

Toward on-line, worldwide access to Vatican Library materials

by F. C. Mintzer
L. E. Boyle
A. N. Cazes
B. S. Christian
S. C. Cox
F. P. Giordano
H. M. Gladney
J. C. Lee
M. L. Kelmanson
A. C. Lirani
K. A. Magerlein
A. M. B. Pavani
F. Schiattarella

The Vatican Library is an extraordinary repository of rare books and manuscripts. Among its 150000 manuscripts are early copies of works by Aristotle, Dante, Euclid, Homer, and Virgil. Yet today access to the Library is limited. Because of the time and cost required to travel to Rome, only some 2000 scholars can afford to visit the Library each year. Through the Vatican Library Project, we are exploring the practicality of providing digital library services that extend access to portions of the Library's collections to scholars worldwide, as an early example of providing digital library services that extend and complement traditional library services. A core goal of the project is to provide access via the Internet to some of the Library's most valuable manuscripts, printed books, and other sources to a scholarly community around the world. A multinational, multidisciplinary team is addressing the technical challenges raised by that goal, including

- Development of a multiserver system suitable for providing information to scholars worldwide.
- Capture of images of the materials with

faithful color and sufficient detail to support scholarly study.

- Protection of the on-line materials, especially images, from misappropriation.
- Development of tools to enable scholars to locate desired materials.
- Development of tools to enable scholars to scrutinize images of manuscripts.

In this paper, we provide an overview of the project, a description of the system being developed to satisfy its needs, and a discussion of how the technical challenges are being addressed.

Introduction

Traditional libraries are, in their essence, services that manage massive quantities of data. They accomplish four primary information-management services for their clients: collection; organization; access (permission to inspect) and retrieval; and analysis, synthesis, and dissemination of information. Librarians and information scientists have developed techniques, procedures, and systems for addressing each of these functions for many kinds of data and presentation. Digital libraries use different methods to accomplish the same things as conventional libraries,

©Copyright 1996 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

0018-8646/96/\$5.00 © 1996 IBM

exploiting digital storage, processing, and communications to enable management of very large numbers of items, searches that would be impractical manually, rapid distribution to or retrieval from afar, and excellent information protection. Like conventional libraries, digital libraries manage massive amounts of information of many media types. However, while digital library services are fundamentally similar to conventional library services, their quantitative characteristics are so different that they allow qualitatively new services to be provided by a library to give the clients quantitatively new abilities. Fuller descriptions of the digital library paradigm, its beginnings, and the breadth of its applications have been thoroughly communicated by Gladney et al. [1], and they are not repeated here.

Since the term *digital library* implies a massive amount of managed information, one could argue that the world has still not truly seen a digital library; however, there have been many prior projects that have contributed to our knowledge of digital libraries. The references in [1] list much prior work. Nevertheless, we would be remiss if we did not mention earlier IBM projects that contributed to the experience of the team working on this project; these include projects with the Indies Archives (Spain) [2, 3], the American painter Andrew Wyeth [4–6], the National Museum of Ethnology in Osaka (Japan) [7–10], and the National Gallery of Art (U.S.A.) [11].

By late 1993, the digital library paradigm had changed from being a relatively obscure topic of interest to a few librarians and computer scientists to being a significant point of interest to every research-university library, to nearly every major library in the United States, and to a gradually increasing number of similar institutions in Europe, the Far East, and elsewhere. Scholarly and public enthusiasm was fanned by the interest of the Group-of-Seven governments in national information infrastructure (NII) programs and by public-press hyperbole about the information superhighway.

On the one hand, the popular expectation that digital libraries will soon be massively deployed seems realistic, because few truly fundamental problems block what is hoped for. On the other hand, this expectation may be unrealistic, because significant engineering challenges remain, and a significant deployment of service will depend on large changes to the infrastructure of institutions that collect, hold, and disseminate information. When one considers all of these factors, with realistic estimates of when the known problems might be solved and of how quickly the infrastructure can be made to evolve, it seems likely that digital library service of significant scale will be accessible to universities in about five years, and to the general public in about ten years.

One class of problems to be overcome is the creation of a significant corpus of valuable digital information,

through the digitization of sufficiently large collections of retrospective material, and through the capture of prospective material from digital source material before it is discarded (after conversion to more traditional media). Particularly for existing corpora, the specifics to be addressed depend on the nature of the materials to be converted; what is needed is different for 19th-century audio recordings, for 20th-century scientific journals, and for works of fine art. In this paper, we focus on manuscripts of the 9th through 18th centuries. The source materials used in this study are among the rarest, most valuable, and most beautiful manuscripts ever created. They represent an incredible diversity of base materials (e.g., paper, parchment, deerskin), coloring materials, sizes, and shapes. Many are bound into volumes. Some are quite fragile. Capturing and preserving their content and beauty is a challenge to our scanning, image processing, and display technologies. These challenges and their solutions are discussed in the section on imaging needs of the Vatican Library system.

Another class of challenges is inherent in the means provided for finding, acquiring, and presenting copies of the documents, pictures, videos, and audio material of interest to the end users. How the materials are to be presented to end users depends not only on the subject matter and its desired appearance, but also on the objectives of the end users and the resources they can bring to bear. For example, presentation of geographic maps to a civil engineer engaged in highway renewal or city planning must be substantially completed within a second if the engineer has to inspect a blowup of a small portion of a map to search for anomalies of immediate interest, but the presentation may be allowed to take many minutes for a full map that might be used for planning future work. For a historian who today must travel long distances and can afford to do so only once a year, having documents delivered within 24 hours can effect major improvements in the quality of his or her work.

The system that we describe in this paper was designed and implemented to meet the needs of the Vatican Library and a community of users desiring access, at a distance, to Vatican Library materials. The system requirements, to a great extent, were determined by interviewing representatives of the Vatican Library and the user community; however, many of the system requirements are not unique, but typical of applications that are described as “digital library” applications.

Although the digital library is an exciting new paradigm, it is relatively unexplored, many questions remaining unanswered. One way to explore these questions is to build an on-line digital library system that meets the real needs of a specific user community, and to use experiences gained in using that system to examine some of those questions. The overall goals of the Vatican Library project

involve exploring many such questions; however, we restrict ourselves in this paper to discussing issues that are directly related to the implementation of the digital library system being developed. The requirements and the implementation of that system are described in the following section.

The Vatican Library system

The inspiration for the Vatican Library system came from the Latin American scholarly community. In Latin America, there are many scholars who desire access to Vatican Library materials for their artistic, historic, theological, and scientific importance, yet their access to these materials is currently quite limited because of the cost of travel to Rome. The desires of the scholarly community were best expressed and ably advocated by faculty members of the Pontifical Catholic University of Rio de Janeiro (PUC-Rio). In brief, the scholarly community desired access to Vatican Library materials via the Internet.

The choice of the Internet as the network of access deserves some small mention. The Internet is ubiquitous in North America and accessible from most places in the world, especially from universities. Although accessible from relatively fewer locations in Latin America than in the U.S., it enjoys a good reputation. The choice was obvious.

Fortunately for the project, the Vatican Library views its mission as providing access to its collections to the worldwide scholarly community. The Library was not only interested in, but also enthusiastic about, using new technology to fulfill its mission. Yet there were concerns that manuscripts might be damaged during image capture. These concerns are addressed as a project requirement, described below.

The targeted user community was broadened, beyond the community of Latin American scholars, to include scholars worldwide. To address this worldwide challenge, a global partnership was formed by the Vatican Library, IBM, and PUC-Rio. Other collaborators have since joined the project team, including technical personnel at Case Western Reserve University.

It was decided to structure the project as a sequence of distinct phases, each having clear technical goals. Development of a fully functional system was one of the goals of the first phase. The Scholars Advisory Committee was formed to advise the project's technical team on the needs of scholars interested in Vatican Library materials. Ten scholars, all of whom were familiar with Vatican Library manuscripts, were chosen to compose this committee; they were selected for their academic distinction, diversity of interests, and geographic diversity. A technical advisory board was formed to guide the technical directions of the project.

Requirements of the Vatican Library system

A number of scholars on the Scholars Advisory Committee were polled by the Worldwide Workflow Consulting Practice, a consulting group with expertise in image systems that is part of IBM's Industry Solutions Organization, to determine their requirements for the system's operation. This led to requirements that the system

- Provide access to "cataloging information" describing Vatican Library materials.
- Provide access to high-quality images of Vatican Library materials.
- Provide scholars with access to this information through the Internet.
- Provide scholars with timely response.
- Provide the information in the most widely used data formats, so that scholars with diverse hardware and software would be able to utilize this information.
- Enable humanities scholars with modest computer literacy to find and use desired materials.

The Vatican Library further required that the system

- Safely capture images of the materials.
- Permit inspection of digitized materials at the Vatican Library.
- Protect the Library's intellectual property rights (especially for reproduction) to the digitized materials.

We believe that these requirements, validated through this project, are typical of what many libraries desire. The primary differentiator for this library was the nature of the source materials: ancient, varied, valuable, and sometimes fragile, illuminated manuscripts.

Of these nine requirements, the final one has had the most profound impact on the system design. Because of this requirement, the "visible digital watermarking" technique, described later in this paper, was developed in order to unmistakably identify images as Vatican Library property and discourage their misappropriation, without diminishing their utility for scholarly study. Because of this requirement, it was decided that unmarked, uncompressed, high-resolution images would not be made available to the Internet or to the servers attached to the Internet, lest they be misappropriated. Consequently, two physically separate systems were implemented in Rio to provide images for local access and remote access. The local-access system provides higher-resolution, unmarked, compressed and uncompressed images to its users. The remote-access (Internet server) system provides lower-resolution, watermarked, lossy-compressed (inexact replicas) images to its users.

On the basis of the nine requirements, a system architecture was devised and the phase-one project was defined. For this phase, six scholars were invited to participate, and it was decided that 20000 images of Vatican manuscripts would be captured, processed, and made available to them through the Internet. This number of images to be scanned was chosen as a compromise between the needs of the user community and the capabilities of the technology (mainly preexisting IBM technology). It was felt that the needs of the project dictated that the system enable meaningful research, and it was correctly conjectured that the scholarly community would require access to entire books, rather than selected individual pages, in order to conduct meaningful research. If we estimate that a typical manuscript contains 400 pages, that we are supporting six scholars with diverse interests, and that each scholar is provided with eight complete manuscripts, we arrive at a total of approximately 20000 single-page images. We initially estimated that 100 images per day could be scanned; this led us to expect that 20000 could be scanned within a year. For the sake of simplicity, we desired to store the entire collection of locally accessible images, in both compressed and uncompressed formats, on a single (40-gigabyte) optical library, and our initial estimates of a compression factor of 10 led us to believe that this was feasible. While our estimates of the compression that could be achieved were imprecise, as described below, they were sufficiently accurate to establish the scope of the project.

● *Overview of the Vatican Library system*

The system designed to satisfy the project's needs includes the following:

- A subsystem, located at the Vatican Library, that is used to scan Vatican Library manuscripts and enter the cataloging information that describes them. This subsystem was located at the Vatican Library to safeguard the manuscripts during the scanning process, as required; shipping them elsewhere to be scanned would jeopardize their safety. This subsystem also provides image storage and display functions that enable scholars at the Vatican Library to inspect the digitized images, as required.
- A subsystem, located at PUC-Rio, that provides access to the images and cataloging information via the Internet. This subsystem also provides additional image display functions to support the needs of local scholars.
- A subsystem, located at the IBM Thomas J. Watson Research Center, in Hawthorne, NY, that is used to examine the scanned images, to archive them, and to replicate and ship them to PUC-Rio for entry into the subsystem at Rio. Routine examination of the scanned images is done to ensure that the scanning subsystem is

operating correctly and the scanned images are of high quality, as required.

While the Internet provides access to the system from many places, the time required to download an image varies. At the beginning of the project, the bandwidth of the U.S.-Brazil Internet link was 64 kilobits per second. Downloading a typical 250-kilobyte image through this link would consume at least 30 seconds in the presence of no other Internet traffic; otherwise it could take several minutes. Therefore, regardless of which side of the link the Internet server occupies, service will be degraded for users on the other side. To provide timely response to the system's U.S.-based and European users, additional servers are also needed in the U.S. and in Europe.

It was envisioned that all of the participating scholars would access the information through the Internet, and that they would access the information from workstations with a broad diversity of capabilities. Some scholars might have excellent image-display capabilities, some might have limited image-display capabilities, and some might have none at all. To enable a subset of the scholars to examine high-quality images of the manuscripts, we developed the *Scholar's Interface Application (SIA)* program for the project. This program enables user-scholars to locate images of interest, download them to their workstations, display them, and magnify portions of them to view their detail, with accurate color.

At the beginning of the project, it was apparent that no preexisting product would provide a unified system to manage the flow of images from scanner to server(s) to storage to display to scholar. For this project, building a unified system involved not only developing new code for the project and integrating it with existing products, but also developing work flow and data-format specifications that permit cooperative operation of the subsystems. In the following sections, we describe both the subsystems developed for the various sites and the work flows that enable unified operation.

● *The subsystem at the Vatican Library*

The subsystem at the Vatican Library, briefly described in the previous section, was designed to capture images of manuscripts, enter cataloging data, and display the collection of scanned images. It is composed of workstations that run three different application programs.

To support scanning, this subsystem includes two IBM PS/2[®] workstations. Each is equipped with the PISA [12, 13] application program, and each supports a Pro/3000 Scanner [14, 15], which is described in more detail later. PISA and the Pro/3000 Scanner were both developed by the IBM Research Division. Working in conjunction, they can preview an image, scan the image, correct the color

of the scanned image, and store it in the bulk memory (DASD) of the attached workstation.

To support image examination, the IBM Color Image Portfolio Assistant (CIPA) [16] application program is utilized. CIPA, which runs on IBM PS/2 workstations that are interconnected with a local area network (LAN), utilizes an IBM 3995 Optical Library [17], connected to the LAN, to store its images. A CIPA-based system may be thought of as an "image database," which associates images with text that is used to describe them. In addition, CIPA provides image services that include faithful-color image display, restoration of the color of faded transparencies, and an image-export feature that prepares the images for printing. To examine an image stored in a CIPA-based system, one queries CIPA's database, the image is retrieved, and CIPA displays the image. In the remainder of this paper, the term *CIPA system* is used for a system composed of the CIPA application program, workstations, LAN, and optical library.

At the Vatican Library, the scanning workstations are also connected to the CIPA system's LAN. The CIPA system receives scanned images from the scanning workstations and stores them in its optical library. Information that describes the images is automatically entered into CIPA's database, as is described below in the section on the flow of images through the system.

A Geac [18] application program running on an IBM RISC System/6000[®] workstation is used to enter cataloging information and store it in the library-standard MARC record format [19] (see also [20]). This system predates the project described in this paper. It operates independently of the PISA and CIPA programs; indeed, the workstation on which the Geac application runs has no physical connection to the PISA-based and CIPA-based systems.

- *The subsystem at the Pontifical Catholic University in Rio de Janeiro*

This subsystem, located at PUC-Rio, provides remote access to the cataloging information and images via the Internet. One component of this subsystem is the project's Internet server, which is implemented on an IBM PS/2 Model 95 running the OS/2[®] operating system. The Rio Internet server is the only Internet server that contains a complete collection of cataloging information and images captured by the project. (The others are discussed in the following subsection.) It serves all of the scholars who are participating in the project (except for those who are supported locally, at the Vatican). Designed and implemented by PUC-Rio, it utilizes the "gopher" protocol [21], a popular Internet access protocol, to provide access to its information. The IBM GoServe [22] OS/2 program is used to provide a gopher interface to the server's two main directories, one containing the retrospective catalog and the other containing image files.

Only authorized users are permitted access to the data stored on this server.

Once scholars receive access to the Internet server, they can navigate hierarchical lists to locate desired images, or they can search the cataloging information. The search capability, still under development, will support free-form text search, which locates and ranks matches in a catalog's data to an unconstrained sequence of input text. Indeed, we believe free-form text search is vital to an easily used interface; the IBM SearchManager/2[™] [23] software product supplies this search capability to the Internet server. A SearchManager/2-GoServe interface is being developed to permit the GoServe software to pass the search request to SearchManager/2 and then supply the search results to the requesting client. Preprocessing of the images stored on the Internet server, as is described below in the section on processing the images to prepare them for the Internet, reduces their resolution, watermarks them, compresses them, and otherwise prepares them for Internet access. To permit scholars to inspect the images being stored on the Internet server, the PUC-Rio subsystem includes several workstations running the Scholar's Interface Application, which is more fully described in the subsection devoted to it.

To permit local examination of high-resolution, uncompressed images, the subsystem at PUC-Rio utilizes a CIPA system.

For security reasons, only compressed, watermarked images are stored on the Internet server. Images that are not intended for access via the Internet, e.g., the uncompressed scanned images, are stored on the CIPA system, which is not physically connected to the Internet server. This approach provides additional security; if you cannot access an image, you cannot accidentally provide it to the Internet.

- *Secondary servers*

To avoid overloading Internet communication lines, primarily outside the U.S., and to achieve better worldwide performance, the project was designed to provide a distributed system of image servers, under the control of the main server installed at PUC-Rio. A secondary image server has been installed at Case Western Reserve University, in Cleveland, Ohio, to serve U.S.-based scholars. While no additional servers are planned for this phase of the project, we hope to add a European server at some future time.

- *The flow of cataloging information through the system*

As asserted earlier, a careful specification of the flow of information is critical to the operation of the system. This flow describes the information formats to be used, the subsets of the formats to be supported, and the physical media to be transferred from site to site. Without such

a specification, we would be unable to freely exchange information between sites which use software components that were written without regard to compatibility with one another.

The catalogers in the Vatican Library enter a description of each work in a manuscript into the Geac system. The manuscript itself is cataloged as a master record with links to each work within the manuscript. The catalog records are exported in MARC Interchange format, written onto magnetic tape, and shipped to PUC-Rio. (This is less expensive than transmitting these very large files by telecommunications.) The cataloging operation proceeds independently of the scanning of manuscripts; this is important to the throughput of both operations.

Selected fields from the records in the catalog (e.g., manuscript name, page) are also imported into the database within CIPA to index images that will be stored within the CIPA system. This database assists researchers physically located at the Library to locate and retrieve images of the pages of manuscripts.

Following receipt at PUC-Rio, the records are imported into an on-line public access catalog that resides on the Internet server. Selected records from the catalog are imported into a CIPA system at PUC-Rio, as they are at the Vatican, for use by local researchers.

• *The flow of images through the system*

The flow of images through the system begins with their capture at the Vatican Library and ends with their storage on the Internet server in Rio and on the CIPA-based image databases in the Vatican and Rio. Along the way, the images are processed to reduce their data volume and transform them into the formats required by the Internet server and the CIPA program; this processing is described more fully below in the section on processing the images to prepare them for the Internet. The images are also visually inspected (at the Vatican and Hawthorne) to ensure that the high image quality desired of the system is being achieved. The steps in the image flow include the following:

1. The images are scanned at the Vatican Library and stored in a subset of the Tag Image File Format (TIFF™) [24].
2. The scanned images are imported into the CIPA system at the Vatican Library.
3. The images are copied to tape at the Vatican Library, and the copies are shipped to the IBM Thomas J. Watson Research Center at Hawthorne, NY.
4. The images are inspected in Hawthorne with a TIFF image viewer written for the project.
5. The (inspected) images are copied onto magnetic tape and shipped to PUC-Rio.

6. The images are imported into the CIPA system at PUC-Rio.
7. The images are processed at PUC-Rio, by software written to process TIFF images stored in the specified subset, and placed on the Internet server.

At the scanning workstation, images are scanned and stored in a subset of the TIFF format specified for the project. This subset is used to describe and store all images transferred from one to another of the system's three subsystems. The specified subset restricts the stored images to 8-bit-per-pixel monochrome images (only shades of gray) or 24-bit-per-pixel color images; it utilizes all baseline tags [24], tags that record colorimetric information, a tag that records copyright information, and a tag that records a brief annotation entered by the scanner operator to identify the image.

After being processed by the PISA application program, each image is displayed on a high-resolution monitor that allows the operator to verify that the image has been scanned correctly; overall appearance, orientation, color, and cropping area (the part of the total image selected for storage) are all visually checked. The images are then used locally (step 2) and prepared for remote use (step 3).

The CIPA system at the Vatican Library supports the import of all images that conform to the specified TIFF subset. During the import operation, the annotation information is automatically extracted from the TIFF header and entered into the CIPA database; there, it can subsequently be used to locate the scanned image. The specification of the format of the annotation permits the automated entry of images into CIPA; this is another part of the work flow defined for this system. The CIPA system at PUC-Rio uses the same import process.

At IBM Hawthorne, the images are again visually examined to ensure overall image quality. Over time, several problems have been detected and corrected. Power fluctuations that the scanner encountered produced "artifacts" (irregular, anomalous patterns) on some images; this was corrected by adding voltage regulators to the scanners. Some images showed ink bleed-through from the other side of the page; software, described below, was developed to deal with this problem. The most significant surprise we have encountered thus far has been the amount of detail in the image content.

The Pro/3000 Scanner is capable of scanning images at up to 3000 × 4000 pixels; as discussed later, they are routinely scanned at a resolution of 2500 × 3000 pixels. We had initially planned to reduce the size of all scanned images to 1000 × 1000 pixels (maximum) prior to compression and storage on the Internet server, as is described more fully below in the section on processing the images to prepare them for the Internet. While this lower resolution has proven sufficient for most

manuscripts, it is inadequate for many maps, architectural drawings, small calligraphy, and marginal notes. Observation of this surprising result enabled us to suitably alter our plans: As described below, some images are not reduced prior to storage on the Internet server.

At PUC-Rio, the files are read from tape and imported into a local CIPA database. This provides secure access to the images for local scholars. In addition, the images are processed, in batch mode, to prepare them for access through the Internet. The JPEG-compressed images produced (discussed below) are then stored on the server, from which they can be accessed, through the Internet, by the remote scholars.

This workflow leads to three identical copies of each scanned image being created. One is stored at the Vatican; one is stored at IBM Hawthorne; and one is stored at PUC-Rio. This triplicate archive provides assurance that scans will not be lost because of a local problem. The project team recognizes a continuing need to recopy the images onto new media every few years. While magnetic tape may endure for many years, tape drives become obsolete within a few years. The images must be copied onto new media before the tape drives become obsolete.

- *The Scholar's Interface Application*

The Scholar's Interface Application (SIA) was designed to be an easily used Internet client that provides a robust set of image-display functions, not present on most WorldWide Web browsers, that permit the scholarly examination of manuscripts. The SIA was implemented as a Smalltalk [25] program with an integrated set of utilities that organize and display files which are stored on either an Internet gopher server or an *image cache* (user-controlled image storage) on the SIA workstation. The SIA supports a two-monitor configuration that allows the scholar to control the program on the *system display*, while presenting high-quality images, possibly to a group of people, on a high-resolution *image display*. A picture of a workstation running the SIA is given in **Figure 1**, which shows the system display on the left and the image display on the right. The image display can show images up to 1000×1000 pixels in size, with a 65536-color palette.

With the SIA, a scholar may navigate a set of gopher menus to find a manuscript page, display the page, and import it into the image cache on his workstation for further scrutiny; the importance of the image cache is described below. The SIA provides several features that facilitate scholarly study of Vatican Library manuscripts. An integrated utility displays JPEG-compressed images on the image display with a variety of magnifications and a variety of screen layouts. Cropping information, used to specify the portion of an image to be magnified, is entered by drawing a crop box on a monochrome *reference image*




Figure 1

Workstation running the Scholar's Interface Application. The image shown on the image display (right-hand monitor) is from the Vatican Library manuscript Vat.Lat. 3811 (fol. 21v,22r, a map of Italy, from a 15th-century copy of Ptolemy's Geography).

that appears on the system display; a reference image is displayed in the lower right-hand corner of the system display shown in Figure 1. Displaying enlarged image portions side by side, for example, enables the visual comparison of details from two manuscripts.

One of the SIA features permits the creation of *snapshots* (uncompressed, display-ready images that can later be displayed very quickly to give pseudo-slide-shows) of images on the image display. Another feature permits a *screen capture* of the image of a cropped region on the system display, so that it may be shown on the image display; this feature enables us to create descriptive text, using an editor and the system display, that may be shown on the image display alongside the image.

The SIA features a tool bar for selecting the layout of the image screen; the supported layouts include full-page display, left/right half-page display, top/bottom half-page display, and quarter-page display. The scholar has the flexibility to place entire manuscript pages, or cropped regions of them, at various locations on the image display. From the system display's menu, the user may adjust display parameters such as brightness and the bleed-through-reduction threshold (described below), or the user may access an editor to enter annotation information.

Displaying an image that is located on the Internet server may require minutes to download the image and seconds to display it. An image that has already been downloaded to the image cache can be displayed within seconds; hence, the image cache is an important performance-enhancement feature of the SIA. The SIA's import utility creates small *thumbnail images* that are used with an integrated *light-table* viewer. The light table displays an array of thumbnail images (of images stored in

the image cache) on the system display; an image in the cache can then be displayed by clicking on its thumbnail. In Figure 1, a light table is shown on the system display, partly hidden behind a reference image.

Preliminary reports indicate that the SIA is easy to use and provides a good set of image functions for examining manuscripts. A few participating scholars and their students, with minimal training, have become competent, enthusiastic users. Further, they report that the displayed images, assisted by the SIA's image function set, are adequate to support their research.

To be user-friendly, the SIA must provide response times that satisfy the scholar. As mentioned above, most images are stored on the server at a resolution of 1000×1000 pixels (maximum), and a few are stored at a resolution of approximately 2500×3000 pixels. If an image had previously been downloaded into the image cache of a PS/2 Model 95 with a 50-MHz 80486 processor, it would typically take 6–7 seconds to decompress and display an 800×1000 -pixel monochrome image; a color image of the same size requires 12–16 seconds. A 2500×3000 -pixel monochrome image typically takes 25–28 seconds to display; a color image requires 30–40 seconds. Although these display times are far from instantaneous, the scholars have found them to be tolerable.

The amount of time required to download an image into the image cache of an SIA workstation strongly depends on the bandwidth available between the server and the workstation. Experiments were conducted using an Internet server at Case Western Reserve University and a workstation at IBM Hawthorne. These revealed that a typical 800×1000 -pixel image, JPEG-compressed to 100–200 kilobytes, would be downloaded in approximately 10 seconds. A typical 2500×3000 -pixel image, JPEG-compressed to 500–1000 kilobytes, would be downloaded in approximately 40 seconds. These times were observed under very favorable conditions; if one were dependent on a 9600-bit-per-second modem for communication to the server, one would expect the download times to be an order of magnitude longer. A comparison of the download times to the display times reinforces the importance of the image cache in achieving tolerable display times.

Imaging needs of the Vatican Library system

The manuscripts of the Vatican Library are treasured for many reasons, including their historical importance and their artistic beauty. Capturing and preserving their visual appearance with images is a considerable challenge. While there are many aspects to image quality, two of the most important are possession of a high level of detail (spatial resolution) and accurate representation of the colors of the original materials.

Keeping in mind the high-image-quality requirements of this project, we summarize the imaging challenges as

- Capturing images with accurate color and with as much detail as possible, while causing no harm to the originals.
- Compressing the images for transmission through the Internet, while preserving as much visual quality as possible.
- Displaying the images, and enlarged details of them, with accurate color.

These are addressed in the following three subsections. Some members of our team had faced many of these challenges in earlier projects [4, 5], but we had little experience with either scanning images from original manuscripts or compressing manuscript images, in this case for access through the Internet; these challenges became the focus of the project's imaging research.

• *Scanning Vatican Library manuscripts*

At the heart of the scanning environment is the Pro/3000 Scanner [12–15], which is based on an IBM-proprietary charge-coupled device (CCD) imaging sensor chip [26] that provides a signal-to-noise ratio greater than 3000:1. Developed by IBM Research and refined through use in earlier projects with the painter Andrew Wyeth [4, 5] and the National Gallery of Art (U.S.A.) [11], the Pro/3000 consists of an integrated copystand and digital camera that together can capture transmissive and reflective originals covering a broad range of sizes. The Pro/3000, shown in Figure 2, captures images at resolutions up to 3072×4000 pixels, with 36 bits of color information per pixel. Its copystand provides diffused quartz halogen illumination to illuminate reflective originals. One of the most important features of the Pro/3000 is its colorimetric filter set, which permits accurate color capture of nonphotographic materials, such as manuscripts, fabric samples, or stamps, that are not composed from photographic dyes. Many scanners are designed to capture the colors of only materials produced with photographic dyes and do poorly on nonphotographic originals.

The Pro/3000's digital camera is supported by a motorized column. A bellows unit is used to focus the digital camera. Because of the limited height of the column, the largest original that can be scanned is 45×60 cm.

The PISA application program works with the Pro/3000 to capture monochrome or color images. Its user interface displays both the proper digital camera height and the bellows position for proper focus for a number of sizes of originals, helping the scanner operator to more quickly position and focus the digital camera when the size of the original is changed. PISA also includes a utility that records the lighting pattern produced by the illumination.

In the process of creating a scanned image, PISA uses this information to correct the image, so that it appears as it would if the illumination were spatially uniform. PISA also includes a utility to analyze the scan of a color-calibrated test chart and determine the color characteristics of the scanner on the basis of that analysis. In the process of producing a scanned image, PISA uses this information to correct the colors of the scanned image so that they will appear correct.

For those versed in color theory, we note that this "color calibration" [27] computes a best-fit 3×3 matrix for mapping the scanner's red, green, and blue signals onto CIE XYZ color coordinates, which may be used for color matching. During scanning, PISA processes the raw scanned data to produce an image with accurate X, Y, and Z coordinates. The CIE XYZ color coordinates are briefly described in the next subsection.

As we noted above, one of the goals of the project was to digitize 20000 images of Vatican Library manuscripts. In planning the project, we based our throughput projections on the expectation that we would be scanning many photographic transparencies of manuscripts and few originals. The physical handling of transparencies is much easier, so the scanning throughput is greater. However, we soon discovered that scanning originals directly led to images having much better quality. When originals are scanned, the color accuracy of the image capture is limited by the color errors added by the Pro/3000; since the Pro/3000 features colorimetric filters, these errors are relatively small. When scanning film, the color accuracy is further degraded by the color errors added by the film; these errors are several times as large as those added by the Pro/3000. The discovery that scanning directly from originals produced superior-quality images led the project to concentrate on the scanning of original manuscripts and the problems inherent therein. Achieving the desired number of digitizations then became a much greater challenge.

Nearly all of the original materials scanned were manuscripts in the original sense—handwritten. Each is unique and irreplaceable; they are, on average, five hundred years old. Hence, our primary responsibility was to do no damage to the manuscripts. Usually written on parchment, the manuscripts are extremely sensitive to environmental conditions such as temperature and humidity. The Pro/3000's quartz halogen lights, filtered by glass diffusers, produce little ultraviolet light, which is known to be damaging to organic materials; however, they do add heat. Most of the heat is directed away from the manuscripts with dichroic reflectors, but it does tend to warm the room containing the scanners. For this reason, air conditioning was added to the scanning area. The scanning area's environment was continuously monitored and maintained within a restricted range. Over the course




Figure 2

The Pro/3000 Scanner.

of time, procedures were developed for scanning the originals. A pane of glass was placed on top of the manuscript being scanned; this tended to both reflect heat and keep the pages somewhat flat. Black curtains were also installed around the two Pro/3000 Scanners installed at the Vatican to prevent stray light from contaminating the scanned images.

Many of the originals we scanned were bound manuscripts, and many problems were introduced by variations in their size and thickness. The sizes of the manuscripts that we have scanned range from 30×42 cm to 39×57 cm for a single page, with a binding thickness up to 13 cm; while this may seem a wide range, it is far smaller than the range of sizes and thicknesses of Vatican Library manuscripts. The manuscripts must be supported during the scanning process in such a way that the binding is not stressed. Capturing the areas between the original writings and the bindings, or gutter areas, of manuscripts is a requirement, since many important notations (added later) are located there. Keeping the whole page in focus is needed for capturing the gutter detail. The pane of glass placed on top of the manuscript being scanned partially flattened the parchment sheets to help maintain focus over the whole page, without flattening the pages so much that folds and wrinkles acquired




Figure 3

The easel used for scanning bound manuscripts.

through the centuries would disappear. Because of the manuscripts' varying thicknesses, the scanner's focus had to be readjusted for each page. Fortunately, the Pro/3000 can be focused by eye through a viewfinder; unfortunately, the scanner focus had to be checked at every page—a process that slows the scanning.

Although scanning from original manuscripts produced superior images, the mechanical positioning of the manuscripts for scanning, as described above, was a limiting factor for scanner throughput. To increase scanner throughput and to better protect the manuscripts from handling wear, a copystand/easel configuration was designed. The copystand holds the camera, lights, and easel, which sits on top of the flat surface of the copystand. The easel, shown in **Figure 3**, is used to support the manuscript. The manuscript is placed in the "landscape" mode, with a maximum page dimension of 45.7×36.6 cm; this represents a scan resolution of 170 pixels per inch when the scanner is operating at a 2500×3000 -pixel resolution. The manuscript is held with its binding and back supported, while the page to be scanned is gently pressed flat against a fixed glass plate at the top of the easel.

In developing the easel configuration, we procured and modified a commercially available book easel.¹ One modification extended the glass plate that covers the manuscript to permit scanning in the gutter areas. Another modification replaced the support for the side of the bound manuscript not being scanned; this permitted us both to scan heavier manuscripts and to support them at a greater variety of angles to the horizontal. A new copystand was

provided, with a lower work area, so that the surface of the easel would be at a comfortable working height; the lighting supports were modified so that they could be directed at the surface of the easel; a longer copystand column was provided, to compensate for the loss of column height caused by the height of the easel; and finally, the surface area of the copystand was increased, which, combined with the longer column height, permitted larger flat original copy (up to 58×76 cm) to be scanned with the easel removed.

One concern, expressed early in the project, was that scans of manuscript pages that are not flat would exhibit curvature. In practice, we have not found this to be a significant problem. A small amount of shadowing and a small amount of distortion due to curvature add to the perceived realism of the image. Both the earlier manual book positioning and the newer easel configuration hold pages acceptably flat.

Still, the scanning area of the Pro/3000 is limited to 58×76 cm. Scanning from photographic reproductions was necessary to capture images of the largest pages. Our knowledge of scanning photography was also refined through the course of the project. We learned that the Vatican Library's preexisting microfilm (both positive and negative) and 35-mm color slides did not provide the quality required by our objectives. The high-contrast film used in microfilming tends to eliminate intermediate gray levels, turning them into black or white. Both the microfilm and the 35-mm color slides were in general too small to permit the transfer to a 2500×3000 -pixel digital image of all the detail present in the original illuminated and handwritten pages. On the other hand, the other photographic format commonly used by the Library, the 5 \times 7-in. color slide, proved to be more than adequate for our purposes in terms of resolution. We have, however, observed a certain variability of color dyes in color slides, depending on the age of the film and the kind of film and development process used.

Some compromises to the quality of captured images were made where the resulting increase in throughput justified the degradation in image quality. The PISA software can scan images at either a 2500×3000 -pixel resolution or a 3000×4000 -pixel resolution. The 2500×3000 -pixel resolution, commonly used, proceeds more quickly and is adequate for capturing the level of detail present on most of the manuscript pages. In practice, the scanned images are somewhat smaller than the maximum image size permitted by PISA because of cropping. Scanning an image in monochrome (shades of gray) instead of in color also proceeds more quickly; monochrome scanning is commonly used for those pages consisting solely of dark ink on a light background (paper or parchment). Finally, we note that the 36 bits of color data captured for each pixel are reduced to 24 bits of color

¹ Information about the Linhof Buchwippe book-copying easel is available from Linhof Präzisions-Kamera-Werke, GmbH D-8000 München, Germany.

data by the PISA software, and the 12 bits of monochrome data are reduced to 8 bits. This color quantization is done by both normalizing the image data and picking quantization levels that are more perceptually uniform; it is more fully described in [13] and [28]. We feel that it degrades the image only slightly. Thus, in common practice, the output of each scan is a 24-bit-per-pixel color image or an 8-bit-per-pixel monochrome image, with an image area slightly smaller than 2500×3000 pixels.

• *Processing the images to prepare them for the Internet*

In preparing the images for access through the Internet, it is essential to reduce their data volume while preserving their color and sufficient detail for them to be adequate for scholarly study. The processing steps used to prepare the images for access through the Internet are as follows:

1. Reduce the image to the desired size.
2. Sharpen the image.
3. Rotate the image to its proper orientation.
4. Transform the image to the desired color space.
5. Apply a *visible digital watermark* to the image.
6. Compress the image, using JPEG.

A batch software process that implements these six steps is executed at PUC-Rio. Steps 1 and 6 are designed to reduce the data volume, while steps 2 and 4 are designed to improve the image quality (by enhancing detail and improving the color rendition). Step 3 is performed for the convenience of the user, and step 5 is designed to prevent the images from being used for purposes other than scholarly study.

The techniques used for the processing of each step were developed by IBM researchers to preserve the color content of the images through each step. While it is beyond the scope of this paper to review the underlying color science, we recommend Reference [29] to the interested reader; in particular, Chapter 8 describes the CIE *standard color observer* that is used for color matching. If two color patches produce the same triplet of standard-color-observer coordinates, and the surrounding conditions are the same, those two color patches will also produce a visual match for most observers. The triplet of color-observer coordinates for a color are called the CIE X, Y, and Z coordinates of the color. The standard color observer is the basis for the image-processing methods that were used to preserve color. We define red, green, and blue *linearized color components* to be proportional to the CIE Y coordinate. Then, a filtered image will appear to have the same color as the unfiltered image if the filtering is applied to the linearized red, green, and blue image components, and the frequency response of the filter for zero frequency is exactly 1. This argument is more fully presented in [30].

The first processing step reduces the size of the image. In all cases, the scanned image is archived. Whenever possible, a lower-resolution image is stored on the Internet server, to reduce the data volume. Many of the images, which are scanned at a resolution of 2500×3000 pixels, are quite usable at a resolution of 1000×1000 pixels (maximum). Let us take as an example the manuscript page shown in **Figure 4**. This page was scanned at a resolution of 2840×1895 pixels and reduced to a resolution of 1000×667 pixels; the reduced image is presented in the figure. This image reduction resulted in decreasing the data volume from 16147368 bytes to 2002964 bytes, as shown in **Table 1**; this is a reduction in data volume by a factor of 8.06. The reduction processing is based on techniques described in [31] as applied to linearized red, green, and blue image components.

The image-sharpening step corrects for the optical blurring that occurs during scanning, and makes the image crisper and easier to read. The combination of a linear filter and nonlinear clipping is used to effect the sharpening. An example illustrating the sharpening is given in **Figure 5**, which shows the unsharpened and sharpened versions of an image. The sharpening produces the appearance of a higher-resolution image than is present, so the utility of the image is greater; however, the sharpening does lessen the amount of compression that can be achieved, as described below.

The output of the linear filter for the pixel in the i th row and j th column, $s_{i,j}$, is related to the input of the filter, $p_{i,j}$, by the equation

$$s_{i,j} = (1 + 4\alpha)p_{i,j} - \alpha(p_{i+1,j} + p_{i-1,j} + p_{i,j+1} + p_{i,j-1}).$$

If this filter is applied without care, it may cause color changes in the image; to preserve image color, we apply it to linearized red, green, and blue color components. When the filter is applied with a large value of α , considerable sharpening occurs, but artifacts appear because of ringing of the filter. The artifacts appear as bright halos on the brighter side of sharp edges and dark halos on the darker side of sharp edges. At the suggestion of one of our colleagues,² the amount by which a pixel can be decreased is limited to 50% of its value, and the amount by which a pixel can be increased is limited to 33% of its value (i.e., $0.5p_{i,j} \leq s_{i,j} \leq 1.33p_{i,j}$). This greatly reduced the visible presence of the artifacts.

Rotation by a multiple of 90° is required if the image was scanned sideways or upside down, so that the image, when made accessible through the Internet, has the correct orientation. The need for rotation is recorded in the TIFF header of the scanned image, where it is entered by the scanner operator at scan time.

² Gerhard Thompson, personal communication, IBM Thomas J. Watson Research Center, March 1994.




Figure 4

Color image with watermark. The compression of the image is reported in Table 1. The image shown has been reduced from 2840×1895 pixels to 1000×667 pixels. This image is from the Vatican Library manuscript Chig. C VIII 234 (fols. 19v, 20r, music for a mass to be sung, from a 16th-century Flemish manuscript).

Table 1 Number of bytes in the data for the image of Figure 4, after each processing step used to prepare the image for transmission via the Internet. The numbers are given for the raw data, and for the data after JPEG compression.

<i>Processing step</i>	<i>Uncompressed data (bytes)</i>	<i>Compressed JPEG data (bytes)</i>	<i>Compression factor</i>
Scanned image (2840×1895)	16,147,368	704,200	22.93
After reduction (1000×667)	2,002,964	147,314	13.60
After sharpening (No rotation)	2,002,964	190,546	10.51
After color transformation	2,002,964	189,000	10.60
After watermarking	2,002,964	197,098	10.16

The image is next transformed so that its colors will appear approximately correct on a typical high-resolution display. We have found that the SMPTE chromaticities [32] and a gamma of 2.2 provide a good description for many displays. This color transformation is essentially a 3×3 matrix applied to each pixel's linearized red, green,

and blue color components, but pixels that lie outside the SMPTE color gamut are mapped to its surface. This transformation is also described in [4].

The next step applies a visible digital watermark to the image. Prevention of unauthorized usage of Vatican Library images was and is a serious concern for the



(a)



(b)

Figure 5

Original version (a) and sharpened version (b) of a portion of a scanned image. The image of the figure is from the Vatican Library manuscript Chig. F VIII 188 (fol. 19v, a discourse on the medicinal use of herbs, from a 15th-century Italian manuscript).

project team. Since the images are accessed through the Internet, which is not secure, we were concerned that the Vatican's images might be used by those who had no right to do so. What the project desired was a method that rendered the images perfectly acceptable for scholarly study yet unacceptable for other usages, such as lithographic printing. Making only lower-resolution images available through the Internet helps accomplish this goal, but not all the images are useful at low resolution. The visible image watermark that was developed for this project is another tool that is used to discourage unauthorized use of the images.

The visible watermark clearly marks an image as belonging to the Vatican Library; a watermarked image cannot be purloined through the Internet and published without acknowledging its ownership, since that acknowledgment effectively appears on the image. However, all of the detail beneath the watermark is readily visible, so it is still quite useful for scholarly study. In Figure 4, the watermark is apparent, yet all detail beneath

the watermark is clearly visible. Visible watermarking of images is also used commercially, so that digital images may be used to advertise photographs [33] without giving away the photographic-quality image.

We have tried to make the watermark as unobtrusive as possible, yet readily visible. When a pixel is changed by our watermarking, the brightness is reduced, while the hue and saturation are held constant. Changing only the brightness, we feel, makes the most visible mark on the image for a given degree of obtrusiveness. The use of a watermark that is thematically related to the materials themselves, in this case the Vatican Library seal, also adds to the unobtrusiveness of the watermark. In applying the watermark, we adjust the watermarking's change of brightness to darken image pixels by the same amount (a change in L^* as defined in [29]), perceptually, whether the pixels are light or dark. This "uniformly perceptual" darkening is only approximate, and it can only be accomplished if the underlying pixels are bright enough to be darkened by the desired amount. It has been our

experience that this technique does produce watermarks that are equally obtrusive on many images.

The watermarking software reads the watermark as a monochrome TIFF image and applies it to the manuscript image. The amount of processing needed to apply a watermark is quite small. Where the watermark image indicates that no darkening is to be applied to the image, the image pixel is unchanged. There is a natural conflict between unobtrusiveness and protection; in this project, we have chosen to use watermarks that are large, nearly centered, and fairly unobtrusive. As the presence of color tends to visually mask the watermark, we tend to use greater darkening for color images than for monochrome images, but this leads to a similar perceived obtrusiveness.

If the watermark could easily be removed, it would not offer sufficient protection. To defeat the watermark, some might postulate that one could simply estimate the watermark image and use this estimate to brighten pixels previously darkened by the watermark program. We have used a number of techniques to thwart this strategy. In the pixel-darkening process, randomness is added to the attenuation factors, so that two pixels are seldom darkened by the same amount. Because of the randomization added to the attenuation factors that darken the watermarked pixels, the watermarked and unwatermarked areas exhibit different textures. In order to restore a watermarked image, one would have to correct each pixel by a pixel-unique amount; with many thousands of pixels typically being altered, this would be a time-consuming task. We also try to make the watermark image difficult to calculate. The size and position of the watermark are modified by random parameters, under program control. When one image is watermarked at different times, watermarks of slightly different sizes are applied at slightly different locations.

The final step in the preparation of an image for storage on the Internet server is compression of the image. Image compression can be lossless, a term which describes compression which produces images that can be decompressed to reconstruct the original image exactly, pixel by pixel. Alternately, image compression can be lossy, a term which describes compression which produces images that, when decompressed, merely look like the original. Lossless compression, with existing techniques, seldom reduces the data volume by more than a factor of 2.5, while lossy compression allows greater reductions, particularly for color images. Lossy rather than lossless compression was chosen for the project in order to reduce the data volume to a level more acceptable for Internet transmission.

A significant compression issue involves the choice of the compression technique and the image file format to be used. For this project, we anticipated that the scholars would have a great variety of hardware and software

for examining the images; therefore, we chose what we believed to be the most commonly used compression technique and format. The images are compressed using the ISO-standard JPEG technique [34–36], which is a commonly used standard and one which produces excellent compression of images. Furthermore, we chose to have the compressed images comply with Version 1.02 of the JPEG File Interchange Format (JFIF), defined by C-Cube Microsystems [37]; this format was chosen because it codifies common practice in applying the JPEG compression algorithm (for example, representing the image in the YCbCr color space) and is in widespread use.

It is beyond the scope of this paper to truly review JPEG compression, but we recommend Reference [36] as a brief review to the interested reader. We do, however, describe some aspects of the baseline option of JPEG, so that the ensuing discussion of JPEG performance will be meaningful. JPEG compression decomposes each color plane (e.g., red, green, and blue) of an image into 8×8 blocks of pixels and transforms (by a discrete cosine transform) each block into 64 frequency components. Then, each frequency component for each block for each color plane is divided by a scalar and “quantized” to an integer value; the smaller the scalar, the finer the quantization of that frequency component. Finally, the quantized frequencies are entropy-coded using Huffman coding. Each frequency component for each color plane has its own scalar, so there are 64 scalars for each color plane; these are collectively referred to as the quantization table (or quantization matrix) for that color. For a color image, JPEG uses three quantization tables, one for each primary color. For a monochrome image (composed only of shades of gray, such as a black-and-white photograph), JPEG uses one quantization table. In practice, JPEG is often applied to images that have already been transformed to consist of one luminance (brightness) component and two chrominance (brightness-independent) components, and the same table is used for both chrominance components; hence, only two tables are used.

JPEG is a variable-rate compression technique. It attempts to preserve each frequency component to a given accuracy level. When an image contains a great deal of detail, there is more frequency content to preserve, and JPEG achieves less compression.

Although we use lossy JPEG compression, we set our quality standard high; the compression should lead to images that are visually indistinguishable from the originals when viewed without magnification, and the artifacts should not be annoying when the images are viewed with a $2\times$ magnification (where the display shows an interpolated image that has four times as many pixels as the image before interpolation). The high quality levels we were seeking are not commonly called for using JPEG, since people often choose higher compression and settle for

Table 2 Number of bytes in the data for the sample monochrome image of Figure 6, after each processing step used to prepare the image for transmission via the Internet. The numbers are given for the raw data, and for the data after JPEG compression.

<i>Processing step</i>	<i>Uncompressed data (bytes)</i>	<i>Compressed JPEG data (bytes)</i>	<i>Compression factor</i>
Scanned image (3064 × 2052)	6,288,145	667,589	9.42
After reduction (1000 × 670)	670,820	116,754	5.75
After sharpening	670,820	159,472	4.21
After rotation	670,820	158,803	4.22
After watermarking	670,820	161,825	4.15

lower quality. A significant issue was determining the appropriate parameters (quantization tables) to provide both acceptable image fidelity and good compression. The quantization tables used for the Vatican Library images were derived by the following process, from tables developed by Peterson et al. [38, 39]. A "representative" sample of images of Vatican Library manuscripts was chosen early in the project. This sample set was compressed and evaluated visually to determine whether the quality criterion had been met. In the earliest experiments, the criterion was not met. The quantization tables were then modified according to methods also described in [38, 39] to provide better image quality; the sample set was again compressed; and the results were again evaluated. This process was iterated until the compressed images met the quality criterion. While these tables have served us well in the project, they did not always perform satisfactorily. As one might imagine, when the project's images differed significantly from the sample set, problems could arise.

When JPEG compression was applied to Vatican Library images, with the tables developed, we observed typical compression factors of 10 to 15 for color images and 4 to 5 for monochrome images. Since JPEG is known to deliver excellent image quality with compression factors in excess of 20 and 10 for color and monochrome images, respectively, the smaller compression ratios actually achieved were a mild disappointment, prompting us to investigate the source of the disparity. Table 1 traces a sample color image (Figure 4) through the six processing steps. After each processing step, the resultant image was compressed using JPEG with the quantization tables developed for the project. We note that the scanned image was compressed by a factor of more than 22. This is a modest amount of JPEG compression, indicating a relatively high level of detail in the scanned image. After the image reduction, the image was compressed by a factor of only 13.6. JPEG attempts to preserve the detail of the image; we see that the reduction step reduces the data volume by a factor of 8.06 but significantly increases the amount of relative detail in the image. The image-

sharpening step further diminishes the amount of compression that can be achieved; this process intentionally increases the amount of detail present in the image. The rotation step, when present, does not increase or decrease the amount of detail in the image; its effect on the compression of the image is very small and incidental. The color transformation step does not inherently increase or decrease the amount of detail in the image; its effect on the compression is also small and incidental. The watermarking step adds fine detail (texture) to the darkened areas, which does diminish the compression. We conclude that the JPEG compression was less than expected because we were working with an image that contains a high level of detail and because several steps in the processing increase the relative amount of detail in the image.

A similar pattern emerges in the data of Table 2, which traces the monochrome image of Figure 6 through the processing steps. With the sample monochrome image, most processing steps again increase the relative amount of detail in the image. For the monochrome image, however, the compression after each step is less than it is for the sample color image. It is well known that JPEG typically produces less compression when applied to monochrome images, so this effect was not unexpected.

Although the first two processing steps reduce the compression (for both monochrome and color images) that can be achieved by nearly two to one, those same steps reduce the data volume significantly. Working in concert, all steps achieve a total reduction of the data volume by a factor of 81.9 for the sample color image and by a factor of 38.9 for the sample monochrome image. Both example images are compressed to less than 200 kilobytes, possess excellent image quality, and, while large, can be satisfactorily transmitted through the Internet.

As noted earlier, some images contain too much detail to be reduced. Often, they feature delicate writing that would be obliterated by the reduction. Except for the image reduction, these images, which represent approximately 1% of the images stored on the Rio Internet server, are processed like the others. The unreduced

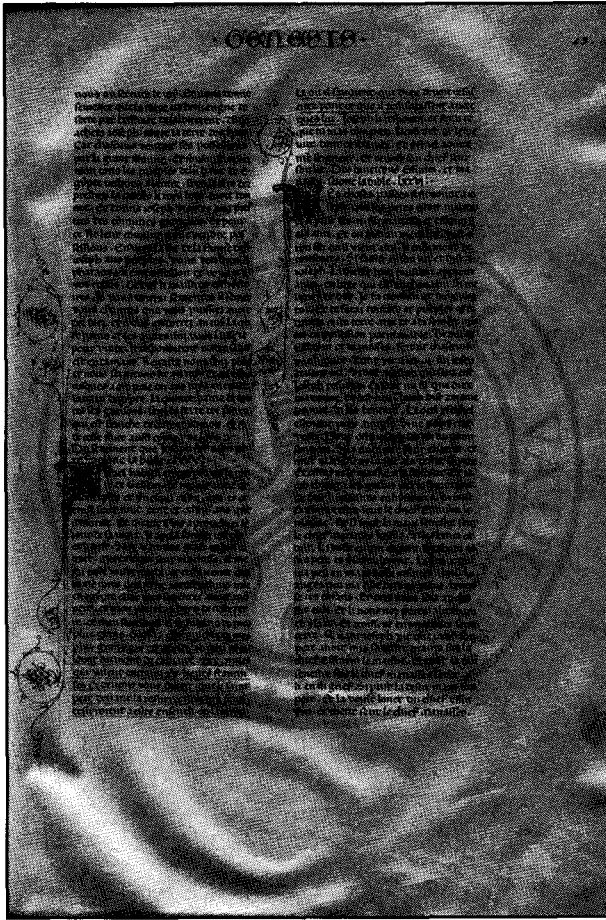


Figure 6

Sample monochrome image, compression of which is reported in Table 2. The image shown has been reduced from 3064×2052 pixels to 1000×670 pixels. The image is from the Vatican Library manuscript Barb.Lat. 613 (fol. 43r, a page of the book of Genesis, from a 15th-century Bible written in French).

images result in much larger data volumes, often in excess of a million bytes. With such high resolution, they are ideally suited to many usages. For these images, the protection offered by the watermarking is even more important.

• *Image display with the Scholar's Interface Application*

The image processing software in the SIA displays images, in a variety of ways, to support visual examination of the source manuscripts for scholarly research. The code consists of two functions: An import-time function is invoked when an image is imported into the SIA's workstation; a display-time function is invoked when an image is to be displayed with high quality on the image display.

The display-time function

1. Decompresses the image.
2. Crops and re-sizes the image.
3. Alters the contrast function to reduce bleed-through.
4. Corrects the color of the image for the display being used.
5. "Dithers" the image to the 65536-color (16-bit) palette of the display adapter of the image display.

The last two steps, described in [40], ensure that the image is displayed with accurate color. Steps two and three alter the size and contrast of an image to assist the scholarly study. The re-sizing is applied to linearized red, green, and blue image components, as described in the preceding subsection, so that the color of the image will not be altered by the re-sizing.

The display-time function supports a variety of cropping and re-sizing options. Any rectangular portion of the image may be selected for display and positioned arbitrarily within any rectangular portion ("window") of the image display, with or without causing the display window to be cleared before the image is displayed. (The SIA defines some particular rectangles that correspond to magnification choices on the image display; these are accessible via the tool bar.) The image may be reduced (by decimation—i.e., reducing the size of an image by discarding unneeded rows and columns) to fit the output window, be displayed at full resolution, or be enlarged by a factor of 2 in each dimension.

For monochrome images, step three of the display-time function may be invoked to reduce the visible effects of ink that has bled through the paper from the reverse side. **Figure 7** shows a portion of an image before and after the bleed-through reduction has been applied. Generally, the images that exhibit ink bleed-through are scanned from manuscript pages composed of essentially black ink on white paper with no illuminations; such images are generally scanned in monochrome. While an analogous bleed-through reduction could be applied to color images, there have not been enough color images exhibiting this problem to justify the effort.

The bleed-through reduction is accomplished by applying a transformation, illustrated in **Figure 8**, to linearized pixel brightnesses. This transformation decreases, by a factor of 2 (the inverse of the slope), the contrast between the brightness threshold (a value specified by the user) and the white value (i.e., the maximum brightness) of the image, while it increases the contrast between the brightness threshold and "minimum black" value of the image. When the brightness threshold is set (by the viewer) to the brightness of ink bleed-through, the visible effect of the bleed-through is indeed reduced. This bleed-through reduction is not intended to eliminate the bleed-through entirely, as this could also

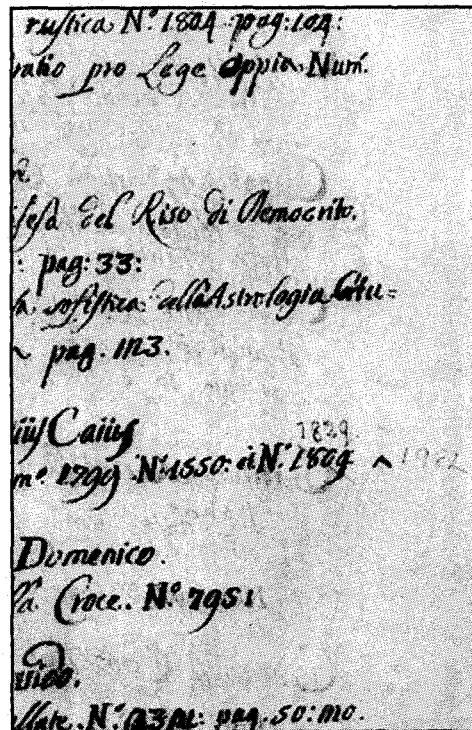
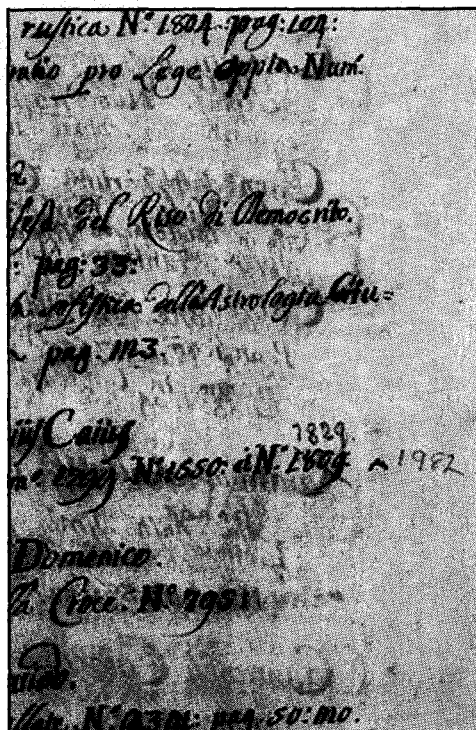


Figure 7

Illustration of bleed-through: (a) before and (b) after bleed-through-reduction processing. The figure shows a portion of a page from the Vatican Library's "Index Codicum Bibliothecae Ottobonianae," the alphabetical index of the manuscripts in the Ottoboni Library, today a portion of the Vatican Library, written in the early 19th century.

remove important information such as faint handwritten annotations. It is intended merely to reduce the visibility of the bleed-through. Note that in the figure the lightly written dates have not disappeared, although the bleed-through is significantly reduced.

The two parameters used by the bleed-through reduction are the brightness threshold, which is entered by the scholar, and the "minimum black," which is calculated as the greatest brightness level that exceeds no more than 1% of the pixels. A side effect of the bleed-through reduction is a decrease in the apparent brightness of the entire image. The application of an inverse display gamma function is used to compensate for this decrease in brightness. The form of the inverse display gamma function is

$$r_{i,j} = (p_{i,j})^{1/\gamma},$$

where $p_{i,j}$ and $r_{i,j}$ are the linearized input and gamma-corrected brightness values, respectively, for the pixel in

the i th row and j th column of the image. The value of γ is entered by the scholar.

We note that this bleed-through-reduction function is fairly unsophisticated. Most of the Vatican Library's manuscripts are in excellent condition, and little more is needed. However, more sophisticated techniques [3] have been developed by our colleagues in Spain to deal with more severe image degradations.

As discussed above, the import-time function creates two "derivative images" from the imported image for display on the system display. One of these, the thumbnail image, is a small version of the image used by the light-table viewer; this image may be either color or monochrome, depending on whether the original image was color or monochrome. Its maximum size is 144×144 pixels. The other derivative image is a larger, monochrome reference image. With the SIA, the scholar selects a subimage to be viewed on the image display by drawing a crop box on the reference image on the system display.




Figure 8

Contrast function used to reduce the visible effects of ink bleed-through.

This reference image is displayed only with shades of gray. The maximum size of the reference image is 512×512 pixels.

Both of the derivative images are prepared for display with a 256-color (8-bit) palette, by reducing the original image to the desired dimensions through decimation, color-correcting the image for the system display, and error-diffusing the result, as described in [40]. At import time, an auxiliary file is created to record miscellaneous data that describe the image: the original image dimensions and the factor by which the image was scaled to produce the reference image, and a brightness histogram of the image (used in the bleed-through reduction described above).

The future

The project's approach to broadening the availability of the Vatican Library treasures has been received with interest and enthusiasm by both the Vatican Library staff and those visiting scholars with whom we have discussed the project. In general, they felt that their research, as well as that of their colleagues, would greatly benefit from the remote availability of electronic copies of the volumes.

Experts and scholars at the Library were very pleased with the quality and fidelity of the digital reproductions of originals. At times, they even felt that our continuous concern about absolute image fidelity was beyond the requirements of most scholars. It is our conviction, however, that the closer the digital replica is to the original, the more often that replica will be sufficient for scholars' needs. The initial concerns about the strain to which manuscripts would be subjected during the scanning

process were addressed with careful handling of the volumes and strict monitoring of environmental conditions. Though information technology was already present in the Library, in the form of electronic cataloging, the introduction of digital imaging and related technologies through our project generated a great deal of interest from various Library departments (e.g., the cataloging and photographic departments) because of the additional possibilities it offers in their areas. On the other hand, the project greatly benefited from the contributions and expertise of the Library staff; their knowledge and experience with the microfilming and photography of manuscripts was invaluable and sped up scanning operations.

The system described in this paper is operational and in daily use. In developing it, the project has achieved many significant milestones:

- A scanning environment has been created within the Vatican Library that is capable of scanning original manuscripts with a high degree of safety.
- The ability to capture images of original manuscripts with high levels of detail and accurate color has been demonstrated.
- The capture of eighty high-quality images of manuscript pages per day per scanner has been achieved.
- JPEG compression parameters (quantization tables) suitable for the high-image-quality needs of the application have been determined.
- A digital watermarking technique that protects the scanned images from misappropriation has been developed.
- Over 21,000 images of manuscripts have been scanned.
- The 21,000 scanned images have been processed and made accessible through the Internet.
- The Scholar's Interface Application has been developed to help scholars find, examine, and study the images made available.
- Scholarly use of the images is in progress.

A small sample of the images produced by this project are currently available for inspection by the interested reader.³

The success of these achievements, given the scope of the project, validates the overall soundness of the design decisions. The choice of file formats (TIFF and JFIF), the subsets of these formats that were used, and the choice of compression technique (JPEG) were foundations of this system. We would make essentially the same choices today.

Were we to begin anew today, however, we would make some different decisions. Because the World Wide Web

³ A sampling of images of Vatican Library manuscripts is available on the IBM home page at URL <http://www.ibm.com/Features/library/manuscript.html>.

has gained great popularity since we began this project, we would undoubtedly choose to implement a Web server instead of a gopher server. The capabilities of commercially available Web browsers have also expanded (e.g., image caching is now a popular feature), so we would likely add features to a commercially available Web browser instead of writing a browser from scratch. Given the time and resources, we would choose to build a more robust image database than CIPA, as we discuss below.

The question of what scanned-image resolution should be used is a complicated one. The answer depends on many factors, including the size of the source materials, the amount of detail in the source materials, the distance from which the materials are normally viewed, and the presence and importance of magnification in the envisioned application(s). Each application must be judged on its own merits; a universal answer will not be forthcoming. Often, today's scanning technology is not capable of resolutions that will support all envisioned applications. In this case, it is often best to capture images at the highest quality that can be afforded—a moving target. Other factors also affect the choice of scanner, including cost, color quality, signal-to-noise ratio, dynamic range, speed, and illumination. The quality of the images captured with the Pro/3000 met the objectives for this project; we believe it was an appropriate choice. It would be our choice today for a project with similar requirements, but it is difficult to generalize this choice to applications with other requirements.

The system we implemented for this project is both a larger image system than we have attempted before and grossly inadequate to deal with the complete holdings of the Vatican Library (which number in the millions of pages) or many other libraries. Although much remains to be done, we have identified and begun to deal with many of the important issues that must be addressed before digital libraries can handle collections of millions of images (or other multimedia objects). These issues include

- Providing a seamless, distributed, network-centric system capable of managing massive quantities of data.
- Providing easily used human interfaces.
- Providing adequate intellectual property protection of the digitized materials.
- Automating the conversion of source materials, reducing the time required, and decreasing the cost of conversion.

The system we have devised for this project is both distributed and network-centric; indeed, giving users access to information through the Internet was a central design objective. Our prototype is not seamless, however, since it requires human intervention to manually move data from subsystem to subsystem. Nor is it able to handle massive quantities of data (objects numbering in the millions). Although we have achieved successful

operation and significant automation, we need to provide both a seamless connection of the numerous dispersed subsystems through an external network, such as the Internet, and an infrastructure that can support millions of objects.

The IBM Digital Library organization, formed after this project was under way, is developing a unifying framework for all IBM digital library applications. It is composed of five functional components: Create and Capture, Storage and Management, Search and Access, Distribution, and Rights Management. Each digital library will contain (at least) one instance of each component; many instances of each component are envisioned. A digital library of color images, for example, would have a different Create and Capture component from that of a digital library of digitized audio. This framework ensures that the different instances of each component can be combined, as needed, to develop unified digital library systems that satisfy diverse needs. At the center of this framework is the Storage and Management component. The IBM VisualInfo™ product [41, 42] is the preferred Storage and Management component. It was designed to manage massive amounts of data; we expect it to become a digital-library-system foundation for IBM and the foundation of the Vatican Library system in 1996.

Many aspects of this prototype are guiding our next steps toward service to museums and libraries with special collections. Recently, IBM joined the National Digital Library Foundation (NDLF) as an advisor and is using experience gained with this project to guide technological elements of NDLF planning.

Providing true ease of use is essential if we are to attract the library-user community to the digital library. We believe this community is not willing to adapt itself to the myriad quirks of technology; it expects the digital library to adapt to its needs. For example, structured queries, which require the user to know the field structure of the database and compose Boolean combinations of search criteria based on that structure, may be acceptable mechanisms for the captive audience or the technophile, but the typical library user neither wishes to know the underlying database structure nor to compose Boolean search criteria. At present, Vatican Library materials can be located in two ways: through a navigation of hierarchical menus or through a search on free-form text. These search aids have the right quality, but much work is still needed to determine how to best organize the underlying information so that location of materials can become truly intuitive.

Similarly, our Scholar's Interface Application has a user interface that was designed to be easy to use. It has a graphical interface whose operation can be learned in less than an hour, and it provides the desired function set. Feedback from the participating scholars has been

positive. The need remains, however, for even greater simplicity to serve library users who may have no training at all. While some system designers aim for a user interface that is as simple to use as that of a VCR, the user community of scholars may need a user interface that is as simple to use as that of a television set.

Protection of the intellectual property of the content owners is a vital problem, which we have only begun to address with the visible digital watermark. In the future, content owners will demand not only the means to protect their information from misappropriation, but also the means to regulate access to it and the means to collect usage fees from those who access it. Already, at Case Western Reserve University, an effort is under way to develop "rights-management" software⁴, which will provide access control, usage monitoring, and royalty assessment for intellectual properties contained in the on-line library; we hope soon to incorporate this software into the Vatican Library system. We recognize, however, that the problems associated with implementing access rules, potentially dependent on user, owner, object, user's operating system, and server's operating system, are many. It may be several years before a comprehensive system of permission management is developed.

Although our image watermarking serves to visibly mark an image and deter misappropriation, it is but one of an ensemble of security measures that one might envision providing. Processing techniques are desired for invisibly marking an image so that its subsequent alteration could be detected. An invisible "audit trail," hidden within an image, is also desired, so that misappropriated images could be traced. There has been much work done recently on invisibly marking images for various purposes [43-49]. Invisible watermarking is properly considered to be part of the field of steganography, the art of hiding information. Another desired technique would add marks that would be invisible when an image is displayed, yet visible when printed. The development of each such multimedia-object security technique would provide important assurance to the content owners that their information would not be misappropriated. Providing multimedia object security is one of the keys to the development of digital libraries. Without the assurance that the content is adequately protected, many owners will be unwilling to make their content digitally available, and the appearance of digital libraries will be greatly delayed. We, and many others, are actively working on devising new techniques for multimedia-object security.

The problem of converting existing materials to digital form is an incredible challenge; we have dealt with one of its most challenging aspects—the conversion of materials

that are varied, difficult to handle, and require conversion at the highest quality levels. Under these conditions, our conversion of roughly 80 images per day per scanner is a significant accomplishment; a capture cost between one and ten dollars per image is acceptable. In terms of capturing collections that involve millions of objects, however, this conversion rate is inadequate and the cost too high.

In the computer industry, many technologies double in performance and halve in price every two or three years. At this pace, it may be a decade before we can accomplish 1000 scans a day per scanner with the necessary quality. Until then, only the most valuable retrospective materials will be captured. There is another trend that offers us hope in seeing significant libraries on line: the digital creation of source materials (for which no conversion is required). Significant creation of digital copy is in practice now. Within a decade, this may be the dominant form of media creation, and digitally created media may be the dominant content of digital libraries.

Acknowledgments

Many people beyond the set of authors of this paper contributed significantly to the technical work described. Richard Cerreta, of the IBM Worldwide Workflow Consulting Practice, has made uncountable organizational and managerial contributions to this project. Lauren Kingman of the IBM Software Solutions Division and Jim Barker of Case Western Reserve University contributed greatly to the project's technical leadership. Howard Sachar of the IBM Thomas J. Watson Research Center provided considerable support and guidance to the project's technical team. Prof. Anthony Grafton of Princeton University chose the theme for the collection to be digitized, provided guidance on content selection, and provided much valuable feedback on the system's operation. At the IBM Thomas J. Watson Research Center, Gordon Braudaway contributed significantly to the development of the system's digital watermarking, Gerhard Thompson contributed significantly to the system's image sharpening and display functions, and Heidi Peterson provided guidance on the choice of the JPEG parameters. Ying Yao, Hon-Sum Wong, and Whan-Soo Kang contributed to the development of the scanner that was so important to this project. Tareq Alrashid of Case Western Reserve University has made many improvements to the CIPA database. Joseph Tarsia of TTI Inc. was instrumental in enhancing the book easel. David Singer and Jon Reinke of the IBM Almaden Research Center provided valued guidance.

Many IBM employees contributed invaluable nontechnical support to this project. Lois Jackson, Robeli Libero, José Schifflini, and Eric Marler, all of IBM Latin America, have been proponents of the project from the

⁴ James A. Barker, Director, Digital Media Laboratory, Case Western Reserve University, personal communication.

time of its initial vision, as have Vincent Yannuzzi, Steven Cutignola, Jean-Paul Jacob, and Richard Abineri.

Other people, not associated with the project, contributed to the paper. We must acknowledge the assistance of the three referees; addressing their insightful comments led to a much improved paper. A list of references on invisible watermarking of images was compiled for us by Minerva Yeung of Princeton University; we thank her for that contribution.

PS/2, RISC System/6000, and OS/2 are registered trademarks, and SearchManager/2 and VisualInfo are trademarks, of International Business Machines Corporation.

TIFF is a trademark of Aldus Corporation.

References and notes

1. H. M. Gladney, E. A. Fox, Z. Ahmed, R. Ashany, N. J. Belkin, and M. Zemankova, "Digital Library: Gross Structure and Requirements" (report from a March 1994 workshop), *Research Report RJ-9840*, IBM Research Laboratory, San Jose, CA, May 1994. This is also available as Virginia Tech. CS TR-94-25 or in NCSTRL via WWW at <http://www.ncstrl.org>. An extended abstract of this paper appeared in Proc. Digital Libraries '94, *Proceedings of the Workshop on On-line Access to Digital Libraries*.
2. *The Documents of the New World*, CD-ROM available from Archivos y Bibliotecas, AIE, Apdo. de Correos 179, 28080, Madrid, Spain.
3. Julián Bescós, Juan Pedro Secilla, and Juan Navarro, "Filtering and Compression of Old Manuscripts by Adaptive Processing Techniques," *Digest of Technical Papers of the 1990 SID Symposium*, vol. XXI, 1990, pp. 384-387.
4. Fred Mintzer and John D. McFall, "Organization of a System for Managing the Text and Images that Describe an Art Collection," *Image Handling and Reproduction Systems Integration Technology, Proc. SPIE 1460*, 38-49 (1991).
5. Fred Mintzer, Ying L. Yao, and John D. McFall, "A Computer System for Scanning and Cataloging the Art of Andrew Wyeth," *Spectra, the Magazine of the Museum Computer Network*, pp. 9-14 (Summer 1992).
6. Fred Mintzer and Ed Pariser, "Porting a Calibrated Color Image Across Corporate Boundaries," *Human Vision, Visual Processing, and Digital Display IV*, B. Rogowitz and J. Allebach, Eds., *Proc. SPIE 1913*, 433-439 (1993).
7. J. K. Hong and S. Sugita, "A Color Image Database for an Ethnology Museum—A Multiwindow System for Electronic Cataloging and Browsing of Ethnological Samples on the PC," *Proceedings, Cologne Computer Conference*, Cologne, Germany, September 1988, pp. 53-60.
8. J.-K. Hong, J. Takahashi, M. Kusaba, and S. Sugita, "A Motion Picture Archiving Technique and Its Application in an Ethnology Museum," *Proceedings, 3rd International Conference on Database and Expert Systems (DEXA92)*, Valencia, Spain, September 1992, pp. 209-214.
9. S. Yamada, J.-K. Hong, and S. Sugita, "Development and Evaluation of Hypermedia for Museum Education; Validation of Metrics," *ACM Trans. Computer-Human Interaction* 2, No. 4, 284-307 (1995).
10. J.-K. Hong, J. Takahashi, M. Takahashi, and S. Sugita, "An Approach to the Digital Museum, Multimedia Systems for an Ethnology Museum," *Proceedings, 3rd International Conference on Hypermedia and Interactive Museums*, San Diego, October 1995, pp. 87-95.
11. The National Gallery of Art's FACT SHEET describing the \MIRA! project may be obtained from Ric Foster, Digital Imaging Specialist, The National Gallery of Art, 4th Street and Constitution Avenue, N.W., Washington, DC 20565.
12. "TDI Pro Scanner Color Scanner System: Operator's Manual, Version 1.3," available from the Image Applications Department, IBM Thomas J. Watson Research Center, Yorktown Heights, NY 10598.
13. "TDI Pro Scanner Color Scanner System: Reference Guide, Version 1.3," available from the Image Applications Department, IBM Thomas J. Watson Research Center, Yorktown Heights, NY 10598.
14. F. P. Giordano, W. S. Kang, T. McCarthy, H. S. Wong, and Y. L. Yao, "A High-Performance, High-Fidelity Color Scanner," *IS&T Symposium on Electronic Photography: Part of the 44th Annual Conference, Advance Printing of Paper Summaries*, pp. 293-296, The Society for Imaging Science and Technology, Springfield, VA, 1991.
15. Y. L. Yao, F. P. Giordano, H.-S. Wong, W. Kang, and J. C. Lee, "Design Considerations for a High Quality Camera Type Scanner," *Proceedings, IS&T Seventh International Congress on Advances in Non-Impact Printing Technology*, Portland, OR, October 1991, pp. 441-450.
16. "Color Image Portfolio Assistant," IBM Product Number 5764-018, available from IBM Mechanicsburg, PA.
17. "3995 Introduction and Planning Guide," Publication No. GA32-0121-09; available from IBM Mechanicsburg, PA.
18. Information about Geac products is available from the Geac Computer Corporation at URL <http://www.geac.com/>.
19. "The USMARC Formats: Background and Principles," *National Information Standards Organization*, Bethesda, MD; available from [gopher://marvel.loc.gov/00/listarch/usmarc/usmarc.pri](http://marvel.loc.gov/00/listarch/usmarc/usmarc.pri).
20. David Bearman and John Perkins, "Computer Interchange of Museum Information Committee: Standards Framework for the Computer Interchange of Museum Information," *Museum Computer Network*, Silver Spring, MD, May 1993, pp. 43-44.
21. Information about gopher is available from the University of Minnesota at boombox.micro.umn.edu in directory /pub/gopher.
22. Information about GoServe is available from the IBM Corporation at URL <http://www2.hursley.ibm.com/goserve/>.
23. "SearchManager/2 Overview," Publication No. GK10-2018-01; available from IBM Mechanicsburg, PA.
24. *TIFF Tag Image File Format Revision 6.0*, Aldus Developers Desk, Aldus Corporation, 411 First Avenue South, Seattle, WA 98104, April 1992.
25. Adele Goldberg and David Robson, *Smalltalk-80: The Language and Its Implementation*, Addison-Wesley Publishing Co., Reading, MA, 1983.
26. E. S. Schlig, "A 3072 x 32-Stage TDI Imaging Device," *IBM J. Res. Develop.* 35, No. 1/2, 283-287 (January/March 1991).
27. Gordon W. Braudaway and Hon-Sum P. Wong, "Color Calibration of the TDI Pro Scanner," *Research Report RC-19240*, IBM Thomas J. Watson Research Center, Yorktown Heights, NY 10598.
28. Fred Mintzer and Helen Delp, "Requantization of Scanned Color Based on the CIE L*a*b* Color Space," *Research Report RC-15578*, March 1990; IBM Thomas J. Watson Research Center, Yorktown Heights, NY 10598.
29. R. W. G. Hunt, *The Reproduction of Colour, in Photography, Printing & Television*, Fountain Press, Tolworth, England, 1987.

30. Fred Mintzer and Gordon W. Braudaway, "Processing Color Images While Preserving Their Color," *Proceedings of the Ninth Workshop on Image and Multidimensional Signal Processing*, Belize City, Belize, March 1996, pp. 106-107.
31. Gerald Goertzel, Fred Mintzer, Gerhard Thompson, and Yi-Hsin Chen, "The Design of Interpolation Filters and Their Application to Image Scaling and Rotation," *Research Report RC-20175*, 1988; IBM Thomas J. Watson Research Center, Yorktown Heights, NY 10598.
32. K. Blair Benson, *Television Engineering Handbook*, McGraw-Hill Book Co., Inc., New York, 1986, pp. 2.31-2.32.
33. Jim Pickerell and Andrew Child, *Marketing Photography in the Digital Environment*, DiSC Co., 100 Frederick Ave., Rockville, MD 20850, 1994.
34. CCITT Recommendation T.81 | ISO IS 10918-1, Information Technology—Digital Compression and Coding of Continuous-Tone Still Images, Part 1: Requirements and Guidelines.
35. William B. Pennebaker and Joan L. Mitchell, *JPEG: Still Image Data Compression*, Van Nostrand Reinhold, New York, 1993.
36. Gregory K. Wallace, "The JPEG Still Picture Compression Standard," *Commun. ACM* 34, No. 4, 30-44 (April 1991).
37. Eric Hamilton, "JPEG File Interchange Format, Version 1.02," C-Cube Microsystems, Milpitas, CA, September 1, 1992.
38. Heidi A. Peterson, Albert J. Ahumada, and Andrew B. Watson, "An Improved Detection Model for DCT Coefficient Quantization," B. Rogowitz and J. Allebach, Eds., *Human Vision, Visual Processing, and Digital Display IV, Proc. SPIE* 1913, 191-201 (1993).
39. Andrew B. Watson, "DCT Quantization Matrices Visually Optimized for Individual Images," *Human Vision, Visual Processing, and Digital Display IV*, B. Rogowitz and J. Allebach, Eds., *Proc. SPIE* 1913, 202-216 (1993).
40. F. C. Mintzer, G. Goertzel, and G. R. Thompson, "Display of Images with Calibrated Color on a System Featuring Monitors with Limited Color Palettes," *Digest of Technical Papers, 1992 SID International Symposium*, Society for Information Display, Playa del Rey, CA, Vol. XXIII, 1992, pp. 377-380.
41. "VisualInfo General Information," Publication No. GC31-7728-02; available from IBM Mechanicsburg, PA.
42. "VisualInfo User's Guide," Publication No. SC31-7670-01; available from IBM Mechanicsburg, PA.
43. D. Caruso, "Two Plans for Watermarks, Which Can Bind Proof of Authorship to Electronic Works," Technology section, the *New York Times Service*, August 7, 1995.
44. "A Digital Notary," *Science News* 147, 135 (March 4, 1995).
45. E. Sapwater and K. Wood, "Electronic Copyright Protection," *PHOTO: Electronic Imaging* 37, No. 6, 16-21 (1994).
46. G. L. Friedman, "The Trustworthy Digital Camera: Restoring Credibility to the Photographic Image," *IEEE Trans. Consumer Electron.* 39, No. 4, 905-910 (November 1993).
47. R. G. van Schyndel, A. Z. Tirkel, and C. F. Osborne, "A Digital Watermark," *Proceedings of the IEEE International Conference on Image Processing, ICIP 1994*, IEEE, Piscataway, NJ, Vol. 2, November 1994, pp. 86-90.
48. W. Bender, D. Gruhl, and N. Morimoto, "Techniques for Data Hiding," *Proc. SPIE* 2420, 40 (1995).
49. S. Walton, "Image Authentication for a Slippery New Age," *Dr. Dobb's Journal of Software Tools for Professional Programmers* 20, No. 4, 18-26 (April 1995).

Received January 31, 1995; accepted for publication June 22, 1995

Frederick C. Mintzer IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (MINTZER at YKTVMV, mintzer@watson.ibm.com). Dr. Mintzer has been a research staff member at IBM since 1978, when he received the Ph.D. degree in electrical engineering and computer science from Princeton University. He is currently the manager of the Image Library Applications project in the Image Applications Department. Since 1981, his work has focused on image processing methods to better process images for display or printing. Methods developed by Dr. Mintzer have been included in several current or former IBM offerings, including the Image View Facility, IBM OfficeVision®, ImagePlus®, Ultimedia® Color Delivery/400, and the Color Image Portfolio Assistant. His current research interests center on methods to more accurately capture, store, and display color images. In recent years, he has led IBM Research efforts to develop image library systems for the painter Andrew Wyeth, the U.S. National Gallery of Art, and the Vatican Library. Dr. Mintzer's work has resulted in eleven patents and more than thirty publications. He is a member of the Society for Imaging Science and Technology and a senior member of the IEEE.

Leonard E. Boyle Biblioteca Apostolica Vaticana, Vatican City, 00120 Italy. Father Boyle has been Prefect of the Vatican Library since 1984. He joined the Dominican Order in 1943; he was a professor of history of theology at the University of St. Thomas in Rome from 1956 to 1961, and a professor of palaeography and diplomatics at the Pontifical Institute of Medieval Studies in Toronto from 1961 to 1984. Father Boyle received his Ph.D. from Oxford University in 1956. He is a fellow of the Royal Historical Society, London, and of the Medieval Academy of America, a corresponding fellow of the British Academy, and a member of the Pontifical Academy of Sciences.

Albert N. Cazes IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (CAZES at YKTVMV, cazes@watson.ibm.com). Dr. Cazes graduated from the City University of New York with a B.S. in physics in 1969. He received an M.S. from Hunter College, New York, in 1975 and a Ph.D. from the Polytechnic University, New York, in 1982, both in mathematics. From 1970 to 1983 he worked in the Office of Data Processing of the City of New York, where he computerized several administrative systems. From 1983 to 1993 Dr. Cazes taught courses in mathematics, statistics, and computer science on a part-time basis at CUNY's Baruch College and SUNY's College at Purchase. In 1983 he joined the IBM Research Division, where he has worked in the areas of videotext, functional programming, parallel processing, virtual reality, and image processing. Dr. Cazes is a member of Sigma Xi and the Society for Imaging Science and Technology.

Brian S. Christian *Digital Media Laboratory, Case Western Reserve University, 10900 Euclid Avenue, Cleveland, Ohio 44106 (bsc4@po.cwru.edu).* Mr. Christian is a manager in the Digital Media Laboratory. He joined the group in 1990 to develop client/server software. He received a B.S. in computer engineering from Case Western Reserve University and is currently an M.S. candidate in a practice-oriented program offered by the school of engineering at Case Western Reserve University.

Steven C. Cox *Digital Media Laboratory, Case Western Reserve University, 10900 Euclid Avenue, Cleveland, Ohio 44106 (scc@po.cwru.edu).* Mr. Cox has been a software developer at Case Western Reserve University since 1990. His past projects have included Smalltalk implementations of multimedia and hypertext authoring and presentation systems. He is currently a team member working on a comprehensive Rights Management system used to control networked access to digitized intellectual properties. Mr. Cox holds a B.S. degree from Case Western Reserve University.

Francis P. Giordano *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (GIORDAN at YKTVMV, giordan@watson.ibm.com).* Mr. Giordano is a senior engineer and project leader of the Image Capture Systems group. He received an associate degree in electrical engineering from the RCA Institute of Technology, New York, in 1973 and a B.S.E.E. from the University of Bridgeport, Connecticut, in 1980. Mr. Giordano joined the Thomas J. Watson Research Center in 1973 and has since worked on various aspects of high-end imaging systems utilizing charge-coupled devices.

Henry M. Gladney *IBM Research Division, Almaden Research Center, 650 Harry Road, Almaden, California 95120 (GLADNEY at ALMADEN, gladney@almaden.ibm.com).* Dr. Gladney has been an IBM research staff member since receiving his Ph.D. in chemical physics from Princeton University in 1963. He has held management positions in physical sciences and in computing services, and is a member of the Association for Computing Machinery and a fellow of the American Physical Society. Dr. Gladney has published in scientific journals in physics, chemistry, laboratory automation, computer security, and database technology. He was a principal contributor to three software projects which became widely used IBM products: RACF™, a component of OS/MVS and OS/2 LAN services which aids administrative control and information security; EDX/1, a multiprogramming-multitasking operating system for process control computers; and ImagePlus VisualInfo, an operational image service. Dr. Gladney's current interest is data management in extremely large systems, with special attention to Digital Library opportunities.

Milton L. Kelmanson *Department of Electrical Engineering, Pontifical Catholic University of Rio de Janeiro, Rio de Janeiro, Brazil (vlpgroup@ludovisi.vlp-ibm.puc-rio.br).* Professor Kelmanson received his Ph.D. from Brown University in 1976; he is a member of Sigma Xi and the Scientific Research Society of North America. Besides his interest in operations research and applications of discrete optimization techniques to the layout of VLSI chips, Professor Kelmanson manages the participation of the Catholic University in the Vatican Library Project.

Jack C. Lee *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (JCLEE at YKTVMV, jlee@watson.ibm.com).* Mr. Lee is an advisory engineer in the Image Applications Department working on color scanning and processing. He received a B.S.E.E. degree from City College of New York in 1974, and an M.S. degree in computer science from the Polytechnic University, New York, in 1987. From 1974 to 1983, he worked at IBM East Fishkill/IBM Kingston on VLSI microprocessor design and software verification. Since 1983, Mr. Lee has been a member of the Image Technologies Department, Computer Science, at the Thomas J. Watson Research Center. He has worked on various projects relating to high-resolution color image capture and archiving, and also image processing for displays and for printing. His current interests include the integration of high-resolution imaging into the Digital Library arena. Mr. Lee is a member of the IEEE Computer Society.

Antonio C. Lirani *IBM Brazil Industry Marketing Education and Research, São Paulo, Brazil.* Mr. Lirani received his B.S. in mechanical engineering from the University of São Paulo, São Carlos, in 1970; he is working toward his Ph.D. in digital systems at the Politechnic School of the University of São Paulo. Since joining IBM in 1976, Mr. Lirani has worked in Industry Marketing Education and Research and has held management positions at the IBM Sumare Plant Information Systems, the IBM Software Institute, and the Software Engineering Institute at the IBM Latin American Technology Institute. He currently manages the IBM Brazil Academic Programs Department. Beginning in January 1994, during the 18 months of its first phase, Mr. Lirani was the manager of the IBM Brazil/PUC-Rio subsystem of the Vatican Library Accessible Worldwide Project.

Karen A. Magerlein *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (ANDRSON at YKTVMV, andrson@watson.ibm.com).* Ms. Magerlein graduated from Duke University with a B.S. in computer science in 1980, joining the Image Technologies Department of the Research Division that same year. She has worked on the development of programming tools, techniques for binary image manipulation, and the JPEG compression standard. Ms.

Magerlein received IBM Outstanding Innovation Awards in 1985 for her contributions to the Image View Facility and in 1991 for her work on bilevel image compression software; she is a co-inventor on sixteen patents. Ms. Magerlein is currently an advisory engineer with the Image Applications Department, working on software for digital library applications. She completed her M.S. in computer science at Columbia University in 1995.

Ana M. B. Pavani *Department of Electrical Engineering, Pontifical Catholic University of Rio de Janeiro, Rio de Janeiro, Brazil (vlpgroup@ludovisi.vlp-ibm.puc-rio.br).* Professor Pavani received her D.Sc. degree from the Federal University of Rio de Janeiro, Brazil, in 1983. Besides being interested in automation in general and in administrative and library automation in particular, she is the technical manager of PUC-Rio in the Vatican Library Project.

Fabio Schiattarella *IBM SEMEA (Southern Europe Middle East Africa) New Markets Unit, Piazzale Giulio Pastore 6, 00144 Rome, Italy (schiatta@vnet.ibm.com).* Mr. Schiattarella joined the IBM Rome Scientific Center in 1988, working in the area of image processing. His other areas of professional interest are multimedia and digital libraries. Since 1993 he has managed local operations in Rome for the IBM-Vatican Library Project. Mr. Schiattarella received his degree in electronic engineering from the University of Rome in 1987.

OfficeVision, ImagePlus, and Ultimedia are registered trademarks, and RACF is a trademark, of International Business Machines Corporation.