ent

NASA Technical Memorandum 4392

Functional Requirements Document for the Earth Observing System Data and Information System (EOSDIS) Scientific Computing Facilities (SCF) of the NASA/MSFC Earth Science and Applications Division, 1992

Michael E. Botts and Ron J. Phillips
The University of Alabama in Huntsville
Huntsville, Alabama

John V. Parker
George C. Marshall Space Flight Center
Marshall Space Flight Center, Alabama

Patrick D. Wright
Universities Space Research Association
Huntsville, Alabama

NASA
National Aeronautics and
Space Administration
Office of Management
Scientific and Technical
Information Program
1992

Acknowledgements

The completion of this study and document would not have been possible without the hard work and persistence of a number of people. Initial thanks should go to Ron Koczor for having the insight to initiate this task and for his assistance and patience during the process. Much praise is reserved for the members of the SCF Working Group who dedicated much time and effort to the numerous, and sometimes demanding, MSFC EOS/SCF Working Group meetings. In particular, the outside efforts of Bill Lide, Karen Butler, Matt Smith, and Mike Goodman are highly appreciated. The time, effort, and cooperation provided to the Working Group by the individual SCF scientists, including Drs. Pete Robertson, Roy Spencer, Hugh Christian, Tim Miller, and Dan Fitzjarrald, were extremely important and greatly appreciated.

Glossary of Acronyms

2D - Two Dimensional
3D - Three Dimensional
4D - Four Dimensional (3D space plus time)
ACE - Advanced Computing Environment
ANDF - Architecture Neytral Distribution Format
ARC - Advanced RISC Computing
ASCII - American Standard Code for Information Interchange
AT&T - American Telephone and Telegraph
ATTSV - AT&T System V Unix
AVS - Advanced Visualization System
BSD - Berkley System Development
CAD - Computer Aided Design
CaPE - Convection and Precipitation/Electrification Experiment
CASE - Computer-Aided Software Engineering
CCM - Community Climate Model
CD - Compact Disc
CFD - Computational Fluid Dynamics
CISC - Complex Instruction Set Computer
COHMEX - COoperative Hunstville Meteorological EXperiment
CPU - Central Processing Unit
DAAC - Distributed Active Archive Center
DEC - Digital Equipment Corporation
DOS - Disk Operating System
DPS - Display PostScript
EADS - Engineering Analysis and Data System (MSFC Institutional System)
ESA - European Space Agency
ESSC - Earth System Science Center at Pennsylvania State University
EOS - Earth Observing System
EOSDIS - Earth Observing System Data and Information System
ESACF - Earth Science and Applications Computing Facility
ESAD - Earth Science and Applications Division (MSFC)
FDDI - Fiber Distributed Data Interface
GB - Gigabytes (10^9)
GCM - General Circulation Model
GE - SGI Geometry Engine
GIS - Geographical Information Systems
GL - Graphics Library
GUI - Graphical User Interface
HIPPI - High-Performance Parallel Interface
HP - Hewlett-Packard Computer Corp.
IL - SGI ImageVision Library
IMSL - International Math and Statistics Library
I/O - Input/Output
IRIS GL - SGI Open Graphics Library
ISO - International Standards Organization
kb - Kilobits (10^3)
kB - Kilobytes (10^3)
LAMPS - Limited Area Mesoscale Prediction System
LAWS - Laser Atmospheric Wind Sounder
LIS - Lightning Imaging Sensor
LOS - Line of Sight

Glossary of Acronyms (concluded)

MB - Megabytes (10^6)
McIDAS - Man-computer Interactive Data Access System
MFLOPS - Millions of Floating-point Operations (10^6) Per Second
MIMR - Multichannel Imaging Microwave Radar
MIPS - Millions of Instructions per Second
MPP - Massively Parallel Processing
MSFC - Marshall Space Flight Center
MSU - Microwave Sounding Unit
MVUPS - MicroVAX II Units of Performance
NFS - Network File System
NQS - Network Queing System
NTI - New Technology, Inc.
OSF - Open Software Foundation Unix
OSI - Open Systems Interconnection
PC - Personal Computer
PDL - Page Description Language
PEM - Pacific Exploratory Mission
PEX - Protocol Extension to X
PHIGS - Programmers' Hierachical Interactive Graphics System
POSIX - Portable Operating System Interface for Computer Environments
PS - PostScript
PVS - IBM's Personal Visualizer System
RAM - Random Access Memory
RAMS - Regional Atmospheric Modeling Systems
RISC - Reduced Instruction Set Computer
RSS - Remote Sensing Systems in Santa Rosa, CA
SCF - Scientific Computing Facility
SCO - Santa Cruz Operation
SGI - Silicon Graphics, Inc.
SNA - System Network Architecture
SPARC - Scalar Processor Architecture
SSEC - Space Science and Engineering Center at the University of Wisconsin
TB - Terabytes (10^{12})
TCP/IP - Transmission Control Protocol/Internet Protocol
TIGA - Texas Instruments Graphics Architecture
TL - Team Leader
TM - Team Member
TRMM - Tropical Rainfall Measurement Mission
UAH - University of Alabama in Huntsville
UI - Unix International
USRA - Universities Space Research Association
VAR - Value Added Reseller
VMS - Virtual Memory System
WYSIWYG -"What You See Is What You Get"
XGL - Sun's X Graphics Library

TABLE OF CONTENTS

Functional Requirements Document
for the
Earth Observing System Data and Information System (EOSDIS)
Scientific Computing Facilities (SCF)
of the
NASA/MSFC Earth Science and Applications Division

1992

Deputy Approval

Approval:

Ronald J. Koczor (signature) Date: 4-6-92

Ronald Koczor
Deputy, Earth Science & Applications Division
NASA/Marshall Space Flight Center

Functional Requirements Document
for the
Earth Observing System Data and Information System (EOSDIS)
Scientific Computing Facilities (SCF)
of the
NASA/MSFC Earth Science and Applications Division

SCF Working Group Approval

Main Group:

Dr. Michael E. Botts (Chairman - UAH)

Karen A. Butler (NTI)

H. Michael Goodman (NASA)

Wilton C. Lide (NTI)

Dr. Jeffrey C. Luvall (NASA)

Dr. Franklin R. Robertson (NASA)

Matthew R. Smith (NTI)

Dr. J. Briscoe Stephens (NASA)

Technical Group:

Danny M. Hardin (UAH)

John V. Parker (NASA)

Ron J. Phillips (UAH)

Patrick D. Wright (USRA)

PART I. INTRODUCTION

A. PURPOSE

The purpose of this document is three-fold: (1) to define the requirements for the five EOS Scientific Computing Facilities (SCF) which will be in the Earth Science and Applications Division (ESAD) at the NASA/Marshall Space Flight Center (MSFC), (2) to recommend general options for the SCF computing environment at MSFC, and (3) to recommend specific short-term options for meeting the computing requirements of these facilities. The intent of this working group is to investigate the potential for development of a "common" SCF at MSFC.

B. BACKGROUND

The ESAD is involved in both NASA-wide and international research activities, involving the government and the scientific research community. The ESAD supports Earth system science and global change research in support of Mission to Planet Earth. The key component of this mission is the Earth Observing System (EOS). The EOS has three major components: EOS Scientific Research Program, EOS Data & Information System (EOSDIS), and the EOS Observatories. The EOS Scientific Research Program is the foundation of EOS - the international scientific effort guiding future development of the EOSDIS and the EOS Observatories. The EOSDIS is functionally the "Researcher" portion of the system. Its goal is to enhance the use of EOS data by the world-wide research community. The EOS Observatories are the actual scientific instruments which will remotely sense many environmental variables.

There are three types of EOS investigations: Those utilizing data from Facility Instruments, those utilizing data from Instrument Investigations, and those in support of Interdisciplinary Investigations. Each EOS investigation will have an associated SCF. A primary role of the Facility Instrument and Instrument Investigation SCFs is to investigate or develop algorithms to process raw instrument data to derived product (Levels 0 through 4). EOS data levels are described in Appendix C. Interdisciplinary Investigations SCF are similar to the other SCFs, except that they will not process Level 0 data, and will employ data from several instruments.

C. SCOPE

An SCF Working Group was established under directive of Ron Koczor, Deputy of the Earth Science and Applications Division, in May 1991. The group's charter was to deliver this document with the three-fold purpose stated above.

The scope of this document is to define the total SCF computing requirements and options for meeting these requirements for the ESAD within the following scheme of "Needs" and "Demands." As defined within this document, "Needs" are requirements that can be met, individually or collectively, by the SCFs, while "Demands" are those which require action outside of the realm of the SCFs:

Individual Needs - These are requirements of the individual SCFs that can only be met by non-shared resources. Candidate options capable of meeting these demands will be recommended within this document.

Common Needs - These are needs which are common to two or more SCFs, but which require non-shared resources. Compatible systems and standards will be recommended for such needs, when the use of such compatible systems and standards enhances the completion of the SCFs tasks.

Overlapping Needs - Overlapping requirements are those which are common between two or more SCFs, and which may be able to be met through the use of shared computer resources. It will be recommended that overlapping needs be met by the sharing of resources between SCFs, in order to provide a higher capability than would be possible for a single SCF. Such an option should enhance, rather than hinder completion of the SCFs' tasks.

Institutional Demands - These are demands that are not met with existing institutional facilities, and which cannot be met with the anticipated SCF funding.

DAAC Demands - These are demands which should be considered during the creation and establishment of the EOS DAAC facility.

D. DOCUMENT ORGANIZATION

Part I of this document consists of the present section which defines the purpose and scope of the document. Part II will define the individual SCF computing requirements, as determined through a combination of questionnaires, interviews, and reviews of relevant documents. Part III will present information regarding the present state and current trends of the computing industry, as well as industry standards and the requirements placed on the SCFs by the EOSDIS. Based on findings in Parts II and III, recommendations for general computing options (i.e., the computing environment philosophy) for ESAD SCFs will be presented in Part IV. Finally, Part V will present recommendations for specific short-term options for meeting the individual and common SCF computing requirements.

PART II. SCF COMPUTING REQUIREMENTS

A. INTRODUCTION

Five scientists at MSFC have EOS Principal Investigator (PI) or Team Member (TM) status. All are located within the MSFC Earth Science and Applications Division (ESAD). In addition, other scientists at MSFC/ESAD have Co-Investigator (Co-I) status, adding further demands on the SCF requirements. The Investigators establishing SCFs at MSFC/ESAD include:

Hugh J. Christian, PI for the Instrument Investigation of the Lightning Imaging Sensor (LIS); Co-I scientists: Rich Blakeslee, Steve Goodman, and Douglas Mach (UAH)

Franklin Robertson, Co-Investigator (CI) for Interdisciplinary Investigation of the Global Water Cycle (GWC); Co-I scientists: John Christy (UAH), Steve Goodman, Dan Fitzjarrald, and Tim Miller

Roy Spencer, U.S. Team Leader (TL) for the Instrument Investigation of the Multichannel Imaging Microwave Radar (MIMR)

Tim Miller, TM with the Facility Instrument Investigation of the Laser Atmospheric Wind Sounder (LAWS)

Dan Fitzjarrald, also a TM with LAWS

Each of these investigators has specific SCF computing requirements. Through a combination of questionnaires, interviews, and a review of relevant documents, the computing requirements of these SCFs were examined on an individual basis by the SCF Working Group. Documenting the requirements of both mission and pre-mission activities of the SCFs was an objective of the SCF Working Group. However, EOS mission activities are not scheduled to begin until at least 1998, whereas pre-mission activities are ongoing or scheduled to begin within the next fiscal year. In light of the rapid changes occurring in computing technology, it is unrealistic for the Working Group to recommend anything more than general directions for computing needs greater than a 5-year timeframe. Therefore, recommendations of the Working Group are primarily concerned with solving the SCF computing needs for pre-mission activities, with the hope that the general directions chosen will be compatible with future requirements and computer capabilities.

The investigations of individual SCF computing requirements primarily concentrated on five concerns: (1) Data, (2) Computation, (3) Software, (4) Connectivity, and (5) Peripherals. One intent of the Data section is to consider any data storage issues which might not be addressed within the EOS DAAC framework. These include the need for storage of temporary, non-archive data on a local workstation or within an ESAD mass storage system, as well as requirements for data storage on special media (e.g., video, CD, etc.). Also under the data requirements, the Working Group examined the relative importance of the different data types that the SCFs would analyze, including raster imagery, vector plots, computer model I/O, and two-dimensional (2D) and three-dimensional (3D) grid point data.

Computational requirements were considered for two categories: (1) the need for raw CPU-intensive computation (i.e., "number crunching"), such as that associated with numerical modeling and analysis, and (2) the need for graphics-intensive performance associated with visualization activities. Software requirements were examined in order to determine which software needs could be met with present application packages and which required future development, and to evaluate whether the development, availability, or application of software packages would restrict or drive options for hardware. Four areas of software needs were examined. These included anticipated importance of (1) McIDAS, an image display and analysis package presently in use within ESAD, (2) vector and scalar numerical models, (3) visualization applications, including data contouring, image processing, time sequencing (i.e., animation), 3D rendering, and Computational Fluid Dynamics (CFD), and (4) numerical analysis and statistics.

The Working Group examined Connectivity requirements related to interactivity with institutes outside of the MSFC domain, network protocols (e.g., TCP/IP, DecNET, SNA), and the number terminals or personal workstations to be interconnected within the SCF. Requirements for peripherals concentrated on special I/O needs for video recording, video frame capture, color printing, and image scanning.

The following section provides a summary of the computing requirements of each individual SCF, as determined from questionnaires, interviews, and review of relevant documents. In the Summary section, these requirements will be considered jointly, in order to highlight common or overlapping needs between the SCFs.

B. COMPUTING REQUIREMENTS OF THE INDIVIDUAL SCFs

HUGH CHRISTIAN
LIGHTNING IMAGING SENSOR (LIS)

Overview

LIS was proposed as an EOS polar platform instrument, designed to detect and locate lightning with storm scale resolution (i.e., 5-10 km) over a large region of the Earth's surface along the orbital track of the satellite. With a 90% detection efficiency within the area viewed by the sensor, the LIS will detect intracloud and cloud-to-ground discharges during day or night conditions, mark the location and time of occurrence of the lightning, and measure the radiant energy. The primary task of the SCF is to develop algorithms to process the lightning data from Level 0 to Level 3.

LIS is scheduled to fly first on the Tropical Rainfall Measuring Mission (TRMM) in 1997. Pre-mission activities are ongoing, and include several field missions such as the Convection and Precipitation/Electrification Experiment (CaPE), STORMFEST, Defense Meteorological Satellite Program (DMSP) Optical Linescan System (OLS), Tropical Ocean Global Atmosphere/Coupled Ocean-Atmosphere Regional Experiment (TOGA/COARE), and the Pacific Exploratory Mission (PEM). The buildup of the SCF facility will begin immediately, with anticipated funding beginning in FY92. Steve Goodman, Doug Mach (UAH), and Rich Blakeslee are also scientists for the SCF.

Data

The mission data will include time, location,brightness, and actual images.This will require a downstream of 6 kB/sec of lightning information plus background and "housekeeping." This data stream will be a continuous stream of 6 kB/sec 3 years for TRMM. The actual readout onboard at the real time events processor is 8 Megapixel/sec but is compressed on the order of 10^5 or 500 bits/sec. The input passes through a single 10 angstrom wide filter. It is expected that 10% of the 0.5 GB/day data coming down or 0.5 GB/week will be retained. This will be stored at the MSFC DAAC after initial processing on the SCF. At present, it is anticipated that the data will be stored in ASCII format. The algorithm processing from Level 0 through Level 3 will be performed at MSFC.

Pre-mission data consist of multiple data sets from airborne, satellite, and ground-based platforms. These include imagery, data from regional lightning networks, 3D volumetric radar data, and scalar and vector parameters measured along the aircraft flight path. These data are expected to be of the same volume as COHMEX. Current data also include Field Mill and LLP lightning data. This data set will include about 15 million lightning flashes/year for the U.S. alone, and may eventually include international data, as well.

Computational Needs

It is not anticipated that there will be a need for CPU-intensive computation outside of that required for graphics, although some processing of radar and OLS data is currently done on the Cray. Any modeling will be an attempt to derive information regarding the physics of lightning as determined from the data, rather than deriving data from the physics.

Visualization

Visualization needs during mission operation will consist primarily of overlaying data onto radar and satellite imagery. Eight-bit color display, with overlay capabilities, is probably all that will be required. It is not anticipated that 3D capabilities will be required for mission activities. However, for pre-mission studies, 3D will be helpful for visualizing charges and lightning locations within clouds using airborne flight-path data and volumetric radar data.

Software

UNIX and X-windows environment are preferred (no DOS or OS/2). McIDAS will handle some of the image display requirements. Radar is presently processed on a Sun SPARCstation 2 using NCAR ROSS/ZEB software. Other software to be purchased or developed as required.

Connectivity

Basic connectivity to DAAC and other institutional facilities required. Ethernet with TCP/IP adequate for most needs. The SCF anticipates a requirement for four graphics terminals connected to one powerful graphics workstation locally, and one less powerful standalone graphics workstation at a remote location.

Data Storage

During mission operations, a minimum of 0.5 GB/week of data will need to be retained at the DAAC, after initial processing at the SCF. Pre-mission activities will require local storage of data equivalent to the volume of COHMEX data. A national global archive of lightning data products will be developed under Wetnet at MSFC, and will be maintained under the SCF. The daily OLS data will consist of about 15 GB to be stored on three 5 GB Exabyte tapes/day. At least 1 day worth of OLS data will need to be readily accessible for processing.

Peripherals

The SCF will require color hardcopy for both mission and pre-mission, and will require video frame grabbing and recording during pre-mission.

Comments

It is anticipated that the SCF computing requirements can be met with a single high-powered graphics workstation locally, and a less powerful graphics workstation at a remote location. Hugh Christian feels that a survey of existing computing technology by the SCF Working Group would be helpful for defining the SCF.

PETE ROBERTSON
INTERDISCIPLINARY SCIENCE: GLOBAL WATER CYCLE STUDIES

Overview

This interdisciplinary investigation is a joint project between MSFC and Earth System Science Center (ESSC) at Pennsylvania State University. The project will use large amounts of data of every type from multiple sources, as well as model output, in order to investigate the global water cycle. The three main tasks are: (1) documentation of the global hydrologic cycle, its related links to the physical climate system, and any detectable changes; (2) study of processes; and (3) integration of conceptual and predictive models. The output will be a mix of Level 1, 2, and 3 products. The investigations are ongoing. Other participating MSFC scientists include Steve Goodman, Tim Miller, Dan Fitzjarrald, and John Christy (UAH).

Data

This project will employ large amounts of data of several types from several EOS and non-EOS sources, including multiband satellite imagery, 2D and 3D gridpoint data from measured and computer model sources, digital elevation data, aircraft data, in-situ measurements, and various vector or overlay type data sets. Because this investigation addresses process studies, global data sets, and model simulations, the capability of handling a variety of data formats is required.

Computational Needs

There is a strong need for Cray-class computation, primarily for Tasks 1 and 3. This is primarily for running the Community Climate Model (CCM) and other General Circulation Models (GCM). There is not yet a justification for a dedicated Cray; however, the investigator suggests that consideration be given to a division-wide Cray while "keeping the S&E (Science and Engineering) lines open." Presently, a 2- to 4-hour run on the existing institutional supercomputer can involve turn-around times from several days to 2 weeks. There is an ongoing effort to put the CCM on the Stardent 2500 and the Silicon Graphics (SGI) 4D/340VGX.

Visualization

The process studies and comparison of data sets (Tasks 2 & 3) will be the greatest need for visualization. This will require a powerful graphics workstation with interactive 2D and 3D capabilities, including image processing, 3D display of time sequences of multiple data sets (i.e., 4D), and statistical analysis (e.g., time series analysis, principal component analysis).

Software

The ability to effectively run the CCM, and other models and diagnostic codes written in FORTRAN, drives the requirements for Task 1. For Tasks 2 and 3, there is a need for software which allows visualization and statistical analysis of multiple global and regional data sets. McIDAS will probably be used initially for data display. However, it is not anticipated that McIDAS will meet all the needs for visualization and analysis. There is a need for software which allows interactive statistical analysis of data sets, combined with visualization, rather than just creating "pretty pictures." Dr. Robertson does not yet know the total software requirements, nor what software is presently available to meet those needs. A survey of what software is available would be helpful.

Connectivity

Connectivity to the Pennsylvania State facility will be a requirement. GCM files are estimated to be about 200 MB. The transfer of files of this size between MSFC and ESSC is expected to ramp up to monthly frequency within the next 2 years. More direct connection to an institutional supercomputer is required. Otherwise, ethernet connections with basic TCP/IP capability is adequate for most networking needs. There is a need for graphics terminals or personal workstations for Tim Miller, John Christy, Steve Goodman, Dan Fitzjarrald, and Pete Robertson.

Data Storage

Local and division mass storage will be important for temporary storage of multiple large data sets. These data sets need to be easily accessed by workstations and supercomputers on the network. Using EADS or the DAAC for temporary storage of these data sets will be inadequate.

Peripherals

Color hardcopy and video recording are required.

Comments

This SCF is probably the most demanding for both CPU-intensive computing and high-powered visualization, as well as the need for local and division-wide mass storage.

ROY SPENCER
MULTICHANNEL IMAGING MICROWAVE RADAR (MIMR)

Overview

The MIMR is a Facility Instrument being built by the European Space Agency (ESA) to retrieve numerous tropical atmospheric and ocean parameters over time scales of many years. These parameters include, but are not limited to, precipitation rate, cloud water and water vapor content, sea surface roughness, sea surface temperature, snowcover depth and water equivalent, and possibly vegetation parameters. The SCF is responsible for developing algorithms to take MIMR data from Level 0 to Levels 1A, 2, and 3.

Data

The project will require very little ancillary data from the other EOS instruments and is considered self sufficient in this respect. The data will consist of multichannel imagery which will be used to generate the various parameters listed above. Data transfer rates during mission operation are expected to be 50-60 kb/sec.

Computational Needs

For mission support, this SCF has a requirement for five processors to be located at the MSFC DAAC. It may be possible that two or more of these processors could reside on one multiprocessor system. The requirement calls for three dedicated 20 MIPS processors to handle near real-time processing of MIMR data from Level 0 through Levels 1, 2, and 3. A fourth dedicated 20 MIPS processor would be required for information management and retrieval. The fifth processor will be dedicated to retrospective reprocessing of the data. It is envisioned that these processors will consist of individual low-end single processor workstations, or one higher-end multiprocessor workstation.

For pre-mission activities, algorithm development will require two clone mid- to high-end workstations. High I/O rates may be more important than CPU power. This activity will continuously be reading large data sets for new runs; however, the computational algorithms are not that intense.

Visualization

During mission operation, the SCF has a requirement for two clone graphics workstations of around 20 MIPS speed. These will be used primarily for image evaluation, with little 3D rendering required.

Pre-mission activities will also require two low- to mid-range graphics workstations for image analysis and comparison. Some time sequencing of images is important, but not as lengthy as that required for present MSU data studies. Two clone graphics workstations at MSFC and at Remote Sensing Systems (RSS) in Santa Rosa, California could probably handle both the visualization needs and computational needs during this phase.

Software

Most algorithms are to be developed at SCF. However, this SCF has a need for a high-quality statistics package with graphics capabilities; Roy Spencer has concerns with the present International Math and Statistics Library (IMSL) package being used and would like to know the availability of alternative math and statistics packages. McIDAS presently handles some of the SCFs image display needs. Otherwise, visualization software requirements are presently undefined.

Connectivity

SCF requires ethernet and high-speed modem connection to the MSFC DAAC. For internal networking needs, ethernet with basic TCP/IP capabilities is sufficient. At the MSFC SCF, 1-2 graphics terminals will be required.

Data Storage

During mission operation, the MIMR experiment team will require the first 2 years of all Level 1A, 2, and 3 data. The data are expected to be on 16 GB, or better, optical platters, for a total volume of 5400 GB. In addition, the MSFC DAAC will require 40 GB of disk storage. The graphics workstations to be used at MSFC and RSS during pre-mission and post-launch will require 8 GB magnetic disk storage.

Peripherals

The SCF will require color hardcopy of higher quality than presently available with the existing Toyo Color printer. Some video recording capabilities will be important, but not as much as has been used with the MSU data sets.

Comments

A summary of hardware needs at the MSFC DAAC during mission operation includes: (1) five 20 MIPS processors, (2) 40 GB magnetic disk storage, (3) 6 optical disk units (16 GB per platter), (4) one 6250 bpi magnetic tape drive, (5) 1 system console and 1 graphics workstation, and (6) miscellaneous peripherals, including terminals, printers, plotters, high-speed modems.

For the purpose of algorithm development during pre-mission and for post-launch calibration and validation, a summary of the two clone hardware requirements includes: (1) One 20 MIPS processor, (2) 8 GB magnetic disk storage, (3) 2 optical disk units (with 16 GB platters), (4) 1 magnetic tape drive, (5) 1 system control console and 1 graphics workstation, and (6) miscellaneous peripherals including terminals, printers, plotters, and high-speed modems.

TIM MILLER
LASER ATMOSPHERIC WIND SOUNDER (LAWS)

Overview

LAWS is an EOS Facility Instrument that is a candidate to fly on EOS B after the year 2000 or possibly as a follow on to the TRMM in 1997. Miller's investigation will use Line-Of-Sight (LOS) wind data rather than derived vector data. The algorithms begin with make an initial estimate on 3D gridded data and then adjust the grid to "match" the LOS data. The project is more of a feasibility study rather than an operating/practical system, and is not expected to require more than a one or two man-year effort.

Data

The instrument is expected to fly at an altitude of 600 to 750 km depending upon the latitude. The area scanned is a 45° cone. The average pulse rate is about 10 pulses/sec, with about 15 vertical levels resolved for each pulse. These data will be assimilated into 3D models with 51-100 grid intervals in each horizontal direction and about 20 levels.

Pre-mission activity will use data derived from the Limited Area Mesoscale Prediction System (LAMPS) and the Regional Atmospheric Modeling System (RAMS) for algorithm development.

Computational Needs

It is estimated that the LAMPS and RAMS models will require about 200 equivalent CRAY XMP hours/year, with 4 M-words of memory per run. Present runs on the CRAY require 2-4 CPU hours, with a minimum of 2 days turnaround on the data. The LAMPS and RAMS are almost totally non-vectorized. Efforts are underway to benchmark these models on the Stardent, Silicon Graphics, and IBM RISC workstations. It is anticipated that a superworkstation in the 100 MIPS range will be adequate for running these models.

Visualization

PI does not anticipate that visualization requirements for this SCF will be very sophisticated, and will consist primarily of images and plots.

Software

RAMS and LAMPS are both presently available. McIDAS may handle some display needs. SCF will employ whatever institutional visualization tools are available.

Connectivity

Basic ethernet with TCP/IP capabilities will be adequate for all networking needs. Two to three graphics terminals will be required.

Data Storage

Data storage needs during mission and pre-mission activities is expected to be about 10 GB. Of this, 3 GB will be required on the local workstation and 7 GB required on mass storage.

Peripherals

There is a requirement for high-quality color hardcopy of images and plots. Video recording capabilities would also be helpful.

DAN FITZJARRALD
LASER ATMOSPHERIC WIND SOUNDER (LAWS)

Overview

The primary data products from this study will be scientific studies related to the improvement of climate models used by the LAWS boundary layer data. No "standard products" or operation data sets will be produced. Algorithm development will be needed to generate the working data sets. Certain of these interim results will be made available to EOSDIS as necessary.

Data

For mission activity, raw LAWS data (Level 0 or Level 1) will be used from the DAAC. Real-time data are not required. The data stream is 5 MB/12 hr and requires 300 MB to be online. The output is expected to be 5 MB/hr, "forever." The data required is essentially the same as that required by Tim Miller's SCF.

Computational Needs

Mission activity requires low-end supercomputing power for GCM modeling, which is vectorized. Presently cannot run high-resolution model on the CRAY because of a 6 MB memory limit. High-end workstation could handle pre-mission needs which employs the CCM.

Visualization

The visualization requirements are similar to what we currently have (vis-a-vis. MSU data rendering). SCF primarily requires ability to view maps and global data.

Software

GCM and CCM modeling programs are presently available. Other software requirements are to be determined.

Connectivity

No special data network required other than the standard for EOS investigators. Basic ethernet with TCP/IP capabilities is adequate. Will require two graphics terminals for PI and a programmer. One or two other researchers may be added to the panel and would require terminals.

Data Storage

The SCF requires online storage for a minimum of 1 month of data or about 300 MB. Remaining data could be archived.

Peripherals

Color hardcopy is required.

C. CONCLUSIONS

The computing requirements of the individual SCFs are summarized in the requirements matrix in Figure II.1. This matrix indicates the importance of various computing requirements for each SCF on a 4-point scale ranging from no importance to low, moderate, or high importance. These values were derived by direct inquiry of the SCF PIs.

As illustrated by the matrix, SCF investigators will be involved with all types of data. The PIs estimated that 1-4 GB of data storage would be adequate on local SCF workstations, while data

PROJECT SCIENTIST	LIS Christian	GWC Robertson	MIMR Spencer	LAWS Miller	LAWS Fitzjarrald
DELIVERABLES					
Algorithms	●	○	●	◉	◉
Science	●	●	●	◉	○
Simulation	◉	●	○	◉	●
SCHEDULE					
Pre-Mission	ongoing	ongoing	ongoing		
Earliest Mission	TRMM, 97	FY92	EOS-A	EOS-B	EOS-B
Earliest Funding	FY92	FY92	FY92	FY93	FY93
DATA					
Types					
Model I/O	○	●	○	●	●
Raster	◉	●	◉	◉	●
Vector	◉	●	○	●	●
Grid Point	◉	●	●	●	○
Other	-	◉	-	-	-
Storage					
Local	1 GB	3-4 GB	GB	1 GB	1 GB
Division	-	100's GB	230 GB	10 GB	10 GB
DAAC/EADS	100 GB	TB	TB	10 GB	10 GB
COMPUTATION					
CPU Intensive	○	●	◉	●	●
Visualization	◉	◉	◉	◉	◉
SOFTWARE					
McIDAS	○	◉	◉	○	◉
Models					
Vector	-	●	○	○	●
Scalar	○	○	○	●	◉
Visualization					
Contouring	◉	●	○	●	●
Image Processing	◉	●	●	◉	○
Time Sequenced	◉	●	◉	○	◉
3D	◉	◉	○	●	◉
CFD	-	●	○	●	●
Other	-	-	-	-	-
Analysis					
Statistics	○	●	●	-	-
Diagnostics	◉	●	◉	-	-
Other	-	◉	-	◉	◉
CONNECTIVITY					
Locations					
MSFC	○	●	◉	●	●
External	◉	◉	◉	○	○
Protocols					
TCP/IP	●	●	●	●	●
DecNET	-	-	-	-	-
SNA	-	-	-	-	-
Terminals					
Text Only	-	-	-	-	-
Graphics	4	3-5	2	3	2
PERIPHERALS					
Video					
Recording	◉	◉	◉	○	○
Frame grabbing	◉	○	○	-	-
Color Hardcopy					
Printing	●	●	●	◉	●
Image Scanning	◉	◉	○	-	-

● High importance ◉ Moderate Importance ○ Low importance - No Importance

Figure II.1. Computing requirements matrix for the EOS Scientific Computing Facilities (SCF) at NASA MSFC.

storage needs within ESAD ranged from 10 GB to 200 GB for various SCFs. Data stored locally or within ESAD facilities represent temporary data, data used to test algorithms, data in intermediate stages of processing, or other data not suitable for archiving within the EOS DAAC facility. Most of these data would need to be readily accessible to the SCF. Individual requirements for data storage within the DAAC or EADS storage facilities ranged from 10 GB to Terabytes of data. Most scientists lacked confidence in estimating total data storage requirements, as well as the hardware capabilities required to meet their needs.

All SCFs except the LIS team had moderate to high requirements for CPU-intensive computing. These needs generally related to the use of both vector and scalar climate models. All SCFs listed the computational needs for visualization as moderate. Unlike the numerical models presently being run, few scientists had specific software in mind for visualization or analysis. However, most knew which visualization and analytical processes might be important for evaluating their data. In those cases, the specific needs for numerical modeling and visualization provided a better indication of hardware requirements for each SCF, than did the scientist's own input on CPU or graphics requirements. In particular, most of the scientists probably tended to underestimate the computer power required to meet many of their visualization requirements. Requirements for analysis software varied from none to high.

With regard to connectivity, all SCFs anticipated that their requirements for communication outside of MSFC could be adequately met with the existing capabilities. Requirements for high-performance networking within MSFC was high, particularly with regard to the MSFC EADS facility, the EOS DAAC, the ESAD Cray, and other ESAD scientists. No SCFs required protocols other than TCP/IP, NFS, and possibly NQS.

Requirements for video recording were low to high, whereas video frame grabbing needs ranged from none to moderate in importance. Although the ability to record video sequences was important for most SCFs, the PIs, in all cases, anticipated that this need would be sporadic. Requirements for high-quality printing of color images and plots were high and less sporadic, while image scanning requirements were moderate to nonexistent.

Although each individual SCF has specific individual requirements, most of the SCFs have common needs, and in many cases overlapping needs. Some of these requirements could be met with shared facilities. With regard to individual SCF needs, none of the SCFs' requirements hindered the possible use of compatible systems.

PART III. CONSIDERATIONS

A. INTRODUCTION

The tremendous speed and magnitude of change in the computer industry makes it very difficult to make proper decisions regarding the purchase of computer hardware and software. The computing power available to scientists is growing by an order of magnitude every 6 or 7 years. Vendors continue to leapfrog one another in power and functionality as new CPU chips are introduced, as computer architectures are redesigned, and as compiler technology is improved. It is important that those responsible for making decisions regarding computing philosophies and actual computer purchases be aware of the directions of the computer industry, in addition to its current state.

In order to determine the best options for meeting the total SCF requirements, the SCF Working Group examined the current offerings and the trends of the computer industry. In addition, the requirements of the EOSDIS and the capabilities of existing computing facilities at the MSFC ESAD Earth Science and Applications Computing Facility (ESACF) were reviewed. The results of these reviews are presented below.

B. REQUIREMENTS OF THE EOSDIS

The functional objectives of a communication interface between the SCF and the DAAC are to perform data quality control and assurance, to provide algorithm updates to the DAAC, to transfer special data products, and to coordinate the testing and integration of new SCF-developed software on the DAAC. The physical interface of this link is provided by the EOS Science Networks (ESN), which is functionally composed of two separate networks, the mission essential network and the mission success network. The mission success network utilizes the NASA Science Internet with gateways to "The Internet." This network is what the EOS science users would utilize to access the DAAC and would not be used by the SCF to routinely transfer data between the DAAC and the SCF. In contrast, the mission essential network serves as the data pipeline in DAAC to DAAC transfers and between the SCF and the DAACs, and is therefore the primary network to be utilized by the SCF in performing it duties with the DAAC.

The requirements levied on the SCF by the Version 1 DAAC are currently summarized in the EOSDIS Core System Request for Proposal (RFP5-13268/307). The DAAC is composed of three subsystems, the Product Generation System (PGS), the Information Management System (IMS), and the Data Archive and Distribution System (DADS). The individual PGS, IMS, and DADS system requirements are detailed in Appendix B using the EOSDIS Core System Requirements Specification numbers. At present, these requirements place few restrictions on decisions regarding hardware purchases for the SCF.

C. INDUSTRY DIRECTIONS

GENERAL DIRECTIONS

Standardization

Standardization of computer hardware and software is a trend that began toward the end of the 1980s. It began as a reaction to the rapid introduction of dissimilar, proprietary systems introduced in the early 1980s from start-up and established companies. This push to standardize has been driven by customer demands for system longevity, as well as industry risks associated with introducing proprietary technology.

The trend toward standardization has manifested itself in the development of open architectures and cooperative alliances between computer companies, many of whom were originally competitors. Further standardization efforts include increased use of off-the-shelf hardware components and the adoption of software standards for operating systems, programming languages, graphics libraries, windowing environments, and networking protocols. These alliances and standards will be discussed in more detail in the following sections on "Industry Alliances" and "Industry Standards." Consideration of developing standards and their acceptance within the computer industry is vital for ensuring longevity, upgradability, and interoperability of hardware purchases, as well as aiding portability and expandability of software development efforts.

CPU Trends

The Central Processing Unit (CPU) is the primary computation engine of a computer. Therefore, its speed and characteristics drive the maximum theoretical performance of any computer platform. In reality, the actual performance of a computer running a given application is dependent on a large number of factors, including I/O rates, the amount of cache memory available, the performance of the graphics subsystem, the data transfer rates between memory and the CPU or graphics subsystem, and the proper balance of the performance and data flow rates throughout the entire system architecture. Because of potential "bottlenecks" within computer architectures, most computers do not approach full CPU potential for a given application. Still, it is important to consider the present and future directions of CPU technology when evaluating possible computing platforms.

Within the last decade, CPU performance for a given cost has increased by an order of magnitude every 6 to 7 years. As an example, CPUs with 1 million instructions per second (MIPS) were prevalent by 1980, only to be replaced by CPUs near 10 MIPS by the mid 1980s. By the end of 1991, CPUs with performance of near 100 MIPS per chip are anticipated. Industry analysts predict that 1,000 MIPS per chip will be available before the end of the decade. This logarithmic increase in computer power places additional importance on the upgradability of computer hardware and on software development with an eye to future power.

Within the last few years, most microprocessor CPU designs began the shift toward the RISC (Reduced Instruction Set Computer) architecture as compared to earlier designs employing a CISC (Complex Instruction Set Computer) architecture [Leibowitz,1991; Hennessy and Joupp, 1991]. Whereas CISC instructions vary in length (i.e., the number of bits in the machine language instruction), all RISC instructions are the same length and operate only on data in the CPU registers, allowing the instructions to be processed in a "pipelined" fashion (similar to an assembly line), effectively allowing the CPU to execute one (or more) instructions per clock tick and thus raising the CPU's MIPS rating. Pipelining, register-based operations, and a small instruction set of same-length instructions are central features of the RISC architecture. Figure III.1 illustrates the logarithmic increase of CPU power used in supercomputers, mainframes, minicomputers, and microprocessor-based computers (i.e., workstations and PCs). The figure shows a sharp break and rapid increase in the power of microprocessors beginning in 1985, and corresponding to the increased power provided by RISC-based CPUs. The performance of CISC-based CPUs has continued growing at the rate exemplified by the pre-1985 segment of the microprocessor curve.

The CISC architecture is exemplified by the Intel 808x and 80x86 family of microprocessors (used in the IBM PC and PS/2 lines) and the Motorola 680x0 family (used in the Apple Macintosh lines). These processors have grown logarithmically in processing power over the years but not at the same rate as the newer RISC architectures. Every new generation in these families has maintained, for the most part, compatibility with previous generations. Unfortunately this compatibility requires that the instruction set grows with each new

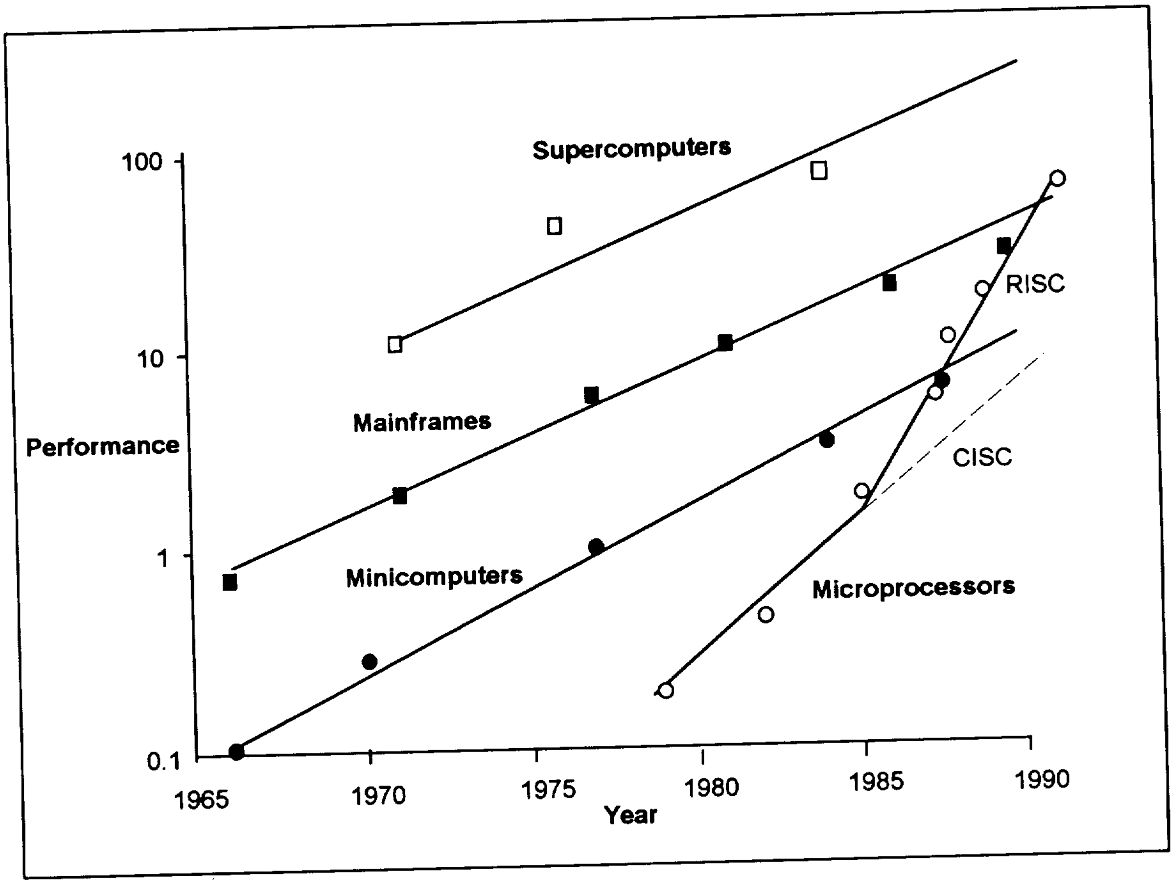

Figure III.1. Trends in microprocessor, mainframe, minicomputer, and supercomputer technology, showing a very rapid increase in performance of microprocessors associated with the RISC technology [adapted from Hennessey and Joupp, 1991].

generation leading to more inefficient instruction execution. To combat this, CISC manufacturers are employing RISC-like features in newer generations.

The RISC architecture is exemplified by the Mips R-series (R2000, R3000, and R4000), the Intel i860 family, the Motorola 88000 chip set, the Sun SPARC family, the IBM RS/6000 chip set, and the HP/Apollo PA-RISC chip set [Corcoran, 1991a; Smith, 1991a, b; Iverson, 1991]. Though most of these newer CPUs actually have an instruction set that compares in size to earlier CISC designs, they are still considered primarily RISC architectures since they try to execute at least one instruction every CPU clock tick. Newer RISC designs attempt to execute multiple instructions per tick by employing "superscalar" and "superpipelining" techniques that add multiple instruction pipelines and more instruction processing stages to each pipeline.

Clearly the choice of CPU that a computer manufacturer uses is critical if they want to maximize performance growth potential and minimize future compatibility problems. Typically, the CPU characteristics (other than speed) are not visible to application users nor to high-level language (FORTRAN, C, etc.) programmers. However, the customer should be concerned with potential obsolescence and incompatibility of their purchased hardware should the computer manufacturer choose a different CPU line in future products. Figure III.2 shows a timeline of past, present, and future CPU utilization for key manufacturers of PCs and workstations.

Parallel Processing and Distributed Processing

Multiprocessing, or parallel processing, involves simultaneous computation on two or more CPUs within a single computer. In theory, multiprocessing increases the computing power by *n* times, where *n* is the number of parallel processors. The major benefits of parallel processing are two-fold. First, on multitasking/multi-user workstations, separate tasks and users may be running on different CPUs simultaneously and thus not competing with each other for CPU time. The second benefit comes from using a parallelizing compiler on application code such that during execution, separate pieces of the code run simultaneously on different CPUs thus giving more processing power to the application.

There are four ways to distribute instructions and data across multiple processors: Single Intruction-Single Data (SISD), Single Instruction-Multiple Data (SIMD), Multiple Instruction-Single Data (MISD), and Multiple Instruction-Multiple Data (MIMD). Single and multiple instruction refers to whether the processors must all execute the same code in lockstep or whether they are free to execute different code. Single and multiple data refers to whether the processors must share the same data or can operate on different data. The MIMD architecture obviously provides the greatest flexibility and is typically employed for most parallel systems with few (roughly < 64) processors. Primarily due to a lack of development tools, MIMD systems have traditionally been more difficult to program than the other architectures; however, the popularity of the MIMD architecture should increase availability of better development tools. The other architectures trade off the flexibility inherent in MIMD for easier programmability and cheaper cost.

For MISD and MIMD architectures, it is important to differentiate between asymmetric multiprocessing, in which each processor basically works independently, and symmetric multiprocessing, in which the processors are tightly coupled by the operating system [Forbes, 1991a, b]. Asymmetric processing is greatly limited in its ability to take full advantage of multiple CPUs, whereas in tightly-coupled symmetric processing, there is a near-linear relationship between the number of processors and the level of performance [Forbes, 1991b]. The general trend in recent years is to provide a growth path for additional CPUs within a given computer platform and to progress from asymmetric to symmetric processing [Baum and Winget, 1990].

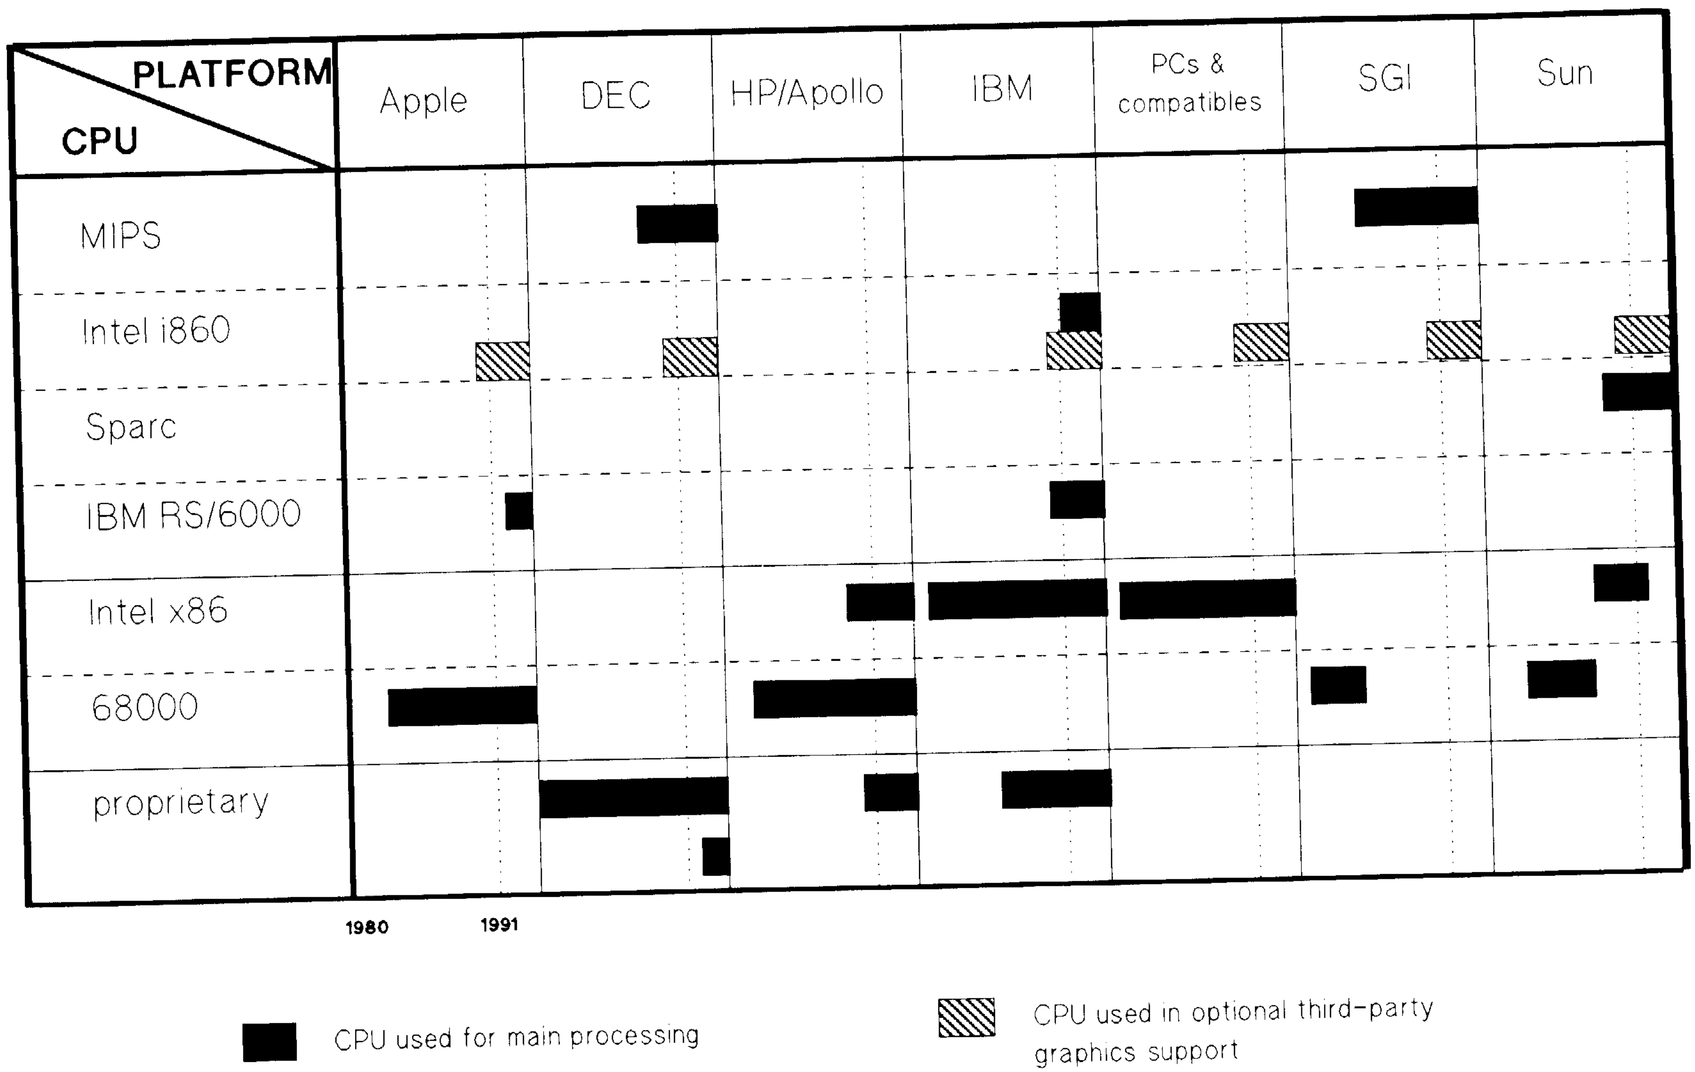

Figure III.2. A CPU history plot showing the usage of various CPUs by workstation vendors. Within each column, the time scale ranges from 1980 to 1994 with the dashed line representing 1991.

Symmetric multiprocessing originated in the mainframe and parallel supercomputer markets. Of the major workstation vendors, only SGI and Stardent have provided symmetric multiprocessor Unix workstations since 1988 [Montgomery, 1991]; Sun introduced a multiprocessor SPARCserver at the end of 1991; however, it will not be symmetric until the release of the Solaris 2.0 operating system in mid-1992 [Corcoran, 1991c]. Also introduced in 1991 were a few first-generation massively parallel processing (MPP) servers, such as the Wavetracer, which consists of thousands of processors and utilizes other workstations for front-end processing [Smith, 1991d]. In August 1991, IBM announced its high-end 32-processor Power Visualizer which utilizes the IBM RISC 6000 as its front-end, placing the platform somewhere between traditional multiprocessors and MPP technology.

The same benefits of multiple CPUs can be achieved even if the CPUs are not in the same physical computer. Distributed processing allows different computers to share their workload as well as their file systems, assuming that the proper connections and software exist on the machines. An additional benefit of distributed processing is the ability to distribute particular subtasks to various computer systems that are best suited for solving that particular task. Unlike parallel processing, distributed processing introduces more compatibility and management issues since networking and dissimilar machines are usually involved. Distributed processing provides far more flexibility regarding upgrades and expansion since the number of computers that can be added is limited only by networking constraints. The network speed is also the biggest bottleneck since it is typically several orders of magnitude slower than the bus speed in a multiprocessor system. Having a mixture of machines with varying performance characteristics ideally allows for a closer coupling between application run-time requirements and hardware availability than does a multiprocessor system.

Importance of Software

The availability of application software and software development tools is a critical concern when making purchasing decisions for computer hardware. Computer power is useless unless it can be directed by application software toward solving the needs of the computer user in a timely and efficient manner. Application software must either be available for the computer platform under consideration or it must be custom developed. Realistically, even if all needs can initially be met with off-the-shelf applications, at some point custom development will be needed to enhance functionality. For custom development, the availability of development and project management tools, tool libraries, reference materials, and development support are important concerns. Many computer manufacturers have invested considerable effort and finances to provide these items, the availability of which can drastically shorten development times while increasing the quality of the final product.

Merging of Personal Computer and Workstation Environments

Workstations were introduced in the early 1980s as a result of the successful introduction of personal computers in the late 1970s. The personal computer set the stage for the "one user - one computer" philosophy that ran counter to earlier uses of mainframes and minicomputers; the workstation was conceived as a minicomputer packaged in a personal computer form-factor for scientific and engineering applications. As CPU performance increased and prices dropped, personal computers approached the computational power of low-end workstations while low-end workstations became comparable in price to personal computers. The current trend in the industry is to provide a single platform which couples the diverse, user-friendly software of a personal computer with the computational power, connectivity, and graphics capability of a workstation.

Significant efforts are underway to allow the cohabitation of DOS, Macintosh, and Unix operating systems on the same platform and to allow the interchange of application software and data between these environments. These efforts include the actions of the Advanced Computing Environment (ACE) consortium and the IBM-Apple alliance, to be discussed below, as well as the release of Sun's Solaris operating system for Intel-based systems. At the present, Unix is available on PCs through Santa Cruz Operation's (SCO) Open Desktop, while DOS applications can be run on Unix workstations at 10 MHz 80286 speeds or better through the use of a DOS emulator from Insignia, Inc. Macintosh applications can be run as is on RISC-based Unix workstations utilitizing a combination emulator/library package from Quorum Software Systems, Inc. [Picarille, 1992].

Merging of Workstation and Supercomputer Technology

As a result of increases in CPU power, parallelization, and compiler technology, workstations are now capable of performing compute-intensive tasks traditionally reserved for supercomputers only a few years ago. This trend toward incorporating supercomputer-class power into the graphics workstation environment began with the introduction of the Ardent Titan and the Stellar GS superworkstations in the late 1980s. These platforms included vector processing hardware similar to that used in supercomputers with the added functionality of high-speed graphics for displaying the computed results. Since that time, other manufacturers, such as SGI, IBM, and HP, have introduced high-performance superworkstations based on scalar hardware, believing that most application code cannot be effectively vectorized.

As shown in Figure III.1, RISC-based CPU performance is growing at a faster rate than CPU performance of traditional supercomputers. The supercomputer architecture of the future will probably not rely on faster processors for increased performance, but will instead gain increased performance through massively parallel processing (MPP) architectures with thousands of interconnected CPUs. MPP technology is exemplified by many available or announced platforms, including the Connection Machine offered by Thinking Machines, Inc., Kindell Square Research's KSR1, Alliant Corporation's Campus, and the MasPar system offered by Digital Equipment Corporation (DEC) [Smith, 1991e]. MPP systems typically employ the SIMD architecture to ease programmability but newer architectures (Thinking Machines' CM-5) are employing a MIMD architecture with the capability to run in a SIMD-like fashion. For full-scale supercomputing applications, analysts predict that MPP systems will "reach the crossover point" - achieve equality with traditional vector systems - around 1994 or 1995 [Smith, 1991c]. Additionally, many "supercomputers" of the future will actually consist of an array of individual workstations networked together and operating under a distributed processing environment.

INDUSTRY STANDARDS

Operating Systems

Unix, DEC's VMS, the IBM-PC operating systems (DOS, OS/2, DR-DOS), and Apple's Macintosh operating system are the prevalent operating systems found on personal computers and workstations today, while Unix, VMS, and IBM's MVS are the prevalent operating systems found on minicomputers and supercomputers machines. Based on the number of ports to different architectures and current trends of manufacturers, Unix has become the de-facto standard for workstation and supercomputer class machines. Unix is also available for the PC environment through SCO's Open Desktop. In a recent move, Sun Microsystems announced the availability by mid-1992 of the latest version of their Unix operating system, Solaris 2.0, on PCs [Burns, 1991b; Bucken, 1991].

Unix currently comes in two major forms: AT&T System V (ATT SV) and Berkeley System Development (BSD). Most manufacturers have developed flavors of Unix based on ATT SV, with Berkeley extensions related to communications and connectivity. Few companies continue to operate strictly under BSD due to the early momentum in Unix support gathered by AT&T. The POSIX standard is a government-imposed Unix operating system specification based on ATT SV with Berkeley extensions. Recently, two major alliances have formed in order to standardize and promote the growth of Unix: Unix International (UI), founded by AT&T, and the Open Software Foundation (OSF), formed by IBM and DEC with the backing of most major system vendors. Both comply with POSIX and offer additional functionality for distributed computing support and distributed resource management support. Although most manufacturers have a different name for their Unix operating system, the functional differences between these Unix flavors have become relatively minor due to standards compliance. In addition, the OSF has recently embraced the Architecture Neutral Distribution Format (ANDF) which promises to extend the advantages of off-the-shelf or "shrink-wrapped" software to architectures that are not binary compatible [Serlin, 1991]. This would allow software developers to create a single version of a Unix application, which would then run correctly on any other OSF Unix machine regardless of the vendor.

DOS was developed by Microsoft for the IBM Personal Computer as a single-user, single-tasking operating system. Due to the success of the PC and the demands of the users, DOS has grown in capability but remains primarily a single-user, single-tasking operating system. The Windows environment, also developed by Microsoft, is a graphical user-interface shell, not an operating system, and functions on top of DOS. OS/2 was developed as a separate effort by IBM to perform true multi-tasking on the upper end of the PC line and uses its own graphical user-interface. OS/2 has not met with the anticipated development and market success, although some analysts feel that OS/2 Version 2.0, due in early 1992, will show great potential by allowing DOS, Windows, and OS/2 applications to coexist within it while Windows will be limited to only DOS and Windows applications. The success of OS/2 or Windows will depend greatly on the availability of useful applications and their perceived efficiency on the millions of existing PCs. The future directions of DOS and OS/2 are uncertain in light of the recent destruction of old alliances, such as IBM-Microsoft, and the creation of new alliances, such as the ACE and the IBM-Apple deal. The Windows interface standard will remain a viable player in the industry due to the large market presence of Microsoft software in PCs and low-end workstations.

In a move that perhaps best illustrates the rapid changes toward RISC and Unix in the computer field, DEC made the startling announcement in September 1991 that they will soon license their once closely-held VMS operating system in a RISC-based form to other computer manufacturers in an effort to recapture the market shift toward open, standards-based systems such as Unix [Vizard and Ballou, 1991; Vizard, 1991b]. The performance of this RISC-based version of VMS running on a new DEC-developed RISC CPU, codenamed Alpha, due sometime in late 1992 or early 1993, is expected to be three to four times greater than the next generation CISC-based VMS system scheduled for release by the end of 1991. An even clearer sign of the gathering strength toward Unix as a standard operating environment is DEC's announced schedule to deliver a POSIX compliant interface for VMS by the end of 1991 so that VMS can operate essentially as Unix in order to ease portability of applications across VMS and Unix environments.

Windowing Environments

Originally developed on workstations at the Xerox Palo Alto Research Center in the early 1970s, windowing environments have spread to most personal computers and all workstations starting commercially with the Apple Lisa and Macintosh in the early 1980s. During the mid-80s, most manufacturers created their own, non-portable versions of a windowing environment. By the late 1980s and early 1990s, the X window environment emerged as the standard for

Unix workstations. The success of X is greatly attributed to its development as a networking protocol allowing graphical connectivity to a wide variety of machines. Since X purposely defines only the protocol for distributed windowing graphics and not the look and feel of the computer display, two standards which define the look and feel have emerged: OSF X/Motif and Unix System Laboratories' (USL) X/Open Look. Table III.1 illustrates the vendor and third-party support for these two standards on different platforms. Clearly, OSF X/Motif has become the window interface of choice among the majority of Unix-based hardware vendors [Burgard, 1991; O'Connell, 1991]. It is important to note, however, that these two standards do not imply incompatibility; any X window-based environment can run programs compiled with either standard. In a network environment, X provides greatly enhanced capabilities for both text and interactive graphics display relative to that available for simple text-only terminals. On PC platforms, X/Motif windowing capabilities are available through SCO's Open Desktop Unix, and within the DOS environment through emulators, such as VisionWare's XVision and Spectragraphics PC-XView [Gill and Hammett, 1992]. X environments exist for the Macintosh as well, available through Apple and third parties.

With the breakup of the Microsoft and IBM team and the creation of a new alliance between IBM and Apple, the future direction of windowing environments on the PC is not as clear as before. Clearly the combination of the undeniable market success of Microsoft's Windows environment, the market failure of OS/2, and the support of Microsoft's future Windows NT (New Technology) environment within the ACE consortium will provide strong impetus for the Windows API. However, more PC users in the future may be enticed over to the Unix/X windowing environment as a result of the incorporation of DOS, Windows, and Macintosh applications into the Unix operating system under efforts of the ACE and IBM-Apple alliances. This transition is being aided by the current existence of user-friendly, "Mac-like" Unix interfaces, such as the icon-based X.desktop from IXI Limited and the Workspace tools from Sun and SGI [Hansen, 1991]. The IBM-Apple alliance will introduce a new object-oriented operating system and windowing interface, codenamed Pink, that can run PC and Macintosh applications as well as a new generation of media-intensive applications. However, this operating system is not scheduled to be released until at least 1993, at which time the ACE X/Motif and Windows NT environments, and Sun's Solaris operating system, should have become well entrenched.

Text: Hardcopy and Screen

The 1980s saw the emergence of "What You See Is What You Get" (WYSIWYG) capability for textual and simple 2D graphic editing and output. By the mid 1980s, the PostScript (PS) page-description language (PDL) from Adobe Systems had emerged as the leading standard for describing resolution-independent text and 2D graphics for hardcopy output in black and white and in full color [Barkes, 1991]. For simple text output (e.g., program or data listings), there are no special requirements for printing to a laser or dot-matrix printer, and many manufacturers, such as Hewlett-Packard, offer more affordable printers with proprietary and simpler PDLs.

With the emergence of the X window standard, the X font capability has become the standard for simple text display on Unix-based workstation screens. For more complex text display and manipulation, Display Postscript (DPS) has emerged as a standard. When DPS is supported on a platform, it generally coexists as an option to the X fonts. There are few other font technologies, most notably TrueType from Apple Computer, that provide resolution-independent scalable fonts that can be generated at the speed required for interactive use on a workstation display. PostScript leads the text presentation market primarily because it has provided a consistent implementation for display on workstation screens and on a wide variety of hardcopy devices.

Table III.1. Vendor and third-party support for Motif and Open Look X window-based interfaces [Burgard, 1991].

VENDOR	MOTIF	OPEN LOOK
IBM	V	T
DEC	V	T
HP/Apollo	V	T
Unisys	V	
Sun	T	V
Solbourne	V	V
Compaq	V	
Dell	V	
Prime	V	
Data General	V	
Silicon Graphics	V	
MIPS	V	
NCR	V	
AT&T	V	V
Wang	V	
NEC	V	
Hitachi	V	
Commodore		V

V - Vendor supported

T - Third-party supported

3D Graphics

No defacto standard has yet emerged for interactive and static 3D graphics. The present X standard was developed as a 2D text and graphics protocol, and has proved inadequate for real-time 3D graphics. This is particularly true for color graphics, or for 3D X display over the network [Hayes, 1991; Hess, 1990]. Protocol Extensions to X (PEX), a protocol for improving 3D graphics over X, is under development. The PEX protocol is based on the Programmers' Hierarchical Interactive Graphics System (PHIGS), a standard developed primarily for Computer Aided Design (CAD) applications, and is not well suited for the rapidly changing geometric information found in scientific visualizations [Jenkins, 1991b]. PEX, like X, requires that each manufacturer support and implement the standard on their hardware, and as of now, it is not clear whether PEX will achieve the industry endorsement anticipated by the developing committee. Although not often used as the primary 3D graphics language for scientific applications, PHIGS is presently available on most workstation platforms. Likewise, Sun's XGL library is based on PHIGS+ [Johnson, 1991].

Renderman is an established standard developed by Pixar that encompasses both a photorealistic renderer and a textual means to express complex 3D scenes. It was designed and is used primarily for generation of static images or non-interactive 3D graphics and not for scientific visualization.

A significant development with regard to 3D graphics standards has been the public release of Silicon Graphics' (SGI) proprietary IRIS GL. Undoubtedly the increased success of SGI as a manufacturer of graphics workstations has been in part due to the endorsement of the GL language by graphics programmers. Unlike PHIGS, IRIS GL is designed for rapidly changing geometries, is specifically tailored to high-speed graphics hardware, and is the most widely used graphics library in the industry with more than 1,400 applications [Siino, 1991]. IRIS GL has been enhanced considerably over the years while maintaining backward compatibility, an important issue when considering a 3D graphics capability for future development use. In the past, the GL included not only graphics presentation capabilities but a windowing library as well. In addition, the GL has in the past been tightly coupled with the SGI Geometry Engine (GE) hardware. These two factors have, in the past, inhibited the acceptance of the GL as a true standard for 3D graphics.

However, in recognition that the GL provides primarily a high-speed interactive rendering capability, SGI has been concentrating efforts on establishing the GL as a 3D graphics rendering standard. These efforts include decoupling earlier GL-based windowing protocols from the GL and developing a software implementation of the GL. In September 1991, SGI announced that it would license the GL to all vendors and has developed an enhanced version that supports the X windows standard [Corcoran, 1991a]. SGI has created an advisory committee, which includes members from IBM, DEC, Compaq, and Intel, with the task of guiding the development and future of the GL. An agreement between DEC and SGI includes plans for incorporating PEX into the GL, and for support of the GL under DEC's VMS and Ultrix operating systems [Grygo, 1991b].

Prior to the September announcement, the GL had already been licensed to IBM for the RS/6000 platform and to Microsoft Corporation for future incorporation into its Windows NT. In addition, the GL has recently made available for SPARC platforms, including Suns, through Du Pont Pixel Systems', PX/GL [Du Pont Pixel Systems, 1991]. Immediately following SGI's Open GL announcement, Compaq Computer Corp., DEC, and Intel Corp. pledged their support for the GL, with Intel announcing its intention to integrate the GL into its i860 chip. DEC dropped development on their own nearly completed graphics library and offered the new technology to SGI [Grzanka, 1991]. The supercomputer firms, Cray and Convex, both resell DGL, a distributed-server version of GL, and have applauded the opening of IRIS GL. By November, six more firms had licensed the IRIS GL, primarily for the PC arena. As will be discussed below, the GL will be made available to all ACE compliant platforms, running either Microsoft Windows NT or SCO Open Desktop [Bruno, 1991]. The IRIS GL is expected to become a major standard for interactive 3D graphics, although probably not at the expense of PEX, which will still serve the CAD market. The support of IRIS GL on a particular computing platform will greatly aid in-house development of high-performance visualization applications, as well as increase the availability of third party visualization software.

Networking

In the workstation domain, the Transport Control Protocol/Internet Protocol (TCP/IP) is a firmly-established standard on Unix-based systems with connectivity available on more proprietary systems such as Crays, VAXes, Macintoshes, and PCs. It is typically used with Ethernet as the connectivity hardware, theoretically providing over 1 MB per second but in practice delivering 300 to 700 kB per second communication capability. TCP/IP support includes standard network file-copying software (ftp) and remote login capability (telnet and rlogin). Higher-level TCP/IP-based tools include the Network Filing System (NFS) for transparent remote file sharing capability and the Network Queueing System (NQS) for remote process capability.

A more recent development in the hardware connectivity domain is the Fiber Distributed Data Interface (FDDI) that increases the theoretical communication performance to over 12 megabytes per second (3 to 7 MB per second in actual use), a factor of 10 over Ethernet [Green, 1991]. The present FDDI standard utilizes fiber optics for connectivity, but recent committees have been formed to produce a standard using traditional twisted-pair wiring in hopes of lowering the connection and wiring cost. A similar development is the High-Performance Parallel Interface (HIPPI) that delivers a theoretical performance of 100 megabytes per second (30 to 70 MB per second in actual use), roughly 10 times faster than FDDI and 100 times faster than Ethernet. HIPPI is an increasingly popular standard for connecting supercomputers and high-speed graphics platforms to deliver real-time distributed visualization capability. In addition to efforts directed toward increasing the bandwidth of the transmitting cable, other development has concentrated on increasing network transfer rates through intelligent switching nodes. Many in the industry feel that this option will provide a better, more scalable solution to increasing network transfer rates than FDDI.

The International Standards Organization (ISO) has defined an Open Systems Interconnection (OSI) model that was intended to become a standard for communication connectivity and protocols but which has yet to be realized on the same scale as TCP/IP and Ethernet [Varney, 1991]. Other proprietary systems have their own connectivity hardware and software such as DEC's DECnet, Novell's Netware and Microsoft's LANman for the IBM PC, IBM's SNA, and Apple's AppleTalk.

INDUSTRY ALLIANCES

Advanced Computing Environment (ACE)

In Spring 1991, over 20 hardware and software vendors announced the creation of the Advanced Computer Environment (ACE) consortium, formed to develop and support an environment which would essentially merge the PC and RISC-based Unix workstation environments. Since then, the consortium has rapidly grown to more than 85 members, including Digital Equipment Corp. (DEC), Compaq Computer Corp., Microsoft Corp., Mips Computer Systems, The Santa Cruz Operation (SCO), Silicon Graphics (SGI), NEC Corp., Prime Computer, Inc., Wang Laboratories, and Zenith Data Systems, among others [May, 1991; Ballou, 1991; Flack, 1991, Smith, 1991a; Porter *et al.*, 1991].

The ACE consortium standard will support two hardware architectures: the Intel 80x86 PC standard and a new definition of RISC-based hardware called the Advanced RISC Computing (ARC) specification, which includes the MIPS R3000/R4000 CPU. In addition, ACE will support two advanced operating systems, either of which will operate on both platforms: one Unix-based, based on the SCO Open Desktop and OSF/1, and the other PC-based, consisting of the Microsoft Windows NT operating system. Furthermore, ACE will be compatible with both PC- and Unix-based networking services, including comprehensive interoperability with Novell, Banyan, and Microsoft networking services, and TCP/IP, SNA, and DECnet protocols. First-generation ACE compliant workstations are already available from SGI and DEC, with more system releases expected from other vendors in early 1992. The ACE compliant SCO Open Desktop Unix and the development version of Microsoft's Windows NT operating system were released in late 1991 [Grygo, 1991a; Johnston, 1991].

SPARC International

SPARC International was founded in 1989 as an independent association of corporations, institutions, and individuals with an interest in the standardization and evolution of the Scalar Processor Architecture (SPARC) technology [Hubley, 1991]. Developed by Sun Microsystems, the SPARC is available at low cost and has been implemented by many international hardware manufacturers creating low-end, affordable RISC computers. One of the main design philosophies for the SPARC CPU was performance scalability with little change to the architecture. This has held fairly true over the years but the design has not lent itself well to current efforts by Texas Instruments and Sun to produce a higher-speed superscalar SPARC CPU, called the Viking, which would be capable of running 50 to 100 MIPS [Corcoran, 1991b; Wilson, 1991]. The relationship between Sun and some SPARC vendors in SPARC International was recently soured by Sun's surprise announcement to its dealers and value-added resellers (VARs), that they could not sell any SPARC-based computers other than laptops and mainframes if they wished to continue business with Sun [Poole, 1991].

Apple-IBM Alliance.

In a surprise announcement in the spring of 1991, former competitors Apple Computer and IBM announced an alliance to develop a new operating system for IBM RS/6000, Intel 80x86, and Motorola 680x0 CPU lines and as-yet unreleased Apple and IBM workstations running with a single-chip implementation of the RS/6000 chip set [Vizard, 1991a; Quinlan and Scannell, 1991; Scannell and Quinlan, 1991; Quinlan *et al.*, 1991; Jenkins, 1991a]. The single chip implementation, dubbed the PowerPC, is to be manufactured by Motorola with assistance from IBM and marketed by Motorola. The operating system, as yet unnamed and to be marketed by the jointly owned Apple-IBM spinoff, Taligent, will be developed primarily with object-oriented technology from Apple's "Pink" project and from Metaphor, a company previously purchased by IBM to develop an object-oriented operating system. A second jointly owned company, Kaleida (as in "kaleidascope"), will provide multimedia capabilities using Apple's QuickTime audio/video technology. Any Kaleida technology will certainly find its way into Taligent's operating system developments and should serve to differentiate the new operating system from current systems. This new system will incorporate backwards compatibility with DOS, IBM OS/2, and Apple's Macintosh OS using "personality" modules. Apple and IBM have also pledged to merge their A/UX and AIX Unix offerings into a single Unix, dubbed PowerOPEN, able to run on RS/6000 and Macintosh platforms. Many analysts view this alliance as a reaction to the threat of the ACE consortium as well as a reaction to the present PC software monopoly held by Microsoft [Vizard, 1991a]. The earliest anticipated release of the new operating system and the single-chip RS/6000 implementation is 1993 [Jenkins, 1991a].

D. MARKET SURVEY

COMPANY PROFILES

AST Research, Inc.

AST Research, Inc. began in the early 1980s as a manufacturer of add-on boards for PCs and compatibles, and moved into manufacturing a full PC compatible line in the mid 1980s. AST has concentrated on providing high-quality, high-performance PCs, with a somewhat unique proprietary Cupid architecture that allows easy future upgrades to newer or higher performance subsystems.

.

Apple Computer, Inc.

Founded in the mid-1970s, Apple popularized the now common window-based user-interface with their Macintosh line, introduced in 1984. The well-rounded ease of use of the Macintosh remains often imitated but as yet unmatched in the personal computer and workstation markets. The Macintosh line will be updated to a RISC-based platform to compete in price/performance by utilizing the single-chip implementation of the RS/6000 CPU, dubbed the PowerPC. Additionally, the Macintosh operating system is undergoing changes to compete with Unix-based systems. These changes will take several years, and in the interim, the Macintosh will be under strong competition from PC and ACE systems running Microsoft's Windows NT operating system. In addition to information management software, an abundance of scientific software, most of it public domain, exists for the Macintosh.

Compaq Computer Corporation

Compaq was created to make a down-sized, "luggable" version of the popular IBM PC in the early 1980s. Since that time the company has moved into desktop personal computers and low-end workstations and has pushed the market toward higher-end (and higher priced) PC-based architectures, including multi-processor designs. Seeking newer RISC-based workstation markets, Compaq helped to create the ACE consortium. In recent years, Compaq allowed lower-priced PC clone companies to undercut their market and has found itself having to reduce prices. Compaq's new management has stated that they will leave their higher-end markets and, aside from supporting the ACE, will instead concentrate on the PC clone market.

Digital Equipment Corporation (DEC)

DEC became popular with their PDP and VAX minicomputer lines. Over the years DEC associated itself largely with the scientific and higher education markets. DEC offers a wide range of computers, from personal workstations to powerful minicomputers but has recently found itself under attack from inexpensive, high-powered workstation vendors that do not have to support such a large installed customer base. Perhaps as a reaction to that, DEC became an original member of the ACE consortium. At the present, DEC workstations are based on the Mips CPU. However, DEC will introduce a new RISC-based CPU design, code-named "Alpha", during 1992 that should offer significant performance over older DEC CPUs and compete with the Mips CPU, as well. Alpha-based systems will initially be offered in larger platforms running at 200 MHz, though workstation-class machines should appear in 1993 or 1994. Compatibility between older VAX-VMS systems and the Alpha-based systems is critical to retaining their customer base and DEC plans to provide that. DEC has also announced intent to support IRIS GL, which should appear on their workstation platforms within 1992, and has joined SGI in incorporating PEX into the IRIS GL.

HP/Apollo

HP/Apollo workstations represent a merging of the Hewlett-Packard lines with the Apollo Computer lines. HP bought out Apollo in the late 1980s and received a strong boost from Apollo's RISC compiler development and RISC workstation design teams. HP/Apollo has recently stormed the low-end workstation market with RISC-based machines that offer the best price/performance of any other in the market. This platform is based on HP/Apollo's proprietary RISC-based CPU, the PA-RISC. HP/Apollo's machines started a price/performance war with IBM and DEC that continues into 1992. HP/Apollo's primary challenge at the moment is to lure a sufficient mass of third party software support since they have only recently become such a competitive player in the workstation arena.

IBM

The world's largest and oldest computer company, IBM, created the RISC concept decades ago yet never achieved financial success with their RISC systems until the RS/6000 family, introduced in 1990. The RS/6000 line provides good price/performance and offers a Unix and X-windows environment along with standard Ethernet connectivity and IRIS GL capability. Rather than develop proprietary graphics hardware for the RS/6000 line, IBM choose to purchase Personal IRIS Geometry Engine (GE) boards from SGI. IBM's RISC workstation line is beginning to compete with their traditional minicomputer market, as well as their PC lines, due to similar price/performance on their high-end and low-end RS/6000s, respectively. Recently, IBM announced a "breakup" of their various divisions, thereby opening up direct competition between various IBM platforms, as well as between IBM and other vendors. IBM recently announced the POWER Visualizer line as a very high-end visualization server based on up to 32 i860 chips. This system offers a highly specialized and somewhat unique approach to visualization solutions, is priced in the millions of dollars, and depends on the RS/6000 for front-end visualization. Whether it will be successful in the market is yet to be determined.

Intergraph

Intergraph has become very successful as a supplier of "turn-key" computer solutions, primarily for vertical markets such as CAD, architectural engineering, urban development, GIS, and image processing. The main forte of Intergraph has been the "turn-key " integration of all of their software and the optimizing of this software to their propreitary platforms. Intergraph workstations have traditionally not been regarded as generic workstations for running a wide variety of third party software. Having originally developed and supplied software on VAX-based systems, Intergraph has within the past few years ported its huge inventory of in-house code to a proprietary platform based on their RISC-based Clipper chip set. However, realizing that their potential software market may be limited by the high price/performance of their workstations, Intergraph has begun to port its software to various other platforms, including Sun, HP/Apollo, Macintosh, and PC workstations.

NeXT, Inc.

Founded in 1988, NeXT provides a family of low-end, low-priced CISC-based workstations running the proprietary, object-oriented NextStep operating system. NeXT distinguishes itself from all other workstation vendors by providing the only object-oriented operating system and integrated development environment. The capabilities of NextStep provide a good picture of the future of operating systems and development environments. Recently NeXT announced an Intel 80486-based version of NextStep, dubbed NextPC, for use by PC clone manufacturers.

This will provide an important market for NeXT, since their workstation hardware is not unique enough to compete with similarly priced, higher performance workstations.

Silicon Graphics Inc. (SGI)

SGI created the low-cost, high-performance 3D graphics market in the mid-1980s. Until recently, SGI has competed primarily in the mid- to high-end workstation and server markets and has become known for offering the best mix of high-performance 3D graphics capability, networking and disk I/O speed, and software development support within a family of binary-compatible machines. SGI is the largest manufacturer of 3D workstations and has the largest base of 3D applications, and is an original member of the ACE Consortium. The company introduced their symmetric multiprocessing operating system, IRIX, along with their multiprocessing workstations in 1988. SGI recently introduced the IRIS Indigo, a low-end workstation that, for the price, offers good CPU speed and unmatched graphics capability and disk I/O. SGI's IRIS GL technology is poised to become a standard for 3D visualization. SGI offers extensive applications, development tools, and source code to allow programmers to quickly utilize the 3D capabilities of SGI systems.

Sun Microsystems

Sun became popular during the 1980s for offering inexpensive CISC-based workstations that could be easily networked using Ethernet and that provided well-rounded traditional program development tools. Due to a large installed base and good development tools, Sun supports the largest number of third party software applications in the traditional workstation market. Until recently, Sun possessed the majority share of the low-end workstation market, but that share has begun to erode due to strong competition in price/performance from other vendor's low-end RISC-based products. Sun is working on their first multi-processor platform, utilizing higher-speed SPARC CPUs, and is scheduled to introduce it in 1992 along with plans for a symmetric multiprocessing operating system to fully utilize their platform.

SPECIFICATIONS

Because of the various ways in which vendors configure platforms, as well as differences in performance measurements, it is difficult to gather and present data that offer meaningful comparisons between different platforms. When comparing platform performance, particularly for visualization requirements, it is important to try to separate CPU performance from graphics performance. In addition, other factors can greatly affect the overall performance of a computer regardless of the capability of the CPU and graphics subsystems. These include I/O performance, the various bus bandwidths, availability of optimized compilers for languages other than FORTRAN, the speed of memory, and the size and speed of cache. Bottlenecks within any of these subsystems can greatly degrade any overall performance. In multiprocessor platform, the availability and performance of a symmetric operating system and the availablity of parallelizing compilers is important.

Tables III.2 - III.6 provide specifications on various aspects of present workstations from the major workstation vendors considered. The specifications include CPU Performance and Base List Price (Table III.2), System Configurations (Table III.3), Graphics Options (Table III.4), Graphics Features (Table III.5), and Graphics Performance (Table III.6).

Table III.2. CPU Performance and Base List Price

Model	Base List Price	CPU & Speed (MHz)	SPECmark	SPECint	SPECfp	MIPS	MFLOPS
AST							
Premium II 386DX	$2,795	80386DX @ 33				8	
Premium II 486SX	$4,495	80486SX @ 33				14	
DECstation							
20 (Personal DEC)	$3,995	R3000 @ 20	16			22	2
25 (Personal DEC)	$4,995	R3000 @ 25	19			27	3
125	$6,495	R3000 @ 25	16			27	3
133	$8,495	R3000 @ 33	26				6
200	$4,000	R3000 @ 25	16	19	18	27	3
240	$11,995	R3000 @ 40	32			43	6
DEC VAX 6000							
610	$247,000	NVAX @ 83	40	30	47		
660	$659,000	6 NVAX @ 83	197	144	243		
HP/Apollo 9000							
705	$4,990	PA-RISC @ 35	34			35	8
710	$9,490	PA-RISC @ 50	50	35	62	57	12
720	$14,990	PA-RISC @ 50	60	40	78	57	17
730	$20,990	PA-RISC @ 66	77	51	101	76	22
750	$48,190	PA-RISC @ 66	78	52	102	76	22
IBM RS/6000							
220	$6,595	POWER @ 33	26	18	34		7
320	$5,500	POWER @ 20	33	16	53	30	9
320H	$12,000	POWER @ 25	41	20	67	37	12
340	$18,895	POWER @ 33	57	29	89		15
350	$26,895	POWER @ 42	71	36	112		19
520H	$28,100	POWER @ 25	44	22	69		12
530H	$31,500	POWER @ 33	57	27	96		20
550	$52,500	POWER @ 41	72	34	120	62	26
560	$64,110	POWER @ 50	89	44	144		31
930	$59,500	POWER @ 25	43	20	72		15
950	$94,500	POWER @ 41	72	34	120		26
Intergraph							
2430	$18,500	C400 @ 40	33			36	10
6450	$24,500	C400 @ 40	33			36	10
6480	$37,900	C400 @ 40	33			36	10
Silicon Graphics 4D							
RPC (Indigo)	$7,995	R3000 @ 33	26			30	4
30 (Personal IRIS)	$15,000	R3000 @ 33	27			27	5
35 (Personal IRIS)	$21,000	R3000 @ 33	31			33	6
Crimson	$27,900	R4000 @ 50	70				16
310	$49,900	R3000 @ 33	25			30	5
380	$154,900	8 R3000 @ 33	128			234	60
420	$94,900	2 R3000 @ 40	56			72	23
480	$194,400	8 R3000 @ 40	166			286	70
Sun SPARCstation							
ELC	$4,995	SPARC @ 33	20			21	3
IPC	$9,995	SPARC @ 25	14	12	11	16	2
IPX	$12,000	SPARC @ 40	24			29	4
2	$18,000	SPARC @ 40	25	21	22	29	4

Note: SPECmark is a current industry standard benchmark measuring a mixture of integer and floating-point tests. SPECint is the integer subset of SPECmark. SPECfp is the double-precision floating point subset of SPECmark. The memory and disk configuration of a base system is listed in Table III.3.

Table III.3. System Configurations: Memory Capacity, Disk Storage Capacity, and Expansion Slots

Model	Base RAM (meg)	Max RAM (meg)	Base Disk (meg)	Max Internal Disk (meg)	Internal & External Disk (meg)	Expansion Slots
DECstation						
20 (Personal DEC)	8	40	-	426		2 TurboChannel
25 (Personal DEC)	8	40	-	426		2 TurboChannel
125	8	128	-	852		3 TurboChannel
133	8	128	-	852		3 TurboChannel
200	16	120	-	852		3 TurboChannel
240	16	120	-	852		3 TurboChannel
VAX 6000						
610	64	128	-	4,000	4,000	
660	64	512	-	8,000	7,000,000	
HP/Apollo 9000						
705	8	64	0	840	9,450	0 slots
710	16	64	0	840	9,450	0 slots
720	16	128	420	840	64,000	1 opt. EISA
730	16	128	420	840	64,000	1 EISA
750	32	384	1,300	2,600	236,000	4 EISA
IBM RS/6000						
220	16	64	0	400	4,400	2 MCA
320	8	128	160	800	8,800	4 MCA
320H	16	128	400	800	8,800	4 MCA
340	16	128	160	800	8,800	3 MCA
350	32	128	160	800	8,800	4 MCA
520	16	512	355	2,500	26,500	4 MCA
520H	16	512	400	2,500	26,500	7 MCA
530H	32	512	400	2,500	26,500	7 MCA
550	64	512	800	2,500	26,500	7 MCA
560	64	512	800	2,500	26,500	7 MCA
930	32	512	670	11,900	53,100	7 MCA
950	64	512	1,700	11,900	53,100	7 MCA
Intergraph						
2430	16	128	426	426	4,430	0 slots
6450	16	256	426	1,000	5,000	5 slots, VME & SRX
6480	16	256	426	1,000	5,000	5 slots, VME & SRX
Silicon Graphics 4D						
RPC (Indigo)	8	96	0	2,200	2,200	2 GIO-32
30 (Personal IRIS)	8	128	0	3,600	82,000	1 VME
35 (Personal IRIS)	8	128	0	3,600	82,000	1 VME
Crimson	16	256	780	3,600	82,000	4 VME
310	8	256	780	14,000	116,000	4 or 10 VME
380	8	256	780	14,000	116,000	4 or 10 VME
420	8	256	780	14,000	116,000	4 or 10 VME
480	8	256	780	14,000	116,000	4 or 10 VME
Sun SPARCstation						
ELC	8	64	-	-	5,200	0 slots
IPC	8	48	207	207	15,600	2 Sbus
IPX	16	64	207	207	15,600	2 Sbus
2	32	128	424	424	20,800	3 Sbus

Table III.4. Graphics Options

Model	List Price	CPU Availability	Monitor Sizes	Resolution
DECstation				
HX	$1,995	All	16", 19"	1280 x 1024
PXG+	$5,000 - $12,000	All	16", 19"	1280 x 1024
PXG Turbo+	$25,000	All	16", 19"	1280 x 1024
HP/Apollo 9000				
GRX	-	All	19"	1280 x 1024
CRX	$6,000	710	16", 19"	1024 x 768, 1280 x 1024
CRX-24	$13,000	720, 730, 750	19"	1280 x 1024
CRX-24Z	$21,000	720, 730, 750	19"	1280 x 1024
TurboVRX	$78,000	720, 730, 750	19"	1280 x 1024
IBM RS/6000				
Grayscale	-	All	16", 19", 23"	1280 x 1024
Color	$1,646	All	16", 19", 23"	1280 x 1024
High-performance 3D color	$1,039 - $1,779	All	16", 19", 23"	1280 x 1024
Gt3	$2,491	All	16", 19", 23"	1280 x 1024
Gt4	$8,043 - $12,000	All	16", 19", 23"	1280 x 1024
Gt4x	$8,328 - $12,100	All	16", 19", 23"	1280 x 1024
GTO	$14,430 - $21,830	All	16", 19", 23"	1280 x 1024
Intergraph				
GT+	standard with 2430	2430	19"	1184 x 884
GT II	standard with 6450	6450	19", 27"	1184 x 884, 1664 x 1248
Edge II+	standard with 6480	6480	19", 27"	1184 x 884, 1664 x 1248
SGI 4D				
Entry Level	-	Indigo, 30, 35, Crimson	16"	1024 x 768
XS	$6,000	Indigo, Crimson	16"	1280 x 1024
XS24	$8,000	Indigo, Crimson	16"	1280 x 1024
TG	$10,000	30, 35	19"	1280 x 1024
Elan	$15,000	Indigo, 30, 35, Crimson	19"	1280 x 1024
VGX	$30,000	Crimson, Power Series	19"	1280 x 1024
VGXT	$60,000	Crimson, Power Series	19"	1280 x 1024
Sun SPARCstation				
IPX	-	IPX	16", 19"	1152 x 900
2GX	$18,495 incl. CPU	2	16", 19"	1152 x 900
2GXplus	$22,495 incl. CPU	2	19"	1280 x 1024
2GS	$26,995 incl. CPU	2	19"	1152 x 900
2GT	$49,995 incl. CPU	2	21"	1280 x 1024

Table III.5: Graphics Features

Model	Base Color Planes / Double-Buffered, Max Color Planes / Double-Buffered	Z-Buffer Bits*	Hardware Support for Alpha-Blending	Hardware Support for Interactive Texture Mapping
DECstation				
HX	8 / 0	-	no	no
PXG+	16 / 8, 48 / 24	24 (opt.)	yes	no
PXG Turbo+	48 / 24	24	yes	no
HP/Apollo 9000				
GRX	8 / 0	-	no	no
CRX	16 / 8	-	no	no
CRX-24	24 / 12	-	no	no
CRX-24Z	24 / 12	24	yes	no
TurboVRX	24 / 12	24	yes	no
IBM RS/6000				
Grayscale	8 / 0	-	no	no
Color	8 / 8	-	no	no
High-performance 3D color	8 / 4, 24 / 12	24	no	no
Gt3	8 / 0	-	no	no
Gt4	16 / 8, 48 / 24	24	yes	no
Gt4x	16 / 8, 48 / 24	24	yes	no
GTO	16 / 8, 48 / 24	24	yes	no
Intergraph				
GT+	8 / 8	-	no	no
GT II	8 / 8	-	no	no
Edge II+	24 / 12	24	no	no
SGI 4D				
Entry Level	8 / 4, dithered to 24	24 (s/w)	no	no
XS	8 / 4, dithered to 24	24 (s/w; opt. h/w)	yes	no
XS24	24 / 12	24 (s/w; opt. h/w)	yes	no
TG	24 / 12	24	yes	no
Elan	24 / 12	24	yes	no
VGX	48 / 24	24	yes	yes
VGXT	48 / 24	24	yes	yes
Sun SPARCstation				
IPX	8 / 4	-	no	no
2GX	8 / 4	-	no	no
2GXplus	16 / 8	-	no	no
2GS	24 / 12	16	no	no
2GT	48 / 24	24	yes	no

(*Z-buffer: s/w - software & CPU-based; h/w - dedicated hardware-based)

Table III.6: Graphics Performance

Model	2D Vectors/sec (x 1,000)	3D Vectors/sec (x 1,000)	3D Gouraud Shaded, Z-buffered Triangles/sec (x 1,000)
DECstation			
HX	621	-	-
PXG+	345	401	70
PXG Turbo+	445	436	106
HP/Apollo 9000			
Grayscale	950	950	-
Color	950	950	-
CRX-24	1,150	1,150	-
CRX-24Z	1,150	1,150	165
TurboVRX	1,360	1,360	317
IBM RS/6000			
Grayscale	72	-	-
Color	125	-	-
High-performance 3D color	90	90	10
Gt3	650	-	-
Gt4	650	400	20
Gt4x	800	800	80
GTO	990	990	120
Intergraph			
GT+	500	300	-
GT II	800	500	-
Edge II+	600	500	5
SGI 4D			
Entry Level	200	200	18
XS	250	250	60
XS24	250	250	60
TG	219	219	40
Elan	1,000	1,000	225
VGX	1,000	1,000	780
VGXT	1,000	1,000	780
Sun SPARCstation			
IPX	480	310	-
2GX	480	310	-
2GXplus	480	310	-
2GS	190	150	-
2GT	500	500	100

PART IV. RECOMMENDED GENERAL OPTIONS

A. RECOMMENDED STANDARDS

OPERATING SYSTEMS

It is recommended that the SCFs be based on the Unix operating system. The portable Operating System Interface for Computer Environments (POSIX) has been established as a Federal standard for operating standards [Federal Information Processing Standards Publications (FIPS PUBS) 151-1]. Although POSIX compliance is not yet mandatory, NASA/MSFC expects vendors that provide hardware and software will either (1) deliver POSIX-compliant capability, or (2) commit to upgrading the delivered capability to POSIX compliance as soon as it is available, but no later than the date set in the future by the National Institute for Science and Technology for mandatory compliance [EADS II RFP No. 8-1-9-AI-00120].

GRAPHICS, WINDOWS, & VIDEO DISPLAY

It is strongly recommended that all SCF workstations have X-window capabilities. The X performance on various computers and X terminals can be highly variable, and it is important that the performance of X on a potential workstation be considered in light of the intended use of that workstation. In particular, the X performance of any workstation to be used for interactive 2D graphics should be very high. However, the X standard, by itself, is not adequate for workstations or terminals intended for 3D applications.

Both PEX and IRIS GL are highly likely to become future standards for 3D graphics. It is recommended that any potential vendor of SCF 3D workstations should have an expressed intent to support PEX in the future. However, PEX is more likely to become a standard within the 3D CAD market and not within highly interactive 3D graphics markets such as visualization. In contrast, the Working Group feels that the acceptance of IRIS GL as a near-future interactive 3D standard is assured, and that the tremendous benefits of IRIS GL for development and application software availability warrant support of the IRIS GL within the SCF and ESACF environment.

A large display screen of 1024x768 resolution has become the de-facto minimum standard for all workstation vendors due to the increased screen "real estate" required by a window-based user-interface. A resolution of 1280x1024 is recommended for visualization workstations. For color display, 8 bits is a minimum, providing 256 colors and can suffice for 2D imaging applications. As demand for more colors has increased over the years, the "true color" capability provided by 24-bit displays (16.7 million colors) has become affordable relative to 8 bit prices and is recommended as a standard for 3D visualization and multi-band or true color 2D imaging. Interactive 3D and image sequence applications benefit greatly from hardware support for double-buffering. Double-buffering provides an off-screen buffer for drawing to take place in. Upon drawing completion, the resultant image is quickly copied to the screen buffer for display. Without double-buffering, all the interim drawing steps occur in the screen buffer and are therefore visible; for sequences of images, it is then difficult to determine what constitutes the intended final image for each frame. Additionally, hardware support for z-buffering of 3D graphics makes the difference between slow and tedious interactivity and real-time interactivity and adds little to the system cost. For 3D applications, a z-buffer associates a "depth" value with each pixel on the screen and any later attempt to draw something at that pixel that has a larger depth value is not drawn, thus simulating how objects closer to an observer obscure more distant objects. The standard z-buffer depth is 24 bits, providing 16.7 million depth values, and this is therefore the recommended minimum z-buffer depth. For applications that composite multiple images or overlay multiple objects in 3D space, the ability to set a transparency value to "see behind" an object or to mix colors is important. Most high-end visualization workstations provide some support for this capability, termed "alpha blending." It is important that a high-end visualization workstation be able to render transparent objects without suffering a significant loss in graphics performance. Therefore the recommended minimum bitplane standard for a high-end visualization workstation is at least 80 bits: 24 bit double-buffered color, 8-bit overlay/underlay, and 24-bit z-buffer. Additional bitplanes for alpha blending and other pixel manipulations are advisable.

CONNECTIVITY

Presently, Ethernet, FDDI, and HIPPI are accepted standards in the marketplace for computer connectivity. Since HIPPI is 10 times faster than FDDI and 100 times faster than Ethernet, these standards provide an efficient means to match computer demands with the most appropriate communication speeds. Cost to support and install these is important, since HIPPI support is far more expensive than FDDI (several tens of thousands of dollars versus one to two thousand), and FDDI is several times the cost of Ethernet. For some locations currently using Ethernet, the cost of adding FDDI optical cable may outweigh the speed benefits. Additionally, HIPPI is a point-to-point standard and is normally used over short distances (~50 meters). By contrast, FDDI and Ethernet provide a "backbone" to which many users can connect.

TCP/IP is by far the most widely used communication protocol in the workstation and supercomputer markets and is used on Ethernet, FDDI, and HIPPI. In practice, TCP/IP utilizes roughly 30% to 70% of the available connectivity bandwidth depending on how well the computer manufacturer has optimized their networking hardware and software. DECnet and SNA protocols are still used at sites with an investment in older mainframe technology; thus, a capability to communicate externally using those protocols is important. ISO/OSI is still not supported as a viable standard by most manufacturers yet. It is important to keep in mind that, for all these standards, the communication protocol is decoupled from the connectivity hardware; thus, a change in protocol does not require an expensive "re-stringing" of the network cabling.

B. GENERAL OPTIONS

COMPUTATIONAL REQUIREMENTS OVERVIEW

All SCFs, excluding the LIS team, indicated a moderate to high need for CPU-intensive computation. High demands for CPU power is related to the requirements for computer modeling, as with the GWC and LAWS teams, or for numerical analysis, as involved in MIMR activities. Between these activities, there were equally-high demands for vector and scalar processing, as illustrated in the Models and Analysis sections under Software in Figure II.1.

All SCF investigators indicated the importance of visualization within their interviews and specified moderate to high ratings for several computationally demanding visualization techniques. Since visualization is most effective when it is highly interactive, the visualization software requirements as specified by the SCF PIs require hardware support for high-performance graphics capabilities. Specifically, contouring, 3D, and CFD visualization rely on high-performance line and polygon drawing. Three-dimensional image manipulation depends on real-time texture mapping capabilities, while 2D image processing and time sequenced images depend on high pixel copying rates. Any interactive use of transparency, including both isolevel and voxel rendering of 3D volume data, relies heavily on hardware alpha blending.

HIERARCHICAL COMPUTE ENVIRONMENT

Based on the requirements provided by the respective PIs, and considering the standards and current trends in the computer industry, the SCF Working Group developed a general option for meeting the total SCF computing requirements. As symbolically depicted in Figure IV.1, the general option is based on a hierarchical environment consisting of a combination of shared and individual resources. This hierarchy is composed of five major system types:

- **Type A** - Vector-Based Supercomputer for high-speed computation of vectorized numerical models
- **Type B** - Scalar-Based Multiprocessor Superworkstation for high-speed computation of scalar numerical models
- **Type C** - File server for high-capacity, high-speed disk storage
- **Type D** - High-end graphics and visualization workstations for advanced visualization applications which might also serve as local file server within one or more SCFs; these platforms will be designed to meet the individual SCFs' visualization needs beyond those solutions available or practical with the Type E personal workstations
- **Type E** - Personal workstations designed to meet the needs of the individual investigators

The Type A and B computers would be shared among all ESAD SCFs and are designed to meet the requirements for vector and scalar climate models, respectively. The Type C computer will provide high-speed file management for the Type A and B computers, as well as for general data storage for all SCFs. Type A, B, and C platforms would exist as part of the ESACF. Type D platforms would be shared among team members of one or more individual SCFs, providing higher performance visualization capabilities than might be possible using the Type E computers that serve as personal workstations for each SCF team member. For some SCFs, the option exists for substituting higher-end personal visualization (Type E) workstations in lieu of the shared Type D platform. This option should be further considered with regard to individual needs, preferences, and budget constraints of each SCF.

Such a hierarchical distribution of platforms allows overlapping needs to be met by higher-end shared systems (i.e., Type A and B platforms), while different configurations of Type D and E platforms allow for special needs of the SCF or individual investigators, respectively. Also Type A, B, and C computers are delegated to tasks associated with numerical modeling and analysis, leaving Type D and E workstations for visualization and other individual information management needs. This minimizes possible duplication of requirements associated with numerical modeling, including very-high CPU demands, high RAM capacity, high I/O bandwidths, and large-capacity, high-speed disks. Conversely, Type D workstations can be configured to best meet the needs of visualization, without concern as to whether they also meet needs for numerical modeling.

In order to best meet the total and individual SCF needs as defined in Part II, the following general specs are recommended for the five type platforms:

Type A - **Vector-Based Supercomputer**: These requirements will be met by the ESACF Cray XMP.

Type B - **Scalar-Based Multiprocessor Superworkstation**: Symmetric multiprocessor with very high CPU performance, parallelizing compilers for C and FORTRAN 77;

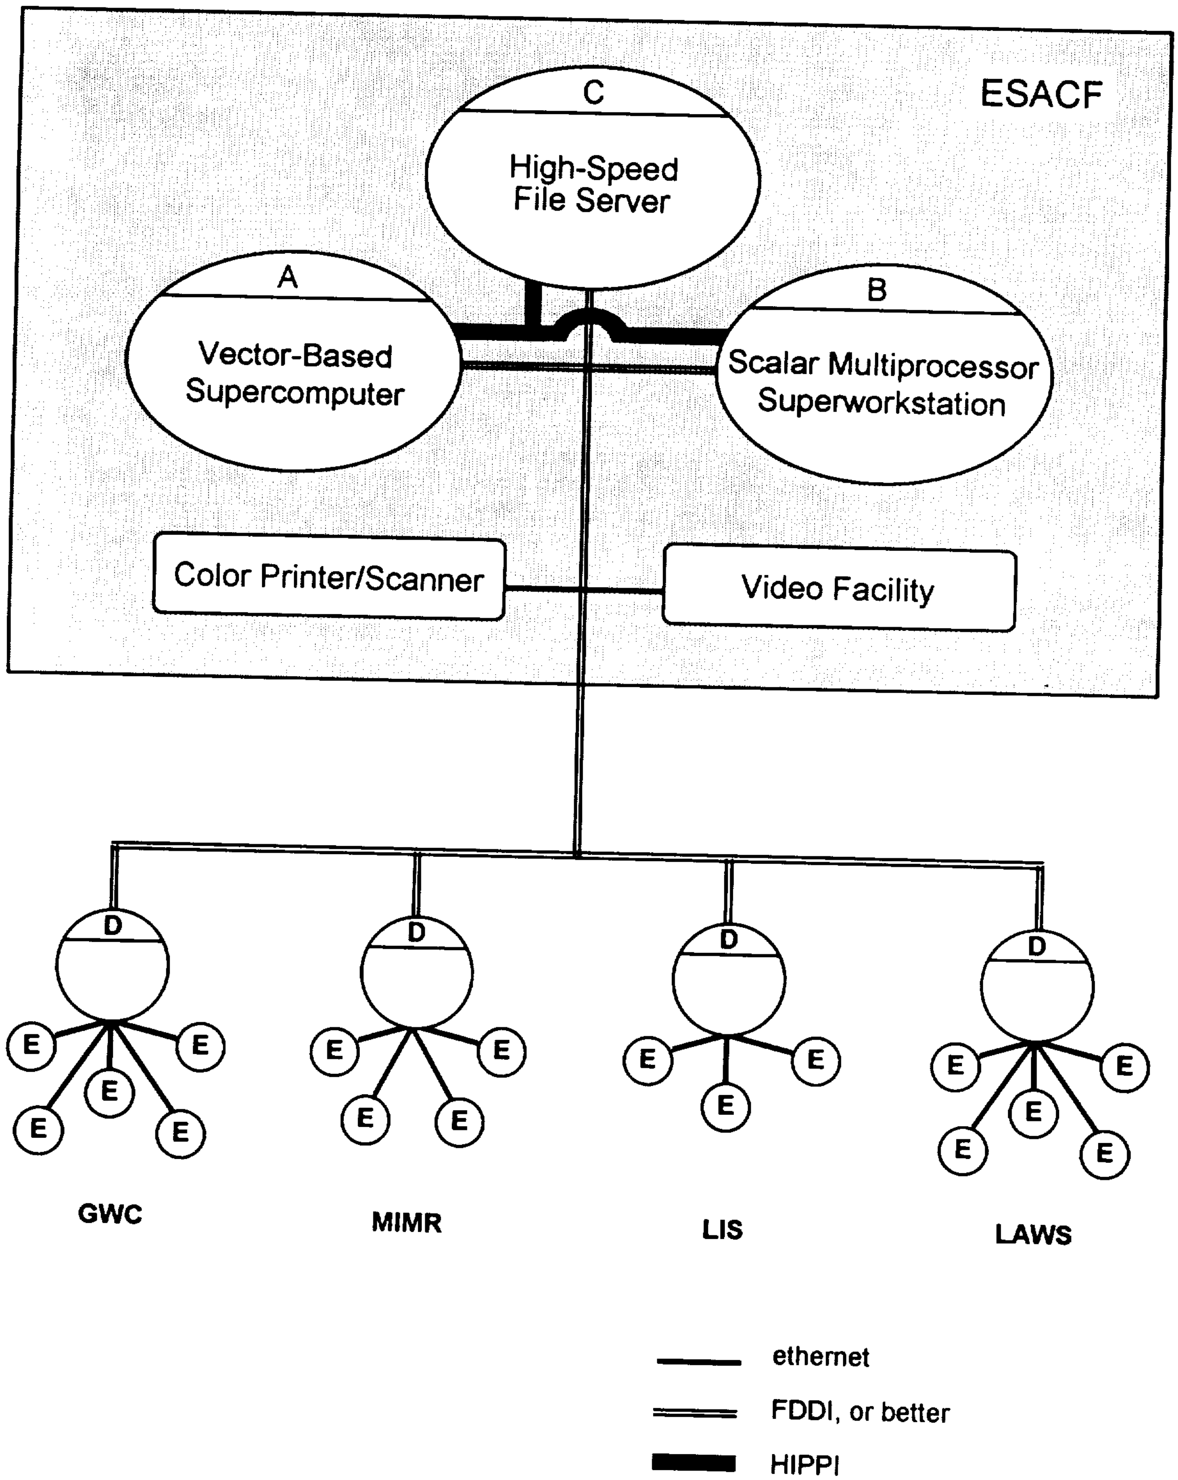

Figure IV.1. Schematic of a hierarchical option for meeting the total computing requirements of the MSFC SCFs. See text for complete description.

high-speed IPI-2 disks or better, large RAM capacity (256 MB minimum); very high I/O speeds; Ethernet, HIPPI and FDDI support.

Type C - **File Server**: High-speed IPI-2 disks or better; high disk storage capacity (up to 96 GB or more); high back-plane I/O speeds with CPU capable of handling all data requests and data flow from HIPPI and FDDI network, and the high-speed disks; Ethernet, HIPPI and FDDI support; available network and file server tools; support for Oracle or other high-end Database Management software; I/O devices: 8 mm Exabyte, CD-ROM, 9 track tape, 1/4" tape.

Type D - **High-End Visualization Workstations**: Moderate to high CPU performance, high pixel throughput and high-performance 3D rendering (1 Million+ 3D vectors/ sec; 150,000+ 3D polygons/sec), 1280x1024 minimum resolution with minimum 32-bit, double-buffered color (i.e., 64-bit color planes), 24-bit z-buffering; hardware support for interactive texture mapping and alpha blending; available visualization software and software development tools (e.g., GUI builders, IRIS GL, CASE tools, ANSI C, C++, FORTRAN 77); Ethernet and FDDI support; I/O devices: 8 mm Exabyte, CD-ROM;

Type E - **Personal Workstations**: Low- to high-performance visualization capabilities; moderate to high X window performance; 1024x768 minimum resolution with 8- to 24-bit color (double-buffered); Ethernet support;

SOFTWARE

The availability of relevant application software and ease of software development are critically important issues when considering hardware purchases. The most common and costly mistakes involve the purchase of high-performance computer hardware with few or no software applications or development tools. It is important that relevant application software exist for meeting immediate requirements, while development tools and support are important for custom applications, as well as for encouraging the future availability of relevant third party software.

Most SCF PIs indicated that McIDAS would be of moderate to low importance for meeting immediate requirements. Most also specified that McIDAS did not meet all of their visualization requirements and they anticipated that other visualization tools would become available which could better meet these needs. Still, it is important to consider the directions of McIDAS development, considering that it plays a major role in meeting present visualization needs. The Space Science and Engineering Center (SSEC) at the University of Wisconsin, the developer of McIDAS, has recently ported McIDAS to Unix-based workstations running X windows. Although OS/2 versions of McIDAS will still be supported for some time, the SSEC does not anticipate continued development under that operating system. Thus far, the SSEC has indicated future support for McIDAS running on IBM RISC 6000, SGI IRIS, and Sun platforms. However, if SSEC complies with Unix and X standards, porting of McIDAS to other platforms should not be prohibitive.

For scientific visualization, there are three types of software products available. The first type includes vertical application packages which are designed to meet the specific needs for a smaller group of scientists. For example, programs such as FAST, ELAS, VoxelView, and ARC/INFO are designed to meet the specific needs of computational fluid dynamics (CFD), image processing, volume rendering, and geographical information systems (GIS), respectively. The second type of visualization software includes application packages which provide a collection of tools for meeting a wide range of general visualization needs, such as image processing, 3D graphics, and volume rendering, but allow limited capabilities for user-customized development. These packages include Precision Visuals PV Wave, Wavefront's

Data Visualizer, Spyglass Transform, and Sun's SunVision software [Kriz, 1991; Mercurio, 1991; Burns, 1991a]. The third type of visualization software includes environments which allow custom design of visualization tools by the user or software developer. Within these environments, several functional modules can be linked by the user using a graphical interface, in order to rapidly create a custom application to meet the specific needs of the scientist. Having begun with SGI's Conman [Upson,1990], this concept has rapidly grown in popularity and is now incorporated in AVS, Inc.'s Application Visualization System (AVS) [Upson, 1990; Mazor, 1991], SGI's IRIS Explorer [Pope, 1991; Gorey, 1991; Mazor, 1991], IBM's Data Explorer [Mazor, 1991], and the Ohio Supercomputer Center's apE [VandeWettering, 1990; Upson, 1990]. Various visualization software packages are briefly described in Appendix C, while the availability of these packages on different platforms is documented in Table IV.1

In particular, the visualization environments offered by AVS and SGI's Explorer are very promising for meeting a wide range of visualization needs when flexibility and timeliness are important. In most cases, existing modules are available with little or no custom programming, allowing the end-user to quickly see their data. For SGI's Explorer, the capability exists to easily bundle the user's collection of modules into a stand-alone application complete with the point-and-click user-interface generated by Explorer. A major benefit of these modular environments is the ability to run modules over a distributed network of machines, while hiding the complexities of distributed processing from the end user. As a general distributed processing rule, it is clearly important to keep the multi-machine dataflow within the bandwidth limits of the network and the respective machines. Explorer allows computationally demanding modules that filter data to run on appropriately powered computers such as the Cray or a scalar multiprocessor, while the graphically demanding output can be interactively viewed on suitably configured visualization workstations without requiring large network traffic for rotating, zooming, clipping, probing, etc.

If it is necessary for the user to develop software, then it is important that one considers availabilty of tools for the particular hardware under consideration. In addition to possible languages requirements such as C, C++, Pascal, Ada, and FORTRAN, development tools might include graphics and image libraries, graphical user interface (GUI) builders, network toolkits, video toolkits, reference materials, Computer Aided Software Engineering (CASE) tools, and optimizing/parallelizing compilers. For development needs, it is also helpful if the vendor supplies example source code illustrating how to access particular features of the system. These development aids become increasingly important considering the added complexities introduced by windowing standards, multiprocessing, distributed processing, and the integration of GUIs with interactive 3D graphics. Many hardware manufacturers have invested considerable effort and finances to provide the necessary software and libraries through in-house development or third-party support, giving computer buyers far more to consider than just hardware capabilities.

A significant tool for high-end visualization development is the IRIS GL, which allows a developer to rapidly create interactive high-speed graphic applications within a windowing environment such as X. Since IRIS GL displays graphics in immediate mode, as opposed to the display-list mode found in PHIGS and PEX, it is far more applicable to the changing geometries associated with scientific visualization. This ease and speed of development using IRIS GL has resulted in the largest library of 3D software applications than any other 3D environment. The opening of the IRIS GL as a vendor-independent 3D standard will assuredly result in an increase in the number of third party applications. Because of its tremendous benefits, it is anticipated that future in-house development at ESAD will heavily employ IRIS GL. It is recommended that at least the Type D platforms support, and are optimized, for IRIS GL. Although IRIS GL is not required for meeting all SCF needs for Type E platforms, these platforms would benefit by supporting IRIS GL since distributed graphics processing and display of GL-based applications running over the network would be possible.

Table IV.1. Support of various visualization software products on different platforms.

	Apple Macintosh	Cray	DECstation 5000	HP/Apollo 9000	IBM RS/6000	PCs & compatibles	SGI	Sun SPARC
AGIS						●	●	
aPE	●	●	●	●	●		●	●
ARC/INFO			●	●	●	●	●	●
AVS		●	●	●	●		●	●
CORE Image Processing			●		●		●	●
ELAS							●	●
ERDAS			●	●	●	●	●	●
ER Mapper								●
FAST (NASA Ames)		●					●	
GENAMAP			●	●	●	●	●	●
Global Imaging				●				
GRASS	●	●	●	●	●	●	●	●
IBM Data Explorer					●			
IDL			●	●	●		●	●
IISS							●	
IMSL Statistics FORTRAN	●	●	●	●	●	●	●	●
IMSL Statistics C			●		●		●	●
IMSL Exponent Graphics			●	●	●		●	●
Intelligent Light Fieldview			○	●	●		●	●
Khoros Image Processing	○	●	●	●	●	○	●	●
Linkwinds (JPL)							●	
Mathematica	●		●	●	●	●	●	●
McIDAS		●			●	●	●	●
METPRO							●	
NCAR Graphics		●	●		●		●	●
NCSA Visualization Tools								
Image, XImage, etc.	●	●			●		●	●
Isolev Visualizer, X DataSlice, Height-Color Visualizer							●	
Contours, DataScope	●							
Paragon Vis. Workbench			●	○	●		○	●
PCI EASI/PACE			●	●	●	●	●	●
PV-Wave:								
Visual Data Analyzer			●	●	●		●	●
Point & Click			●	●	○		○	●
SciAN					●		●	
SGI Explorer		●					●	
SoftImage					●	●	●	●
Spyglass	●						●	
Sun's SunVision								●
Terramar Image Processing				●	●	●	●	●
UNIRAS agX/Toolmaster		●	●	●	●		●	●
Vertigo Data Slicer							●	
Vis5D					●		●	
VisAD					●		●	
Vital Images Voxelview	●				●		●	
Wavefront Data Visualizer		●	●	●	●		●	●
Whip (Hannaway & Assoc)				●	○		●	●

● Application supported ○ Future support announced

For Type E Platforms, one should consider the availability of information management tools such as word processing, spreadsheet and low-end database management software. At present, some software packages to meet these needs are available for Unix-based platforms. Other DOS and Macintosh applications can be supported on Unix-based through the use of emulators and special operating systems, including Insignia SoftPC and Quorum Equal [Picarille, 1992]. On PC platforms, SCO Desktop provides the Unix operating system while allowing access to more familiar DOS information management tools. With the growth in the number of low-end PC-priced Unix workstations and the current trend toward providing traditional PC applications in Unix, the availability of common information management tools on Unix platforms should grow quickly.

PERIPHERALS

Video Recording

While the SCF requirements indicate a need for video recording and video frame-grabbing, these needs are, in most cases, low to moderate, and are generally sporadic. SCF requirements for video recording and video frame-grabbing could be adequately met with an upgraded ESAD video facility, which would be accessible over the network and compatible with SCF workstations. This facility should have the capability to record sequences of images stored in files, as well as directly record screen output from a visualization workstation. The ability to record animation sequences of about 25 sec, or 700 images, is generally adequate for output of scientific animation; however, the ability to easily record longer sequences would, at times, be advantageous.

The needs for video frame-grabbing require a facility which is capable of accurately digitizing individual frames or frame sequences. The facility should be frame-accurate in both record and frame-grabbing modes and should allow triggering of frames based on video Time Code. The facility should support VHS, 3/4" U-matic, and possibly Betacam formats for recording, and should be capable of digitizing frames from 3/4" U-matic or Betacam video.

Color Hardcopy

Most SCFs indicated a high need for color printing of images, text, and plots. It is recommended that one or more high-quality color printers be configured into the ESACF and possibly within individual SCFs. The ESACF printer should be directly accessible and capable of receiving and printing files from all workstations on the network. It is recommended that for accurate reproduction of computer imagery, these printers be capable of accurately reproducing 16.7 million continous-tone colors at a minimum of 150 dpi (300 dpi preferred) on both paper and overhead transparency. At present, the best printers for high-quality photorealistic hardcopy employ the color sublimation technique or direct photography. For applications that require both the printing of text and color imagery, these printers should be capable of printing color PostScript files, as well as other image files generated by the SCF's visualization workstations. Video printers are generally not recommended, since they tend to have low resolution and poor color reproduction. Image scanning requirements are low to moderate and could be adequately satisfied with a high-quality ESAD image scanner which is easily accessible over the network.

DATA STORAGE

Most SCF PIs had difficulty estimating total data storage requirements. However, immediate requirements for SCF data storage can probably be met by a file server with around 40 GB of rapid access disk storage and an expandable capacity of up to 96 GB. In order to maximize the speed of disk I/O of very large data sets, it is important that these disks be IPI-2 (6 MB/sec) or better. In addition to requirements for rapid access data storage, future SCF requirements for online access to larger amounts of data might be satisfied by attaching on-line mass storage devices, such as the 1 TB capacity Epoch Renaissance Infinite Storage. Since the Type C file server should be capable of meeting all data storage requirements associated with numerical modeling, requirements for data storage on local SCF machines should primarily be driven by needs for data visualization and local manipulation.

CONNECTIVITY

At present, TCP/IP protocol is adequate for most SCF requirements. SCF systems should have standard capabilities for ftp and telnet. In order to allow transparent management of files residing on various machines, NFS capabilities should exist between the file server and all SCFs machines, as well as within each SCF environment. This capability is important for allowing rapid access to all relevant data regardless of location within the network, and should minimize the need for duplicate data files on various machines.

All SCF workstations should have Ethernet connectivity as a minimum. In order to adequately support the increased network traffic resulting from distributed processing, X protocol, and the transfer of large data files, it is highly recommended that the "backbone" network for the SCFs have at least the transfer rate of FDDI, and, as illustrated in Figure IV.1, that the Type A, B, and C servers be directly connected with FDDI-equivalent, or better, speeds. Further research should be conducted to evaluate whether FDDI or new network switching devices would best meet these needs. HIPPI connectivity is required between the Type A, B, and C computers, in order to support very high-speed transfer of large data files during compute-intensive numerical modeling.

All SCF workstations should be allowed direct connectivity to Type A, B, and C platforms, using telnet, rlogin, or ftp protocol. In addition, distributed processing capabilities with these platforms should be developed and encouraged, particularly with regard to real-time visualization or steering of numerical models.

PART V. SPECIFIC SHORT-TERM OPTIONS

A. HIGH-COST ALTERNATIVES

One goal of this investigation has been to investigate cost-effective solutions for meeting the SCF requirements. There are high-cost alternatives for meeting the SCF requirements, including supercomputers, minisupercomputers, and mainframes, as well as the IBM POWER Visualization System (PVS). Particularly with the availability of the ESACF Cray XMP, these alternatives do not, at present, add significant benefits to warrant their high costs.

MINISUPERCOMPUTER OPTIONS

There are several supercomputer and minisupercomputer vendors which could meet some of the SCF requirements, including Alliant, Convex, Intel, and NEC. These vendors offer peak vector-based double-precision performance from hundreds of MFLOPs to the GFLOP range, large RAM support (1-4 GB), high disk capacity, high bus speeds (up to 500 MB/sec), and second or third generation architectures. These systems are priced in the multi-million dollar range and require similarly expensive support. Although all these systems provide support for scalar and vector operations, the high processing speeds given are for highly vectorized code; scalar code will not achieve similar rates. As an example of price/performance, the Convex C3240 (the four processor version of their lowest-end system) offers a peak double-precision vector speed of 200 MFLOPS, a 64-bit architecture, a maximum of 2GB RAM, support for high-speed disks and networking, and well-established vectorizing and parallelizing compilers. A base C3240 with 256 MB RAM lists for just over $1,000,000 without disk storage or compilers. At the high-end, the Convex C3840 (offering just over twice the peak double-precision vector performance of the C3240) in a similar configuration lists for over $3,400,000. For roughly 10% of these costs, there are scalar-based multiprocessing superworkstation systems available (considered below) which offer just under half the double-precision speeds of the C3240, comparable disk and networking performance, similar compiler technology, and a far cheaper support cost.

These more expensive systems have an architecture that is well suited to greater expandability and higher utilization and control of those larger resources. However, these benefits come at a cost that is not proportional to the increased benefits when compared with what is available on present and soon-to-be-released superworkstations. Recent advances in 64-bit microprocessor design, standard high-speed bus architectures, standard multiprocessing support hardware, and increased RAM memory densities at a lower costs, will soon allow several workstation vendors to provide equally capable computers at a substantially lower price. Traditional minisupercomputer manufacturers that cannot compete will either completely change their architectures to become cost-competitive (leaving the early buyers with an expensive, minimally supported system) or leave the market.

Of the two 64-bit CPUs announced as of March 1992 (DEC's Alpha and Mips' R4000), only Mips' R4000 is in production, most notably in SGI's Crimson, the first single-processor machine to use the R4000. SGI, which has an aggressive history of product development based on standard technology, is known to be developing a multiprocessor machine based on the R4000 which should offer comparable performance and expandability to most minisupercomputers at a strongly competitive price. Similarly, DEC will offer a range of systems utilizing Alpha that should offer similar price/performance. Mips and DEC have been open in selling their CPUs to other computer manufacturers thus increasing the odds that major price/performance drops will occur within the next year in the minisupercomputer and high-end workstation markets.

Based on these considerations, minisupercomputers and mainframes are not considered cost-effective solutions to the scalar modeling and file server requirements of the SCFs at MSFC. Present and near-future superworkstations can meet the SCF requirements at a very significant reduction in purchase and maintainance costs, and without a significant decrease in performance. As the MPP architecture matures and becomes more stable in the next few years, these platforms should be considered as possible solutions for meeting supercomputing applications for all models.

IBM'S POWER VISUALIZATION SYSTEM (PVS)

In August 1991, IBM introduced the PVS server, which consists of 8 to 32 i860 RISC processors with 256 to 768MB of system memory, and a wide-band (256-bit) data path. The PVS can be connected to 8 disk array substems for a storage capacity of up to 170 GB. For visualization applications, the PVS renders complete images using IBM's Data Explorer visualization software, and transports these images through a Visualization Video Controller to a RS/6000 workstation for display. The PVS is capable of high-resolution rendering of very large data sets at over 30 frames per second. However, in order to view these images and to interact with the visualization process, these images must be compressed and transferred over HIPPI lines to a RS/6000 with a video controller board which is capable of decompressing the images in real-time.

Although the PVS is a unique approach to demanding visualization needs, the PVS is not, at the present, considered to be a cost-effective alternative for meeting the SCF requirements at MSFC. The reasons for this are as follows: (1) The cost of an appropriately configured PVS is around $3 million; (2) For interactive visualization needs, the PVS is capable of supporting only one user at a time; (3) The PVS relies almost entirely on the Data Explorer package for meeting programming needs and is not designed for general program development, compiling, or file access; (4) The requirement for HIPPI support to visualization workstations is costly; (5) The PVS solution must be accepted as a whole, requiring IBM Rs/6000's as visualization front-ends and the IBM Disk Array as a file server; and (6) The functionality that the PVS provides for visualization can be obtained by less costly solutions using, for many cases, a standalone high-end graphics workstation, or, in cases where more compute power is required, through distributed processing between such workstations and the ESACF Cray XMP.

B. COST-EFFECTIVE OPTIONS

TYPE A. VECTOR-BASED SUPERCOMPUTER

The requirements for the vector-based supercomputer can be satisfied with the ESACF Cray XMP. The SCFs do require that the Cray have direct access to large-capacity, high-speed data storage and that interactive access to the Cray be available from Type D and E workstations. The Cray should support distributed versions of various visualization and analysis software, including for example, AVS Inc.'s AVS, SGI's Explorer, the SGI ImageVision Library (IL), NASA Ames' Plot3D, and the IMSL statistical package.

TYPE B. SCALAR-BASED MULTIPROCESSOR SUPERWORKSTATION

As presented in Part IV, the requirements for the Type B platform are for a symmetric multiprocessor with parallelizing compilers for C and FORTRAN 77; high-speed IPI-2 disks (6 MB/sec capability) or better, large RAM capacity (256 MB minimum); very high I/O speeds, including Ethernet, FDDI, and HIPPI support.

The two platforms which, from a hardware perspective, meet or come near the Type B requirements are the DEC VAX 6000-600 and the SGI 480S.

Based on the following considerations, the SGI 480S is the recommended platform for meeting the needs of the Type B computer:

- The VAX 6000-600 runs with the VMS operating system and does not support the current version of DEC Ultrix (i.e., Unix).

- The VAX 6000-600 does not currently support a HIPPI network connection. The SGI does.

- The SGI 480S is a 236 MIP computer with a SPECmark rating of 166 and a double-precision MFLOPS rating of 70. The VAX 6000-600 has a SPECthruput rating of 196.9, which is a similar specification for parallel processing platforms.

- In late 1992 or early 1993, the SGI 480S, currently utilizing the MIPS R3000 processor, will be field upgradable to the MIPS R4000 CPU, increasing its performance by at least three times, and converting it to a 64-bit architecture. The VAX 6000-600 is a 32-bit architecture.

- The list price for an SGI 480S with 256 MB RAM and 3 GB disk space is half the price of a VAX 6000-600 with 256 MB RAM and 3 GB disk space. Adding in the cost for the field upgrade of the SGI to the R4000 processors brings the list price to roughly the VAX price.

- The SGI 480S has very fast I/O for both small and large-sized data files, as reported by Montgomery [1991]: "Even the fastest I/O subsystems we've evaluated (those of the Hewlett-Packard Apollo's 9000 Model 720 and Sun Microsystems' Sparcserver 490) are no match for the striped IPI-2 disks. ... One of the most remarkable aspects of the Silicon Graphics I/O subsystem is the flatness of its performance lines. Whether we worked on a 1MB file or a 200 MB file, we generally saw the same incredibly high throughput." Being a new machine, the DEC VAX 6000-600 has no disk I/O ratings published. Both DEC and SGI platforms support disk striping, effectively parallelizing disk throughput.

- AVS, SGI's Explorer, and IRIS GL are available for the SGI 480S providing easy scientific visualization use within a distributed environment; Explorer and IRIS GL support are standard with no additional cost.

- If required to meet future needs, the SGI server can be upgraded with high-performance graphics capabilities. The VAX 6000-600 has no graphics expansion capability.

- Both platforms offer capable network and disk management tools. The SGI system management tools, leveraging off of years of SGI expertise in the graphics market, are primarily graphical-based applications providing an intuitive and easy-to-use interface. The DEC VAX systems are still primarily text-based systems and would have only recently, if yet, had graphical management tools ported to them for use via an X-window terminal.

TYPE C. FILE SERVER

The requirements for the Type C file server include high-speed IPI-2 disks or better; high disk storage capacity (up to 96 GB or more); high back-plane I/O speeds with CPU capable of handling all data requests and data flow from HIPPI and FDDI network and the high-speed disks; Ethernet, HIPPI and FDDI support; available network and file server tools; support for Oracle or other high-end Database Management software; I/O devices: 8 mm Exabyte, CD-ROM, 9 track tape, 1/4" tape. Platforms potentially capable of meeting these needs again include the VAX 6000-600 and the SGI 480S. For the same reasons presented for the Type B platform, the SGI 480S is the recommended platform for meeting the needs of the Type C file server. In addition, a suite of file server and network management tools are available for the SGI servers, allowing easy maintenance and backup of large volumes of data storage.

TYPE D. HIGH-END VISUALIZATION WORKSTATION

The recommended requirements for the Type D visualization workstations include moderate to high CPU performance, high pixel throughput and high-performance 3D rendering (1 Million+ 3D vectors/sec; 150,000+ 3D polygons/sec), 1280x1024 minimum resolution with minimum 24-bit, double-buffered color (i.e., 48-bit color planes), 24-bit z-buffering; hardware support for interactive texture mapping and alpha blending; available visualization software and software development tools (e.g., GUI builders, IRIS GL, CASE tools, ANSI C, C++, FORTRAN 77); Ethernet and FDDI support; video I/O support; I/O devices: 8 mm Exabyte, CD-ROM.

The potential platform candidates include the DECstation 5000 model 240PXG Turbo+, the HP/Apollo series 9000 model 730 TurboVRX, the IBM RS/6000 model 560 GTO, the SGI 4D Crimson VGX or Power Series VGX, and the Sun SPARCstation 2GT. The CPU and graphics performance of these candidates is summarized in Table V.1. The list price configuration includes 64 MB RAM memory and roughly 800 MB disk memory. The prices are rounded up to the nearest $5,000.

Based on the following reasons, the SGI Crimson VGX and the SGI Power Series (i.e., 300 and 400 series) VGX are highly recommended for the Type D platform:

- The CPU and graphics performance of the Sun and DEC platforms are significantly lower than the other platforms and are unacceptable for the Type D platform.

- The HP/Apollo platform does not provide 24-bit, double-buffered color (i.e., 16.7 million colors), but only 12-bit, double-buffered color, or 4,096 colors. This would greatly hinder the rendering of volumetric data and shaded surfaces by producing unwanted banding and artifacts, and is unacceptable for the Type D visualization platform.

- The SGI platform offers comparable graphics performance for 2D and 3D vectors, while greatly outperforming the other platforms in 3D shaded triangular surfaces, which are the basis for all 3D surfaces:

	3D triangles per second	3D vectors per second
SGI VGX	780 K	1,000 K
HP/Apollo TurboVRX	317 K	1,360 K
IBM GTO	120 K	990 K

- Only the SGI platform offers specific hardware support for texture mapping, important for mapping images onto 2D and 3D surfaces (e.g., satellite imagery mapped to elevation data).

- The SGI Crimson offers comparable single CPU performance to the other platforms, while only the SGI Power Series allows expandable higher performance through parallel processing.

- The MIPS R4000 CPU is the only true 64-bit architecture, providing 64-bit registers, 64-bit internal data and instruction paths, and a flat 64-bit memory layout which is important for large (i.e., >4 GB) models and datasets.

Table V.1: CPU and Graphics Performance for Type D Platform Candidates

	DEC	HP/Apollo	IBM	SGI	Sun
CPU Performance					
MIPS	43	76	~80	~100 - 286	29
MFLOPS	6	22	31	18 - 70	4
SPECmarks	32	`77	89	70 - 166	25
Graphics Performance					
2D Vectors/sec (x 1,000)	445	1,360	990	1,000	500
3D Vectors/sec (x 1,000)	436	1,360	990	1,000	500
3D Gouraud Shaded, Z-Buffered Triangles/sec (x 1,000)	106	317	120	780	100
Double-Buffered 24-bit Color	yes	no	yes	yes	yes
24-bit Z-Buffer	yes	yes	yes	yes	yes
Hardware Support for Alpha-Blending	yes	yes	yes	yes	yes
Hardware Support for Interactive Texture-Mapping	no	no	no	yes	no
Graphics Programming Interface	X, PHIGS	X, PHIGS	X, GL, PHIGS	X, GL, PHIGS	X, PHIGS
List Price	$75,000	$90,000	$90,000	$90,000 - $220,00	$70,000

- SGI has the longest history of providing high-performance graphics on workstations, has consistently provided leading edge graphics technology and performance, and has traditionally offered an abundance of software development tools and sample code to aid graphics programming.

- SGI supports and is optimized for the IRIS GL allowing faster development of visualization applications, and thus has access to the largest available library of 3D applications. In addition, much on-going in-house development, as well as outside development of meteorological and visualization applications, relies on the IRIS GL.

- The SGI IRIS Explorer visualization application is free with every system, providing a quick and intuitive means for end-users to visualize their data without having to write (or have written for them) custom programs. Similar visualization environments, such as AVS, cost around $5,000 for the SGI and other platforms.

- As seen in Table IV.1, SGI provides the most complete platform for visualization software, particularly for interactive volumetric rendering.

- As noted above for the Type B and C platforms, SGI has excellent disk I/O performance compared with the other Unix platforms.

- The SGI Power Series VGX platforms are available through the NASA Ames WKSII contract at a discounted price and quicker delivery.

TYPE E. PERSONAL WORKSTATIONS

Type E personal workstations should be configured around the specific needs of the SCF and the individual user. Therefore, specific recommendations for these platforms are minimal. The Type E workstations could range in price from less than $10K to around $30K, with the higher-end personal workstations being a consideration for SCFs that might choose to forego Type D platforms. Several candidate platforms are described below:

Digital Equipment Corporation (DEC) DECstation 5000

The DECstation 5000 family uses a RISC-based MIPS R3000 CPU and has a SPECmark range of 13.9 to 32.4. Base prices run from $3,995 for a diskless model 20 with 8 MB RAM to $13,495 for a model 240 with 16 MB RAM and a 336 MB disk. Disk-based systems run with the Ultrix operating system, DEC's version of Unix. Graphics options range from the low-end 2D-only HX option for $1,995 to the high-end 3D PXG Turbo+ option with 96 bitplanes for $25,000. The graphics are based on a DEC-modified version of PHIGS. Expandability options include Turbochannel I/O slots and SCSI. All models support Ethernet connections.

IBM RS/6000

The RS/6000 family uses a proprietary RISC-based RS/6000 CPU and has a SPECmark range of 25.9 to 89.3. Base prices run from $6,500 for a diskless model with a grayscale monitor to $21,500 for a model 520 with 8-bit color. The RS/6000 line has good CPU price/performance, particularly for code compiled with the newer optimized FORTRAN compiler. Software is binary compatible within the RS/6000 family, but not with any higher-end platforms. For graphics support, IBM purchased SGI's Personal IRIS geometry engine (GE) hardware, and licensed and slightly modified SGI's IRIS GL. With slight modification, this allowed easy porting of a large number of 3D applications from the SGI to the RS/6000, and surely aided the success of this platform. Considering that the RS/6000's graphics performance relies on the low-end Personal IRIS GE, there is some question as to how IBM will choose to upgrade the graphics performance on this platform. Outside of buying higher-end boards from SGI, dropping support of the GL, or introducing a completely new architecture, IBM will probably rely on future implementations of the IRIS GL on Intel i860 co-processor boards.

HP/Apollo 9000 Series (425 & 700 family)

The HP/Apollo 9000 model 425 series uses an older CISC-based Motorola 68040 CPU and has a SPECmark of 11.5. Base prices run from $15,490 for a 425e to $33,490 for a 425s. In contrast, the HP/Apollo 9000 model 700 family uses a RISC-based PA-RISC CPU and has a SPECmark range of 34.0 to 77.5. Base prices run from $4,990 for a diskless, grayscale model 705 with lower-speed CPU to $43,190 for the model 750 with color capabilities and a higher-speed CPU. The HP/Apollo 9000 model 700 family currently has the best CPU price/performance and the fastest mid-level (8- to 24-bit single-buffered) color capable X-windows and PHIGS-based vector graphics performance of any workstation vendor.

Silicon Graphics Inc. (SGI) Indigo and Personal IRIS

The SGI Indigo and Personal IRIS use a RISC-based MIPS R3000 CPU and have a SPECmark range of 26.0 to 31.0. Base prices run from $7,995 for a diskless 8-bit color Indigo to $18,000 for a Personal IRIS Model 35 with 8-bit color. Both platforms support upgradable graphics options ranging from 8-bit, single-buffered base and XS graphics to the high-performance 24-bit, single-buffered Elan graphics. IRIS GL is standard on all platforms, providing access to the largest base of 3D applications, as well as providing ease of graphics development and the ability to run GL applications in a distributed environment. The windowing environment, with the release of IRIX 4.0, is X/Motif. Software is binary compatible with all SGI platforms, including the high-performance multiprocessor servers. SGI Explorer visualization software is free on all SGI platforms; AVS is also supported through AVS, Inc. for roughly $5,000. The CPU in both systems will be upgradable in the near future to the RISC-based MIPS R4000 with a three times performance increase. The Indigo line comes standard with audio processing of CD quality sound; this is optional on the Personal IRIS family.

Sun Microsystems SPARCstation

The Sun SPARCstation has a RISC-based SPARC CPU and a SPECmark range of 13.5 to 25.0. Base prices run from $4,995 for a diskless monochrome ELC to $18,000 for the SPARCstation 2 with monochrome graphics. The ELC model is limited to black and white capabilities, while the IPC and IPX provide 8-bit color. The SPARCstation 2 graphics performance is upgradable from the 8-bit color GX to the higher-performance 48-bit color GT. The SPARC-based platforms have binary compatibility throughout the platform, and boast a large installed base of platforms and an abundance of 2D, 8-bit color application software, much of it developed for the science community. The GS and GX models are available in small "pizza-box" configuration allowing for easy desktop use but limited expandability; the GT model has an additional tower for under the desk. The windowing environment on the Sun platforms is X/Open Look. Sun supports a version of PHIGS, as well as their XGL graphics library which is based on PHIGS.

PC Compatibles

The personal computer has gained in capabilities and power primarily through the introduction of higher powered CPUs, increased software functionality, and the availability of add-on boards for memory management, communications, and graphics. However, in order to preserve backwards compatibility, major architectural changes which might significantly improve performance and functionality have been limited. These limits have not been present for the newer low-end RISC workstations which now offer a better price/performance than equivalently configured high-end PCs. However, for traditional MIS needs, or for users with a significant investment in a PC system, there are options for hardware and software upgrades and enhancements.

DOS, OS/2, and Unix are supported on PC systems. X-windows is supported under the Unix operating system as well as under the Windows-DOS environment. Only the highest-end PCs can handle the interactive data flow required for running the operating system, the X-windows client software, and any graphical-based or heavy text-based X applications. Except for very simple X-terminal sessions, a graphics accelerator is essential for running X-clients on a PC.

Most PCs currently rely on VGA graphics for display. The capabilities of extended VGA allow higher resolutions of 1024x768 and 256 colors. However, "VGA is inherently slow, so in applications where you're doing a lot of bit-mapped color stuff, a graphics accelerator can really help" [Farris, 1992]. Options for graphics accelerators include IBM's 8514/a, at the lower end, and, at the higher end, IBM's XGA and the Texas Instruments' Graphics Architecture (TIGA) utilizing TI's 34010 or 34020 graphics coprocessors. These boards offer resolutions from 1024x768 to 1280x1024 resolutions and are priced from $600 to $3,000. For the TIGA, the X server software is resident on the board itself, freeing the PC CPU for compute-intensive tasks. Recently, IBM announced that it was dropping support for the 8514/a card to concentrate on XGA. Since 8514/a is a subset of XGA, 8514/a-based applications should run without change under XGA.

With the opening of IRIS GL as a vendor-independent 3D graphics standard and the decision by Intel to support the GL on its i860 processor, low-cost PC coprocessor cards based on the Intel i860 chip running standard GL are expected by mid-1992. This should serve to rapidly create a viable market for low-end 3D visualization packages for PCs since GL-based applications have existed for years on SGI and IBM workstations and can be easily recompiled for PCs running under Unix.

However, even with add-on coprocessor boards, memory expansion, and Unix and X-windows functionality, the architectural bandwidth limitations found on most PCs will curb the overall effectiveness of rapid graphical interaction with medium to large datasets. In addition, for new purchases, the cost of configuring a high-end PC with Ethernet connectivity, higher memory, the Unix operating system, and graphics coprocessor boards will increase the cost of the PC to that of a lower-end RISC-based graphics workstation with much greater performance than the PC. At present, the main advantages of PCs is the lower entry-level price and the abundance of DOS and Windows-based information management software. However, the first advantage is only applicaple to requirements that don't demand much data bandwidth or graphics performance, such as word processing and spreadsheet applications. The second advantage is rapidly becoming less of an issue with the migration of DOS and Windows-NT to RISC-based architectures running Unix.

PART VI. CONCLUSIONS

CONCLUSIONS

The purpose of this document is three-fold:

- Define the requirements for the five EOS SCFs at NASA/MSFC ESAD
- Recommend general options for the SCF computing environment at NASA/MSFC ESAD
- Recommend specific options for meeting the computing requirements of these facilities

The requirements of the MSFC SCFs have been discussed in Part II and summarized in a requirements matrix in Figure II.1. These requirements are summarized as follows:

- **CPU Intensive Computation**: moderate to high needs driven by both vector and scalar modeling
- **Visualization**: moderate to high needs for contouring, image processing, time-sequenced data, 3D, and CFD visualization
- **McIDAS**: low to moderate needs initially; less or no need as other tools become available;
- **Analysis tools**: varying needs from none to high
- **Connectivity**: networking within MSFC important; no special external needs; TCP/IP only protocol required
- **Video**: moderate sporadic needs for video recording and frame-grabbing
- **Color printing/scanning**: high needs for frequent high-quality color printing; low to moderate sporadic needs for image scanning
- **Data storage**: need for rapid access data storage, both locally and within the division; estimated: 1-4 GB required on local workstations, 10-200 GB required for rapid access

Part III discussed several important considerations, including the requirements of the EOSDIS, the current directions of the computer industry, presently accepted industry standards, and the configurations and specification ratings of the workstations being offered by major vendors. These considerations are summarized below:

REQUIREMENTS OF THE EOSDIS

- There are very few requirements imposed on the SCFs by the EOSDIS at this moment;

INDUSTRY DIRECTIONS

- **Standardization**: trend toward standardization of software and hardware, and use of off-the-shelf hardware components

- **CPU Technology**: trend from CISC to RISC-based CPUs, with an order of magnitude increase of performance every 5-7 years (see Figure III.1)

- **Parallel Processing and Distributed Processing**: increased use of symmetric parallel processing to increase CPU performance and to handle multi-tasking and multi-users; increased emphasis on distributed processing to distribute an application's workload over different computing platforms

- **Importance of Software**: Availability of software and software development tools of high importance; hardware vendors investing considerable effort in developing tools and attracting third party software developers

- **Merging of Personal Computer and Workstation Environments**: significant effort underway to allow cohabitation of DOS and Unix on all platforms, and to allow interchange of software between DOS and Unix operating systems

- **Merging of Workstation and Supercomputer Technology**: increase of workstation power to supercomputer levels; trend toward replacing vector-based supercomputers with Massively Parallel Processing (MPP) technology

INDUSTRY STANDARDS

- **Operating Sytems**: Unix AT&T System V with Berkley Extensions & POSIX Compliance

- **Windowing Environments**: X/Motif or X/Open Look

- **Text**: X fonts and Display PostScript (DPS) for screen; PostScript (PS) for hardcopy

- **3D graphics**: PEX/PHIGS for static geometries (e.g., CAD); IRIS GL for interactive and dynamic geometries (e.g., visualization)

- **Networking**: TCP/IP protocol (including telnet, ftp, NFS, and NQS) with ethernet, FDDI, and Hippi connections

MARKET SURVEY

- **Vendors Considered**: AST, Apple, Compaq, DEC, HP/Apollo, IBM, Intergraph, NeXT, Silicon Graphics (SGI), and Sun

- **SPEC Tables Available**: CPU performance, System configurations, Graphics Options, Graphics Features, Graphics Performance

- **Highest CPU Performance**:

	SPECmarks
Multiple Processor:	
DEC VAX 6000/660	197
SGI 480	166
Single Processor:	
IBM RS/6000/560	89
HP/Apollo 9000/750	78
SGI Crimson	70

- **Highest Graphics Performance**:

	3D shaded triangles per second	3D vectors per second
SGI VGX	780 K	1,000 K
HP/Apollo TurboVRX	317 K	1,360 K
IBM GTO	120 K	990 K

In Part IV, general options for meeting the MSFC SCF requirements were presented. These included recommended standards and, as presented in Figure IV.1, a hierarchical cost-effective option for meeting the total SCF requirements. The hierachical compute environment consists of various shared and individual workstations as described below:

HIERARCHICAL COMPUTE ENVIRONMENT

- **Type A** - Vector-Based Supercomputer for high-speed computation of vectorized numerical models
- **Type B** - Scalar-Based Multiprocessor Superworkstation for high-speed computation of scalar numerical models
- **Type C** - Shared File Server for high-capacity, high-speed disk storage
- **Type D** - High-end visualization workstations for advanced visualization applications beyond the capabilities of the Type E personal workstations; might also serve as a local file server within one or more SCFs
- **Type E** - Personal workstations designed to meet the needs of the individual investigators

PERIPHERALS

- Video Recording/Frame Grabbing: Requirements could be met by a high-quality video facility within ESACF and accessible over the network
- Color Printing/Scanning: A high-quality color printing and scanning capability should be accessible by network within the ESACF; Individual SCFs may require local color printing capabilities, as well

CONNECTIVITY

- All SCF platforms should have ethernet connectivity with TCP/IP protocol; all should have telnet and ftp capabilities, while most should have NFS capabilities, as well
- HIPPI connections should exist between Type A, B, and C platforms for high-speed I/O of large data files
- An FDDI network should serve as the backbone for the SCF and ESACF facilities, and should have direct connectivity into type A, B, C, and possibly D, platforms; This is required to handle increased network demands arising from transfers of large data files, increased distributed processing activity, and increased use of X window protocol

SOFTWARE

- Efforts should be expended in parallelizing and optimizing scalar numerical models
- AVS and SGI's Explorer are important visualization tools; SCFs should look seriously at the purchase of these packages for the visualization platforms and for the Cray
- Off-the-shelf visualization software exists for meeting many of the needs of the SCFs, and the availability of these tools should be examined when considering candidate hardware platforms (see Table IV.1 for listing of visualization software on various platforms)
- General and specific software will need to be developed within the division in order to meet many needs of the SCFs, and much importance should be placed on the availability of software development tools on candidate platforms

- The ability to run DOS-type Information Management software (e.g., word processors, spreadsheets, etc.) on Unix workstations is rapidly becoming less of an issue

Part V discusses specific short-term options for meeting the immediate requirements of the MSFC SCFs. In particular, the following platforms are recommended as candidates for meeting the types above; detailed support for these recommendations are provided within the document:

TYPE A: VECTOR-BASED SUPERCOMPUTER

- The ESACF Cray XMP will meet this requirement

TYPE B: SCALAR-BASED MULTIPROCESSOR

- Candidates: DEC VAX 6000/600 and the SGI 480S
- Recommended: SGI 480S

TYPE C. FILE SERVER

- Candidates: DEC VAX 6000/600 and SGI 480S
- Recommended: SGI 480S with a mix of IPI-2 and SCSI disks

TYPE D. HIGH-END VISUALIZATION WORKSTATIONS

- Candidates:
  - DECstation 500 model 240PXG Turbo+
  - HP/Apollo 9000/750 TurboVRX
  - IBM RS/6000 model 560 GTO
  - SGI 4D CrimsonVGX or SGI Power Series VGX
  - Sun SPARCstation 2GT
- Recommended: SGI Crimson VGX or SGI 400 series VGX

TYPE E. PERSONAL WORKSTATIONS

- Candidates:
  - DEC DECstation 5000 family
  - IBM RS/6000 family
  - HP/Apollo 9000 Series (425 & 700 family)
  - SGI Indigo and Personal IRIS families
  - Sun SPARCstation family
  - PCs and compatibles
- Recommendation: Type E workstations should be configured around the specific needs of the SCF and individual user

PART VII. REFERENCES

REFERENCES

Ballou, M-C. (1991), *ACE formation rocks desktop system architecture alliances,* **Digital Review**, June 17, 1991, pg. 1.

Barkes, K.G. (1991), *Clear sailing for PostScript,* **DEC Professional, Vol. 10**, No. 9, September 1991, pg. 64.

Baum, D.R. and J.M. Winget (1990), *Parallel graphics applications,* **Unix Review, Vol. 8**, No. 8, September 1990, pg. 50.

Bucken, M. (1991), *Sun targets Windows NT,* **Software Magazine, Vol. 11**, no. 12, October 1991, pg. 40.

Burgard, M. (1991), *Who's winning the GUI race,* **UnixWorld**, August 1991, pg. 80.

Burns, M. (1991a), *Scientific visualization wars,* **Supercomputing Review**, October 1991, pg. 31.

Burns, M. (1991b), *Sun's new operating system straddles PC/workstation border,* **Supercomputing Review**, October 1991, pg. 13.

Bruno, L. (1991), *Six more firms license SGI's GL: Vendor's addition to help promote GL as a standard for 3-D graphics,* **UNIX Today!**, November 25, 1991, pg. 3.

Coopee, T.C. (1992), *Speeding into the RISC power zone,* **Digital Review,** January 6, 1992, pg. 23.

Corcoran, C. (1991a), *MIPS launches R4000 RISC chip,* **InfoWorld, Vol. 13**, Issue 40, October 7, 1991, pg. 25.

Corcoran, C. (1991b), *SGI licenses IRIS library,* **InfoWorld, Vol. 13**, Issue 38, September 23, 1991, pg. 21.

Corcoran, C. (1991c), *Sparcservers await Solaris 2.0 to go symmetrical,* **InfoWorld, Vol. 13**, Issue 39, September 30, 1991, pg. 6.

Du Pont Pixel Systems (1991), *Insight: GL on SPARC,* **Computer Graphics World, Vol. 14**, No. 10, October 1991, pg. 48+.

Farris, R. (1992), *Maximizing your PC for graphics,* **UnixWorld, Vol. IX,** No. 2, February 1992, pg. 61.

Flack, D. (ed.) (1991), *ACE reshuffled into DEC,* **UnixWorld, Vol. VIII,** No. 10, October 1991, pg. 15.

Forbes, J. (1991a), *Multiple processors: are two heads better than one?,* **Personal Workstations, Vol. 3,** No. 2, February 1991, pg. 32.

Forbes, J. (1991b), *The Silicon Graphics approach,* **Personal Workstations, Vol. 3**, no. 2, February 1991, pg. 38.

Gill, P.J. and J. Hammett (1992). *Fast-Forward to UniForum,* **UNIX Today**, January 6, 1992, pg. 5.

Gorey, K. (1991), *IRIS Explorer: The example/the product,* **IRIS Universe**, No. 17, pg. 34.

Green, P.E. (1991), *The future of fiber-optic computer networks*, **Computer, Vol. 24**, No. 9, September 1991, pp. 78-87.

Grygo, G. (1991a), *DEC pledges to market NT operating system; MIPS launches R4000*, **Digital Review**, October 7, 1991, pg. 1.

Grygo, G. (1991b), *DEC backs Iris library in graphics workstation war*, **Digital Review**, September 23, 1991, pg. 1.

Grzanka, L. (1991). *Silicon Graphics opens GL licensing*, **Supercomputer Review**, November 1991, pg. 50.

Hansen, A. (1991), *The desktop of the future arrives*, **UnixWorld**, May 1990, pg. 97.

Hasler, A.F., K. Palaniappan, and D. Chesters (1992). *Visualization of multispectral and multisource data using an interactive image spreadsheet (IISS)*, Preprint Volume of the **8th International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology**, January 1992, Atlanta, pg. 85.

Hayes, F. (1991), *X rated in 3-D*, **Unixworld, Vol. VIII**, No. 10, October 1991, pg. 54.

Hennessy, J.L. and N.P. Joupp (1991), *Computer technology and architecture: An evolving interaction*, **Computer, Vol. 24**, No. 9, September 1991, pp. 18-29.

Hess, M. (1990), *Flexing PEX*, **SunTech Journal, Vol. 3**, No. 1, Winter 1990, pg. 72.

Hibbard, W., C.R. Dyer, and B. Paul (1992), *A development environment for data analysis algorithms*, Preprint Volume of the **8th International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology**, January 1992, Atlanta, pg. 101.

Hibbard, W. and B. Paul (1992), *Progress with the VIS-5D System*, Preprint Volume of the **8th International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology**, January 1992, Atlanta.

Hubley, M. (1991), *SPARC International: guiding the market*, **Personal Workstation, Vol. 3**, No. 6, June 1991, pg. 77.

Iverson, W. (1991), *Allies rally around new MIPS chip to confront Sun, IBM, and HP; Also support Microsoft OS/2 3.0 and Intel x86 series microprocessors*, **Supercomputing Review**, May 1991, pg. 11.

Jenkins, A. (1991a), *IBM, Apple detail roles in joint desktop system plan*, **Digital Review**, October 7, 1991, pg. 1.

Jenkins, A. (1991b), *PEX takes X into the third dimension*, **Digital Review**, April 22, 1991.

Johnson, S. (1991), *XGL dynamic acceleration*, **SunWorld**, October 1991, pg. 71.

Johnston, S.J. (1991), *NT to run 16-bit software on ACE systems*, **InfoWorld, Vol. 13**, Issue 38, September 23, 1991, pg. 8.

Kriz, R.D. (1991), *PV-Wave Point & Click*, **Pixel**, July/Aug. 1991, pg. 28.

Leibowitz, M.R. (1991), *The new generation of RISC*, **UnixWorld**, August 1991, pg. 70.

May, C. (ed.) (1991), *SGI joins forces in ACE to build industry standard,* **Silicon Graphics World**, July/August 1991, pg. 10.

Mazor, B. (1991), *Imaging at SIGGRAPH Las Vegas: Environments in spotlight*, **Advanced Imaging**, September 1991, pg. 26.

Mercurio, P.J. (1991), *The Data Visualizer*, **Pixel**, July/Aug. 1991, pg. 31.

Montgomery, J.I. (1991), *Premier parallel performer*, **Digital Review**, September 9, 1991, pg. 35.

O'Connell, B. (1991), *GUI update*, **DEC Professional, Vol. 10**, No. 3, March 1991, pg.60.

Pepke, E., K.W. Johnson, J. Murray, and P.S. Ray (1992), *The SciAn visualization package and applications to problems in meteorology*, Preprint Volume of the **8th International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology**, January 1992, Atlanta, pg. 93.

Picarille, L. (1992), *API to let Mac software run on RISC systems*, **InfoWorld, Vol. 14**, Issue 3, January 20, 1992, pg. 1.

Poole, G.A. (1991), *Sun booms, SPARC vendors fume*, **UnixWorld, Vol. VIII**, No. 7, July 1991, pg. 15.

Pope, D. (1991), *Visualization news; Tools in the news*, **Pixel**, July/Aug. 1991, pg. 10.

Porter, S., B. Robertson, and A. Doyle (1991), *News in Brief*, **Computer Graphics World, Vol. 14**, No. 10, October 1991, pg. 14.

Quinlan, T. and E. Scannell (1991), *Apple, IBM ready to finalize joint venture*, **InfoWorld, Vol. 13**, Issue 38, pg. 1.

Quinlan, T., E. Scannell, and K. Scott (1991), *IBM, Apple ink historic deal*, **InfoWorld, Vol. 13,** Issue 40, October 7, 1991, pg. 1.

Scannell, E. and T. Quinlan (1991), *IBM, Apple set to launch deal*, **InfoWorld, Vol. 13,** Issue 39, September 30, 1991, pg. 1.

Siino, R. (1991). *Iris Graphics Library: The Standard for 3D*, **IRIS Universe**, No. 18, pg.18.

Serlin, O. (1991), *Unix inches closer to off-the-shelf software*, **UnixWorld**, September 1991, pg 33.

Smith, N.P. (1991a), *Allies rally around new MIPS chip to confront Sun, IBM, and HP; Also support Microsoft OS/2 3.0 and Intel x86 Series microprocessors*, **Supercomputing Review, Vol. 4**, No. 5, May 1991, pg.11.

Smith, N.P. (1991b), *Intel's new 2.5 million-transistor i860XP - big impact on massively-parallel market, weaker prospects for workstations*, **Supercomputing Review**, June 1991, pg. 6.

Smith, N.P. (1991c), *New NERSC Director pursuing CRAY-3 program, seeking active role in Cray research MPP plans*, **Supercomputing Review**, October 1991, pg. 12.

Smith, N.P. (1991d), *Wavetracer: Low-cost MPP power*, **Supercomputer Review, Vol. 4**, No. 10, October 1991, pg. 18.

Smith, N.P. (1991e), *Important newsystems from Alliant, Cray Research, Intel, Kendall Square and Thinking Machines: Will MPP reality catch up with its potential at last?,* **Supercomputer Review, Vol. 4**, No. 12, December 1991, pg. 12.

Upson, C. (1990), *Tools for creating visions*, **Unix Review, Vol. 8**, No. 8, September 1990, pg 38.

VandeWettering, M. (1990), *apE 2.0,* **Pixel**, Nov./Dec. 1990, pg. 30.

Varney, S. (1991), *Global networking to push OSI momentum*, **Digital Review**, October 7, 1991, pg. 47.

Vizard, M. (1991a), *IBM signs up Apple to trump ACE,* **Digital Review**, July 8, 1991, pg.1.

Vizard, M. (1991b), *Posix Interface, RISCserver tools further Unix integration*, **Digital Review**, September 23, 1991, pg. 1.

Vizard, M. and Ballou, M. (1991), *Open VMS to come as DEC licenses RISC-based version to rival vendors*, **Digital Review**, September 23, 1991, pg. 1.

Wilson, R. (1991), *Sun unveils Viking superscalar design*, **Electronic Engineering Times**, Issue 657, September 2, 1991, pg. 1.

APPENDIX A.
DATA LEVELS OF THE EOSDIS

DATA LEVELS

Level 0 - Reconstructed unprocessed instrument data at full resolution.

Level 1A - Reconstructed unprocessed instrument data at full resolution, time referenced, and annotated with ancillary information, including radiometric and geometric calibration coefficients and georeferncing parameters (i.e., platform ephemeris) computed and appended, but not applied, to the level 0 data.

Level 1B - Level 1A data that have been processed to sensor units (e.g., radar backscatter cross section, brightness temperature). Not all instruments will have level 1B equivalent.

Level 2 - Derived environmental variables (e.g., ocean wave height, soil moisture, ice concentration) at the same resolution and location as the level 1 source data.

Level 3 - Variables mapped on uniform space-time grid scales, usually with some completeness and consistency (e.g., missing points interpolated, complete regions mosaiked together from multiple orbits).

Level 4 - Model output or results from analysis of lower level data (i.e., variables that are not measured by the instruments, but are derived from these measurements).

APPENDIX B.
REQUIREMENTS OF THE EOS DAAC-SCF INTERFACE

CONSIDERATIONS

The functional objectives of a communication interface between the SCF and the DAAC are to perform data quality control and assurance, to provide algorithm updates to the DAAC, to transfer special data products, and to coordinate the testing and integration of new SCF developed software on the DAAC. The physical interface of this link is provided by the EOS Science Networks (ESN), which is functionally composed of two separate networks, the *mission essential network* and the *mission success network*. The mission success network utilizes the NASA Science Internet with gateways to "The Internet." This network is what the EOS science users would utilize to access the DAAC and would not be used by the SCF to routinely transfer data between the DAAC and the SCF. In contrast, the mission essential network serves as the data pipeline in DAAC to DAAC transfers and between the SCF and the DAACs, and is therefore the primary network to be utilized by the SCF in performing it duties with the DAAC.

The requirements levied on the SCF by the Version 1 DAAC are currently summarized in the EOSDIS Core System Request for Proposal (RFP5-13268/307). The DAAC is composed of three subsystems: the Product Generation System (PGS), the Information Management System (IMS), and the Data Archive and Distribution System (DADS). The individual PGS, IMS, and DADS system requirements are detailed below using the EOSDIS Core System Requirements Specification numbers. The overall DAAC-SCF interface is summed up in the following DAAC functional requirement:

DAAC0110 The DAAC shall interface with the SCFs to support the development of data product algorithms and quality assurance of data products.

THE PGS-SCF INTERFACE

The PGS supports the integration of new and updated science algorithms into the production environment through an interactive link to the scientists located at the SCF. The SCF transmits new or revised algorithms along with associated documentation and test data to the PGS. Test products generated by the candidate algorithms are sent to the SCF. Reviews of the test products are sent to the PGS. The PGS also sends requests for scientists to assess the data quality of its products and receives the resulting quality assessments. All algorithms will be placed under software configuration control.

PGS-0610 The PGS shall accept from the SCFs new or modified calibration coefficients to be validated in the test environment.

PGS-0640 The PGS shall accept from the SCF new or modified Standard Product algorithms to be tested at the processing facility.

PGS-0900 The PGS shall send test products to the SCF for analysis.

PGS-1030 The PGS shall provide a toolkit to the SCF containing version so the routines specified in requirements PGS-0970 to PGS-1025.

PGS-0970 The PGS shall provide file access subroutines that enforce compliance with the adopted standard ECS formats.

PGS-0980 The PGS shall provide job control routines that provide all required task parameters to the Standard Product software.

PGS-0990 The PGS shall provide error logging subroutines for use by Standard Product software in notifying the system operators of conditions requiring their attention.

PGS-1000 The PGS shall provide error logging subroutines for use by Standard Product software in notifying the users of conditions requiring their attention.

PGS-1010 The PGS shall provide mass storage allocation subroutines that provide algorithms for dynamic allocation of storage for temporary files.

PGS-1015 The PGS shall provide ancillary data access subroutines that provide Standard Product software access to ephemeris data (e.g., solar, lunar, and satellite ephemeris), Earth rotation data, and time and position measurement data.

PGS-1020 The PGS shall provide mathematical libraries including linear algebra (e.g., LINPAC, IMSL) and statistical calculations (e.g., SAS, SPSS).

PGS-1025 The PGS shall provide a science processing library containing routines such as image processing, data visualization, and graphics.

PGS-1130 The PGS shall receive product quality assurance from the SCF, which shall describe the results of the scientist's product quality review at an SCF.

THE DADS-SCF INTERFACE

The interface between the DADS and the SCF shall provide for an exchange of documents, special data products, metadata, correlative data, calibration data, ancillary data, and algorithms.

DADS0010 The DADS shall receive from the SCF updated metadata products that have been quality assured by the SCF.

DADS0190 The DADS shall receive from the SCF the following:

a. Special products (Level 1-Level 4)
b. Metadata
c. Ancillary data
d. Calibration data
e. Correlative data
f. Documents
g. Algorithms

DADS1250 The DADS shall, in response to a quality assurance product request from the SCF, transmit the specified products and associated metadata to the SCF for quality assurance approval.

DADS1260 The DADS shall send special algorithms to the SCF.

DADS2380 The DADS shall send to the SCF the following:

a. Level 0-Level 4
b. Special products (Level 1-Level 4)
c. Metadata
d. Ancillary data
e. Calibration data
f. Correlative data
g. Documents
h. Algorithms

THE IMS-SCF INTERFACE

The Information Management System provides the interface between the SCF and the EOSDIS Core System information management functions. The primary role of the IMS is to give the users (e.g., SCF users) efficient access to the DAAC data products, tools, and information to search, locate, select, and order data products. A Virtual IMS data base management system (provided by the DAAC) will manage local data at the SCF.

IMS-1400 The Virtual IMS Information Management software shall operate with a local data base using a data base management software provided by the SCF which conforms to a set of determined standards, thereby facilitating the process of importation of the local data base into EOSDIS Core System.

IMS-1410 The Virtual IMS Information Management software shall provide metadata management services for the local SCF metadata.

IMS-1420 The Virtual IMS Information Management software shall provide the capabilities to search the local SCF data base.

APPENDIX C.
AVAILABLE VISUALIZATION SOFTWARE

DESCRIPTION OF VISUALIZATION SOFTWARE

AGIS - A GIS package with some image processing capabilities handling both raster and vector-based data; developed and distributed by Delta Data Systems [(601) 799-1813].

aPE - A "pipeline" dataflow visualization package, allowing a variety of general-purpose visualization capabilities, including volume rendering and slicing of 3D data, and colormap processing of images; originally developed at the Ohio Supercomputer Center; distribution and support now through TaraVisual Corporation [(800) 458-8731].

ARC/INFO - A vector-based GIS analysis, database management, and display package, allowing limited image processing as well [ESRI: (714) 793-2853].

AVS - A "pipeline" dataflow visualization package, allowing a wide variety of visualization capabilities through the use of distributed and user-defined modules; capabilities include 2D/3D slicing and isolevel surface rendering, trajectories, image processing, and plotting; developed originally by Stardent; development, distribution, and support now through AVS, Inc. [(508) 287-0100].

CORE Image Processing - An image processing package with a well designed X/Motif GUI interface; a collection of modules for image processing, and 2D/3D visualization can be purchased and used in user-developed applications through **HARD CORE** [Image Data Corporation: (818) 796-9155].

ELAS - NASA developed, command-line driven, image processing package with a large collection of data I/O and data analysis modugles [NASA Stennis: (601) 688-1920].

ERDAS - Image processing and raster GIS software with a menu and point-and-click graphical user interface and a variety of data import tools [ERDAS, Inc. (803) 242-6109].

ER Mapper - General Image processing package marketed by Earth Resources Mapping in Australia [+61 9 388 2900].

FAST (NASA Ames) - A point-and-click and menu-driven package for the calculation, processing, and 4D visualization of Computational Fluid Dynamics (CFD) data, including isolevel rendering, contouring, and trajectories; developed at NASA Ames and Sterling Software; available through NASA distribution [(415) 604-4463].

GENAMAP - A vector-based (**GENAMAP**) and raster-based (**GENACELL**) package for GIS modelling, analysis, and display [GENASYS: (800) 447-0265].

Global Imaging - Command-line driven image processing package [Global Imaging (619) 481-5750].

GRASS - A robust command-line driven GIS package developed by the U.S. Army Corp of Engineers; public domain (?).

IBM Data Explorer - A "pipeline" dataflow visualization package that includes image processing, 2D/3D contouring and isolevel surface rendering, etc.; handles time-sequenced data by rendering and storing individual images which are then played back in sequence [(205) 830-6118].

IDL - Image and signal processing, general scientific visualization, data reduction and analysis, and data plotting package; developed and marketed by Research Systems, Inc. [(303) 786-9900].

IISS - Interactive Image Spreadsheet being developed at NASA Goddard SFC which allows image processing and 3D slicing using a spreadsheet interface, and geared toward very large, multichannel data sets [NASA GSFC (Hasler *et al.*, 1992)].

IMSL - Statistics and data analysis libraries, including **IMSL statistics FORTRAN** and **IMSL statistics C**; also available is **IMSL Exponent Graphics** for graphically displaying 2D and 3D data sets [IMSL: (800) 222-4675].

Intelligent Light Fieldview - GUI-driven 3D CFD visualization tool that includes slice planes, isolevel surface rendering, and particle tracing [Inteligent Light: (201) 794-7550].

Khoros Image Processing - An integrated software development environment for image processing, numerical analyses, and data visualization; developed at the University of New Mexico; Khoros is available through anonymous ftp at no charge [(505) 277-6563].

Linkwinds (JPL) - A visualization program developed at NASA JPL that allows multi-variable spatial data to be analyzed and correlated by linking relevant linked windows [FTS 792-0693].

Mathematica - A general purpose data processing and plotting package, supporting a wide range of mathematical functions and allowing 2D and 3D graphs and color rendering [Wolfram Research: (217) 398-0700].

McIDAS - A meteorological data analysis and display system developed at the SSEC at the University of Wisconsin [(608) 262-0783].

METPRO - An operational and research-oriented meteorological information processing package, supporting real-time data ingest [General Sciences Corporation: (301) 953-2700].

NCAR Graphics - A collection of FORTRAN callable routines that allows plotting of data on a variety of display devices [NCAR: (303)497-1000].

NCSA visualization tools - A collection of public domain packages, including basic image processing (**Image, XImage**, etc.), isolevel surface rendering (**Isolev Visualizer**) and slicing (**X DataSlice**) of 3D data, surface generation based on 2D data values (**Height-Color Visualizer**), and contouring (**Contours**) and probing (**DataScope**) of 2D and 3D data.

Paragon Visualization Workbench - Integrated image processing package and visual programming environment; consists of image processing modules and an algorithm development environment [Paragon Imaging, Inc.].

PCI EASI/PACE - An image processing/GIS package that allows integrion of raster and vector GIS and image data [PCI: (202) 785-8281].

PV-Wave - General visualization package for 2D/3D data with features including color and line contouring, plotting, images on mesh surfaces, and minor image processing. Comes in two versions, including **Visual Data Analyzer**, a command line-driven program, and **POINT & CLICK**, a mouse- and icon-driven version. [(800) 447-7147].

SciAN - A public domain 2D and 3D visualization and animation package, allowing isolevel surface rendering and color contouring, with animation recording support [Florida State University: ftp.scri.fsu.edu (Pepke *et al.*,1992)].

SGI Explorer - A "pipeline" dataflow visualization package that provides image processing, and 2D/3D visualization capabilities, including data slicing, isolevel surface generation, true volume rendering, trajectory paths and particle tracing, and data plotting; free with all purchases of SGI hardware [(205) 830-5400].

SoftImage - The 4D Creative Environment is an animation package often used for enhancing scientific visualization.

Spyglass - Developed using NCSA software, Spyglass **Transform** uses a spreadsheet concept for analyzing and visualizing 2D and 3D data, using color contours, trajectories, and plots, while **Dicer** allows analyses and visualization of 3D data, through color contouring, plots, isolevel rendering, and data slicing [Spyglass: (217) 355-1665].

Sun's SunVision - A general visualization package allowing basic data and processing, as well as slicing and isolevel surface rendering of 3D data [Sun Microsystems].

Terra-Mar Image Processing - General purpose, menu-driven image processing software, in two versions: MICROIMAGE for PC-based systems, and IDIMS for Unix-based workstations [Terra-Mar: (415) 964-6900].

Uniras - **agX/Toolmaster** is a development tool allowing visualization functionality to be integrated into existing application software, while **agX/Volumes** allows 3D gridded data to be sliced and rendered with isolevel surfaces.

Vertigo - A software package which has traditionally been used to create 4D animation, is now adding modules such as a 3D data slicer and isolevel surface renderer for scientific visualization.

Vis5D - Interactive public domain application allowing isolevel rendering, contouring, and trajectory plotting of time-sequenced 3D data [SSEC, University of Wisconsin: (608) 262-0783; Hibbard and Paul (1992)].

VisAD - Public domain application for VISualization of Algorithm Development that directly links an algorithm script to the intermediate and final visual results [SSEC, University of Wisconsin: (608) 262-0783; Hibbard et al (1992)].

Vital Images Voxelview - 3D volume rendering package, allowing true voxel rendering, as well as isolevel and slice plane rendering [Vital Images: (515) 472-7726].

Wavefront Data Visualizer - General data visualization package primarily for isolevel rendering, slicing, and probing of 3D data [Wavefront Technologies: (805) 962-8117].

Whip - A collection of image processing routines, developed by GW Hannaway & Associates [(303) 440-9631].

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (*Leave blank*)

2. REPORT DATE: June 1992

3. REPORT TYPE AND DATES COVERED: Technical Memorandum

4. TITLE AND SUBTITLE: **Functional Requirements Document for the Earth Observing System Data and Information System (EOSDIS) Scientific Computing Facilities (SCF) of the NASA/MSFC Earth Science and Applications Division, 1992**

5. FUNDING NUMBERS

6. AUTHOR(S)

Michael E. Botts,* Ron J. Phillips,* John V. Parker, and Patrick D. Wright**

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

George C. Marshall Space Flight Center
Marshall Space Flight Center, Alabama

8. PERFORMING ORGANIZATION REPORT NUMBER: M-692

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

National Aeronautics and Space Administration
Washington, D.C. 20546

10. SPONSORING/MONITORING AGENCY REPORT NUMBER: NASA TM-4392

11. SUPPLEMENTARY NOTES

***The University of Alabama in Huntsville, Huntsville, AL **Universities Space Research Association, Huntsville, AL. Prepared by Earth Science and Applications Division, Space Science Laboratory, Science & Engineering Directorate.**

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Unclassified—Unlimited

Subject Category: 60

12b. DISTRIBUTION CODE

13. ABSTRACT (*Maximum 200 words*)

Five scientists at MSFC/ESAD have EOS SCF investigator status. Each SCF has unique tasks which require the establishment of a computing facility dedicated to accomplishing those tasks. A SCF Working Group was established at ESAD with the charter of defining the computing requirements of the individual SCFs and recommending options for meeting these requirements. The primary goal of the working group was to determine which computing needs can be satisfied using either shared resources or separate but compatible resources, and which needs require unique individual resources. The requirements investigated included CPU-intensive vector and scalar processing, visualization, data storage, connectivity, and I/O peripherals. A review of computer industry directions and a market survey of computing hardware provided information regarding important industry standards and candidate computing platforms. It was determined that the total SCF computing requirements might be most effectively met using a hierarchy consisting of shared and individual resources. This hierarchy is composed of five major system types: (1) a supercomputer class vector processor, (2) a high-end scalar multiprocessor workstation, (3) a file server, (4) a few medium- to high-end visualization workstations, and (5) several low- to medium-range personal graphics workstations. Specific recommendations for meeting the needs of each of these types are presented.

14. SUBJECT TERMS

EOS, Scientific Computing Facilities (SCFs), Computing, Computer Specs, Workstations

15. NUMBER OF PAGES: 104

16. PRICE CODE: A05

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	**Unclassified**	**Unclassified**	**Unlimited**

NSN 7540-01-280-5500

Standard Form 298 (Rev 2-89)
Prescribed by ANSI Std Z39-18
298-102

NASA-Langley, 1992