range of interests, and as the model evolves to an earth system model and is applied to help address critical societal issues, such as global warming, the user base will expand. Data storage, distribution, and support will require enhancement to meet the needs of the community. The current data storage requirements for CCSM are conservatively estimated to be around 120 terabytes. As CCSM is ported to more computational centers, data management will become an even more complex issue.

The estimated resource needs to address data requirements are:

- mid-level software engineer to oversee CCSM data storage needs, migration and cataloging of data, interfacing with SCD and DOE data centers
- junior-level software engineer to carry out data migration and storage and address day-to-day data requests

Figure 9. Number of individuals who have downloaded either the CCSM model or data from the CCSM control simulation from May to November 2002.

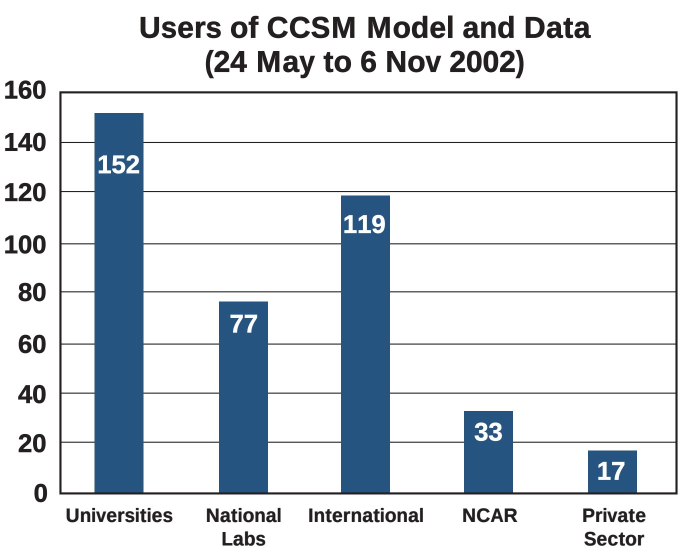

e. Educational Outreach Goals

A high-priority concern of the SSC is the education and training of the new generation of climate system modelers. A comprehensive program should be established to facilitate the training of new scientists in this area. This interaction should include:

- NCAR or other national laboratory scientists teaching summer courses in specific topics of climate system modeling
- graduate students spending extended time at NCAR to work with scientists on a modeling project
- NCAR scientists visiting universities to offer lectures or teach on modeling topics, and serve on Ph.D. committees of graduate students
- university faculty and CCSM scientists developing a CCSM-based curriculum on climate/earth system modeling
- the CCSM community creating a postdoctoral fellowship program, ideally multi-agency, that would allow next-generation climate modelers flexibility to work at a number of climate modeling centers

Finally, CCSM must continue to expand its Web-based outreach. Data acquisition and communication of model structure and development will benefit from enhanced capabilities in the area of Web technology outreach.

Accomplishment of these outreach goals will require a substantial enhancement of support. Funding to support summer courses, student visits, travel by CCSM scientists to universities, and coordination of curricula development all require funding that currently does not exist within the CCSM structure. Maintaining and developing the Web technologies will require a full-time Web manager.

V. Funding Strategies

The existing directed funding for the CCSM program enables the support of a number of liaison activities. Current funds primarily support software engineers, the annual workshop and SSC, CAB, and working group meetings, and administrative support. These existing directed funds for CCSM must continue to be available and grow over the next five years to ensure that the basic CCSM program can continue. Note that with only these existing funds, the program cannot commit to substantial growth in any new areas. The program has reached a point of *stable fragility* in terms of what it can support and develop. Thus, the funds that currently exist will continue to support model development, code infrastructure, distribution of the model code, and data archiving. The history of the current funding is described in Section III.

The focus of this section is on defining funding trajectories that will enable accomplishment of part or all of the proposed activities described in the previous section. Two funding strategies are presented: the sustainable and the optimal path to growing the program. The sustainable funding level will allow for moderate growth in the efforts to meet the needs of the community and to further develop the CCSM. The optimal funding trajectory will allow for the full evolution of the program into a fully interactive physical, chemical, and biological system model. The evolution of hires for the two trajectories is given in Table 1. Note that these trajectories favor the near-term buildup of staff to address scientific development and software infrastructure support. This strategy will enable the development of the full system model over the five-year time period, since much of the model

Table 1. Sustainable and optimal scenarios for CCSM staff hires from 2004 to 2008. The tasks of these hires are presented in Section IV.

Sustainable

Year	Ph.D.s	Proj. Sci.s	S.E.s	Admin. Assts.
2004	1	1	2	0
2005	0	2	1	0
2006	1	2	1	1
2007	0	1	1	0
2008	1	1	0	0
Total	**3**	**7**	**5**	**1**

Optimal

Year	Ph.D.s	Proj. Sci.s	S.E.s	Admin. Assts.
2004	1	2	2	0
2005	2	2	2	1
2006	1	2	2	0
2007	1	2	1	0
2008	0	0	0	1
Total	**5**	**8**	**7**	**2**

development should be completed before the end of the next five years. The early hire strategy also ensures a more stable and robust code infrastructure than currently exists within the CCSM.

Using current mean salary figures for each of the hiring levels—Ph.D. scientist, project scientist, software engineer, and administrative assistant—and the time series of hiring projections given in Table 1, two funding trajectories result. These two pathways for funding are shown in Figure 10. Again, it is important to note that these funds are solely for new hires to ensure model stability and scientific growth. These pathways do not include the assumed modest increase in existing funding. When the 2006 plan is created, new needs may exist, and a readjustment to the trajectories may be required. The present plan is focused on achieving the explicit goals outlined in Section IV. The total staff funds grow to slightly over $7 million a year by the end of the five-year period.

The CCSM Strategic Business Plan describes both sustainable and optimal growth scenarios for staff. The assumption for the sustainable growth is the addition of 16 full-time employees (10 scientists, 5 software engineers, 1 administrative assistant) over the five-year period. For the optimal scenario, another 6 full-time employees (3 scientists, 2 software engineers, 1 administrative assistant) are added. The total budget from FY04 to FY08 also includes the FY03 program, plus annual inflation to sustain the program. Each year a 4% inflation factor was used. Included in the current FY03 program (approximately $2 million) are about 12 full-time employees (2 scientists, 4 associate scientists, 5 software engineers, and 1 administrator), and budget to support the annual workshop and SSC, CAB, and working group meetings ($160,000 per year). Salaries, benefits, materials and supplies, purchased services, travel, computer equipment, costs for computer technical support, and overhead are included in the budget.

To begin a CCSM Educational Outreach Program, the estimated funding needs would be to support 5 visiting scientists, 5 postdoctoral fellows, and 10 students either at NCAR or elsewhere (universities and national laboratories) to work on CCSM issues. The estimated cost for this part of the Educational Outreach Program is $2 million. The program also would include support for 15 grants through NSF at $125,000 per grant for CCSM research by university researchers and support for 7 postdoctoral fellows to do research on CCSM issues at universities as well. The estimated cost of this part of the program is $2 million. In all, the total cost of the optimal scenario is $11 million.

At present NSF is the major funding source for the CCSM program. As the program continues to grow in areas of more comprehensive system modeling, other sources of funding may need to be sought. For example, the DOE may be approached to identify any new sources of funds for this activity. Also, the new Climate Change Research Program (CCRP) may open further opportunities for funding in the future.

We envision that starting in September 2006, a new five-year strategic business plan will be developed to maintain continuous funding for the program. The new plan can address new issues that arise for CCSM, and hopefully address any shortfalls that hinder the community effort.

Figure 10. Funding trajectories to accomplish sustainable and optimal levels of staff hires for CCSM.

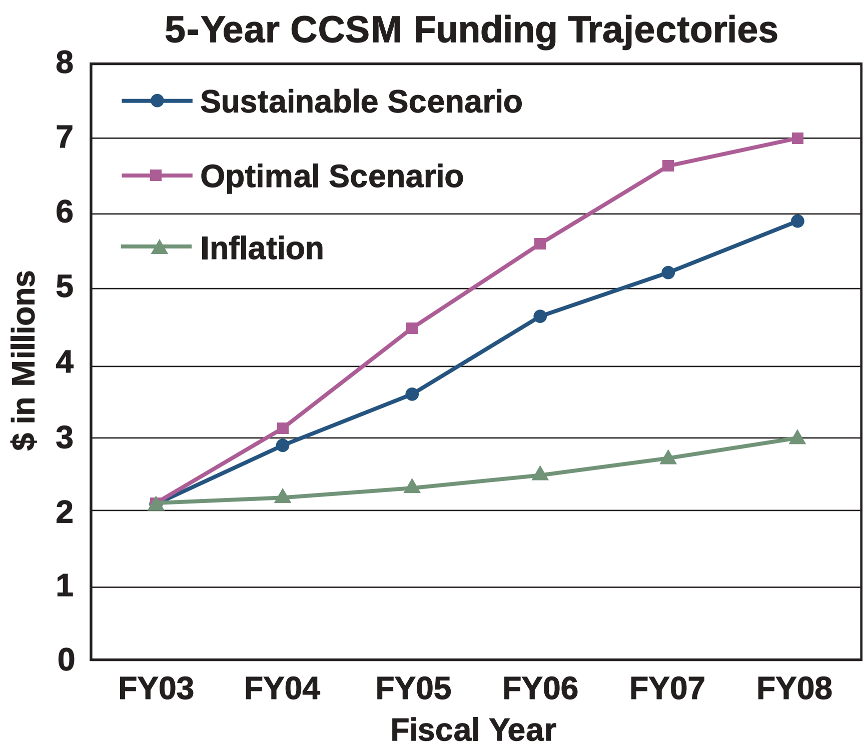

VI. Conclusion

This strategic plan provides a description of the scientific and infrastructure goals of the CCSM program for the next five years, i.e., 2004 to 2008. It is argued that the present CCSM effort is untenable for the next five years. The existing structure is stable, but fragile, and cannot continue to grow at past rates. The user numbers reflected in Figure 9 indicate both the great success and the tremendous demands of this program. To continue to meet these demands and develop a stable, robust, and scientifically sound model will require the implementation of the current CCSM Strategic Business Plan. The optimal funding strategy is strongly recommended for a tenable future of the CCSM effort.

Accountability should be an integral part of this strategic plan. The SSC commits to hold semiannual reviews of the plan's implementation and to make any midcourse corrections necessary. The SSC will report progress of the implementation to the CAB and directors at the annual Washington, D.C., meeting.

An exact implementation strategy will depend on the funding trajectory CCSM follows. Once the funding situation is known, the SSC will map out a yearly implementation plan, and the details of this implementation plan will be presented to program managers, the director of NCAR, the director of CGD, and the president of UCAR at the CAB meeting in June.

It is hoped that this CCSM Strategic Business Plan provides a sound basis for ensuring the continued growth of a remarkable program. Sociologically, the CCSM program has been a unique experiment, which has provided other disciplines with an example of how to organize a large community effort around a central scientific theme. That these communities are turning to the CCSM program as the paradigm for such an activity is yet another indication of the program's success. It is now essential to secure this success for the future, and that is the main purpose of this plan.

VII. List of Acronyms

AMWG	Atmosphere Model Working Group
CAB	CCSM Advisory Board (formerly CSM Advisory Committee)
CCWG	Climate Change Working Group
CCRP	Climate Change Research Program (NOAA)
CCSM	Community Climate System Model
CGD	Climate and Global Dynamics Division (NCAR)
CRB	Change Review Board
CSL	Climate Simulation Laboratory (NCAR SCD)
CSM1	Climate System Model version 1
CVWG	Climate Variability Working Group
DOE	U.S. Department of Energy
ESMF	Earth System Modeling Framework (NASA)
IPCC	Intergovernmental Panel on Climate Change
ITCZ	Intertropical Convergence Zone
NASA	National Aeronautics and Space Administration
NCAR	National Center for Atmospheric Research
NOAA	National Oceanic and Atmospheric Administration
NSF	National Science Foundation
OMWG	Ocean Model Working Group
SCD	Scientific Computing Division (NCAR)
SciDAC	Scientific Discovery through Advanced Computing (DOE)
SSC	Scientific Steering Committee
UCAR	University Corporation for Atmospheric Research
WACCM	Whole Atmosphere Community Climate Model