



Research Institute for Advanced Computer Science
NASA Ames Research Center

IN-60
43066
P38

Final Report on the Workshop on NASA Workstation Technology

Robert L. Brown

(NASA-CR-188866) WORKSHOP ON NASA
WORKSTATION TECHNOLOGY Final Report
(Research Inst. for Advanced Computer
Science) 38 p

CSCL 09B


63/60

N91-32800

Unclas
0043066

RIACS Technical Report 90.30

NASA Cooperative Agreement Number NCC2-387



Final Report on the Workshop on NASA Workstation Technology

Robert L. Brown

JULY 1990

The Research Institute of Advanced Computer Science is operated by Universities Space Research Association, The American City Building, Suite 311, Columbia, MD 21044, (301) 730-2656.

The work reported on herein was supported by Cooperative Agreement Number NCC2-387 between the National Aeronautics and Space Administration and the Universities Space Research Association.

Final Report on the Workshop on NASA Workstation Technology

Robert L. Brown

*Research Institute for Advanced Computer Science
NASA Ames Research Center
Moffett Field, CA*

RIACS TR 90.30

ABSTRACT

On March 13-15, 1990, RIACS hosted a Workshop on NASA Workstation Technology. The workshop was designed to foster communication among those people within NASA working on workstation-related technology, to share technology, and to learn about new developments and futures in the larger university and industrial workstation communities. This report documents the workshop and its conclusions. Briefly, the workshop was a success; many people asked that it be repeated regularly. New collaborations were established as a result of it. We learned that there is both a large amount of commonality of requirements and a wide variation in the modernness of in-use technology among the represented NASA centers.

1. Introduction

This report documents the results of the "Workshop on NASA Workstation Technology" held by the Research Institute for Advanced Computer Science on March 13-15, 1990 near NASA Ames Research Center. The purpose of the workshop was to bring together people working for or with NASA in the area of workstation technology development to discuss their projects and future needs. Also invited to the workshop were industrial and academic researchers and visionaries to present their directions for the future.

The workshop was the result of a request by the Chief of the Human Factors and Information Sciences Division within the Office of Aeronautics and Space Technology (now Office of Aeronautics and Exploration Technology) for the purposes just described and to provide input for managing a diverse set of workstation-related projects. The Telescience Consortium at NASA Ames Research Center helped support the effort of creating and running the workshop

The intent of this report is to convey the results and findings of the workshop and to make recommendations concerning where additional emphasis is needed in the future. The intended audience of the report is principally program managers and other government personnel concerned with establishing and implementing research and development funding priorities. The goal is to leave the reader with a good sense of the state of workstation technology within NASA relative both to mission needs and the state of the technology outside NASA, and where emphasis should be placed in the future. In this sense, this document can be used as a reference for the near future, but because the technology is advancing so quickly, it may become obsolete within one year's time.

1.1. Summary of Results

Before describing the details of the findings of the workshop, it is worthwhile to examine a few of the most obvious results. No survey was conducted to verify that most of

these items were universally agreed upon, but in private conversations and in small-group discussions, these were most often mentioned.

- 1) There is a lot of commonality of interest in workstation technology within the agency, and a fair amount of replication of effort. However, the replication is not necessarily harmful since, because the technology is advancing so quickly, a large amount of effort is required by all parties just to stay current.
- 2) There is a wide variation among the NASA centers in the level of technology currently in use. The variation nearly matches the nearness to operational missions for each center, that is, centers that focus on research and technology development are generally more advanced and those that focus on missions, particularly operations are less so. This was not a surprise.
- 3) There is an intense interest in continuing the type of information exchange that took place in the workshop. This was one of the first opportunities for some of the attendees to learn about others' work and to share experiences.
- 4) The area of "computer graphics" or "scientific visualization" uses and produces technology that can be used in other, traditionally non-visual, applications. Currently, the best computer graphics in the agency are produced from numerical simulations on super-computers. The same techniques can be applied to help visualize observational data.
- 5) Standards are more important than ever. Workstation-based systems usually involve a variety of equipment interconnected by a network of some kind. The equipment is usually supplied by more than one vendor and is a mix of old and new. To make such a varied system work and to take advantage of developments from other sources and projects, close adherence to a well-defined set of standards is necessary. Choosing the "right" set of standards, however, is not a trivial task.
- 6) "Virtual reality" technology (defined later) will become increasingly important in

NASA missions, though its direct application today is not clear. We need to pay strong attention to virtual reality and watch for ways to bring that technology to bear.

These results will be elaborated upon in later sections, and some will result in specific recommendations.

2. Vision

It is important to separate, in discussions of workstation technology, what workstations *are* or *do*, from what types of applications they *support*. These two are not entirely orthogonal; once a model of a workstation has been developed, then each variable in that model can be emphasized or de-emphasized for instances of workstations that support a particular activity. This section attempts to define a model of a workstation that can help structure thinking about how to create an optimal workstation (at minimal cost) for a particular activity.

A workstation is a form of computer, but one that emphasizes an input/output relationship with its user. In the ultimate, as William Bricken stated in his keynote speech, a computer is not simply a machine to compute with numbers, but instead should be considered to implement an abstract reality. The workstation then becomes the interface between that reality and the human. As an example, a simulation model of a physical model in almost all cases attempts to model the physical universe, so that experiments can be performed using the model rather than the real world (which often is not economically feasible). Then, because the computer is simulating a reality, the presentation of the information it computes about that reality plays a pivotal role in the understanding of its results. This is the role of the workstation.

We did not attempt to define the lower bounds of the hardware, software, or capabilities associated with the term "workstation." Such an exercise is nearly meaningless; some may argue that a Teletype device is a workstation with a very low "presentation bandwidth" -- the volume of information it is capable of presenting to its user -- and with a very low "control bandwidth" -- the volume that the user can transmit through the device to the computer. Few users of workstations, though, would call a Teletype a "workstation." To be useful as a workstation, a device must have a moderate to

high presentation and control bandwidth between itself and the user. The evolution of workstations depicts an ever increasing capability for both.

2.1. Workstation Models

There are several ways to develop a model for workstations. One is a “functionality” model which describes the basic capabilities of the hardware and software that constitute a workstation. Another is a “capability” model, describing what particular types of activities a workstation enables its user to engage in. Here we present both.

2.1.1. Functionality Model

Most forms of computations can be described as having three parts: some input, some processing, and some output. This is often referred to as the IPO model -- input, processing, output. Another way to describe it, symbolically, is

$$y = f(x)$$

where x is the multivalued input, f is the computation, and y is the multivalued output. Though this is a gross simplification of computation, it serves as a useful model in describing workstation responsibilities.

This model applies to entire computations, not just that activity that takes place within a workstation, a supercomputer, or the network, but the combination of all resources and activities involved. Hence, the f in the description can be arbitrarily complex and involve other subcomputations, each with their own inputs and outputs.

One view of the role of a workstation is that it is only responsible for those parts of the overall computation that directly involve the human user -- principally presenting output for visual and audio inspection and gathering input, generally but not limited to hand manipulations (key presses, mouse movements, etc.). We call this the *minimal function*, or I/O, model because it describes the workstation as only an I/O device, only supporting

human-computer interaction (HCI). This provides a good starting point for workstation models because the HCI is an exclusive task for workstations.

Many people consider local processing power a fundamental attribute of workstations. It is true that local processing power is necessary to drive sophisticated HCIs. Having dedicated processing power for other tasks, such as local analysis, is not a part of the minimal-function model. However, if local processing power does not exist, then the f in the model above must be completely performed by a remote computing resource. Often the remote processor is shared and cannot guarantee an upper bound on response time. Since low response time is critical for effective HCIs, if it is not achievable by relying only on remote processing, local processing power must be made available. Hence, the second-level, or *IPO*, workstation model includes sufficient processing power to perform those computations that are critical to maintaining a usable HCI. A problem with this model, however, is that it does not provide an upper bound on the amount of power needed in the workstation. The processing power to put into a workstation is a function of the demands of the application, the effectiveness of the network for remote resource access, the responsiveness of the remote processing resource, and the desired upper bound on response time.

The IPO model is more descriptive of the current state of the art in workstations than the I/O model is, and is better as a prescriptive model because it allows for more customization of the workstation to the task-at-hand. However, the IPO model is low-level; it describes the machinery -- input, processing, and output devices. It does not address what types of activities the workstation enables.

A pictorial representation of the IPO model is given in Figure 1. Here, input, output, and processing are represented on separate axes, each independent. Each instance of a workstation defines a point in that space. Figure 2 shows two examples: a scientific RISC-

based workstation with graphics capabilities, and a virtual reality station for viewing the relatively simple graphics databases (such as the Ames VIEWS system).

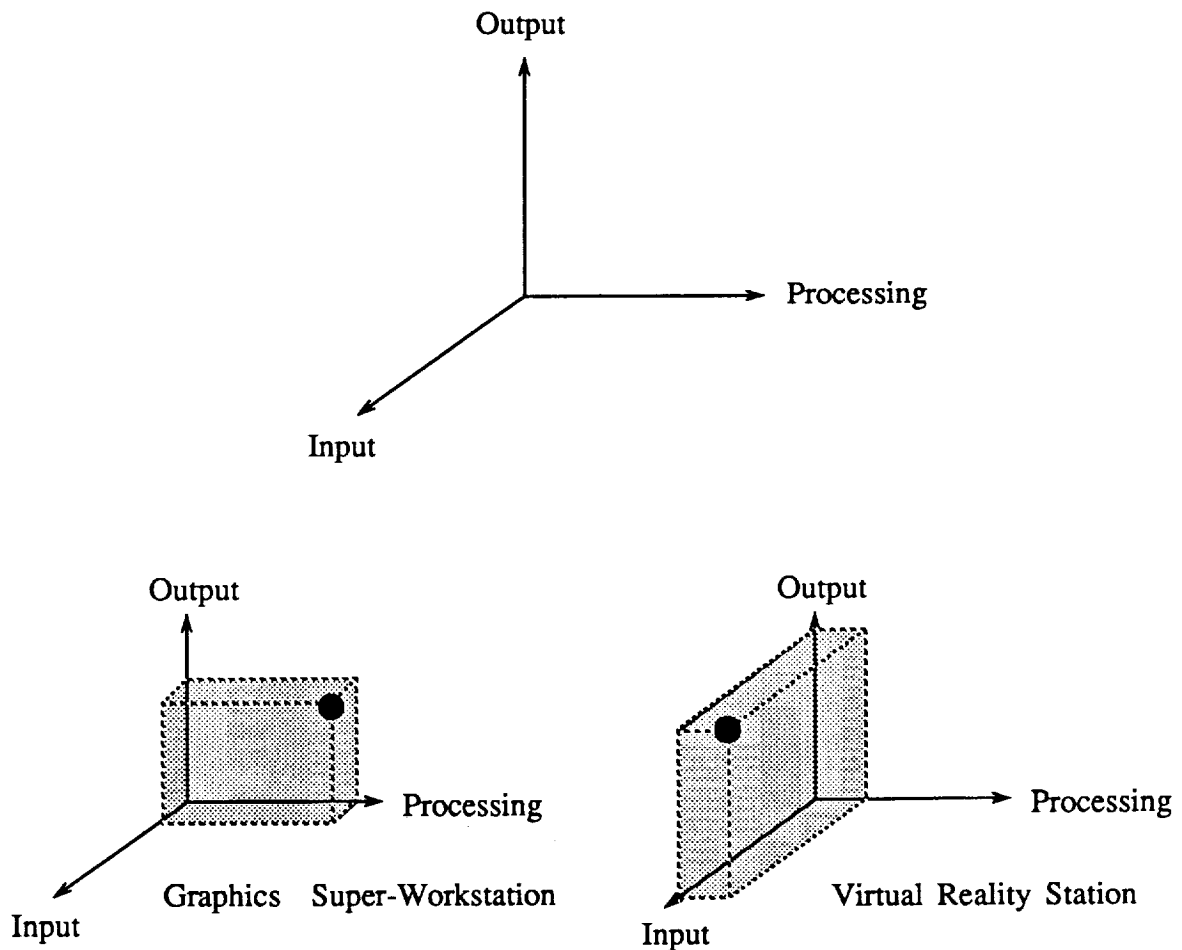


Figure 1. The IPO Functionality Model and Examples

2.1.2. Capability Model

What does a workstation do? It is worthwhile to examine how workstations are used and, hence, why they are deemed important. A typical application that employs workstations also employs much more in the way of hardware and subsystems interconnected by a network. Powerful computing resources, sensing instrumentation, large databases, and process control actuators are often included. Often neglected but of tremendous importance is other people; workstations and their networks can provide a convenient and powerful way

of interacting with colleagues and coworkers. Because the workstation needs to interact in a consistent and predictable and timely fashion with its human user, it must have dedicated processing power to drive those interactions. This dedicated processing power can also be put to use for local analysis of remotely obtained data.

Thus, the workstation is a tool that one uses to access information, to perform analysis, to monitor and control abstract or real processes, to perform design activities, to develop new software more readily, and as an interface to collaborate with others. We call these the “capabilities” of a workstation, and the application to which a workstation is applied drives the extent to which each of these functions is emphasized. None of these items are necessarily useful in isolation, they merely provide capability that enable a workstation to be applied to a particular task.

The reports about workstation usage within NASA, as presented at the workshop, back up this model. There are, however, groups who use workstations in a laboratory as stand-alone computers. These systems can be viewed as a microcosm of the larger, and more typical usage scenario. Stand-alone workstations are either applied to narrowly focussed projects, are used as testbeds for prototyping, or are used as vehicles for learning about workstation technology.

The multipart model here is not meant to describe or prescribe the hardware and software components of a workstation; that is discussed in the “Technology” section. Rather, these parameters describe higher-level functions that workstations are capable of supporting. Each real workstation will employ some combination of these parameters in varying degrees, depending on the task to which they are applied.

2.1.3. Access

Access, apropos to workstations, refers to the ability of the user to locate and acquire information. This information need not be local, since inherent with workstations is their

ability to communicate with remote resources by way of networks. Because a workstation is not a self-contained computing system, but instead is an interface to a much larger computing environment, its ability to access remote resource can become a critical attribute.

As networking of computers and workstations becomes more pervasive, the need for high-quality methods for accessing remote resources becomes more critical. In particular, remote data resources are becoming increasingly important. Managing and traversing the variety of databases available via networks is a form of exploration, and the workstation is responsible for producing the best presentation and control. These databases can be arbitrarily complex and contain information in multiple media, such as text, imagery, video, audio, and others. Workstations can be created to accommodate any number of these.

The ability of a workstation to access remote information is a function of the level of technology employed in its networking interfaces and the level of technology in the network itself (e.g. Ethernet vs. FDDI vs. modems). For accessing massively large databases, the time required is a function of the throughput of the network and network interface. Time is the critical factor, since a workstation is an interactive device, information should be accessible at interactive rates. However, for the most part, the time required to gain access to a remote data resource is not a function of the level of workstation technology used, but rather of the level of networking technology.

Because of the way they are integrated into larger systems, workstations need the ability to readily access remote resources, both computational and data.

2.1.4. Analysis

The local computing power of a workstation allows its user to control the upper bound on the time it takes to perform a particular analytical task. In this sense, the workstation can become a "personal computer" for the user. Indeed, many people think of the traditional, un-networked personal computer as a workstation, which matches our model because of its

ability to perform analysis, but fails matching entirely because of its inability to access remote resources.

At the low end, a workstation provides only a modicum of local analysis capability; many applications do not need much more. On the high end, workstations may have power equivalent to that of supercomputers of just a few years ago. The advantage of localizing such power is that the user can be guaranteed exclusive use of it.

The ability to perform local analysis is derived from the requirement on workstations to have sufficient local processing power to communicate with its user at a sufficiently high bandwidth for the particular application. Often, that leaves spare processing power that is made available to the user for those computations that would otherwise be shipped to a more powerful, centralized, computing resource.

The local processing power necessary to drive the user interaction is sufficiently high as to leave reserve for local analysis. Also, oftentimes the need for quick response from computationally intensive tasks mandates a high degree of local processing power for analysis tasks..

2.1.5. Monitoring & Control

Workstations present the ability to create “software control panels” and displays that can present information about real or abstract processes. For monitoring, workstations with advanced displays can present images that emulate well-know physical gauges, such as meters and strip chart recorders, or images that are nearly photographic in quality depicting the object, or a metaphor for it, that they are monitoring. The information used in monitoring a system may come from a variety of sources; the workstation is then responsible for fusing these multiple data sources into one or more highly understandable presentations.

For control, workstations allow highly configurable software control panels that do not require hardware modification to implement changes. This reconfigurability allows a single machine to be used for many different monitoring and control tasks simply by running

different software. This is an important point for NASA missions; it enables the creation of a sophisticated mission operations center that can be applied to a wide variety of tasks and reused for a new mission when the current one ends.

Control devices, that is, the hardware the human uses to communicate to the workstation, are highly varied. Traditionally, commands typed at a keyboard have been used. However, there is an increasing interest in new and novel human interface devices that allow the user to communicate control commands in more natural ways such as by voice or gesture.

Workstations are the focal point of interaction; they need the ability to fuse and display information as well as the ability to allow the user to manipulate either the process or the interface itself.

2.1.6. Design & Development

Design is a complicated process, being a mix of creative thinking and successive refinement. Though it is not clear how the workstation can augment creative thinking (perhaps by automating mundane tasks), it has been very clear for many years that computer-aided design (CAD) tools are extremely powerful for augmenting the refinement process. Once the conception of a design exists and has been entered into a CAD system, the designer can "play" with alternatives and, depending on the power of the CAD system, the computer can check the design against a set of rules, or constraints, to determine if a particular design is manufacturable and if so, potentially its quality, based on a criteria set.

The design process is highly manipulative, and workstations are the focus of manipulation, or interaction, between the user and the computational power available. Design is also very visual, relying on high-quality renderings of the object being designed. In a sense, the designer is exploring the object, finding places where enhancements or corrections can be made, and then implementing those changes. This requires a high degree of visualization and control capability in the workstation hardware and software.

Because of its ability to present potentially large amounts of information, the workstation can be a powerful platform for development activities, not necessarily limited to software development. Computer-aided software engineering (CASE) tools require the ability to present information in multiple media to the user, various views of the software under development, both statically and dynamically. These views increase the understanding of the program under development and have the potential of reducing development time.

Computer-aided design and software engineering can provide a highly interactive, manipulative, method for designing a physical object or piece of software. The workstation is the vehicle for implementing that interaction, whether by presenting high-quality images of the object under design, or by offering a high degree of control over that object.

2.1.7. Collaboration

Workstations provide the opportunity for multiple dispersed individuals to engage in computer-supported collaborative activities beyond what was previously possible. Because the networks that interconnect workstations are capable of carrying a variety of media, multimedia workstation-based collaboration tools can be developed. Even though text-based electronic mail can be an effective first step in supporting collaboration among geographically dispersed colleagues, the ability to send images, video, sound, and structured data such as spreadsheets, graphs, and graphics, can enhance the quality of the collaboration just as these same media have done for collocated collaborators.

Multimedia collaboration is a relatively immature area; no standards yet exist for media representation and transmission. Such standards will become critical for widespread use of the workstation as a collaboration tool, just as standards for telephony are required for large telephone systems to work together.

Collaboration can be either batched or real-time (conversational). With the former, information is shared and exchanged at a relatively low granularity which ranges from hours

to days or even weeks. The latter entails a much higher granularity of information exchange, on the order of seconds. Workstation-supported multimedia conferencing enables a group of people to teleconference using their workstations and networks to exchange information and work together on a shared set of information.

Workstations and their networks can enable a group of dispersed collaborators to share information readily.

2.2. Model Summary

Just presented is a model of the capabilities of workstations without specific mention of the application areas in which they are used. For example, a workstation used by a mission planner may need to perform well in access, design, development, and collaborations, but perhaps less so in analysis and monitoring. The workstations used in mission operations will require high capabilities in monitoring and control, analysis, and perhaps collaboration, but little in the way of design and development capabilities.

The model can be depicted as in Figure 2. Each application will require different levels of each of these capabilities.

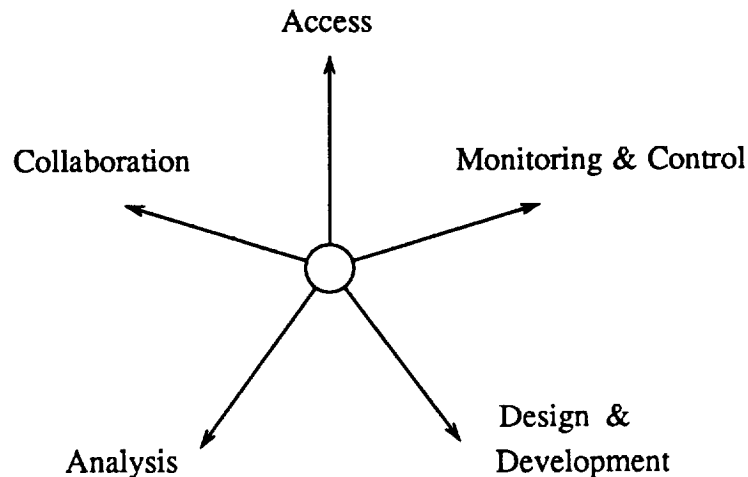


Figure 2. Depiction of the Model of Workstation Capabilities

3. Applications and Payoffs

The benefits of workstation technology are clear. Workstations enable their users to more efficiently access, organize, and understand data.

3.1. Management

Management users of workstations principally use them to access information about what they manage, to help with budgeting and forecasting, and to study alternatives using simulation or optimization techniques. These are traditional management functions, however, and most can be performed on a stand-alone computer. The payoff in workstation technology for management is that this advanced technology will enable managers to more precisely and effectively communicate with their group within the organization, regardless of geographic dispersment. The "group" can be either a group of peers (*i.e.*, other managers) or those being managed. Existing tools, such as electronic mail, can have a strong positive impact on project management, for example, because they provide a medium for regular, non-intrusive, reporting of project status. Using its presentation capabilities, the workstation permits managers to track progress through a project by visualizing status on a PERT, or similar, chart. Using its access power, the workstation can provide an automatic mechanism for updating that chart.

Communication is a key aspect of management. Hence, a workstation for management purposes requires a high degree of collaboration capability to allow the managers to engage in multimedia conferences with others in the organization. Face-to-face meetings and conferences are the mechanism in use today for much of the communication that takes place in management. Effective use of present and near-future workstation technology will permit managers to engage in meetings and conferences without leaving their offices. The workstation will augment the possibilities for meetings by providing shared spreadsheets, charts, and graphics. Multimedia real-time conferencing and multimedia electronic mail systems are beginning to appear [Acev85], not just for the traditional research and

commercial sector [Sari85, Reyn85], but for military management and information control as well [Pogg85].

Management of the future will rely on advanced workstation technology for management of information and projects, and as a medium, using real-time multimedia conferencing, for communication.

3.2. Science

The value of a scientific workstation cannot be overstated, and has gained wide acceptance inside the agency. The volume of information involved in a typical scientific application can be quite large, regardless of whether that data is computer generated, as with numerical simulation, or sampled from nature. Understanding the data and identifying patterns or features is nearly impossible without computer assistance for visualization.

Science objectives are major drivers for NASA missions. The Great Observatories will be generating extremely large quantities of data; EOS will overwhelm our capability to store and analyze scientific data. It is conceivable that any scientist with a modest research grant will have access to most of this data and will require a great deal of processing and visualization power to analyze it. There are insufficient supercomputers in the nation to be used by this corps of scientists to perform their analyses or to process the data to verify their models. With the cost of powerful scientific workstations (in 1990, in the 10-20 MFLOP range) within the budget capability of a large number of scientists, it can be safely assumed that much post-observational data processing will migrate from centralized data analysis centers to the researchers' desktops. Because of the input/control capabilities of workstations, the researchers will be able to control and adjust their analyses as they progress; a capability not possible with centralized data analysis.

Visualization is perhaps the most important capability that workstations bring to the scientific research community. NASA deals both with image data and numerical data resulting from computational simulations. Each type of data has unique requirements for

processing and visualization, but both require a stable and powerful set of software for these tasks. Image processing packages abound (e.g. AIPS and IRAF) but must be extended to accommodate scientific domains of interest to space science and earth science. Similarly, there are numerous post-simulation scientific visualization packages (e.g. GAS, Rivers, etc.). In many cases these packages are freely available and well-supported. This fact has resulted in great progress within NASA in the areas of scientific visualization.

Finally, NASA is in an era of multidisciplinary missions, EOS is a good example. The data must be shared by scientists in several disciplines and compared and fused with data from other sources. The community of researchers using NASA data will grow dramatically over the course of the next decade. Hence, NASA must pay attention to standards for data representation and provide convenient methods of accessing the data, preferably directly to and from the scientists' workstations.

Scientists in the future will need to process vastly greater quantities of data and have direct control over the manipulation and visualization of that data. The workstation provides the visualization and control functions, and the cost of adding sufficient processing power for local analysis is approaching the range of most budgets.

3.3. Operations

Consoles for mission operations have been commonplace within NASA since the beginning of the space program. In the past, these consoles were custom-built for each mission at high cost. NASA has entered an era of launching many smaller missions, and the need for a retargetable mission operations capability has been recognized by many. Operations consoles present displays of operations parameters for the purpose of managing a mission. Hence, in a way, operations and management use of workstation technology is very similar, except that operations requires more sophisticated real-time interfaces to instrumentation relaying the health and status of the spacecraft.

As in management, communication is important in mission operations. Modern workstations can augment the operational capability by providing a means for the operations team to communicate among themselves (whether collocated or not) and to interact with the mission science team (usually at a different site) as applicable. Just as managers can share spreadsheets and PERT charts, operations teams can share visualizations of the spacecraft attitude and health, and engage in group conferences.

Planning and scheduling can be supported by workstation technology. Understanding a complex schedule requires visualizing a graphical representation, which may be created within the workstation. A mission plan is very much like a spreadsheet; there are interdependencies among the events in the plan. During operations, plans usually must be altered to accommodate delays and opportunities. The workstation can maintain a constant up-to-date view of the mission plan for the mission operators.

On Space Station Freedom, workstations will be the focus of information for the astronauts. They will present displays of station health, payload status, resource utilization, consumable inventory, and non-consumable item location. However, today's most powerful workstations are too heavy and generate too much heat for SSF. In the future, more powerful on-board computing will be needed as operations expand, hence, packaging powerful workstations into lightweight and low-heat units will be necessary for manned missions. This will become especially critical for future lunar and Mars missions.

For operations, future workstations will allow mission operators to work in close collaboration with each other and with scientists, and will provide advanced, sharable displays of mission status, including spacecraft and payload health, and mission schedules. Packaging will be critical for future on-orbit workstations.

3.4. Development & Design

As the focal point for design and development in the context of projects, workstations enable their users to access a wide variety of resources to put to bear on the task. In

software development, having access to software archive servers provides a wealth of source code, all of which is free and much of which is reusable (a recently published list of archive servers on the internet listed 644 sites). Hence, the workstation as an interface to remote resources provides a capability for software developers heretofore not available.

Visualization in language environments has evolved from inherently textual displays to graphical interfaces. The Smalltalk environment is an early example, providing a visual browser allowing the programmer to “wander” through the types hierarchy. Later examples of visually-oriented programming environments include Cedar, Pecan, Garden, the Cornell Program Synthesizer, and Aloe [Ambl89].

Parallel program development can benefit from workstation-based visualization techniques. There are visual languages that aid in the design of parallel systems. Similarly, understanding the behavior of a running parallel program requires the distillation of a large quantity of information. Performance debugging tools exist and can greatly simplify the debugging process [Lehr89].

Many design operations are inherently visual, and the workstation provides the interface in many cases between the designer and a software system that provides the design elements and manages design constraints. With proper use of workstation technology and tools, a designer can design, and, where applicable, test or validate a design against a set of rules. Often advanced input technology is incorporated in design workstations; devices such as tablets, styli, touch screens, and data gloves are now or are becoming commonplace.

More visualization will be incorporated into software development environments in the future because of the increasing size and complexity of software systems. Input technology is a critical area for future non-software design systems.

4. Virtual Reality as a Goal

The workshop began with an inspirational keynote speech from William Bricken, the Chief Scientist at the University of Washington's Human Interface Technology Center. His talk was titled "A Vision of Virtual Reality." Dr. Bricken presented virtual reality as a way of thinking about how computers can impact our understanding of models. The talk spurred much conversation; numerous individuals immediately recognized how virtual reality technology could directly benefit their project. Yet even Dr. Bricken admitted that the technology necessary for a true virtual reality (often called "cyberspace" in the literature) could be 20 years in the future.

As previously mentioned, computers are not just fast calculators of arithmetic processes; they have the ability of creating new realities that either directly model known realities or transcend those realities, creating new ones. Though the description often sounds as if it comes from the literature of science fiction, the notion has concrete benefit to NASA. Indeed, separating the image of fantasy from practical implementation can be a limiting perceptual factor in any description of virtual reality. The concept has direct applicability to many NASA missions, yet the perception is that virtual reality is too futuristic to be considered seriously.

Virtual reality requires metaphors. Xerox invented one of the first highly useful workstation metaphors with the creation of windowing systems -- the desktop metaphor. Apple Computer popularized the idea. Xerox has moved on to newer office metaphors, the "rooms" metaphor, for example, which permits the workstation user to establish screen configurations (based on window layout) and easily switch from one room to another.

At one point in the post-workshop sessions, a small working group concluded that virtual reality is the ultimate goal of all workstation technology. However, we concluded that to depict the ultimate virtual reality system as the end-point of technology development was too narrow-sighted. Instead, the models described in Section two of this report are more

realistic; each key aspect of workstation technology can expand continuously outward, and the fidelity of a virtual reality system is a function of the state value of each parameter in the models.

Virtual reality systems are as important to NASA missions as artificial intelligence systems, perhaps more so. So long as humans are involved in the process of operations or exploration (mission science), the degree to which those humans can immerse themselves in the environment of study, the more effective they become in their understanding. For operations, if the human operator is in the control loop, that person must have as much understanding of the situation (of the spacecraft) as is possible. A useful metaphor for this type of operations is to allow the operator to “be” the spacecraft, and have direct access to all its instrumentation. Precisely how this maps into a control interface, however, is a research topic. For exploration, providing as much spatial information as possible, and allowing the explorer to manipulate, in some restricted way, the environment results in a metaphor wherein the explorer is translocated to the exploration site.

These, and other, metaphors, require research in order to translate them into systems that we will accept as easily as we accept the desktop metaphor for office automation computer systems.

5. Technology

The technology, both hardware and software, of computer workstations is advancing so quickly that the typical time between new product development to announcement in the industry is six months. The workstation industry, with a couple of exceptions, is engaged in hardware competition, primarily to increase the speed of the processors in order to drive higher quality displays.

5.1. Hardware

The hardware of a typical contemporary workstation consists of a processor, memory, screen, keyboard, pointing device, network interface, I/O bus, and sometimes specialized processors (graphics, signal processors, etc.). All of these features can be categorized into the elements of the functionality model -- input, processing, and output. These categories are treated separately.

5.1.1. Input

The bandwidth of information from the workstation to the user has always been considered to be more important than the bandwidth in the other direction. One reason is plain; the quantity of information that the workstation may need to present is potentially very large, typically measured in kilobytes or even megabytes. Presentation data is quantifiable, and the quantities are much, much larger than just a few years ago. But, there is no consistent quantification of the number of bytes of informations that a user can direct at the workstation, and it is generally considered to be very low. Metrics such as "keystrokes per second" are commonplace, but are not accurate indicators of potential bandwidth from the human user to the computer.

As a result of the input bandwidth being considered low, the interfaces that are designed to accommodate it typically cannot handle anything but slow input. When these interfaces are confronted with even moderate speed (*e.g.* 9.6kb/sec) data, they become a bottleneck. When presented with multiple sources of moderate speed input, they collapse

into ineffectiveness. A typical input interface interrupts the workstation processor on every character. At 19,200 baud input, a character is received approximately twice a millisecond. The overhead in processing a character interrupt is typically on the order of a hundred or so microseconds, often as high as 500 microseconds. Receiving 19,200 baud input can saturate many workstation processors simply because the interfaces and operating systems are not designed to handle the load. New interfaces to high-bandwidth devices must be designed and standardized to accommodate such devices.

Apropos to input technology, or those devices that improve the "control bandwidth" of the workstation, are:

- Mouse -- 2D positioning
- Space ball -- 6 degree of freedom specification
- Tracker -- 6 degree of freedom movement
- Data glove -- complete hand orientation and gesture
- Data suit -- complete body orientation

5.1.2. Output

Workstation output technology traditionally has meant cathode ray tube screens. However, at least two presentations at the workshop described projects that are investigating the use of new visual display technology. In one case, a stereoscopic head-mounted display is being used in order to present spatial information. In the other, a viewer built like eyeglasses is being used for portability. Visual display technology offers many possibilities today, including:

- Polarizing stereoscopic display
- Flat panel LCD display
- Head-mounted stereoscopic display
- See-through head-mounted display
- True 3-D flexible display

The fidelity of a display and its field-of-view govern how well it can create a visual virtual reality. Higher fidelity results in more realistic images; higher field-of-view prevents distractions from "leaking" in. However, high fidelity imagery is not necessary for a convincing vir-

tual reality; low fidelity imagery combined with proper field-of-view and head tracking is sufficient.

5.2. Software

More than anything, the software within a workstation defines its capabilities. In almost every presentation at the workshop, software issues dominated hardware issues. It is the case that software technology has not kept abreast of hardware technology. The speed of raw processing and graphics performance has been increasing at a rate greater than all informed predictions in the past four years, yet software technology remained basically unchanged in that period.

5.2.1. User Interface Software

Workstation application developers all mostly use the same language (the "C" language) but there is not yet a commonly popular graphics user interface (GUI) system or language binding. Indeed, two major factions in the GUI arena continue to feud (Open Software Foundation, OSF, and Unix International, UI) and there is not yet a clear common alignment of the major workstation vendors to these factions. Specifically, Sun Microsystems and AT&T are aligned with UI, and IBM, DEC, and HP are aligned with OSF. Each belittles the others' efforts.

Each of these toolkits are based on a common workstation window system platform, the X Window System from MIT's Project Athena. X is derived from an earlier windowing system developed at Digital Electronics for the V Kernel developed at Stanford University. The protocol that defines what an application can do with the workstation screen is limited in X, constrained to describing operations in terms of individual dots, or "pixels" ("picture elements") on the workstation screen. With an ever changing display technology, confining applications to describing operations in terms of pixels is too limiting. Additionally, the X protocol only allows applications to send a stream of drawing commands to the screen. Hence, highly repetitive operations, such as drawing a regular grid mesh, result in a very

high amount of communication. In a similar vein, fonts in the X Window System are fixed in size, defined by the number of bits on the screen they occupy.

5.2.2. Operating System Software

Operating system software is less of an issue. The UNIX operating system is a clear winner, though there exist two major families of UNIX. One is derived from the original UNIX systems, developed at AT&T Laboratories in Murray Hill, New Jersey. These versions are now called "System 5." The other family derives from an earlier AT&T version and was developed at the University of California at Berkeley, and are called "BSD" (for "Berkeley Software Distribution") and are known for their superior capabilities in interfacing to TCP/IP-based networks. Combining the two families has been a goal of the industry for at least six years. Now a major proponent of the BSD faction, Sun Microsystems, and the owner of the System 5 faction, AT&T, have joined forces to create a merged system.

Yet UNIX and its derivatives are limited in capabilities, with respect to research operating systems such as Mach and Amoeba. Berkeley UNIX added virtual memory and networking capability to the popular UNIX system; Mach and Amoeba add the capability of crafting truly distributed applications, those whose components exist on several dispersed computers across a network.

6. Enablers and Inhibitors

Many presentations in the workshop discussed both positive and negative influences to infusing workstation technology into NASA missions. We tagged these “enablers,” or those things that were positive influences, and “inhibitors,” or those that were negative. The major influences included a notion of NASA “culture,” technology transfer concerns, ever shifting yet unifying standards, and the NASA procurement process.

6.1. Culture

At the workshop, there was a fair amount of discussion about the NASA “culture” and its impact on infusing new workstation technology into missions. In all cases, it was considered an inhibitor; attitudes about advanced workstation and computer technology often blocked progress. An inherent factor in NASA culture is the length of project lifetimes. NASA engages in long-term projects. The Apollo mission was nearly ten years long from conception to first major success, and was built on the technology of the Mercury and Gemini projects. As a result of the length of typical missions, as one workshop presenter described it, there exists a “technology frustration gap” between the workstation technology frozen into the design and that which is currently possible. In Figure 3, the distance between the lines describes this gap, and the “frustration” is that of those engineers responsible for the workstation technology applied.

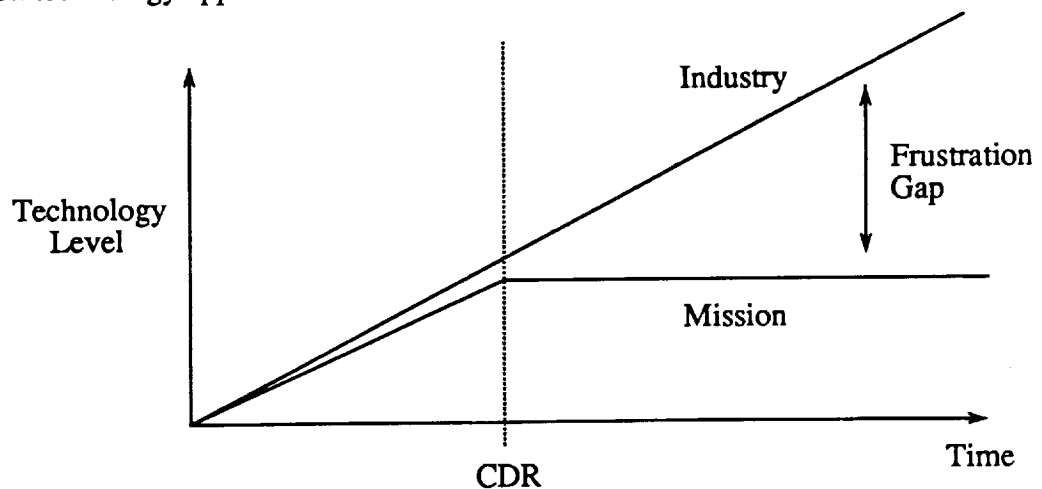


Figure 3. The Technology Frustration Gap

Another important inhibiting factor related to culture is the need to retrain personnel to use new technology. This retraining cuts across all levels of staffing -- from upper management to the engineers in operations. An effort must be made to demonstrate the benefit of advanced workstation technology before effective training can occur.

Yet another inhibiting factor is the multiple contractor approach to missions. Without standards defined for the type of workstation technology to be used, or if those standards are not sufficiently advanced as to make good use of the technology, the contractors are likely to invent their own standards, resulting in incompatible systems which exacerbates the attitude that advanced workstation technology is hard to learn and hard to use.

6.2. Standards

The diversity within NASA is vast. For technology to be shared efficiently among projects, adherence to standards is critical. However, the proper selection and use of standards is even more critical. Apropos to workstation technology, the following categories of standards are important:

- Operating system
- Programming language
- Windowing system
- Graphical user interface toolkit
- Graphics language
- Network interface
- Network protocols
- Application program interfaces

For each of these categories, there is more than a single standard available. There is little guidance from the standards themselves as to which to choose; each standard only defines itself so that two projects using the same standard will be able to exchange pieces of their work.

Standards are complex. To make intelligent choices about which standards to use requires the choosers to become experts in each of the categories listed above because an

advanced workstation application may need to choose one standard from *each* of those categories. The possible permutations of selected standards from those groups is enormous, but in order for two disjoint development efforts to enjoy the exchangeability that standards permit, the same set of standards must be chosen by each. The set of standards for a particular project is called an "application environment profile" [Isaa90].

Different domains (*e.g.*, scientific, commercial) may need different application environment profiles, but within a single domain (if narrowly defined), there may be a single set that best fits the applications. The profile must be complete, that is, all needed pieces are in place. The profile must also be coherent, that is, all pieces must interrelate (for example, if two standards are needed from a single category, they must both be able to operate with the same standards selected from the other categories).

Standards are created and selected as much by popularity as functionality. At the workshop there was a strong favoring of varieties of UNIX as the operating system of choice. Both an IEEE (1003.1) and FIPS publication exist describing a portable standard for UNIX, and the major workstation vendors are close to compliance. For windowing systems, the X Window System from Project Athena at MIT was the most popular, but no single GUI toolkit to use above the X Window System has emerged as the most popular in the community. Indeed, at least two presentations at the workshop were concerned with internally-developed toolkits.

6.3. Procurement

Because advanced workstation technology has potential benefit for almost everyone within NASA, the quantity of workstations that must be acquired are very high. However, procurement procedures can make it difficult to initiate volume purchases. Workstations have become, in many groups, a necessary piece of office equipment, as common as a telephone. Yet the procurement mechanism for workstations has lagged behind this change.

In fact, in many cases, purchases of ADP equipment are actively discouraged because of special regulations and procedures.

The typical cycle from design to market for new workstation products has shortened to six months in the industry. Hence, a protracted procurement procedure that begins with a capability specification, proceeding through competitive bidding and selection, and ending in delivery can easily result in the installation of a previous-generation workstation. If NASA is to stay abreast with the technology, procurement procedures for workstations must mirror the fast-moving industry development cycle for hardware and software.

Ames Research Center has instituted a procurement procedure entitled the "Interactive Systems" procurement, managed by Ames Code RC. The systems allows Ames branches to order workstations from any of three vendors by completing a Service Request form. This form requires about two weeks of internal processing before an order is placed with the vendor, and then approximately one week of post-delivery check-out is required before installation. This is a lightweight procedure for workstation acquisition and is representative of what is needed at all centers.

7. Recommendations

These are specific recommendations derived from the workshop. These recommendations are domain-specific in some cases, because, as previously presented, different domains have different demands on workstation technology. Specifically, the domains of workstation technology represented here are as follows:

- Mission operations consoles
 - Spacecraft health & safety monitoring
 - Payload status and control
- Numerical simulation visualization
 - Algorithm computation status
 - Results visualization
- Observational data visualization
 - Data integrity validation
 - Physical object rendering

7.1. Technology

Workstation technology is advancing very rapidly, faster than NASA's standard procurement mechanisms can track. Yet, in many cases, the technology is insufficient for our needs. Following the categories listed above, the key technology areas for NASA to track and advocate are as follows:

Mission Operations. Standards are critical. Application Environment Profiles describe a clustering of standards. To merely select a single standard, or even two, does not invoke the full benefit of standards. A complete suite must be selected in order to assure exchangeability and interoperability of operation systems. Also, virtual reality, in the form of telepresence, has great potential for remote operations.

Numerical Simulation. This is a two-fold arena; for those who have access to the NASA supercomputer facilities, such as the Numerical Aerodynamic Simulator at Ames, visualization hardware is critical. The Workstation Applications Office at Ames has defined a new standard for visualization systems, and requires continued support. Graphics power

is key for this class of individual. For those without supercomputer access, sufficient local processing power, in the form of superworkstations, is required for their work. Processor technology is key for this class of scientist.

Virtual reality can play a major role in the understanding of numerical simulation results. Display technology, and control technology for exploring a simulated space becomes critical.

Observational Systems. Though not limited to workstation systems, storage technology will be a critical factor for those engaged in exploring the data sent back by future observing systems. Though if the storage issue were to be solved, the ability to pull the mass of information needed into the workstation would continue to be critical. We must focus attention on workstation interfaces to high-speed networks. Yet each NASA center has a heavy investment in 10megabit/second Ethernet center-wide networks. These networks are insufficient for the needs of future workstations used in understanding new observational data.

7.2. Enablers and Inhibitors

The “culture” issue discussed above can be partially or wholly overcome by education. There was both a strong desire expressed at the workshop to repeat it, and a skepticism that even continued open forums will result in improved interactions among groups engaged in similar work.

Repeat the workshop. An annual workstation technology workshop and a regular reporting mechanism that cuts across all centers would maintain a level of information exchange that would enable less replication of effort by allowing all to build on the efforts of others.

Infuse modern workstation technology into all levels. By placing the technology into the hands of managers, those managers will better understand and accept the technology.

Export scientific visualization tools into a broader community. Those groups responsible for the state-of-the art work in scientific visualization should be chartered with finding ways to apply that technology to other arenas, such as mission operations and observational data visualization.

Support and publicize standards efforts. NASA does support national and international standards efforts, but there is no single reference source that projects using workstation technology can access to understand the agency's stand on standards.

8. Bibliography

- Acev85 Aceves, J. Joaquin Garcia Lunas and Andrew A. Poggio, "Guest Editor's Introduction: Computer Based Multimedia Communications," IEEE Computer, 18(10), October, 1985, p. 10.
- Ambl89 Ambler, Allen L., and Margret M. Burnett, "Influence of Visual Technology on the Evolution of Language Environments," IEEE Computer, 22(10), October, 1989, p. 9.
- Isaa90 Isaak, Jim, "Applications environment Profiles: A significant tool for simplifying and coordinating standards efforts," IEEE Computer, 23(2), February, 1990, p. 69.
- Lehr89 Lehr, Ted, Zary Segall, Dalibor F. Vrsalovic, Eddie Caplan, Alan L. Chung, and Charles E. Fineman, "Visual Performance Debugging," IEEE Computer, 22(10), October, 1989, p. 38
- Sari85 Sarin, Sunil, and Irene Greif, "Computer-Based Real-Time Conferencing Systems," IEEE Computer, 18(10), October, 1985, p. 33.
- Reyn85 Reynolds, Joyce K., Jonathan B. Postel, Alan R. Katz, Greg G. Finn, and Annette L. DeSchon, "The DARPA Experimental Multimedia Mail System," IEEE Computer, 18(10), October, 1985, p. 82
- Pogg85 Poggio, A., J. J. Garcia Luna Aceves, E. J. Craighill, D. Moran, L. Aquilar, D. Worthington, and J. Hight, "CCWS: A Computer-Based Multimedia Information System," IEEE Computer, 18(10), October, 1985, p. 92.

9. Agenda

Day One - March 13, 1990

- 8:30 AM** **Registration desk opens.**
Coffee and pastries served in the Ballroom lobby.
- 9:00 AM** **Welcoming and opening remarks.**
Barry Leiner, Peter Denning, and Robert Brown
- 9:15 AM** **Charge and goals for the workshop.**
Lee Holcomb
- 9:30 AM** **Keynote speaker.**
William Bricken, Chief Scientist, Human Interface Technology Center, University of Washington
- 10:10 AM** **Mission Operations Workstations.**
Jay Costenbader, GSFC, *Three Mission Operations Workstation Projects*
Jim Jeletic, GSFC, *Workstation Technology Used for Flight Dynamics Mission Support*
Mike Wiskerchen, Stanford University, *Application of Advanced Workstation Technology to Shuttle Operations*
- 12:00 PM** **Break for lunch.**
Location: *Ballroom C*
- 1:00 PM** **Application Development Environments**
Eric Hardy, CMU, *The SEI User Interface Project*
Randy Davis, University of Colorado, *Oasis: Present and Future*
Jay Costenbader, GFSC, *TAE+*
- 3:00 PM** **Coffee break.**
- 3:15 PM** **Andrew Potter, MSFC, *The SoftPanel Prototype***
Mission Science Workstations.
Tim Castellano, ARC/RIA, *SHOOT*
Patricia Liggett, JPL, *The VNESSA System*
- 5:00 PM** **Technical sessions end.**
- 6:30 PM** **Vendor show, poster and demonstration session.**
Hors d'oeuvres and no-host bar
Ballroom C & D
- 9:30 PM** **Day one ends.**

Day Two - March 14, 1990

- 8:00 AM** **Coffee and pastries.**
Ballroom lobby
- 8:30 AM** **Mission Operations (continued)**
Gaius Martin, JPL, *Sharp*
Tom Engler, MSFC, *MSFC Workstation Lab:SSF and AXF*
Gregory Blackburn, JSC, *SSF Multipurpose Applications Console (MPAC)*
- 10:30 AM** **Coffee break.**
- 10:45 PM** **Science Data Visualization.**
Eric Hibbard, ARC, *Code RCD Graphics*
Leo Blume, JPL, *Linkwinds -- A Prototype Scientific Visualiztation System*
- 12:00 PM** **Break for lunch.**
Location: *Ballroom A*
- 1:00 PM** Val Watson, ARC, *Workstation Applications Office Projects*
Stephen Coles, JPL, *EASE: An Engineering Analysis Subsystem Environment
for Real-Time Spacecraft Control*
- Productivity and Collaboration Tools.**
Keith Lantz, Consultant, *Multimedia Workstations for Collaboration*
- 3:00 PM** **Coffee break.**
- 3:15 PM** Barry Leiner, RIACS, *The National Collaboratory*
- Other Workstation Projects.**
Joe Hale, *MSFC, Workstation projects at MSFC.*
Michael McGreevy, ARC, *Ames Virtual Reality Projects*
- 5:00 PM** **Workshop wrap-up.**
Robert Brown
- 5:15 PM** **Break for the day.**