

Imaging and the Internet

By Michael Greenhalgh

gremarth@fac.anu.edu.au or Mike.Greenhalgh@anu.edu.au

<http://rubens.anu.edu.au>

The paper is an introductory "how to do it" briefing for creating and using computer-accessible images with WWW/Mosaic. with appendices discussing in detail how I proceed with my ArtServe server in Canberra (Appendix I for description). It discusses the techniques and the hardware and software required to make large quantities of images transparently available to the Internet. Areas covered include input (scanners & video cameras. storage (laserdisks and computer hard disks) specialised hardware (framegrabbers and image compression boards). software (freeware shareware write-it-yourself and WWW/Mosaic (how to generate pages automatically from a set of data records with or without inline images). Strategies for dealing with higher-quality images in a WWW/Mosaic networked environment are described in Appendix II. A brief survey of Australian initiatives with WWW/Mosaic is given in Appendix III.

INTRODUCTION

To people who deal with images (scientists as well as art historians, and now artists) the web, together with recent advances in reasonably-priced digitizing equipment, memory, storage and machine speed, offers opportunities which simply did not exist even eighteen months ago. Indeed, cast your minds back to what was needed three years ago for any project involving images: one platform, a dedicated and expensive commercial package, probably, a platform-specific image format, and no networking to speak of. But with the appearance of network-transparent software, freely available, the whole scene has changed: we need no longer think, and we should no longer think, of one-user one-machine setups (expensive "fortresses" which are difficult to update), but a client-server setup which works just as well between Bangkok and Canberra as between central Bangkok and Chiang Mai.

Yet even with free and flexible software, developing systems which will replay the considerable investment they require takes time in thought, experimentation, and

user-testing. We must reasonably expect any imaging work to survive for at least two or three years, before higher-quality output devices make today's images look; very inferior in quality.

This paper discusses the advantages and the disadvantages of various technologies, and describes how I have set up my ArtServe server, and why I went the path I did. This might help other people, or at least explain where I have gone wrong. I make several assumptions: (a) that any worthwhile setup will require several thousand images; (b) that allowance must somehow be made for people with slow lines; (c) that flexibility in the quality of image delivery is therefore a good idea; (d) that the quality required of images will rise as display technologies get cheaper and better; (e) that network speed will (Oh vain hope!) catch up with the things people want to do on the networks, including transmitting big images:

1: INPUT & STORAGE

1. 1: Input

Inputting large quantities of images to make them computer-addressable requires a production-line facility, and some speed, not to mention a methodology for naming the image files which will last some years and, for preference, work on all platforms. Note that "addressable" does not necessarily mean "digitized": see the section on laserdisks, below.

SCANNING seems to me problematical for large tasks. Not only is it slow, and can produce very large files, but it is operator-intensive, since automatic hoppers are few and far between, and expensive. Assuming each scan takes five minutes to set up and execute, then twelve images are processed in one hour, which is too few for my purposes. Then again, taking a detail from a scan requires even more intervention: since a scanner has no zoom, the whole work is scanned, presumably at a higher resolution, and the required detail excised for further treatment.

VIDEO INPUT is to be recommended, always assuming the user does not mind an image the highest resolution of which can be no higher than that of a video scan, which is about 760 by 525 pixels. If this is acceptable, two devices recommend themselves. A S-VHS VIDEO CAMERA with a macro facility mounted on an enlarger stand allows the capturing of photographs, documents, and book illustrations. The **Sony PHV-A7E Photo Video Camera** does the same thing for slides: with positive/negative, focus, colour correction, and 6x zoom facilities, this device works very well for offering the video image, again at S-VHS quality, for digitizing. The

device consists of a "light box" surmounted by a carrier for a negative or slide holder, and with a lens above that. The focusing can be automatic (button), or the user can disengage this and do it by hand. Controls on the side allow: (a) a change from negative to positive (i.e. the machine will handle negatives as well as transparencies, whether these latter be monochrome or colour); (b) an iris to allow changes to the amount of light thrown on the image; (c) white balance (read the manual); (d) a colour adjustment, which gives overall more or less colour to the image. On the front is a colour balance wheel (with an on/off switch) allowing the user to counter-adjust for images which are too blue, red or whatever. Very important is the motor-driven (therefore smooth and accurate) zoom control, which allows anything up to six-times magnification. If autofocus is on whilst you are zooming, then focus should be preserved.

Getting the video image into digital form requires a computer with a framegrabber, which can convert the analog video signal into a bitmap, preferably one with 24 bits of information, eight each for the red, green & blue channels. The make of computer doesn't matter: I use both an Amiga and a DECStation, but anyone which could produce - say - a TIFF or a PICT or a JPEG would do.

FRAMEGRABBERS provide the alternative digitizing hardware to scanners, and they are machine-specific, usually with their own software. Little of use can be said about them here, except that they are getting faster all the time, with real-time JPEG cards (which can convert full-frame incoming live video, at 25 frames per second) now a reality, assuming a suitably speedy bus to the hard disk.

1.2: Storage: Hard Disks

Storing large quantities of images is getting cheaper every day, but strategies are still needed to make the best use of any available space, bearing in mind not only that materials must be backed up to safely on tape (so the more compact the better), but also that the work you do on images should ideally last for several years, so other media such as CDROM and laserdisk can look attractive.

1.3: Storage: CDROMs

Certainly, the CDROM is a useful storage medium - cheap to reproduce, capacious, and not too difficult to write yourself, especially if you are a masochist. But the procedure requires prior storage on a large hard disk: a CDROM cannot conveniently be written direct from a framegrabber. How many images will a CDROM hold? Assuming the images are compressed from video resolution, in 24 bits, then each

resultant JPEG image will be under 100Kb - which means that 5,000 will fit easily on a CDROM, in 500Mb. One tip here if you are writing CDROMs yourself is to make sure and use numeric filenames which abide by the old DOS "8+3" rule, and capital letters only: the ISO9660 standard for CDROMs does indeed in theory allow much more flexibility than this; but practicalities dictate that this restriction must be obeyed if files are to record on an ISO CDROM to be played in any machine. Thus "12234.GIF" and "14643.JPG" are licit filenames: whereas "plc.bangkok", "palace.bangkok" and "PALACE.BANGKOK" are not.

But CDROM is not the only viable storage medium for people interested in images, because its capacity is relatively puny. This applies to the Kodak PhotoCD which, with 100 images, seems pointless except for archiving, for which it is very good and relatively cheap.

1.4: Storage: Laserdisks

I am not one of those who think the day of the laserdisk has gone. Antique computer equipment from, say, 1983) would be a masochist's joy; but one of my Department's laserdisk players dates from that year. It still works perfectly, and it is still serviceable. For many of the gurus, of course, the laserdisk is obsolete technology, because they have only ever used one for viewing movies (or in karaoke bars?); their first reaction is - be digital! Encode the 54,000 (or 108,000 on the two sides) onto a large hard disk, and proceed from there. Not everybody has such capacities (over 5Gb for one side of a videodisk, and over 10Gb for both sides) at their disposition.

Each commercial laserdisk platter in CAV format (Constant Angular Velocity) can hold up to 54,000 frames per side, at video resolution and these can be stills, or a linked series of images (i.e. a movie), with stereo sound thrown in. It's not digital, but who cares? The technology interfaces well with a computer, and frames can be digitized as they are needed, perhaps via a JPEG board (see below, the section on commissioned software for laserdisk).

The laserdisk is a recordable medium as well. Sony, for example, market a VHS recorder the removable platter for which holds some 36,250 frames on each side (other brands are available at S-VHS quality). Recordability is a decided strength, because the speed at which images can be encoded is much higher than is possible with a scanner or, indeed, with any kind of digitizer. I've got 13,000 of my research slides encoded, and I still have a lot left to do!

But is the technology networkable? yes. Why not a computer-controlled laserdisk player, its frames recorded in a simple database, which responds to requests from over the network by offering the required frames to a framegrabber, which the computer then sends on their way? Simple, and it keeps your hard disk (and network) unclogged, because the laserdisk is the storage, not the hard disk.

This is what I have done on my ArtServe server. The following are required: (a) a framegrabber talking to software which automates its operation; (b) a strategy for handling the images on the computer side, preferably with scripts for writing HTML documents. What I have done is to set up a directory called laserdisk, which contains within it the HTML indexes and pages (both generated by a perl script called "salami" see Appendix II). Under this directory are the image directories, and these hold only the small GIF images which act as "inlines" in HTML documents.

Placing such material on CDRom is a distinct possibility, but the material probably needs storing on hard disk first: and some twelve CDRoms (and hence a jukebox) would be needed to hold 54,000 images at video resolution. This is peanuts these days in storage terms (our StorageTek Tape Robot currently holds one terabyte of data, and this capacity will increase rapidly), but the capacities are easier to talk about than to fill.

1 .5: Commissioned Software for the Laserdisk

How does material get from the laserdisk to the hard disk? This is where specially commissioned software is used, written some in C, some in perl, which talks not only to the laserdisk but also to the framegrabber program on the computer. For example. I have a program which allows the user to specify the following parameters:

sony IN OUT /where / FORMAT SCALING

Thus "sony 12200 12299 /path/ GIF 10" would place GIF images at 10% their original size in the specified directory and give them .GIF suffixes; whilst "sony 12200 12299 /path/ JPG" would do the same for full-sized JPG images.

Thus dumping full-sized images to the hard disk (i.e. bypassing the framegrabber or the laserdisk for day-to-day operation) is easy. This adds flexibility: if you have the hard disk capacity, dump everything onto the disk; if not, service Internet requests from the laserdisk, which feeds the requested frames one-by-one to the framegrabber.

2. APPLICATIONS

2.1: An X-Window Image Database in tcl/tk

The first stage in making our materials computer-accessible was to construct a short database file for a selection of images of prints from the mid-15th century to the end of the 19th century, and to digitize the related images. The database was simple, not to say crude: artist-name, born, died, work-title, date-of-work, comments, path to filename of image. This exercise taught me that: (1) image databases with up to 100 images are dead easy, but useless; (2) image databases with 1,000 images and more are a struggle, and very exacting to design, to prepare and then to manage.

To convert the "naked" database to a real application, I employed a programmer (David Clarke, Faculties Computer Unit, Australian National University, David.Clarke@anu.edu.au) to write an interface for me in tcl/tk for my DECStation running Ultrix. The user can interrogate the Prints Database by filling in a template. This may be done "raw", as it were (fill it in, and hope for the best), or by selecting from the data on offer in one or more of the fields via scrolling lists:

The selected records are then displayed in their own scrolling list, or "selected records window", and clicking on a record brings up a thumbnail of the associated image; clicking on the thumbnail brings up the image itself, which is displayed using xv at half the side (i.e., one quarter the area of the stored image. This is to conserve memory, because we envisage a user wanting several images on the screen at once. Any images of special interest can be enlarged as required:

The "selected records window", once finished with can be closed, and another query begun; alternatively, it can be treated as a database, and further queries run upon it. Since much of the material of the Prints Database concerns classical, mythological and historical subjects, a button offer the user the chance to execute a string-search through one or more books (from the Bible and Shakespeare to historical and philosophical works). When time and money allow, we shall scan some out-of-copyright editions of texts of interest to Art Historians, such as Giogio Vasari, and to interrogate them using the TextSearch program which we have developed (see below: "WWW/Mosaic Text Queries"). This excellent and flexible application (it will work with any collection of data, images, and text files) is network-accessible only in

the sense that a user with the appropriate permissions can log into my DECStation and run a windowing session on their terminal/machine. It is not, and cannot be, available under WWW/Mosaic, for it serves a more focused teaching purpose.

2.2: Enter the Network: WWW/Mosaic

The next stage in the process of broadcasting images, as it were, was to look at Internet access, and I was fortunate to do so when WWW/Mosaic had just appeared on the scene. I cut my teeth on an XMosaic client, and used its features (especially VIEW SOURCE and its help screens) to study how presentations were put together. David Green (of the ANU's Server called <http://life.anu.cdu.au>) gave me some additional tips and an initial push. I decided to try WWW/Mosaic out.

What constitutes a trial? I started with some 70 images, mostly of prints and classical architecture, using WWW/XMosaic, and put these on servers in the USA, Europe and Australia in June 1993. The reaction was positive, so decided it was time to do something serious rather than simply pretty. In July 1993 I therefore prepared what was intended to be a one-hour student "tutorial" on *The Palace of the Emperor Diocletian at Split in ex-Yugoslavia*. and placed that on the same servers. This took me some 18 hours to construct, including digitizing some 40 of the circa 80 images (the others were already done). You can already see how long it might take to construct a whole "computerised course", although one would probably get a little faster with practice.

However, these trials were like fragmentary illustrated books: you could turn the pages, admire and download images, but you couldn't choose what you wanted to see. Clearly, some database work was in order. I already had the 2800-record datafile for the Prints Database. The question was: how to make the database available to WWW/Mosaic, without inserting everything by hand, like filling up a stamp album? Writing HTML **documents** is not difficult. but it is time-consuming and exacting, for they need to be correct. Obviously. a program was needed to automate the task - a program flexible enough to take care of the various "flavours" of HTML document and audience I had in mind.

2.3: Salami: an HTML record slicer

With programming help, I therefore developed a filter which, for obvious reasons, I call "salami": the data file goes in at the top, and various switches are set, for example: length of the Mosaic index page(s); length of the display pages when no inline images are included; number of inlines across the page when images ARE included; and

number of inlined records per page. The program assumes the inlines (GIFs) are in the same directory as the JPEGs (i.e. the big images) unless you specify otherwise. The records appear on the page (with or without the inlines above them) with basic information: for the Prints Database, the artist's name, birth and death dates, title of work, and date of work. The title of the work is hot-spotted, and clicking on it brings up the JPEG image. Different applications (such as classical architecture, where there are no or few artist-names) naturally require a different setup.

The program works just like a salami-slicer. Although there are obvious problems with differing screen resolutions (i.e. what looks good on an X-terminal bulks too large on a small-screen Macintosh), I find that about 17 lines per index page and per no-inline-images page works well; and that, for inline images, five images across with the associated records directly underneath, and with three such blocks down the page (i.e., 15 records (and their inline images)) per page looks reasonable:

Doing this kind of thing by hand would be impossible: to give an example, the prints Database, which has some 2,800 records, generates in its inline-images configuration, 12 index pages and 209 pages to which they refer.

2.4: WWW/MOSAIC Database Queries

The database setup described above is quite slick, but it suffers from two main disadvantages.

The first is that the database appears in an order or orders I determine, except that there are index pages so that a user can jump to the main target area. The two "flavours" currently operating for the Prints Database are: (1) alphabetically by artist, with the work title alphabetically within each artist-list; and (2) alphabetically by subject-matter (from allegory to war), with the artist-list alphabetically within that. Because of the largish number of records involved, the user cannot often go directly to a known image: getting Durer's "Rhinoceros", for example, would require looking through several pages of records either classified under "Durer" or classified under "animal".

The second is that the user struggling against a slow network has to wait for a lot of unnecessary material to float past the eyes before the target image (we trust!) appears. This can of course be mitigated by switching off the "autoload images" function, but not all users know about this.

Clearly, some kind of database-like true query facility is required, so that the user can specify "Durer" and "rhinoceros" and get straight down to animal fetishism without wading through all the classical mythology. Such sophisticated searching is not yet a fully-supported feature of WWW/XMosaic, but the provided forms support is a help, and thereby a query can be formulated for passing to a database, and for the response to be shown to the world. Currently, the FORMS interface works only under XMosaic, but it is promised soon for MacMosaic and Windows.

I therefore commissioned from the United Kingdom a search engine so that the user could query my database according to scrolling lists of criteria, or user-entered parameters. Incidentally, this process taught me about what distributed data processing really is, and just how useful WWW can be in the development of such an interface. Rather than the programmer writing the code 12,000 miles from me, and sending it to me by ftp, etc. etc., I have sent him the data sets. To test it during the development process. I simply accessed a "private" directory downstream of his httpd server, and tried out what he was been doing, reporting directly to him in a box on the XMosaic screen. The images, needless to say, resided on my machine, 12,000 miles distant from the data files manipulating them.

The FORMS interface is working well, and offers querying not only on the Prints & Print History Database, but also on a database of Classical Architecture of the Mediterranean Basin, with some 2,500 records and associated images. Here is a forms page, which the user can fill in or, alternatively, use the scrolling lists provided. Such lists take away some of the guesswork, since the database is effectively listing the possible hits for each field:

2.5: WWW/MOSAIC TEXT QUERIES

But if Art Historians work with images, they also work; - as do most disciplines! - with words. Database files are fine for one particular kind of searching - but I now wanted a WWW/Mosaic-capable version of the string-searching program already implemented under tcl/tk (see above). This offers the user the chance to search for a string or strings through designated directories or sets of files. This can be used for ordinary reference purposes (e.g., *what does Homer say about Achilles? what s that quote from Shakespeare that Delacroix is illustrating in this print?*). The "hits" are marked by boldening, and movement between them can be by clicking on the icon provided, or by using the application's scrollbar.

An important feature of TextSearch is that it works with ordinary ASCII files, and also with HTML documents; a further feature is that any retrieved HTML document still has its hotlinks intact, and so can be used as a springboard for further work.

Another reason for incorporating text searching in my suite of WWW/Mosaic scripts is because I have mounted a 120,000 word book on the net, and text-searching must be an essential tool for any networked reader. The exercise may help demonstrate whether the technology is capable of sustaining "power" applications, and whether an "electronic" book can be more flexible, better illustrated and more useful than a printed book. Here is a TextSearch window:

How should a book be handled on the Web? The first assumption must be that it will "work" differently from a conventional paper book, or we are simply wasting our time. For example, implementing what is no more than an electronic page-turner would be silly, because we would be ignoring the "hyper-" part of the opportunity the Web provides. I'm sure I haven't yet got the format correct, but equally sure that a suitable and flexible format will emerge which will give electronic books many advantages over paper ones, not least their undateability, flexibility with media such as still and moving images, and sound, as well as their "outreach" via hotlinks to other resources on the Web. But for the moment I was conscious that many people have slow lines, and that offering a long book; as one file might make people wait too long. So I adopted two strategies. The first was to slice the book into discrete HTML documents, and provide the top and bottom of each document with hot buttons which, reading from left to right, offer access to a Help Page, to the Table of Contents (itself a hotlist), to the Database Forms Interface, to TextSearch, to the Main menu of ArtServe, and to images relevant to the particular section:

2.6: ArtServe in the University:

Images for Students

In my Department, we have decided that the display and networking technology is now sufficiently advanced to make the accessing of decent-quality images across the network a possibility in a university teaching environment, where network access is now common and computer provision equally so. I have therefore proceeded to do some trial digitizing, in order to test whether the technology will help with teaching and learning. After all, the technology is but a shop-window, and the good displayed might range from exciting to useless!

Mindful of the arguments about "better in the future", and sceptical (but not cynical) about the actual value of the Internet in teaching and learning, we have begun with a "dry run", by making all images used in first-year classes available in digital form over the network using WWW/Mosaic, and accessed via a simple database. We envisage that such a setup will be used by students for private study: as well as being able to print out the images as aides memoire should they so wish, the annotation feature might well help them augment their lecture notes, and their preparations for presenting their own seminars. Such a study aid is an improvement on maintaining free access to the slide collection, and may restrict damage to the slides, and mis-filing. We understand that there are no copyright restrictions on copying for strictly pedagogic use and, of course, access to copyright material will be restricted by software.

As for the use of such materials in lectures, this is certainly possible with some preparation. The main lecture theatres at the Australian National University are equipped not only for tape-fed video but also for network-fed video derived from computers, so the lines of communication are already there to allow a lecturer to show a succession of images from a pre-prepared script. However, such images will currently be of video quality, which is a great contrast with 35mm slides. A video image is about one-third of a megapixel, whereas the resolution of a 35mm slide perhaps approaches sixteen megapixels. So for lecturing purposes with current video resolution, there is something extremely perverse about digitizing a 35mm slide and degrading its quality, and then using computer technology to display the image at something like one-fortieth of the resolution easily attainable through the "old technology" of a slide projector.

3: Conclusion

Why do my own contributions to network teaching & learning (see above) take the form of "neutral" materials, rather than that of "teaching packages"? Because (1) I remain to be convinced that packages (whether in the form of textbooks, or automated on a computer) are suitable for any but the imparting of the most basic skills: (2) teaching and learning are two-way processes involving human interaction, whether one-to-one or in a larger group, and a computer should be used as a tool to help not to dominate; (3) the provision of data (the copyright for which resides with me, and which I waive) backed by a rudimentary database is intended to help teachers by providing them with a network-available bank of images - we'd have 20,000 or more if ten other people did the same.

From my own experience of teaching and of using and providing materials for the network I suggest, as an avowed sceptic, the following minimalist approach for the immediate future:

(a) We should concentrate on providing a range of materials for learning, rather than materials for teaching - that is, passive data files (rather like books in a library - but multimedia books where possible) rather than programs which attempt to interact with and leach the student, which is what human beings have been doing quite well since at least the time of Plato. WWW/Mosaic can help here, because of the ease with which threads can be provided by the teacher, Ariadne-like, through the amazing maze of cyberspace;

(b) a modest number of base-level course units in which some elements are available over the network. In my own field, I could imagine my own tutorial on "The Palace of Diocletian at Split" occupying a course on Roman Art & Architecture for perhaps one hour - so the scope for industry on the part of our colleagues is large! I would expect such elements to be referred to (and perhaps actively promoted) in various universities across the world;

(c) we should promote the idea of referring students to information resources (beyond simple library catalogs) available over the network. These would include not just online books or encyclopaedias, but databases fronted by Mosaic/WWW, such as that of Australian Flora at the CSIRO Botanical Gardens in Canberra;

(d) we should encourage academics to become network-conscious, and publish materials on the net, and consequently for University librarians to become more or less worried about the integrity of their empires depending on whether they view networks as a help or a threat;

(e) we should expoint the submission not only of undergraduate papers over the network (which already happens in some subjects), but the use of network resources in the submission of higher degrees (just as some multimedia higher degrees are already available on CDROM).

APPENDIX I

A Description of the WWW/Mosaic Server

in the Department of Art History,

Australian National University

(<http://rubens.anu.edu.au>)

ELECTRONIC IMAGING PRE- 1994

Several of the materials which are now on ArtServe were made available on servers in the Netherlands and the USA in the second half of 1993: (a) tutorial on *The Palace of the Emperor Diocletian at Split* 15Mb of text and images; (b) presentations totalling 21Mb of material related to Art History, principally on Prints & Print-making;

INSTALLATION OF A WWW SERVER

ArtServe went live on 4th January 1994, and currently gets about 2,800 accesses per day from all over the world. The material on offer consists of: (a) the *Diocletian tutorial* mentioned above; (b) 2,800 images of *Prints & Print History*. together with their associated data files, available both with inline thumbnail images and via a fill-in-the-form database-query interface. The images date from the beginning of printmaking to about 1900, so as to avoid copyright problems; (c) 2,600 images of *Architecture of the Mediterranean Basin*, currently available only via a fill-in-the-forms database-query interface; these are all taken by myself; (d) a *TextSearcher* (commissioned by me), which operates with plain ASCII text files and also with HTML files; (e) a 120,000-word book. written by me, on *The Greek and Roman Cities of Western Turkey*, "sliced up" into HTML format. This is published here as my contribution to the credibility (in teams of scholarship) of the Internet: i.e., we must use the Internet for publication, otherwise its status will be that of an emailing or publicity billboard or toy. When I have time, the book will be linked to over 2,000 images (some of them counted in the 2.600 art & architecture images mentioned above);

USING THE TECHNOLOGY

FOR TEACHING & LEARNING

We are gearing up to make sets of images seen by our students network available. Two of our units in 1994 (*Introduction to Art History* 870 images; and *Introduction to Modern Art*; 900 images) have all images network available. For 1995, at least four of our standard Art History course units ("*Introduction to Art History*", *Introduction to Modern Art*, *Art & Architecture of the Italian Renaissance*, and *The Art of the Print*) will have access within the Department to banks of images (about 1,000 per unit)

made available to them under WWW/Mosaic, and served via rubens.anu.edu.au either from hard disk or from computer-driven laserdisk. To the best of my knowledge, this setup and combination represents a world first in IT teaching in the Humanities. We estimate that 120 students will be enrolled on Art History IA, 95 on Art History 1B, 60 on *Art & Architecture of the Italian Renaissance*, and 70 on *The Art of the Print*.

During first semester 1994, all images seen in our first-year semester unit (Art History IA) were made available with inline GIFs under a WWW/Mosaic interface. These 1000+ images are only available internally for copyright reasons. The same setup was also used to present a selection of these slides for visual test revision, as it was for a later-year Australian Art unit. Preliminary tests indicate that the students like looking at electronic images, because of their constant availability (instead of struggling to get a look at the ONE set of slides). Already available for The Art of the Print unit for 1995 is a core of some 2,800 images on CDROM, network-available as http://rubenslpnnts_form.html.

Also available is a preliminary store of 840 images for the 1995 first-semester unit on Michael Greenhalgh's *The Art & Architecture of the Italian Renaissance*, again with a WWW/Mosaic front-end, on *rubens*. The technology used here is exciting, for it uses the immense reservoir of laserdisk storage. The user currently sees a sorted listing of records, and selects the required image; the computer addresses the laserdisk, which presents the image to the framegrabber, after being digitized. the image is sent over the network in the usual way. Soon the user will have the same system enhanced by (a) inline thumbnail GIFs, as in the standard ArtServe setup already described; and (b) a switch allowing downward scaling of the retrieved "full-size" image, to make the setup more useful over a busy network or with a slow and small machine. This course unit will, as far as possible, be completely digital (except for the lecturer, who remains analog; and the lectures, which will use slides): apart from a one-page initial handout giving information about accessing the system, almost everything else (schedules, tutorial assignments, lecture topics, essay questions, bibliography, assessment questionnaires, and especially the images shown during lectures and tutorials) will be digital.

WORK IN PROGRESS

1. Development of the *Salami Slicer* which takes a data file, and slices it up painlessly into HTML format, automatically writing the HTML code and the image references required. Nearly finished, this will give complete flexibility in the size of pages formatted, the number of images they contain (they can have just records - no images),

and the quantity and order of the fields in the data record. This is an essential script when dealing with large quantities of data records and their associated images;

2. Writing of a data file for 11,000 images of European art and architecture, all of them taken by me, which currently reside on a laserdisk; currently, some 900 data records with inline GIF images are only available internally for copyright reasons.

3. In connection with (2) I have commissioned and can demonstrate code which (1) has WWW/Mosaic request a frame from the laserdisk; (2) framegrab and digitize the frame; (3) send it out across the network; (4) write the image either to a Mosaic document or to an external image file; and (5) apply any percent of scaling to the operation. This technology means that (subject to copyright) commercial laserdisks can be fed across a network, and be used by more than one person in more than one location. The last bit of code to write is some kind of stacking. so that 2+ accesses in close proximity will get what they want, rather than some brush-off message:

APPENDIX II:

PLANS FOR IMAGE QUALITY

IMPROVEMENT

Introduction

For some people, the resolution available to all video devices (laserdisks, video cameras, still video cameras) means that images captured by such means and then framegrabbed are of insufficient quality to stand the test of time. It is certainly a fact that higher resolutions are possible, as may be seen by viewing a PhotoCD image in all its five available resolutions (including video resolution). Certainly, to preserve all the detail in a 35mm slide requires more than video resolution.

The matter boils down to three questions. Are we satisfied with video resolution? Are there ways to mitigate its inadequacies? Will video resolution seem satisfactory in the future?

On today's monitors (I have a 21-inch 24-bit megapixel Trinitron) I find video resolution fine, since the X-Y resolution does not affect the quality of the image depth-wise: 24-bit JPEG'd images at video resolution can look very good. To mitigate its inadequacies I use a video slide copier (at S-VHS quality) which generates an image of the whole slide, and I can then zoom in (at up to 6X) and take as many

details as I wish. This patchwork-quilt approach isn't quite the same as viewing the whole image at a high resolution, but the whole package will, I think, prove perfectly adequate for small-group and individual learning.

Predicting what will seem satisfactory in the future is very difficult, but I imagine more resolution for less money will always be the order of the day. For many purposes, video resolution will remain adequate. But this leaves hanging the question of what people who deal in images, such as Art Historians, should do about their computerization. Should we go ahead at video resolution, and risk (in a decade's time) to be seen to have worked in vain? Should we all use PhotoCDs, and store them in juke-boxes, all online at once? My department possesses 180,000 35mm slides, so we would need 1800 CDROMs in line - which is a tall order, although not as tall as getting the databasing done so that they could be accessed. I conclude that the PhotoCD, whilst workable, deals in quantities of images too small to interest even moderately sized enterprises such as our's.

A possible solution: scanning

I believe that a useful device for our 180,000-slide dilemma might be a hopper-fed slide scanner (so that it could work unattended overnight, when the network is also quiet), spooling the main images (each how big after compression? perhaps 5Mb? or perhaps two or three times that?) directly to tape across the network, and leaving a small-scale "visiting-card" version of each on the hard disk. However, scanning is slow, and expensive in personnel irrespective of whether a hopper is used.

A better solution: digital camera & laserdisk

Computer imaging is today capable of generating very high quality images, and they can easily be stored in small quantities on hard disks, which are getting cheaper all the time. The snag is that current display technologies cannot easily manipulate images over - say - 1Mb in size. Anyone needing to deal electronically with large quantities of images (e.g. over 20,000) clearly wants to do the work involved only once, to be able immediately to display the images using current technology, but also to have to hand higher-quality images which will not look frankly silly in a few years when display technology has improved and come down in price - i.e., for the next generation of CDROM projects.

I have proposed the installation of a system which meets the above requirements, and makes use not only of high-quality analog and digital technologies, and of the Tape Robot, but can also serve the images transparently over the Internet (or any restricted

subset of it), using WWW/Mosaic. The data records are written for BOTH systems at the SAME time. The setup, which costs about \$100,000, involves:

(1) a digital camera which can write digitized images at 1.6 million pixels (compare video resolution, which can never offer more than a maximum of 400 thousand pixels); this dumps the results to a hard disk on a Macintosh, via a Photoshop plug-in;

(2) a Sony recordable laserdisk player (which I already have), which captures up to 36,000 archival-quality analog images on each side of its removable cartridge;

(3) an autochanger for a Sony recordable laserdisk player which, when associated with the player in (2) above allows computer and hence network access for up to 15 cartridges (at over 72,000 frames per disk) and therefore to over one million images on line;

The setup envisages the processing of large quantities of images, each one of which is written AT THE SAME TIME to a writeable laserdisk and to the Digital Camera system; the common data record is written at the same time. Scanning (which can produce huge files) is not an option, because it is too slow for dealing with anything but a small quantity of images. The equipment will be used for archival storage for research and teaching-related imaging in the Department of Art History, and all the images will (subject to any copyright restrictions) be available under WWW/Mosaic

The work will proceed in three stages:

STAGE ONE: I can already demonstrate network-transparent access to frames on a laserdisk - i.e. video-quality images. The datafile is regimented under WWW/Mosaic, with or without inline thumbnail images, and a request for a full-size image has the computer accessing the laserdisk frame, offering it to the framegrabber, digitizing it, and sending it off across the network to the recipient. Still to be developed for this setup are (a) resolving the contention when a frame is requested and the framegrabber and laserdisk player are busy; (b) implementation of a scaling option for the WWW/Mosaic pages, which already works with specific requests to the URL, but not yet with Mosaic pages: the user could ask for any proportion of scaling for an image. This will help people with slow lines, or even modems; it could be very useful indeed if something similar were to be applied to Tape Robot material (see below);

STAGE TWO: Using the SEPS Digital Camera, large quantities of 35mm slide material, and photographs and prints, will be digitized and stored on a local hard disk. Software needs developing for sending such files (Each 1.6-megapixel file is some

4.5Mb in size) to the Tape Robot for storage, for checking them, and then deleting them locally.

STAGE THREE: Building on Stages One and Two, software needs developing to give WWW/Mosaic clients flexibility in the quality of images they retrieve. Remembering that there will be ONE datafile which will point both to ONLINE material under WWW/Mosaic and OFFLINE material stored on the Tape Robot, the network user should be able to request (a) video-quality images, which are typically 80Kb full-size in JPEG format, or indeed at any percentage of this size; and (b) make requests for the higher-quality 4.5Mb images, or any percentage of this size, which will be stored not on laserdisk (remember that video resolution produces small files - but files lacking in this kind of "resolution") on the Tape Robot.

Software will therefore be necessary: (a) to retrieve image files from the Tape Robot; (b) to scale them if and when requested by the user; (c) to send them directly to the user, or leave them in some "pool" where they can be downloaded by the user, perhaps in dead of night using ftp; and finally (d) to perform (a) and (b) through the WWW/Mosaic interface, and perhaps set (c) in trail as well.

APPENDIX III:

AUSTRALIAN INITIATIVES WITH

WWW/MOSAIC

Even though there are questions needing answering about the use of WWW/Mosaic on "power" projects such as large databases or whole sets of online books, there are several ways in which this ingenious front-end can immediately be used in the everyday teaching and learning processes of teachers and undergraduates, so I shall now sketch out other initiatives under way in my university in Australia, without any suggestion that they are necessarily different (let alone better) than those under way elsewhere. After all, one great advantage of the Internet is that the windows opened up on the computer monitor could derive from Canberra, Caracas, or Chicago, albeit with a different flavour from each.

The initiatives are at two levels. At the strategic level, the Australian Department of Education, Employment and Training (DEET) will support WWW/Mosaic by specifying it as the front-end for a series of "Clearing Houses for Teaching and Learning" which it will fund from 1994. Under this scheme, a Central Clearing House will have responsibility for general standards and publicity, clerical assistance,

programming support and advice on machine, etc., compatibility; like the other three or four Clearing Houses, it will also take charge of one or more discipline areas. The Clearing Houses, using UNIX boxes with plenty of storage and memory, will accept moderated materials for teaching and learning, encourage electronic journals and discussion groups, and offer advice to those who wish to begin or improve their own use of electronic media. An important point is that the whole Clearing House system will encourage and foster the notion of publication: materials will NOT be allowed to reside in the Clearing Houses without prior peer assessment (from the world-wide net community), and users will be encouraged to count materials which are accepted as true publications. The Advisory Group o DEET on the Clearing Houses (and of which I was a member) believe this is an essential step in raising the standards and hence the credibility of networked information materials - countering the dross factor, as it might be called.

At the University, you will probably be familiar with our http servers - Tony Barry in the Library (info.anu.edu.au) Matthew Ciolek in the Social Sciences (coornbs.anu.edu.au) and David Green in BioInfonnaticsand electronic publishing (life.anu.edu.au): between us, we receive over 20,000 accesses per day, from all over the world.

\$\$\$\$\$\$\$\$\$\$\$\$\$\$\$\$

\$\$\$\$\$\$\$\$\$\$\$\$\$\$\$\$

About the Author: Michael Greenhalgh sometime Senior Lecturer in Art History at the University of Leicester (UK) has been the Sir William Dobell Foundation Professor of Art History at the Australian National University Canberra, Australia since 1987. His first interest in computers was when trying to index a book on a Cyber 72; since then he has studied the digitizing storage and manipulation of graphical material and how to make this available to his students, most recently over the Internet, on which ArtServe is run. Subject to funding he hopes to develop ImageServe as described above in Appendix II.