CHARLES E. CATLETT

Manager, Networking Development
National Center for Supercomputing Applications (NCSA)
University of Illinois at Urbana-Champaign

and

JEFFREY A. TERSTRIEP

Project Lead, Networking Development
National Center for Supercomputing Applications (NCSA)
University of Illinois at Urbana-Champaign

The Use and Effect of Multimedia Digital Libraries in a National Network

ABSTRACT

The Internet has supported information archives for some time. These archives have traditionally allowed users to retrieve text and image data as well as software to their own computers for examination. As the Internet grows in scale and in performance and services, more sophisticated information archives and access modes are possible. This paper reviews the growth of the Internet with its current information archive services and proposes methods for providing interactive access to multimedia data. Various information types and their access modes are discussed in terms of their role in defining advanced digital library and network services. A prototype digital library system and user interface developed at the National Center for Supercomputing Applications is examined.

BACKGROUND: THE GROWTH OF THE INTERNET

The term *internet* means a network of networks. Our national network today is composed of a number of national backbone networks (e.g., NSFNET, ESNet), mid-level (regional, consortium) networks, and campus networks. "The Internet" is a network of networks that includes our national network as well as other connected networks in many countries throughout the world. The common thread among all Internet components is that they operate based on the same network protocols and share a common addressing scheme, message forwarding (or "routing") schemes, etc.

In looking at the growth of the Internet, it is helpful to look closely at a major component of the Internet, the NSFNET. In 1986, the National Science Foundation (NSF) established the NSFNET to interconnect six supercomputer centers at 56 kilobits per second (kb/s). At each backbone node, mid-level networks were established. The NSFNET architecture consists of a *backbone* network, *mid-level* networks to extend the backbone connectivity to institutions, and *campus* networks to extend the mid-level connectivity to individual local area networks (LANs). By 1988, these mid-level networks were providing backbone access to over 500 individual sites. In late 1988, the backbone was expanded to thirteen nodes, and the links were upgraded to 1.5 megabits per second (Mb/s). By the fall of 1990, the network had grown to over 2,000 sites, and the backbone was again upgraded to sixteen nodes interconnected at 45 Mb/s.

During the past several years, a number of international links have been established as well, including extensive connectivity to Europe and the Pacific Rim. Campus networks have matured, providing access to many more individual computers so that now the Internet connects over 300,000 individual computers. Figure 1 shows the rapid growth of the Internet. Hosts on the network are shown from the original, centralized host registration at the Network Information Center (NIC) as well as the current decentralized registration system called the Domain Name System (DNS). "NSFNET Backbone Traffic" refers to the number of data packets that are transported across the NSFNET backbone network monthly. Note in particular the growth in foreign (non-U.S.) networks connected, the number of individual hosts, and the growth in the amount of data being passed over the NSFNET backbone (Smarr & Catlett, in press).

It is not at all clear where or when this growth will level out. A large emphasis is being seen now in connecting K-12 institutions to the network, and the various mid-level networks are beginning to concentrate on marketing the network to a number of sectors including industry and education.

"DIGITAL LIBRARIES" ON THE INTERNET TODAY

A number of academic library catalog search systems are accessible from the Internet today. CICNet interconnects Big Ten universities and several others. The library catalog search facilities of the libraries of most of these universities are accessible from the network. These services deliver only information about where data exist but do not provide access to the actual data.

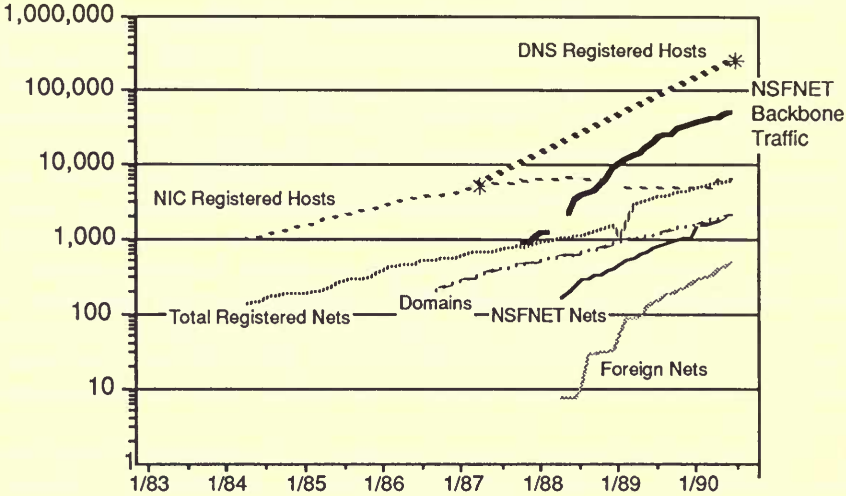

Figure 1. Internet growth indicators (©1991 Catlett, Terstriep)

Internet services that provide access to information are currently limited to archives reachable via file transfer and bulletin boards. An advantage to these services is that they are globally accessible, and the fact that they are heavily used in spite of their shortcomings indicates the demand for these services, which provide access to text, software, and images. Images and software as well as large text files are generally compressed, and the user must decompress them (once they have been retrieved) before using them. This makes the file opaque from the user's point of view. The user must also have access to the facilities required to run any of the available software.

The archives that allow file transfer have a number of limitations. First, there are no universal indexing or naming conventions; therefore, it is difficult to locate items. Second, an item must be retrieved in full before it can be examined beyond reading the file name. Given that file names do not provide any significant indication of the contents of the item, users do not have a viable means to screen items before transferring them over the network to examine them. Third, there are few directories telling users where to find these archives or what is contained in any of them. These archives, then, are useful if one knows what one needs, where it is, and how it is stored.

Bulletin boards are slightly better in that the information is organized in categories. There are several hundred categories generally accessible. However, these categories consist mainly of computer science and popular culture topics. Within a category, there is no index of contents beyond the title (subject) of the entry, the author, and the date. There are a number of software packages available to read these bulletin boards. Many are difficult to use, requiring extensive experience even for simple filtering of information such as searching for entries with a particular subject. Some of the recent access packages that run on Apple Macintoshes, however, provide a fairly straightforward user interface.

Most computer centers operate a storage archive for their users. These large private collections are not generally accessible for general users. Those that provide general access use anonymous file transfer protocol (ftp) as described above. These archives cannot be ignored, if for no other reason than their size and growth rate. A typical major computing center's storage archive can grow at between 50 and 200 Gbytes per month.

Large-scale projects that are underway and require more sophisticated digital access include the Human Genome Project, the Hubble Space Telescope, the Earth Observing System (EOS), and the BIMA (Berkeley-Illinois-Maryland Array) radio astronomy imaging consortium. In the case of the Hubble Space Telescope and the EOS, it is estimated that up to 1 terabyte per day will be collected.

CONTENTS OF A MULTIMEDIA DIGITAL LIBRARY

The scientific multimedia digital library will contain a variety of information types. Table 1a shows the approximate size of various types of data. This table includes individual items, such as images or journal articles, and their average sizes. From the size, one can calculate the network throughput required to retrieve them in a fixed amount of time. Sequences of data such as image sequences, audio, etc., are shown in Table 1b along with the approximate network throughput required to transmit them.

In addition to standard types of data such as those shown in Tables 1a and 1b, the scientific multimedia digital library also contains data sets generated by applications such as numerical models. Direct access to data sets would allow scientists to verify the conclusions of their colleagues by examining the data firsthand. Also, for such applications as global climate models that require hundreds of hours of supercomputer time to run, a number of users will want to "mine" or explore the data.

TABLE 1A
VARIOUS ITEMS (AND THEIR SIZES)
FOUND IN A MULTIMEDIA DIGITAL LIBRARY*

Item	*Size Calculation*	*Size in Bits*
Journal articles, papers, etc. (avg. 5 pages)		
Plain text	5 kbyte/pg. × 5 pg. avg.	25 kbytes
Formatted text	10% overhead to text = 5.5 kbyte/pg. × 5 pg. avg.	28 kbytes
Scanned page images	300 dpi × 7.5 in. × 10 in. = 1 Mbyte/pg. × 5 pg. avg.	5 Mbytes
Single images		
Color NTSC	512 × 512 × 8 bits	.26 Mbytes
G4 FAX	1.7 k × 2.2 kbits	.5 Mbytes
Gray-scale	2 k × 2 k × 8 bits	4 Mbytes
Color	2 k × 2 k × 24 bits	12 Mbytes

*Calculations on the average size of each item are shown as well as the size.

TABLE 1B
SEQUENCES OF DATA AND APPROXIMATE NETWORK
THROUGHPUT TO TRANSMIT

Sequences	*Heading*	*Required Throughput*
Audio		
Low fidelity	Sampling rate	.064 Mb/s
High fidelity	Sampling rate	.64 Mb/s
High-definition TV*		
Production quality	Minimal compression, 30 frames/second	1-2 Gb/s
Post-production quality	Modest compression	200 Mb/s
Distribution quality	Compression with information loss	20 Mb/s
NTSC quality	Compression with information and visual loss	5 Mb/s
VCR quality	Compression with significant information, visual loss	1.5 Mb/s

*HDTV compression rates are from Glenn Reitmeier, Director, High-Definition Imaging and Computing Laboratory, David Sarnoff Research Center.

The scientist will want to examine the data in a number of ways, including extracting portions of information at the byte level. It is critical that the data be stored in such a way that their format and contents can be later ascertained. There are a number of data file formats that are generally used in the computation science community. These formats incorporate a standard header describing the contents of the data file as well as access software for reading, writing, and interpreting the headers. Self-describing data formats might also contain references to data analysis software or perhaps copies of appropriate access and analysis subroutine object code and source code.

The multimedia digital library might also store programs that generate the data rather than the actual data. For example, periodic complete state information (checkpoints) of a long global climate model might be more convenient to examine than the multiple terabytes of data that the model could generate. In this case, the user will generate the data "on-the-fly" by starting up the model at some point in the model's cycle.

Images and sequences of images will be stored in the multimedia digital library as shown in Tables 1a and 1b. Note that the output from a data generator application could also be a sequence of images such as this. Scientists require at least distribution quality imagery for serious examination, although lower quality may suffice for cursory examination or observation of large-scale phenomena (e.g., weather patterns in a climate model).

THE SCIENTIFIC DATA MANAGEMENT FACILITY

A Prototype Multimedia Digital Library

The National Center for Supercomputing Applications (NCSA) is developing multimedia digital library services for a number of projects, including the implementation of a central archive for the BIMA project, storing scores of data sets and images collected by the Hat Creek millimeter array radio telescope. The intent of much of this work is to explore the provision of interactive access to the types of objects that a scientist would find useful in a multimedia digital library.

A prototype has been designed based on several fundamental components of a multimedia digital library aimed at providing access to information used by computational scientists. The data involve multiple formats and media types. The data will be distributed, will in many cases be pre-existing, and thus will have a set format and storage type and must be accessed in that way.

Two major components make up the digital library: directory services and data access. The digital library can be accessed using a variety of applications, including, for example, user interfaces with browsing and examining capabilities and data analysis packages to examine data.

An indexing system or directory service is needed to provide a catalog of location and, preferably, format/type information for the distributed data archives. This function is essentially a database with information about the location of data items, the type of data, and the format of the data. This component is a database.

A mechanism for locating data is needed to access the digital library; this will query the directory database. Mechanisms for browsing data and for examining data are necessary. The mechanisms will differ for the various data formats and media types. In the case of data generators (programs) in the digital library, index entries include information about where the data generator will be executed.

Figure 2 illustrates the functional components of the digital library as implemented in a prototype that was demonstrated during March 1991 at National NET'91 in Washington, DC. This includes the user interface with browsing and examination applications as well as the directory and data archive components. Figure 2 also includes multiple archives with multiple item types, including data generators and the use of data analysis filters.

The scientific digital library prototype has several indexes and several data archives, and some indexes reference multiple archives. The user interface sends queries to one or more indexes. The queries result in lists of relevant items sent to the user, each with one-line description, author, creation date, data type, and a pointer to the actual location of the item on the network.

Depending on the item type, the user is given a choice of examination/browsing options. For example, text can be examined with a text editor, and scientific data sets can be examined using a number of data analysis tools. When the user chooses one of the tools, the interface automatically starts up the analysis tool for the user and informs the tool of the location of the data set. In the case of data generators, a choice of data analysis tools is given for use as the user interface for control and viewing of the process. Users can also elect to transfer a copy of the item to local disk; however, many examine options involve use of the item at its original location. The list of items can also be saved to local disk.

NETWORK ARCHITECTURES, PROTOCOLS, AND MULTIMEDIA DIGITAL LIBRARIES

The nature of access to the data in a multimedia digital library greatly affects the network architectures and protocols required. At the same time, network architecture and also resource billing schemes will determine what are the most cost-effective access methods and thus will affect the way users access data.

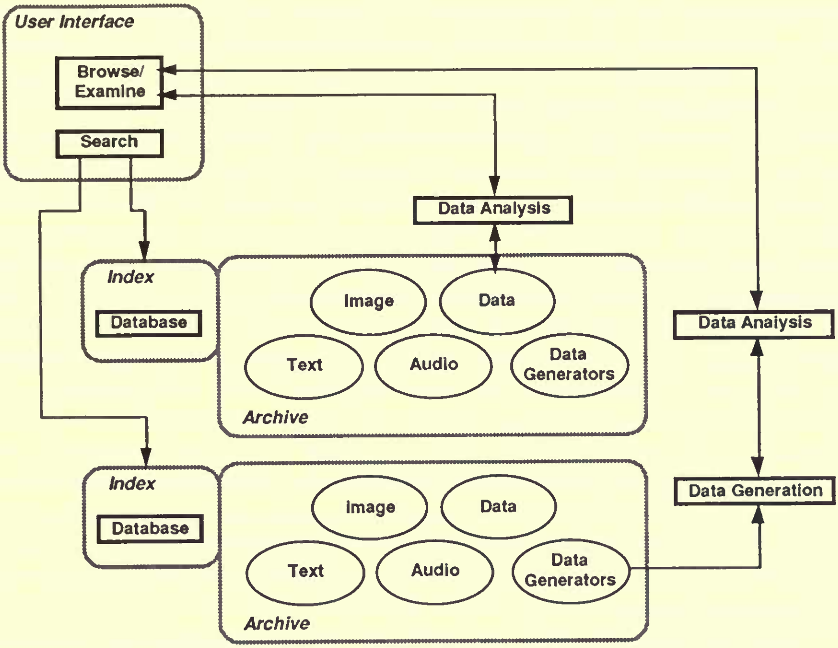

Figure 2. Model used to build the NCSA multimedia digital library prototype (© 1991 Catlett, Terstriep)

For access to remote data over the network, existing Internet archives allow for limited browsing and examination of text (using bulletin boards) but only for retrieving most data (images, large text files, software) to local disk to examine or browse. The multimedia digital library must support remote access to data as many of the items will be too large for the amount of disk available locally to the user and many of the items will require access to facilities (e.g., supercomputers) that do not reside with the user.

To access data remotely without retrieving the data to local disk in total, a number of types of access can be supported. The popular Network File System (NFS) protocol from Sun Microsystems allows files on remote disk to appear as though they are on local disk. This is done by allowing the user to access the files in subunits (transparent to the user) called blocks. However, if the user only requires a particular record or byte (perhaps every *n*th byte), NFS still sends the data in blocks, and the subunits are extracted from the blocks by the user's workstation.

In some cases, it may be more suitable to allow the user's application to extract information from the remote file by individual bytes rather than blocks. In other cases, a user will want a whole group of items (e.g., files) to be sent to the local workstation so that they can be manipulated locally.

Each of these remote access methods—groups of files, whole files, blocks, bytes—dictates requirements on the network supporting the access. For example, where text or images are being browsed, the user should be able to "flip" through several pages per second. This may be most easily done by retrieving in total the entire file or group of files to be browsed. For image sequences, however, the entire image sequence may be too large to fit on the user's disk and must be accessed in parts. Images could be sent over the network in sequence, requiring isochronous, in-order delivery. In this case, lost or damaged images or portions of images will be better skipped than retransmitted. To provide isochronous viewing, groups of images could be sent over the network and played locally, with a number of images waiting to be viewed at the local workstation at any given time. This would allow the workstation to deliver the images at a constant rate, relaxing the requirement of the network to do so. In addition, a queue of images at the local workstation may provide enough time to retransmit lost or damaged images.

It is important to note that delivery of information at a constant rate in the presence of any errors will come at a cost of reliable delivery of all information. This is because error correction measures such as retransmission come at a cost of delays, and these delays may cause more disruption in the image stream than the error they are meant to correct in the first place.

When viewing a sequence of images, part of the information that is contained is the development of features in time. Thus delivery of the information at a rate that distorts the time element will deliver incorrect information to the user.

The billing algorithms of the network will also drive the way that the data are accessed. An example of this effect can be seen in the delivery of electronic mail. When mail is delivered using dial-up phone circuits such as with UUCP (UNIX-to-UNIX copy protocol), the cost of delivery is dependent upon a circuit setup cost and a time-sensitive usage cost. Therefore, to minimize the number of calls and the length of calls, electronic mail is queued and sent periodically. Internet mail delivery systems such as SMTP (simple mail transport protocol) assume that there is no cost to setting up circuits or in usage. Therefore, to minimize delay in mail delivery, each mail message is delivered as soon as it is submitted. The result is that Internet mail is much more interactive than UUCP because of the cost structure of the underlying network services rather than because of any technical considerations.

The current Internet cost structure is a fixed cost, not sensitive to usage. The fixed cost generally involves the cost of equipment at installation time, the cost of leasing telecommunications circuits, and some cost for maintaining an operations staff locally and/or at a central network operations center. In this environment, a multimedia digital library might download small items for local examination and access large items remotely. The difference between large and small will be determined by the capacity of the network, the amount of local storage space, and the amount of time the user is willing to wait while information is retrieved. With time-sensitive network connections such as a circuit-switched (dial-up, ISDN) connection, the trade-off will also include the cost of keeping the circuit up for large retrieval and the cost of keeping the circuit up for long sessions of remote data examination.

DATA GENERATORS: HIGH-PERFORMANCE APPLICATIONS

Several high-performance applications are described below. These applications have intensive network requirements. The multimedia digital library prototype described above allows users to access these types of applications; therefore, they must be taken into account in assessing the effect of multimedia digital libraries on a network.

Radio Astronomy

The Hat Creek radio telescope collects information at 2,048 frequencies. The telescope data must be converted into visual images using computational image-processing techniques. Supercomputers are used for this, acting as the image-forming element of the telescope. The conversion involves a calibration calculation to filter out much of the interference caused by atmospheric anomalies, then a FFT (Fast Fourier Transform) to convert the raw telescope output data into images. For each frequency, a two-dimensional image is produced. Thus the image output to a radio telescope is a spectral cube, with two spatial dimensions and one spectral dimension. In the case of the Hat Creek array, this cube is 2,048 frequencies by up to 4,096 horizontal and 4,096 vertical pixels with each pixel being 16 to 24 bits. For example, a 2,048 by 2,048 spatial size with 24 bits per pixel would involve the following amount of data:

2,048 frequencies • 2,048×2,048 pixels/frequency • 24 bits/pixel • 1 byte/8 bits = ~26 gigabytes.

Reconstruction of these image cubes from the data requires real-time interaction by a scientist who observes roughly two to five images per second being displayed. Typically, one spectral image is used in this process, and the scientist will watch the image reconstruction as the nonlinear deconvolution either converges on an image or begins to diverge, indicating the need to stop the process and restart after adjusting gain parameters. Analysis of the resulting spectral cube involves traversing both the spatial dimensions and the spectral dimensions.

Data are collected continuously by the telescope and are integrated over time to increase the signal-to-noise ratio. The integration time, generally measured in tenths of seconds to tens of seconds, is determined as a function of the signal strength of the object being observed. Where very long integration times are used, the telescope would not necessarily be steered in real time. Short integration times are generally desired for real-time observation.

Several classes of observation require these images to be produced in real time for interactive steering of the telescope. The integration time would be on the order of tenths of seconds to several seconds. These classes include observation of time-variable phenomena such as solar activity, a technique called "mosaic-ing" (where short observations are made on a number of small regions and then reconstructed into a larger image later), and in cases where the atmospheric changes, which happen on the order of seconds, are kept to a minimum.

A prototype that was demonstrated by BIMA scientists at NCSA recently involved the functional decomposition of this type of system, using both the CRAY Y-MP for the baseline calculation and the massively parallel CM-2 for the FFT. By spreading the computation across several supercomputers, the speed of the computation increased significantly.

Atmospheric Sciences

Interactive visualization systems involve both analysis of precomputed data and analysis of running simulations. For the analysis of precomputed data, a supercomputer is used to render thunderstorm data using surfaces, contour plots, massive particle releases, and slices. The supercomputer simulation involves calculating the evolution of a weather system for a region of the atmosphere. For example, a region that is 100 km long by 50 km wide and 30 km high is subdivided into a grid of zones, each zone perhaps 1 km by 1 km by 500 meters. Several variables are associated with each of these zones, including temperature, pressure, and velocity vectors. The supercomputer simulation involves using the laws of physics to compute the evolution of these variables over a period of time from some beginning state. A typical simulation as described above has over 1 million zones, each with nine variables. The variables are stored in 8-byte fields, thus the amount of data required to represent one moment in the storm evolution is

$\sim$1,000,000 zones • 9 variables/zone • 8 bytes/variable = 72 megabytes.

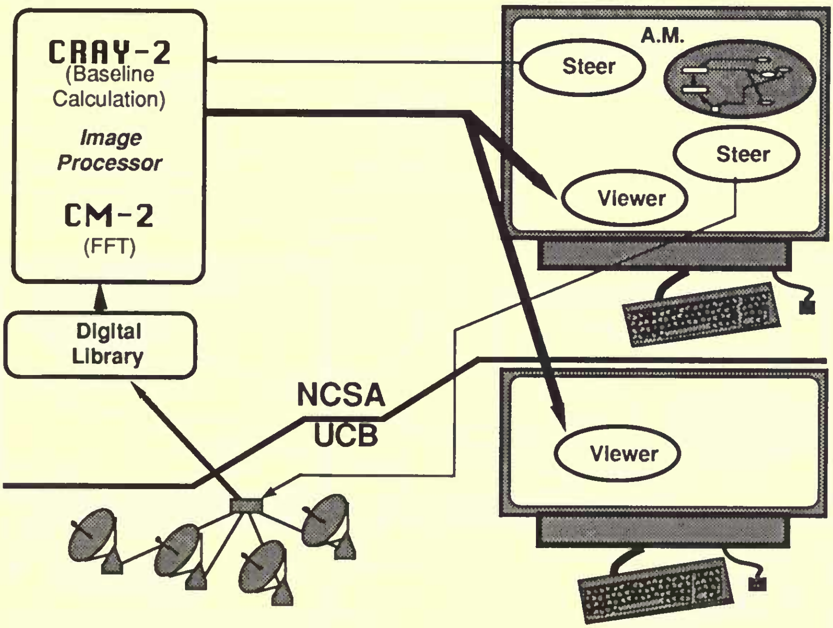

Figure 3. Prototype radio astronomy application that was developed using both the CRAY-2 and the CM-2 to interactively evaluate the frequency response of a given radio telescope antenna configuration

The resolution shown above, with each zone representing an area 1 km by 1 km by 500 meters, is not high enough to study small-scale phenomena such as tornados. In order to reduce the zone size for this scale of activity, the number of zones would increase beyond the capacity of any available supercomputer memory, and the amount of time it would take to calculate all of the variables for even a single moment in time would exceed the compute power of even the fastest current supercomputers. To address the need for higher resolution, interactive systems are being developed to allow the scientist to intervene in the running simulation and request a higher or lower resolution in portions of the simulated storm system. This will allow for increasing resolution in those areas with high activity without increasing the overall size of the simulation beyond feasible limits.

These types of applications are also being distributed over multiple computers to increase the computation rate. A current project involves the use of multiple RS/6000 workstations at NCSA to compute the model. Early studies have yielded a three-fold decrease in turnaround time when comparing one RS/6000 to using six in parallel, even for relatively small model sizes.

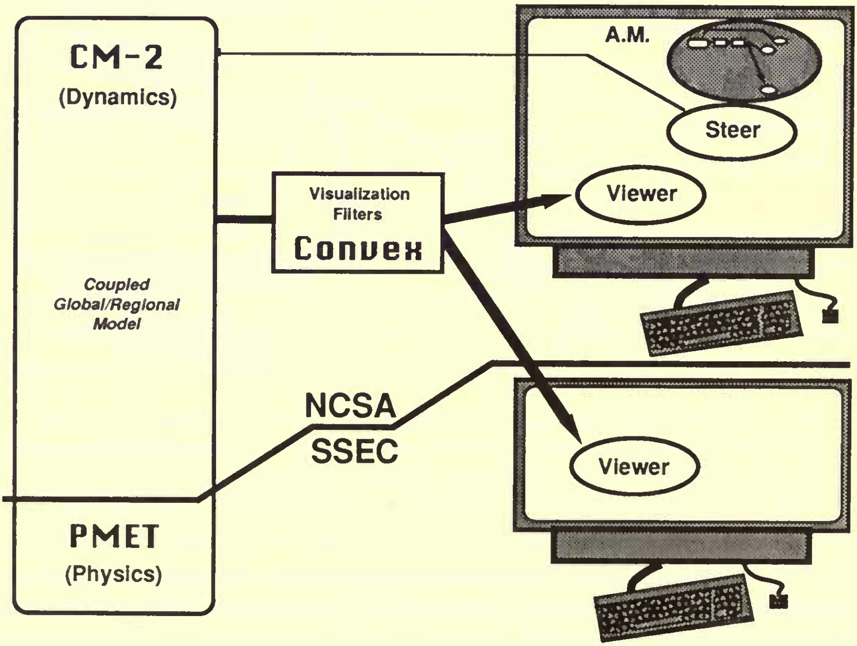

Figure 4. Coupled global/regional climate model using the PMET at Wisconsin and the CM-2 at Illinois

Biomedical Imaging

The Distributed Biomedical Imaging Laboratory (DBIL) is a testbed to integrate imaging instrumentation used in biomedical research with remote high-performance computing environments. DBIL is a joint project between NCSA and the University of Illinois Biomedical Magnetic Resonance Laboratory, with remote collaborators at Lawrence Berkeley Laboratories.

One application that has been demonstrated is the use of a CRAY-2 and CM-2 for 3D image reconstruction simultaneously during data acquisition from a nuclear magnetic resonance imaging spectrometer.

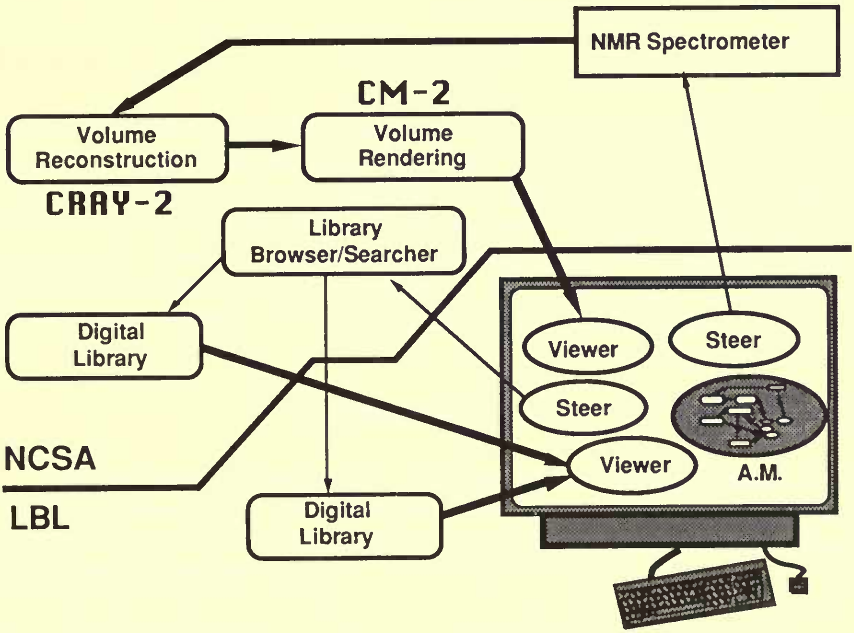

Figure 5. Distributed Biomedical Imaging Laboratory application

The reconstructed volumetric image is then sent to a CM-2 for volume analysis and visualization in a distributed display environment. The CRAY-2 image reconstruction takes roughly 0.05 seconds for each projection, and after one to one hundred projections are calculated, the volume of data (up to several megabytes) is sent to the CM-2. The processes run continuously for the duration of the experiment (10 to 60 minutes).

This original system allowed 3D rendering of a static image—a 3D snapshot. The system that has recently been demonstrated allows 3D rendering of a dynamic image—a 3D movie. Using this system, a frog egg will be observed over a 24-hour period to yield a 3D movie of cell multiplication.

REFERENCE

Smarr, L. L., & Catlett, C. E. (in press). Life after Internet: Making room for new applications. In B. Kahin (Ed.), *Building an information infrastructure.* New York: McGraw-Hill.