Graphics on the WWW

Dr. Anne Mumford

Abstract

This State of the Art Report addresses graphics on the World Wide Web. The WWW is now a major vehicle for the transfer of graphical data. In order to be successful we need to have an agreed set of standard file formats which enable predictable transfer of the information we require. The report looks at the different types of format which need to be transferred including: raster images; 2D line drawings; 3D models including engineering and chemical data; moving images. Commonly used examples of each type of format are discussed in outline. These include: TIFF, GIF, PNG raster formats; the CGM vector format; MPEG movies; VRML models. The report suggests that there is a need to think as graphics as information and not just enhancement of text and for adding interest. The report notes the need to consider indexing and metadata if we are to store and access graphics resources and notes the potential for new tools such as Java in the use of graphics in the future.

Contents

Contacting the author

Dr Anne Mumford
Computing Services
Loughborough University
Loughborough, LE11 3TU
UK

email¸a.m.mumford@lboro.ac.uk

tel¸¸+44 (0) 1509 222 312

Introduction

In an early paper about the WWW presented at the 3rd Joint European Networking Conference in May 1992, Tim Berners-Lee, Robert Caillau and Jean François Groff stated:

> "The aims of the W3 initiative are twofold: firstly to make a single, easy user-interface to all types of information so that all may access it, and secondly to make it so easy to add new information that the quantity and quality of inline information will both increase."

It is certainly true that much of this vision has been achieved. Vast amounts of information are available, although organising this information and quality assurance are still challenges to be fully met. The information is still mostly textual and the images simple and still and serve little purpose in providing information. This report notes the current status and the potential for new forms of graphics to emerge on the WWW.

A recent special edition of IEEE Computer Graphics and Applications (Gershon and Brown, 1996) looked at the use of *Computer Graphics and Applications in the Global Information Structure*.

This issue commences with a statement by Al Gore, Vice President of the USA who says:

> "With more information and data becoming accessible and the wide availability of commercial versions of graphics Web browsers via online services and software producers, the Web is likely to continue its explosive rate of growth. Continued work by the information visualization and computer graphics communities has the potential to make the World Wide Web and the Internet even more accessible and useful to citizens of all ages and from all walks of life. I would like to challenge you to take a critical next step. We need to move beyond Web browsers into the world of collaborative work. The information needs of the government, industry and the public are becoming more complex and demanding new ways to link people across organisational, geographic and educational boundaries."

This is our challenge and why we need to look at more complex forms of graphical data than the familiar GIF and JPEG files.

The WWW, through its graphical users interfaces, has given us a visual window onto a range of linked information sources. The emergence of new versions of HTML as well as non-standard extensions have given a range of methods of displaying information in a way which aims to capture the reader. The development of standard style sheets for the WWW will allow better practice.

The pictures associated with the text on WWW pages are intended to capture the reader's attention through the use of logos, photographs and diagrams. Rarely does this enhance the *information* being presented, it is intended as an attention-catching device.

The ability to display still images with the intention that these will add interest - and possibly aid explanation - is important on the WWW. This does not really take us further than a book, journal or newspaper in that the text and images being presented are still and thus do not progress beyond the page paradigm. In fact, the relative difficulty of looking at a screen, compared to paper, the lower resolution, and the inherent portability of paper, may make paper a preferable medium for many purposes. Of course, such an argument is counterbalanced by the ability to search and access information quickly and to find links to other related data.

We need to have a wider vision of graphics on the WWW for the following reasons:

- graphics can indeed be used to illustrate a point and add interest. We can however move beyond this and have graphical interfaces to information. In the same way as we can display text in a range of styles, we have the technology to allow graphical interfacing to, for example numeric, information to allow it to be displayed in different ways and to be manipulated by the user.

- it is also necessary to think of the information being transported across the network and to consider the relationship between the size of that transferred information and the processing carried out locally. We currently have a situation where a great deal of network traffic is being taken up with simple pictures transferred using a format which minimises local processing needs. Would we be better reducing file size and increasing the processing required of the server? The demands on the network do require some serious thinking about a range of other possibilities.

In order to consider the potential for this wider vision, this report looks at the nature of graphical information. It then moves on to look at a range of example file formats which illustrate different types of data. Some scenarios are then considered. The final part of the report considers the needs for graphical resources to be organised with suitable indexing and metadata. Some examples of this are given.

It seems the aim of any company involved in information provision to make their technology available via the WWW. This report reflects the current (May 1996) technology and gives a snapshot of some of the current movements in the business. While this will change, it is hoped that the general considerations will hold as the market inevitably changes.

Thinking About Graphical Data

It is useful to commence a discussion of file formats with some conceptual models which put a framework on the ideas being presented. Too often we compare file formats as if they are the same, when they are in fact very different. Files composed of image data, for example a satellite image, consist of very different information from a picture of a car engine which might require a fine level of detail for display. This may be different again from an artist's sketch of a car design for an advertising campaign.

The Computer Graphics Reference Model (CGRM) which has been developed to describe data which may have a graphical representation is a good place to start. This model gives a basis for comparing the different file formats. It is clear that these file formats do not follow the model and many cover several potentially different conceptual ideas. However, by looking at the CGRM the role, or roles, for these, and other, file formats become clearer.

Computer Graphics Reference Model

The CGRM recognises five different environments for computer graphics: Construction; Virtual; Viewing; Logical; Realization. Most of the information that is generated by an application and is to be rendered, will have progressed through a number of specific transformations and attribute bindings in order to produce a displayable image. The data moves forward through the stages from the Construction to the Realization environment. At each stage, the graphical information, or composition, can be stored as a metafile of information. the diagram below shows the environments and the names of the related compositions which can be stored.

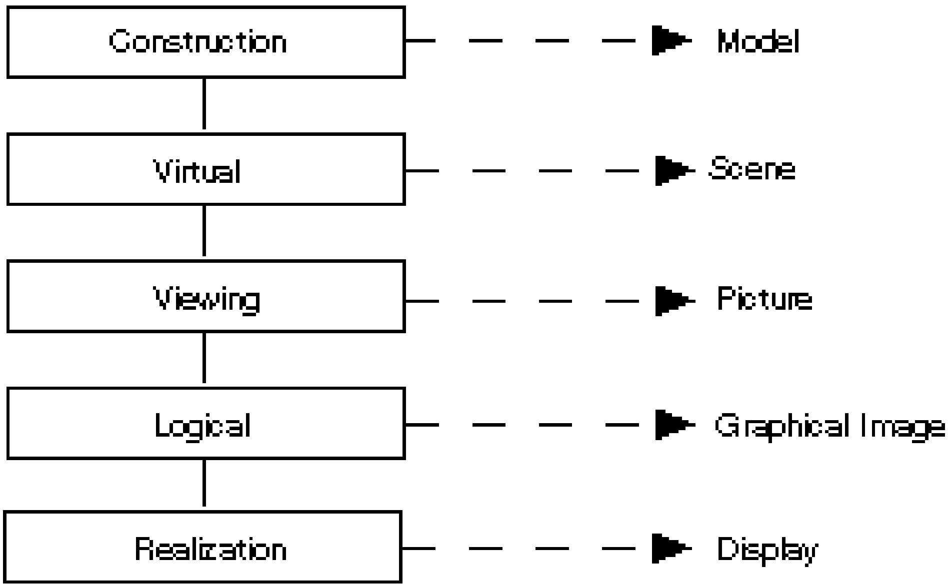

Environments and Compositions in the Computer Graphics Reference Model

The CGRM recognises that data can flow between the different environments. The process of data moving up from a higher to a lower environment is known as "absorption" and the process of data moving from a lower to a higher environment through input of graphical data, for example through digitization, is known as "emanation".

Here we are concerned with file formats and as such we will look at the nature of a typical data capture metafile from each of the environments in the CGRM.

Construction Environment

This is the environment that interfaces to the application. The application data to be displayed is prepared as a "model" from which specific graphics scenes may be produced.

The data stored at this level generally relates to the application and is not readily transferred between dissimilar applications. Examples include CAE interchange, such as STEP and IGES, the PHIGS archive file and the Image Interchange Facility if it is associated with application information.

Virtual Environment

In this environment, a "scene" of the model to be projected is produced from a set of virtual output primitives. As the geometry of these virtual primitives is completely defined, scenes are necessarily geometrically complete.

Within GKS:1985, for example, the production of primitives in Normalized Device Coordinate (NDC) space with attributes bound corresponds to the virtual environment. The Workstation Independent Segment Store (WISS) in GKS is a segment store to which all workstations have access.

Viewing Environment

In this environment a specific view of the scene is taken and the "picture" to be presented is projected. Output primitives in the viewing environment may have a lower geometric dimensionality than in the virtual environment, for example a 3D picture produced by a PHIGS application may be stored as a 2D Computer Graphics Metafile. The Computer Graphics Metafile (CGM) is an example of storage at the viewing environment. GKS:1994 allows the NDC Picture to be stored as a CGM at this level in the pipeline. VRML version 1.0 provides a 3D scene which is also probably at this level. VRML 2.0 moves the format up the pipeline.

Logical Environment

In this environment, the picture is rendered as a graphical image ready for presentation. Associated with each graphical output primitive is a set of properties to be ultimately realized by the rendering. All the attributes are bound to the primitive at this stage but the image itself is not actually rendered. For example, if we consider an output file destined for portrayal on a plotter, although we have defined the conceptual pen to be used the actual colour of the pen and the line thickness have not been resolved; everything is logically bound but the relevant device tables have been looked up. Examples of storage at the logical environment include TIFF, GIF and PNG

Realization Environment

In this environment, the graphical image is presented as a conceptual display for subsequent output to a specific device. This device need not directly correspond to the physical display. At this level all of the attributes (e.g. colours) are tightly bound. The storage at this environment is known as the "lexeme store". Examples of storage at the physical environment and files of compressed image data using perhaps the JPEG and MPEG techniques.

CGRM gives us the basis for thinking more clearly about file formats and to describe more accurately particular formats. It is important to know what we can achieve with particular file formats.

Formats, Transfer and Processing

There is a tendency to only be concerned with simplicity of interpretation and to be less concerned with the file size or with the amount of information which is being exchanged. This has resulted in the graphics on the WWW being lowest common denominator raster graphics. What we see is what we get, but we might actually want more than this and have the ability to modify the graphics and to interact with it.

The selection of formats for exchange is a balance between a number of factors including:

- file size
- processing required to interpret the information
- information in the file

The early history of the WWW has seen a predominance of raster formats being used, mainly either GIF or JPEG files.

It seems likely that this will change as the requirement for exchange of graphical information (and not just "pretty pictures") becomes a more major requirement.

File formats

The discussion of specific formats is split into 3 sections: models, scenes/pictures and graphical image/display formats. This reflects the fact that a number of formats do not fit comfortably into any one level in the Computer Graphics Reference Model.

Model

Engineering — STEP

The STandard for the Exchange of Product Model Data (STEP) is an ISO standard (or more accurately a set of related standards). STEP provides a neutral file format definition for the exchange of CAD data. The aim of the standard is to cover all aspects of engineering design. It is concerned with 3-D information, lifecycle information and assembly information. It uses a language called Express to define the information models.

Some work at Rensselaer Polytechnic Institute (Hardwick, 1996) is concerned with the exchange of STEP models across the WWW. The work also uses CORBA which allows applications to use one anothers resources by supporting message calls between objects through a network. The project uses STEP units of functionality to build the model which is then stored in a database. Information about that database is available via the WWW. The project then uses the data definition language of CORBA — IDL — to describe the interface to applications.

CORBA (Common Object Request Broker Architecture) is a standard for distributed objects being developed by the Object Management Group (OMG). The OMG is a consortium of software vendors and end users. Many OMG member companies are then developing commercial products that support these standards and/or are developing software that use this standard. CORBA provides the mechanisms by which objects transparently make requests and receive responses, as defined by OMG's ORB. The CORBA ORB provides interoperability between applications built in (possibly) different languages, running on (possibly) different machines in heterogeneous distributed environments. It is the cornerstone of OMG's CORBA architecture.

Greenough, 1996 describes the MIDAS project which aims to develop an integrated engineering environment bringing together a number of aspects of engineering. The aim is to develop an open system with a central database objects being described using the Express language adopted within STEP. The data are accessed through programs which use a set of function calls modelled on the STEP Data Access Interface.

Engineering — AutoCAD DXF

This format was developed by Autodesk as an interchange format for AutoCad 3D drawings between applications. It is often used as a 2D vector format. It is widely used. The DWF format has been developed for this purpose by Autodesk who are promoting their Netscape plug in (see below).

Chemistry

The limitations of paper for describing molecular structures have provided a motivation for the work described by Casher and Rzepa, 1996. They take output from popular molecular modelling packages and output the results as VRML which they are building into a library. They are developing a stand alone molecular VRML authoring environment called MOzART which is based on the emerging Molecular Inventor from SGI. Problems with VRML 1.0 are likely to be solved with VRML 2.0. The results are being published as part of an electronic journal project hosted at Imperial College.

Other Model Formats

There are a range of formats which people are exchanging but are less well used and there are currently no (easily found) examples of novel use across networks. Relevant acronyms include: CDF, Net-CDF, HDF, NITF.

A really useful starting point for information about these formats and the links to relevant specifications is the home page for the EC sponsored Open Information Interchange Initiative (OII) which is listed at the top of the set of references in this paper.

A set of 3D models can be accessed through the UK VRSIG which offers a WWW interface to a collection of freely available 3D object files which have been compiled for their applicability to real-time graphics applications, and virtual reality (VR) in particular. All material is freely-available.

Scenes and Pictures

VRML

VRML 1.0 is at the level of a scene or picture in the CGRM. VRML 2 however does move towards being a modelling description and thus its inclusion at this point may not be accurate.

As there is so much progress in this area, the author of this report decided to extract the text below from the VRML Architecture Group (VAG) pages on the WWW. By the time you are reading this it will have moved on, so please follow the reference. Some useful background is hopefully included below.

<start extract from VAG pages>

Introduction

The Virtual Reality Modeling Language (VRML) is a language for describing multi-participant interactive simulations -- virtual worlds networked via the global Internet and hyper-linked with the World Wide Web. All aspects of virtual world display, interaction and internetworking can be specified using VRML. It is the intention of its designers that VRML become the standard language for interactive simulation within the World Wide Web.

The first version of VRML allows for the creation of virtual worlds with limited interactive behaviour. These worlds can contain objects which have hyper-links to other worlds, HTML documents or other valid MIME types. When the user selects an object with a hyper-link, the appropriate MIME viewer is launched. When the user selects a link to a VRML document from within a correctly configured WWW browser, a VRML viewer is launched. Thus VRML viewers are the perfect companion applications to standard WWW browsers for navigating and visualizing the Web. Future versions of VRML will allow for richer behaviours, including animations, motion physics and real-time multi-user interaction.

VRML Mission Statement

The history of the development of the Internet has had three distinct phases; first, the development of the TCP/IP infrastructure which allowed documents and data to be stored in a proximally independent way; that is, Internet provided a layer of abstraction between data sets and the hosts which manipulated them. While this abstraction was useful, it was also confusing; without any clear sense of "what went where", access to Internet was restricted to the class of sysops/net surfers who could maintain internal cognitive maps of the data space.

Next, Tim Berners-Lee's work at CERN, where he developed the hyper-media system known as World Wide Web, added another layer of abstraction to the existing structure. This abstraction provided an "addressing" scheme, a unique identifier (the Universal Resource Locator), which could tell anyone "where to go and how to get there" for any piece of data within the Web. While useful, it lacked dimensionality; there's no there there within the web, and the only type of navigation permissible (other than surfing) is by direct reference. In other words, I can only tell you how to get to the VRML Forum home page by saying, "http://www.wired.com/", which is not human-centred data. In fact, I need to make an effort to remember it at all. So, while the World Wide Web provides a retrieval mechanism to complement the existing storage mechanism, it leaves a lot to be desired, particularly for human beings.

Finally, we move to "perceptualized" Internetworks, where the data has been sensualized, that is, rendered sensually. If something is represented sensually, it is possible to make sense of it. VRML is an attempt (how successful, only time and effort will tell) to place humans at the centre of the Internet, ordering its universe to our whims. In order to do that, the most important single element is a standard that defines the particularities of perception. Virtual Reality Modeling Language is that standard, designed to be a universal description language for multi-participant simulations.

These three phases, storage, retrieval, and perceptualization are analogous to the human process of consciousness, as expressed in terms of semantics and cognitive science. Events occur and are recorded (memory); inferences are drawn from memory (associations), and from sets of related events, maps of the universe are created (cognitive perception). What is important to remember is that the map is not the territory, and we should avoid becoming trapped in any single representation or world-view. Although we need to design to avoid disorientation, we should always push the envelope in the kinds of experience we can bring into manifestation!

This document is the living proof of the success of a process that was committed to being open and flexible, responsive to the needs of a growing Web community. Rather than re-invent the wheel, we have adapted an existing specification (Open Inventor) as the basis from which our own work can grow, saving years of design work and perhaps many mistakes. Now our real work can begin; that of rendering our noospheric space.

History

VRML was conceived in the spring of 1994 at the first annual World Wide Web Conference in Geneva, Switzerland. Tim Berners-Lee and Dave Raggett organised a Birds-of-a-Feather (BOF) session to discuss Virtual Reality interfaces to the World Wide Web. Several BOF attendees described projects already underway to build three dimensional graphical visualization tools which inter-operate with the Web. Attendees agreed on the need for these tools to have a common language for specifying 3D world description and WWW hyper-links -- an analog of HTML for virtual reality. The term Virtual Reality Markup Language (VRML) was coined, and the group resolved to begin specification work after the conference. The word 'Markup' was later changed to 'Modeling' to reflect the graphical nature of VRML.

Shortly after the Geneva BOF session, the www-vrml mailing list was created to discuss the development of a specification for the first version of VRML. The response to the list invitation was overwhelming: within a week, there were over a thousand members. After an initial settling-in period, list moderator Mark Pesce of Labyrinth Group announced his intention to have a draft version of the specification ready by the WWW Fall 1994 conference, a mere five months away. There was general agreement on the list that, while this schedule was aggressive, it was achievable provided that the requirements for the first version were not too ambitious and that VRML could be adapted from an existing solution. The list quickly agreed upon a set of requirements for the first version, and began a search for technologies which could be adapted to fit the needs of VRML.

The search for existing technologies turned up several worthwhile candidates. After much deliberation the list came to a consensus: the Open Inventor ASCII File Format from Silicon Graphics, Inc. The Inventor File Format supports complete descriptions of 3D worlds with polygonally rendered objects, lighting, materials, ambient properties and realism effects. A subset of the Inventor File Format, with extensions to support networking, forms the basis of VRML. Gavin Bell of Silicon Graphics has adapted the Inventor File Format for VRML, with design input from the mailing list. SGI has publicly stated that the file format is available for use in the open market, and have contributed a file format parser into the public domain to bootstrap VRML viewer development.

This is a clarified version of the 1.0 specification. No features have been added or changed from the original 1.0 version of the spec. This is a 'bug-fix' release of the spec, correcting misspellings, vague wording and misleading examples, and adding wording to better define the semantics of VRML.

VRML 1.0

VRML 1.0 is designed to meet the following requirements:

- Platform independence
- Extensibility
- Ability to work well over low-bandwidth connections

As with HTML, the above are absolute requirements for a network language standard; they should need little explanation here.

Early on the designers decided that VRML would not be an extension to HTML. HTML is designed for text, not graphics. Also, VRML requires even more finely tuned network optimizations than HTML; it is expected that a typical VRML world will be composed of many more "inline" objects and served up by many more servers than a typical HTML document. Moreover, HTML is an accepted standard, with existing implementations that depend on it. To impede the HTML design process with VRML issues and constrain the VRML design process with HTML compatibility concerns would be to do both languages a disservice. As a network language, VRML will succeed or fail independent of HTML.

It was also decided that, except for the hyper-linking feature, the first version of VRML would not support interactive behaviours. This was a practical decision intended to streamline design and implementation. Design of a language for describing interactive behaviours is a big job, especially when the language needs to express behaviours of objects communicating on a network. Such languages do exist; if we had chosen one of them, we would have risked getting into a "language war." People don't get excited about the syntax of a language for describing polygonal objects; people get very excited about the syntax of real languages for writing programs. Religious wars can extend the design process by months or years. In addition, networked inter-object operation requires brokering services such as those provided by CORBA or OLE, services which don't exist yet within WWW; we would have had to invent them. Finally, by keeping behaviours out of Version 1, we have made it a much smaller task to implement a viewer. We acknowledge that support for arbitrary interactive behaviours is critical to the long-term success of VRML; they will be included in Version 2.

VRML 2.0

Moving Worlds VRML 2.0 is the second release of the VRML Specification. The specification is currently under development (2nd draft), and is scheduled for functional freeze (Draft #3) on June 5th 1996 and final document on August 4, 1996.

The specification was originally developed by Silicon Graphics in collaboration with Sony and Mitra. Many people in the VRML community have been involved in the review and evolution of the specification (see credits page in the specification). Moving Worlds is a tribute to the successful collaboration of all of us. Gavin Bell, Chris Marrin, and Rikk Carey have headed the effort at SGI to produce the final specification.

The VRML Architecture Group (VAG) put out a Request-for-Proposals (RFP) in January 1995 for VRML 2.0. Six proposals were received and then debated for about 2 months. Moving Worlds developed a strong consensus and was eventually selected by the VRML community in a poll. The VAG made it official on March 27th.

To start using VRML 2.0 you must install a VRML 2.0 browser. See San Diego Supercomputer Center's list of browsers for what's available. Note however that since VRML 2.0 is still a working document, these browsers are in a beta phase. At this point, Sony's CyberPassage is the only browser that supports VRML 2.0 Draft #1. Watch the Silicon Graphics VRML site for news on Cosmo Player for Windows95 coming soon.

VRML 1.0 provided a means of creating and viewing static 3D worlds; VRML 2.0 will provide much more. The overarching goal of Moving Worlds VRML 2.0 is to provide a richer, more exciting, more interactive user experience than is possible within the static boundaries of VRML 1.0. The secondary goal is to provide a solid foundation that future VRML expansion can grow out of, and to keep things as simple and as fast as possible -- for everyone from browser developers to world designers to end users.

Moving Worlds provides these extensions and enhancements to VRML 1.0:

- Enhanced static worlds
- Interaction
- Animation
- Scripting
- Prototyping

<end extract from VAG pages>

There are moves to standardise VRML 2.0 (or a subset of it) within ISO. This proposal is being discussed at an ISO SC24 meeting in June 1996.

One of the other proposals for VRML 2.0 was Active VRML from Microsoft which they are continuing to develop.

CGM

Pictures can be stored using raster formats such as GIF and PNG. There are however severe limitations with using this approach as the diagrams can have "jagged" edges and may not be as detailed as one might need due to poor resolution. The use of vector graphics can result in much smaller files and better representation. There are current moves to get the CGM standard incorporated into the standard WWW tools for these reasons.

The Computer Graphics Metafile (CGM) is the International Standard for storage and exchange of 2D graphical data. Although initially a vector format, it has been extended in 2 upwardly compatible extensions to include raster capabilities and provides a very useful format for combined raster and vector images.

A metafile is a collection of elements. These elements may be the geometric components of the picture, such as polyline or polygon. They may be details of the appearance of these components, such as line colour. They may be information to the interpreter about how to interpret a particular metafile or a particular picture. The CGM standard specifies which elements are allowed to occur in which positions in a metafile.

CGM also has profile rules and a Model Profile to attempt to solve the problem of flavours of standards. 4 Internationally Standardised Profiles (ISPs) are being developed for CGM. These are being used as the basis for defining the way that CGM will be used within MIME compliant email and within WWW. CGM has been accepted as a MIME data type. There are a number of activities concerned with increasing the use of the CGM on the WWW. It has been debated by the WWW Consortium and has received support there. The FIGleaf inline plug in for Netscape, for example, supports CGM as well as other formats including PNG. A viewer for CGM is also being developed as part of the RALCGM project. InterCAP Graphics Systems, Inc., have announced InterCAP InLine which is a Netscape API-compliant graphics viewing tool that operates as a plug-in to Netscape Navigator 2.0. InterCAP InLine supports inline viewing, zooming, dynamic panning and magnification, and animation of intelligent, hyperlinked Computer Graphics Metafile (CGM) vector graphics within the Navigator 2.0 Web browser.

SVF

The Simple Vector Format has been developed by SoftSource and NCSA as a vector format suitable for the WWW. It allows hyperlinks to be included and layer information. A Netscape plug in is available. It is not clear that this format will be widely accepted. A well defined and ISO standard alternative exists in the CGM and SVF seems to be reinventing that particular wheel.

DWF

Autodesk and Netscape announced in April 1996 that they are to work on a format called the Drawing Web Format (DWF). The arguments made in the press release are the general ones for use of vector rather than raster format — compaction resulting in improved performance and accuracy which is generally not obtainable through use of raster formats. The WHIP! plug in for Netscape Navigator is available from the Autodesk WWW address and enables creation, viewing of DWF files. A future version will enable transfer from the AutoCAD DXF format. The format will also have the ability to embed URLs providing links to other locations. This is a format worth watching. There is a need for a widely accepted standard for vector graphics on the WWW. Will it be this, CGM or some other format?

Graphical Images and Display

Introduction

A range of formats is described in this section. The even description here does not reflect the relative use of the formats. By far the most used format is GIF. JPEG is also well used for still images. The formats for moving images described here (MPEG, AVI, Quicktime) are all used as each has a user base on different platforms. It seems likely that MPEG will become the format of choice as hardware supporting it becomes more widely available. The inclusion of PostScript and PDF at this level might be debated, but the author sees the formats as very much a page image containing presentation level information.

GIF

The GIF format defines a protocol which supports the hardware independent, online transmission of raster graphics data (i.e. images). It uses a version of the LZW compression algorithm for its compression.

GIF is defined in terms of data streams which in turn are composed of blocks and sub-blocks representing images and graphics, together with the essential control information required in order to render the resultant image on the target output device. The format is defined on the assumption that an error-free transport level protocol is used for communication i.e. no error detection facilities are provided.

GIF utilities include an encoder program used to capture and format image and graphical data as a GIF data stream and a decoder program capable of re-interpreting a stream. Data streams are encoded such that the decoding process is optimized. The decoder is able to process the data stream in a sequential manner, parsing the blocks and sub-blocks, using the control information to set hardware and process parameters and interpreting the data to render the graphic image.

Although the Graphics Interchange Format (GIF) is the copyright property of CompuServe Inc., they have granted a limited, non-exclusive, royalty-free license for its use in computer software.

Compuserve developed GIF in 1987 and promoted it as a royalty free standard for bitmaps. It is widely used on the WWW and has the advantage of being simple and is a lossless encoding. Unisys hold the patent for the LZW compression algorithm used in GIF and at the end of 1994 announced that they were going to enforce the patent and to charge royalties. This has caused a lot of debate on the WWW (see Wegner, 1995) and has increased the urgency with which PNG was developed.

PNG

The Portable Network Graphics format (PNG - pronounced "ping"). It has been developed by a group supported by W3C and, following the GIF patent issues, by CompuServe. The GIF hiatus speeded the development of the format, though this does not seem to have in any way reduced the quality of the result which is widely regarded as a well defined format. Due to the fact that it was designed with the WWW in mind, it allows a progressive display option and also allows the storage of keywords which can be extracted by search engines. PNG supports a colour look up table (like GIF) as well as true colour with a colour depth of up to 48 bits. It goes beyond GIF by supporting a full alpha channel and image gamma indication, allowing contrast correction for different input and output devices.

Will it take off? The quality of a format has never been a measure of the success of its take-up. It is becoming supported by a number of browsers, though not all have fully implemented the specification — a traditional problem with file formats! Browsers with some capability include: Chimera, Internet Explorer, Mosaic 95, NCSA Mosaic, Netscape Navigator. Clearly the need is for native support in Netscape Navigator rather than an unofficial plug in. Meanwhile, keep an eye on the PNG home page for more details.

TIFF

Aldus Corporation designed and made public the Tagged Image File Format (TIFF) in 1986. TIFF is a raster format. Although initially targeted at desktop publishing applications, it has been widely implemented on all sorts of computing platforms and has become a de-facto industry standard format. There is no general purpose ISO standard for raster interchange although there have been moves to standardise a version of TUFF (TIFFIT) in ISO. The ODA standard has a useful specification for tiled compressed raster, but this is one of 8 different parts defining various content portions and aspects of the overall architecture. There are standards for compression of black-and-white images (e.g., the CCITT/ISO facsimile standards), and compression of colour data (e.g., the ISO JPEG compression standard), but TIFF goes further in offering a complete format for general raster interchange.

The TIFF definition is based on the concept of "tags". Tags simply provide information about the raster image (one of the tags is a pointer to the compressed content of the image itself). Examples range from such critical information as the compression type, size, and bit order of the compressed image, to purely information items such as author, date and time, source software, etc.

This is a well used format though GIF is more widely used on the WWW. There have also been version problems going from version 5 to version 6.

JPEG

This abstract was taken from the JPEG_FAQ written by Tom Lane, organiser of the Independent JPEG Group.

JPEG (pronounced "jay-peg") is a standardized image compression mechanism. JPEG stands for Joint Photographic Experts Group, the original name of the committee that wrote the standard.

JPEG is designed for compressing either full-colour or grey-scale images of natural, real-world scenes. It works well on photographs, naturalistic artwork, and similar material; not so well on lettering, simple cartoons, or line drawings. JPEG handles only still images, but there is a related standard called MPEG for motion pictures.

JPEG is "lossy," meaning that the decompressed image is not quite the same as the one you started with. (There are lossless image compression algorithms, but JPEG achieves much greater compression than is possible with lossless methods.) JPEG is designed to exploit known limitations of the human eye, notably the fact that small colour details are not perceived as well as small details of light-and-dark. Thus, JPEG is intended for compressing images that will be looked at by humans. If you plan to machine-analyze your images, the small errors introduced by JPEG may be a problem for you, even if they are invisible to the eye.

A useful property of JPEG is that the degree of lossiness can be varied by adjusting compression parameters. This means that the image maker can trade off file size against output image quality. You can make extremely small files if you don't mind poor quality; this is useful for applications like indexing image archives. Conversely, if you aren't happy with the output quality at the default compression setting, you can jack up the quality until you are satisfied, and accept lesser compression.

There are two good reasons for using JPEG: to make your image files smaller, and to store 24-bit-per-pixel colour data instead of 8-bit-per-pixel data.

Making image files smaller is a big win for transmitting files across networks and for archiving libraries of images. Being able to compress a 2 Mbyte full-colour file down to 100 Kbytes or so makes a big difference in disk space and transmission time! (If you are comparing GIF and JPEG, the size ratio is more like four to one.)

If your viewing software does not support JPEG directly, you will have to convert JPEG to some other format for viewing or manipulating images. Even with a JPEG-capable viewer, it takes longer to decode and view a JPEG image than to view an image of a simpler format such as GIF. Thus, using JPEG is essentially a time/space tradeoff: you give up some time in order to store or transmit an image more cheaply.

It is worth noting that when network or phone transmission is involved, the time savings from transferring a shorter file can be greater than the extra time needed to decompress the file.

The second fundamental advantage of JPEG is that it stores full colour information: 24 bits/pixel (16 million colours). GIF can only store 8 bits/pixel (256 or fewer colours). GIF is reasonably well matched to inexpensive computer displays — most run-of-the-mill PCs can't display more than 256 distinct colours at once. But full-colour hardware is getting cheaper all the time, and JPEG images look much better than GIFs on such hardware.

A lot of people are scared off by the term "lossy compression". But when it comes to representing real-world scenes, no digital image format can retain all the information that impinges on your eyeball. In comparison with the real-world scene, JPEG loses far less information than GIF. The technical meaning of "lossy" has nothing to do with this, though; it refers to loss of information over repeated compression cycles, a problem that you may or may not care about.

Photo-CD

Photo CD is a format developed by Kodak which has been adopted by many software manufacturers on a range of platforms. It enables images to be stored at a range of compression levels. It is highly regarded as a good way of storing images and viewing them. It is a proprietary format.

Kodak have announced they are working on the following enhancements which are very relevant to network provision of images:

- ¸On-the-fly watermarking of images. That is, existing non-watermarked Photo CD images can be served over the web with a watermark applied
- ¸Controls to limit access to high resolution image data.
- ¸A URL locking parameter to turn interactivity off for specific instances of images.
- ¸The ability to source Photo CD images from anywhere on your system for distributed web sites.

Their WWW site notes the fact that the use of Java will make all browsers able to use the format in the future.

FIF

Another company joining the move to making its technology available over the WWW is Iterated Systems who market the technology for fractal image compression developed by Michael Barnsley. The approach uses fractals as the image compression method rather than the discrete cosine transformation used in JPEG and MPEG. FIF (Fractal Image Format) produces images which can be zoomed in a way which lossy compressed images cannot and is also is smaller than equivalent JPEG and GIF files. The software associated with it has been expensive and the formats proprietary and closed. Users of the format have tended to be large companies, such as Microsoft who licensed the technology for its Encarta CD-Rom.

Like other companies described in this report, Iterated Systems are repositioning themselves to provide technology for the WWW. The result is 2 fractal viewers free of charge (plug in for Netscape and an AVI viewer), an image conversion tool for converting images to FIF and a program developers kit.

It remains to be seen whether this can take off in a WWW world dominated by GIFs and JPEG files.

PostScript and PDF

The most popular format is the PostScript language which is an output option in very many packages, and is supported in firmware in numerous output devices such as laser printers. This is more flexible than raster storage in that the scale can be changed without loss of information. It offers the advantages of potentially high resolution colour output - that is, it is close to being as good as a printed paper copy.

PostScript is a page description language (PDL) designed by Adobe Systems Inc. PDLs are designed for presentation of complete, formatted, final-form (non-revisable) page images on output printing devices. "Virtual paper" is a good metaphor for PDLs. Most PDLs, PostScript included, are oriented toward presentation of pages on laser printers. PostScript is the most successful of the commercial PDLs (though others do exist, for example Interpress from Xerox and QuickDraw from Apple Computer), and has had a heavy influence on the final appearance of the Standardized Page Description Language (SPDL, an ISO standard).

As the "language" part of PDL suggests, PostScript is a true interpretive programming language. Its Level-1 definition includes over 270 operators, which go far beyond basic graphics presentation (definition and maintenance of "dictionaries", boolean and arithmetic operators, etc). The recently released Level-2 definition contains over 420 operators.

PostScript uses programming language constructs and paradigms: procedures, variables, conditional logic, etc. This creates powerful expressive capabilities. The trade off is that, compared to more object-oriented graphics formats, a PostScript graphics file is very difficult and impractical to edit or modify. Although device-independent, the PostScript imaging model demands raster devices for presentation. The language is implemented on powerful onboard micro processors on many raster devices (PostScript demands a lot of memory and computational power to interpret).

Encapsulated PostScript (EPS) is a (large) subset of PostScript which allows storage of information in the PostScript language but excludes any size or positioning information. This means that a part of a page can be brought in to another document. This is most frequently used for the inclusion of graphics within documents where these have been produced a different package than the one used for producing the text.

Adobe have further developed the PostScript concept to define their Portable Document Format (PDF) which links with a suite of software called Acrobat. PDF extends PostScript Level 2 to allow the addition of links within and between documents, annotations, thumbnails of pages, and chapter outlines which link to specific pages for access. The basic metaphor again is the page. This can be very attractive to publishers who wish to define a house style or who wish to have an online version of a paper journal. One such example is the Electronic Publishing journal from Wiley which is described in Smith et al (1994).

Moving Images

Introduction

3 formats dominate the exchange of moving images and associated audio. MPEG has tended to dominate on Unix platforms, Quicktime on Apple and AVI on PCs. Iterated Systems are also promoting fractal compression. It seems likely that the dominant format in the future will be the MPEG-2 format developed by ISO and ITU.

MPEG

MPEG is an international standard for the encoding of moving pictures. The name comes from the Moving Picture Experts Group (of ISO and CCITT - now ITU) who developed the standard. MPEG-2 builds on the original MPEG-1 specification and looks likely to dominate standards in this area. It includes both video and audio and future versions are likely to cover 3-D images. Images encoded in this format are becoming available and hardware to support the encoding and decoding of MPEG is coming down in price.

Quicktime

This was developed by Apple for storage of audio and video information. It can however be used on a range of platforms, though other formats tend to dominate there.

AVI

This format has been developed by Microsoft as part of their Resource Interchange File Format. The format compression and decompression software is bundled as part of Video for Windows.

FIF

Fractal compression technology (discussed above for still images) can also be used for moving images.

Publishing on the Internet

This has been included in this report because the end result we often require is a published report of some kind with a mix of text and graphics. Three products are addressed here. The Acrobat software and associated file format PDF seems likely to dominate for those wishing to have page description files of this kind. The products listed all have the look and feel of printed pages as their model for online publishing.

Acrobat

Adobe have signalled a commitment to the WWW with various moves recently. Their PostScript file format has always provided a format which is largely reliable and portable between machines and software and printers. They built on the success of PostScript a few years ago in producing their Portable Document Format (PDF). Adobe had learned from their PostScript experience that it is possible to have an open standard for exchange and still obtain revenue from those products which used the format - in the case of PostScript these were the laser printers with PostScript capability.

PDF through the Acrobat suite of software tools is a page description language which allows navigation within documents, the ability to annotate, collaboration tools for document development and indexing. Scanned documents can be turned into PDF. PDF can be converted to PostScript for exchange and printing. It has gained popularity with publishers experimenting with electronic journals. The reasons for this include the fact that it allows the publisher to maintain integrity of documents which are stored in a page format thus allowing continuity of style as well as document fidelity.

Adobe and Apple have recently announced PageMill and SiteMill which allow creation and management of pages on the WWW which have been created using Adobe software. A future version of the Acrobat software has been announced, and this will include what is described as being more tightly integrated with the WWW technology. There will no longer be a need for the separate (free) Acrobat Reader. As part of this strategic direction, Adobe are working with suppliers of browsers, including NetScape, to build PDF support into their browsers. Other recently added extensions include the ability to embed AVI video and audio files.

The real problem, in the view of the writer of this report, is whether the page paradigm is the one we wish to follow when using online documents. It seems to have attractions for people for the wrong reasons, which include: maintenance of paper-based styles; maintaining information fidelity (is there a better way?); maintaining equivalence between paper and online forms of documents. There is however a lot of interest in PDF and Acrobat tools, particularly from those investigating the possibilities of the electronic library.

The Acrobat software is available for Windows, Apple Mac and Unix systems.

Envoy

Envoy is marketed by Novell, having been originally launched by WordPerfect and is part of the Perfect Office suite. This gives it a large market share, but it certainly has less visibility than its competitor Acrobat. Recent extensions to Envoy have been made by a company called Tumbleweed and Tumbleweed Publishing Essentials was released in December 1995. Indexing, hyperlinks and conversion from PostScript have been added. An OLE Control object for Microsoft's Internet Studio publishing tool has also been produced. It is supported on Windows 95 and NT. An Apple version is expected.

Common Ground

This product has a relatively small share of the market and is available for Windows only at the time of writing. An Apple version is expected.

Graphics as Information

In the introduction it was noted that much of the graphics on the WWW is there to add interest (and reduce speed!). Yet there is more to graphics than that. It can be a window onto other information, it can be a way of manipulating a central database, it can provide information by its very nature through image processing and analysis tools. It has the potential to offer a secure access to data, for example census data which needs to be anonymised. The user can extract the information through a graphical representation rather than through a numeric table. The ARGUS project described below has such potential illustrated.

We look here at two UK based projects which are working to use graphics as an interface to data.

The KINDS Project

The project is run jointly by the Department of Geographical and Environmental Sciences of the Manchester Metropolitan University, Manchester Computing of Manchester University and the IT Institute of the University of Salford.

It is closely linked to the MIDAS service of Manchester Computing. MIDAS (Manchester Information Datasets and Associated Services) provides on-line support for large and complex data sets. Its aims include the provision of effective, efficient and the widest possible access to these datasets by the UK academic community. KINDS is concerned with large, spatially referenced data sets available to the academic community through CHEST deals.

KINDS is directed to developing an efficient and effective interface to the datasets. The efficiency of access depends on the potential user understanding the datasets and their methods of access. Effectiveness depends on the user knowing what they want and on knowing what the data will provide. For many users and potential users achieving these is a process of education not only about the data but also about their own objectives, the use to which it can be put and the limitations to use.

Developing good practice in using national data sets involves linkages between objectives, data processing methods, software and hardware specifications and data parameters. Understanding and specifying these linkages are complex processes but are necessary if the best and widest use of national data sets are to be achieved.

KINDS will produce an intelligent user interface which sets up the linkages. It will take the form of an interactive system to solicit information from users in order to decide what data are appropriate and then inform how to access them. The interplay will make use of text and visual explanatory materials. These will provide information and ideas which will refine the users understanding not only of the MIDAS service and the data sets it supports but also of the use of spatial data sets in their own area of work. The material will be appropriate for a wide range of users from complete novices to experienced researchers.

The KINDS project will develop an on-line or stand-alone system which will provide as output; instructions on how to access data explanatory materials and quality reports.

The prototype KINDS system provides some user friendly and powerful interfaces for you to browse and search for spatial information about Great Britain. The Bartholomew digital map (GB) dataset has been used as the major spatial information source for the demonstration. The recent development (WWW interfaces to Arc/Info) is also presented here for MIDAS users to experiment with WWW as an alternative platform to handle Bartholomew data and generate feature maps.

Project Argus

ARGUS is joint project between Birckbeck College at the University of Leicester in the UK. Examples of the work to date are:

- Cartographic Data Visualiser (cdv) which is a large collection of tcl/Tk scripts for highly interactive, exploratory spatial data analysis for area aggregate and network data and is available to the academic community in the UK via WWW.

- 3dv which uses readily-available extensions to a public domain GIS called GRASS to illustrate innovative ways of visualizing socio-economic data as well as other abstracts and real surfaces

- a series of WWW pages which illustrate and expand on these developments

Profiles

The transfer of graphical file formats between applications and between machines has caused many users a great deal of frustration. It often "doesn't work". This may be because the software used does not use the file format in the same way. This may be for a number of reasons which include:

- the software may have implemented the specification incorrectly
- the format specification may be ambiguous and open to different interpretations
- the generating software and the interpreting software may have implemented different subsets of the format specifications and thus some parts of the file may not be understood on interpretation

For these (very common) reasons, there has been a move to create profiles (or application profiles as they are sometimes known) for file formats, or combinations of formats. This work has been progressing for some years for graphical file formats with work by MAP (Manufacturing Automation Protocol initiative led by General Motors), TOP (Technical Office Protocol initiative led by Boeing) and CALS (the US Department of Defense initiative). The CALS work has provided the major thrust in the production of profiles for some years and these have included vector and raster formats as well as formats for product data exchange and document formats.

The Computer Graphics Metafile standard has been extended to include the requirements for defining profiles and a pro forma and model profile to assist in the definition of profiles. Profiles are very useful because they can be used as the basis for conformance testing. Profiles for CGM are currently being used by the Air Transport Agency (ATA), the petroleum industry (through the PIP profile) and by CALS. International Standardized Profiles are also being developed for CGM through the European Workshop on Open Systems (EWOS) and these will be harmonised for international standardization. CGM profiles are being used to define a MIME type for use on the Internet and within the WWW.

Other standards are also adopting the profile concept. This is due to the comprehensive nature of standards which try and offer a wide range of facilities not always needed by all applications. The Image Interchange Facility (described below) also uses the concept of profiles.

Bringing Order to Graphics Resources

The pictures on the WWW mostly come about because of their association to text and their use in adding interest or aiding explanation. It is likely that there will be more and more resources of a graphical nature which are stored for their own sakes, perhaps associated with linked text, but maybe simply as resources for inclusion in other documents, such as teaching materials. It is essential that if these resources are going to be useful they need to be indexed to allow them to be found.

Images can be associated with keywords, but often the requirement for the search is more complex than at the keyword level — all the pictures with meadows on sunny days might be accessible by keyword but what about all the images with a certain shade of blue, or with objects that look like "this".

Metadata to allow searching and access across resources is also needed and again the needs of graphical data are more complex than those of textual data.

Much of this work is at the research level still but there is a real requirement to deliver suitable technology soon.

A workshop is being organised by the Coalition for Network Information on the indexing and metadata needs for images (where these are defined a still raster images) in September 1996. See:

http://www.cni.org/CNI.homepage.html

There are also archiving issues. There is a need for service providers to archive material at the highest possible quality and to see archiving and delivery formats as being separate. We can store at high quality and send an adequate copy of an image across the network.

An initiative between the commercial sector and the education sector called the "Knowledge Gallery" has been launched in the UK recently. This aims to provide an interface to image related data which are distributed. It will face the many challenges outlined in this report.

Looking Ahead

If we are to realise the vision of the authors of the WWW as quoted in the introduction applied to graphical information and to take up the challenge of Al Gore, we have to move beyond the lowest common denominator graphics we see as standard on the WWW. We need the standard versions of the browsers to have a wider range of formats accepted by their native form. We need to go beyond GIF to vector formats, 3D models and application-related data.

In realising this it is important to recognise the need for tools for informal and formal testing to help analyse any problems which may occur.

Indexing and metadata issues need to be addressed with regard to the more complex data types involved in graphical applications. The related topics of archiving and quality assurance are also issues.

Looking ahead it is possible that some of these problems may go away with the use of Java. The graphical data can come with appropriate software to handle the format. Maybe browsers will not be the issue, though the data types and file formats in use will still need to be broadened.

There are plenty of challenges!

Acronyms

AGOCG¸Advisory Group On Computer Graphics

ATC¸Advanced Technology Center

AVI¸Audio/Visual Interleaved data

CDF¸Common Data Format

CGM¸Computer Graphics Metafile

CGRM¸Computer Graphics Reference Model

CORBA¸Common Object Request Broker Architecture

DWF¸Drawing Web Format

DXF¸Data eXchange File

FIF¸Fractal Image Format

GIF¸Graphics Interchange Format

GKS¸Graphical Kernel System

HDF¸Hierarchical Data Format

JPEG¸Joint Photographic Expert Group

MPEG¸Motion Picture Experts Group

NetCDF¸Network Common Data Format

NITF¸National Imagery Transmission Format

OII¸Open Information Interchange

OMG¸Object Modelling Group

ORB¸Object Request Broker

PHIGS¸Programmers Hierarchical Interactive Graphics System

PNG¸Portable Network Graphics

STEP¸STandard for the Exchange of Product model data

SVF¸Simple Vector Format

TIFF¸Tagged Image File Format

VRML¸Virtual Reality Modelling Language

W3C¸WWW Consortium

- References
- Contents
- Looking Ahead

Return to AGOCG Home Page

Back to AGOCG Technical reports

References

A very useful WWW address for information on file formats is currently found through the EC Open Information Interchange Initiative which is part of the IMPACT 2 programme. This contains useful briefings on various standards and updates of various workshops. At the time of writing this seems to be up to date and the challenge will be to ensure this is maintained. As with all information on the WWW, sustaining the accuracy of the information in a changing world is not an easy task. The dominance of the information is towards ISO standards though a number of industry standards are also included.

The reference page is:

http://www2.echo.lu/impact/oii/oiistand.html

A further source of information to read and to contribute to is a resource which was initiated through AGOCG. This is found on the AGOCG WWW pages. This allows searching through the database and allows readers to add to the information in the database.

Index to Multimedia Information Sources edited by: Simon Gibbs and Gabor Szentivanyi can be found at:

official home: http://viswiz.gmd.de/MultimediaInfo

A useful general book on file formats is:

Graphics File Formats, James D Murray & William van Ryper, O'Reilly Associates Inc, 1994, ISBN 1-56592-058-9

Information on other formats and resources quoted in this paper can be seen in the references below.

Acrobat	see: Adobe; Smith, 1993
Adobe	http://www.adobe.com/

AGOCG	/agocg/ note that the ":8080" may be removed in the future or contact the author of this paper for more information
ARGUS	http://www.geog.le.ac.uk/argus/
Autodesk	http://www.autodesk.com/
Casher & Rzepa, 1996	Casher, O. and Rzepa, H.S., The Molecular Object Toolkit: A New Generation of VRML Visualization Tools for use in Electronic Journals, Proceedings of the Eurographics UK Conference, 1996 Also available at: http://www.ch.ic.ac.uk/rzepa/eg/ see other papers by the same authors at this site on the use of WWW related tools in chemistry
CGM	Information Processing Systems - Computer Graphics Metafile for the storage and transfer of picture descriptive information', ISO/IEC 8632, Parts 1-4, 1992 and Amendments 1 & 2 see also Henderson & Mumford /agocg/CGM.html http://www.intercap.com
CGRM	Information technology - Computer Graphics - Reference Model, ISO 11072, 1992.
CORBA	http://www.acl.lanl.gov/CORBA/corba.html
DWF	see Autodesk
Envoy	http://www.twcorp.com/viewer.htm
FIF	http://www.iterated.com/
Gershon and Brown	Gershon, N. and Brown, J., Computer Graphics and Visualization in the Global Information Infrastructure, special issue of IEEE Computer Graphics and Applications, March 1996, p60-75
GIF	Graphics Interchange Format Programming Ref., Ver. 89a, CompuServe Inc., 1990.
GKS	Information Processing Systems - Graphical Kernel System (GKS)', ISO/IEC 7942, 1985 and 1994 (new version of GKS)
Greenough, 1996	http://www.cis.rl.ac.uk/proj/midas.html
Hardwick 1996	Hardwick, M. et al, Sharing Manufacturing Information in Virtual Enterprises, Communications of the ACM, Feb 96, Vol 39, no 2, p 46-54 also see: http://www.rdrc.rpi.edu/

Henderson & Mumford, 1993	Henderson, L.R., Mumford, A.M., The CGM Handbook, Academic Press, 1993. ISBN 0-12-510560-6
JPEG	jpeg-info@uunet.uu.net
KINDS	http://cs6400.mcc.ac.uk:80/kinds/
Kodak Photo-CD	http://www.kodak.com/
MPEG	MPEG-2 FAQ: http://www.tc.cornell.edu/Visualization/Education/cs418/MPEG/mpeg2.html
PDF	see Adobe
PHIGS	Information processing systems - Computer graphics - Programmer's Hierarchical Interactive Graphics System (PHIGS), ISO/IEC 9592, 1989.
PNG	http://quest.ipl.nasa.gov/PNG/ http://www.w3.org/pub/WWW/TR/WD-png.html Wegner, T., A PNG Format Primer, EXE: The Software Developers' Magazine, Vol 10,3, August 1995, p 21-27
PostScript	Adobe Systems Inc., PostScript Language Reference Manual, Addison-Wesley, 1985
Quicktime	downloadable viewer: http://www.graviton.com/help/quicktime.html
Smith, 1993	Smith, P.N. et al, Journal publishing with Acrobat - the CAJUN project, Electronic Publishing, Vol 6(4), 1-13, Dec 1993, John Wiley.
STEP	Standard for the Exchange of Product Model Data (STEP, ISO/IEC 10303
SVF	http://www.stat.ucla.edu/develop/web/languages/svf/
Testing	http://www.w3.org/hypertext/WWW/Test *WWW testing* http://www.uark.edu/~wrg/ *test patterns* http://speckle.ncsl.nist.gov/~lsr/cgm.htm *CGM testing*
TIFF	ftp://ftp.sgi.com/graphics/tiff/TIFF6.ps.Z
translators	http://info.mcc.ac.uk/CGU/SIMA/convert.html
UK VR-SIG	http://www.crg.cs.nott.ac.uk/ukvrsig/
Vitrus	+1 919 467 9700

VRML	http://vag.vrml.org/ http://vrml.wired.com http://www.sdsc.edu/vrml