THE TECHNOLOGICAL SETTING OF GIS

M F GOODCHILD

The development of GIS has been driven at least in part by technology, particularly the specific technology required to support spatial and graphic applications in computing. This chapter reviews the degree to which technology has constrained or driven GIS, and looks at prospects for the future given likely trends in the computing industry. GIS is becoming less dependent on technology as factors such as data volumes and the need for trained staff become more critical.

INTRODUCTION

GIS is a late bloomer among applications of computing technology in part because it is so demanding, and simply could not be supported in any useful fashion by the resources available in the typical computer system of, say, 1960. In addition, the spatial nature of geographical data is not easily accommodated within the essentially linear structure of conventional computers, and early input and output devices lacked the spatial resolution necessary to deal with these kinds of data. Indeed, it is often noted that the human eye and mind are still superior to the best digital technology in such spatial tasks as pattern recognition. The previous chapter (Coppock and Rhind 1991) has already argued that the history of GIS development can be linked to more general advances in hardware and software, although other factors such as education and awareness and the actions of key individuals have also been important. The purpose of the present chapter is to review the technological setting of GIS in the early 1990s, to identify the key technologies on which GIS depends, and to discuss certain significant trends that will likely influence this technological setting in the coming decade. The chapter is divided into two major sections: the four parts of the first section discuss the current setting from the perspectives of input, storage, manipulation and output; and the second section looks at the state of the GIS industry in the early 1990s in the context of several key issues.

Looking back over the past three decades of GIS development, it seems that the application imposes certain specific requirements, each of which had to be met before GIS could really blossom. Although developments such as CGIS (the Canada Geographic Information System; Tomlinson, Calkins and Marble 1976) took place when many of these were not available, or available only at high cost, their primitive technological environments certainly presented their developers with enormous problems. The following list is undoubtedly far from complete, but serves to emphasize the demanding nature of the GIS application (see also Maguire, 1989, for an overview of computers and geography):

- *Interactive.* The user must be able to interact with the computing system, continually issuing instructions and receiving responses. Widespread availability of interactive computing dates from the late 1960s, and has now largely replaced the earlier batch mode.
- *Multiuser.* Many users must be able to access the geographical database simultaneously. Multiuser operating systems date from the late 1960s. More recently, technology has become available to support distributed databases, avoiding the need to concentrate all data in one central location.

- *Graphic.* The system must be able to input and output data graphically, as it is otherwise very difficult for the user to work with geographical information. Primitive graphic output devices date from the mid-1960s, while the first digitizer for map input became available in the early years of the same decade.
- *Volume and speed.* Geographical data sets are often large and complex, and therefore require large digital storage devices. At the same time the system must be capable of processing large volumes quickly, and providing immediate answers to queries. Magnetic tape, which was the primary storage medium for computing until the mid-1960s, is too slow for interactive use.
- *Virtual memory.* Until recently the central, random access memory of computing systems was very expensive. The development of virtual operating systems (e.g. DEC's VMS) in the late 1970s allowed applications to process large volumes of data using comparatively small central memories.
- *Database management systems.* Contemporary GIS are enormously complex software systems. The GIS itself must be supported by an operating system, and will likely also depend on the presence of a graphics package for input and output, routines supplied by the programming language in which the GIS is written, and numerous other software products. Many of the more powerful contemporary GIS are designed to rely on a data base management system (DBMS), relieving the GIS developer of many of the common data housekeeping functions. DBMSs were adopted widely in the computing industry in the 1970s (Tsichritzis and Lochovsky 1977). At the same time the marriage between GIS and DBMS which began to emerge in the software products of the 1980s is not perfect. GIS require the rapid display of large quantities of data, and access to information through location as well as attributes, and it has proven difficult to accommodate the multidimensional nature of geographical data (two or three dimensions of location plus attributes) within the framework of many DBMSs.
- *Cost.* GIS is a minor application compared to the needs of accounting, mathematical modelling, word processing, etc. A resource management agency or utility company is willing to adopt GIS only when its perceived benefits exceed its costs (Dickinson and Calkins 1988). The explosion of interest in GIS in the 1980s is due at least in part to the steady and spectacular drop in the costs of computing technology over the past three decades.

THE CURRENT SETTING

Input

Most of the contents of geographical databases continue to be derived from paper maps, and the technology of the early 1990s offers two main methods of input: digitizing and scanning. Digitizing requires a human operator to position a map on a table, and move a cursor over it, thus capturing the positions of points, and building up digital representations of line and area features. Digitizing is error prone, tedious, time consuming and expensive – issues that are dealt with in more detail by Fisher (1991 in this volume). Estimates of the current cost of successfully inserting a single area, such as a land parcel, into a data base range from $1 to $40 (Goodchild and Rizzo 1987; Goodchild 1987). In essence, digitizing technology remains at a stalemate. Map-size digitizers were developed in the 1960s, and there have been few if any advances in the past two decades (see Jackson and Woodsford 1991 in this volume). Software for intelligent management of the digitizing process and interactive editing of the results was developed in the late 1970s. These functions require a small, dedicated workstation with simple graphics capabilities, and can be readily provided by the current generation of personal computers and workstations.

Scanning captures the entire contents of the document automatically. However, the process of interpreting features from the scanned image of the map is error prone, requiring frequent human intervention, and it is often necessary to redraw the document to provide a sufficiently clean image. Map-size scanners are large and expensive, and unlikely to become much cheaper. The most important factor in the development of scanning technology is the software for feature detection. Simple problems such as the recognition of characters in feature labels, or the interruption of contours by elevation annotation remain major challenges for software developers. If they can be overcome, or removed through the appropriate redesign of map formats (Shiryaev 1987) then scanning will inevitably win in the cost effectiveness battle with digitizing, but in the early 1990s the contest is still even. Jackson and Woodsford (1991 in this volume) present a detailed discussion of the current state of data capture technology.

Maps are a form of analogue geographical database in their own right, and a very important source of digital geographical data, but they are certainly not the only source. Maps are often designed for purposes that have little to do with geographical databases, a primary purpose of map design being the communication of information to the reader through visual perception (Robinson 1953). For example, the purpose of a set of elevation data in a GIS is likely to be to provide accurate estimates of the elevation of any point on the earth's surface in the area covered by the data set; the purpose of a map showing contours of elevation is likely to be to provide the reader with an impression of the general form of the surface (Imhoff 1982). Maps use methods of representing geographical variation, such as contours, that may be far from optimal for digital representation.

Thus an increasing amount of GIS data is now being provided by non-map sources. Davis and Simonett (1991 in this volume) review the current state of technology in remote sensing (see Ehlers, Edwards and Bedard, 1989 for a review of GIS/ remote sensing integration), and explain the tension that exists between two paradigms: remote sensing as a method of collecting and compiling information on geography through image interpretation; and remote sensing as a source of data to support studies of environmental processes and earth system science. For map data, the key problem is the high cost of conversion to digital form, but for remote sensing, the problem is just the opposite. The Landsat technology of the 1970s created an avalanche of images from space that greatly exceeded the capacity of the data analysis technology of the time. Despite the enormous advances in image processing and GIS since then, the technology of remote sensing instruments is moving rapidly. The Eos systems proposed for the 1990s will include much higher spatial and spectral resolution than the current generation (10 m and 256 bands are proposed, respectively, leading to estimates of 1 Tb or 10^{12} bytes of information per day) and it is likely that storage and analysis capabilities will lag far behind.

A major advance in data capture technology is occurring with the development of GPS (Global Positioning System: the magazine *GPS World* is a useful source of information on this technology). Hand-held GPS receivers will soon be able to determine position anywhere on the earth's surface in latitude/longitude coordinates to an accuracy of fractions of a second of arc (1 second of arc latitude is approximately 30 m). As well as being an accurate source of raw geographical information, GPS has potential in such applications as tracking vehicles and shipments, emergency response and resource management. Its accuracy is already greater than that of the best base mapping commonly available in most countries, including the United States.

The accuracy of GPS has interesting implications for the future of mapping. For a relatively small investment (around $100 000), it is already possible to equip a vehicle to make accurate surveys of road or railway alignments. Since the vehicle can travel at speeds of up to 100 km/h and still maintain adequate accuracy (defined as the accuracy of currently available topographic mapping, e.g. 1 : 25 000), the costs of this method of map making may be an order of magnitude less than conventional methods based on photogrammetry. GPS is likely to play a major role in the creation and updating of future road and street network databases. In addition, there is strong interest in the merging of GPS-surveyed road and rail alignments with images, for instance, of road surface quality or road signs. This will require multimedia technology, where the GIS is capable of handling both structured network data, to represent the road alignment, and unstructured video images.

GPS and remote sensing are useful illustrations of the difference between data and information, as neither provides any intelligent interpretation of the data they capture. Just as the retina provides limited interpretation of images rather than send them uninterpreted through the optic nerve to the brain, the proposed data rates of Eos are so high that some sort of intelligent filtering will be almost inevitable before the data are transmitted or stored. Digitizing and scanning technologies both concentrate on capturing the contents of a map accurately and completely, leaving the more demanding tasks of data interpretation and compilation to the cartographer. Only recently has it been possible to begin developing more intelligent data capture systems, using digital representations of knowledge about geography. The recent developments in scanning systems described by Jackson and Woodsford (1991 in this volume) fall into this category, and research is currently under way in a number of organizations to apply digital technology to the task of compiling and interpreting geographical information, as distinct from capturing geographical data. Hutchinson (1988), for example, has developed improved methods of building elevation databases using knowledge about hydrology.

Storage

Modern GIS store data using a variety of technologies, some of which are generic and some of which are almost uniquely appropriate to this application. In general, geographical databases demand large amounts of storage, and have only become feasible with the decline of storage costs and dramatic improvements in storage technology of the past three decades. Paradoxically, the paper map is a highly efficient store of data, and the problem of converting its contents to a manageable digital representation has been at least partly responsible for the comparatively slow pace of GIS development.

The primary storage medium is currently the fixed magnetic disk, or 'hard disk', with capacities ranging from tens of megabytes (1 Mb $\approx 10^6$ bytes) to tens of gigabytes (1 Gb $\approx 10^9$ bytes). For comparison, a megabyte is sufficient to store the contents of a book of roughly 200 000 words. The contents of maps are so variable that it is difficult to give meaningful estimates of their storage requirements in digital form, but experience in agencies such as the US Geological Survey and the Defense Mapping Agency suggests that a typical topographic map's contents can be stored with acceptable precision in roughly 10 Mb, to within an order of magnitude, which is also the approximate volume of one band of a Landsat scene; and the US Bureau of the Census TIGER database which includes digital representations of most streets in the United States, occupies roughly 10 Gb. Any item of data on a hard disk is accessible within milliseconds.

Magnetic tape (capacity roughly 100 Mb) remains a useful storage medium for geographical data and the primary method of back-up for fixed disk. However, long times are required to 'seek' data by winding through tape serially. Optical disks, such as the CD-ROM (Compact Disk–Read Only Memory) and WORM (Write Once Read Many) with capacities from 250 Mb to 2 Gb can seek in fractions of a second, and are particularly suitable for archives because of their long-term stability, which matches the nature of much geographical data. Once a master CD-ROM has been manufactured, it is possible to print and distribute copies at very low cost, so this medium is becoming increasingly popular for such data sets as TIGER. But against the advantages of long-term stability of the medium, it must be added that this technology is in a state of rapid flux.

Despite their high capacity, none of these media offers solutions to the problems of storing geographical data sets of terabyte size (1 Tb $\approx 10^{12}$ bytes), such as the US Geological Survey's Digital Cartographic Database, or the databases of many major utilities or cities (For a discussion of very large spatial databases see Crain 1990.) A number of automated magnetic tape stores have been built, capable of mounting tapes from a library of perhaps 10^4 volumes in minutes, but these are cumbersome and expensive. Optical platters can be mounted in 'jukeboxes' offering seek times of tens of seconds, but it is difficult to build stores using fixed disks for more than about 10^{11} bytes.

Because of these constraints, which effectively limit the amount of data available for interactive use at any one time, most large geographical databases are partitioned in one way or another. The term 'tile' is commonly used to refer to geographical partitions, often on the basis of map sheets or political units; although the database may be partitioned internally, the user may nevertheless see it as 'seamless'. Partitions by subject matter are referred to as 'layers', 'themes' or 'coverages', and the system may impose constraints on the different classes of features that can be stored in any one such partition. In addition, many GIS databases are partitioned temporally, maintaining maps of an area at different points in time.

Within any one partition, it is desirable that the user be able to retrieve any feature as rapidly as possible. Rather than search sequentially through the features in the partition in response to a request, many systems use some form of index for more rapid access. Indexes may be based on the characteristics of features, or on their locations, and used to support queries based on characteristics and locations respectively. Buchmann *et al.* (1989) include several papers on recent developments in indexing large spatial databases, and Samet (1989) gives a comprehensive review.

The most critical form of memory for a computer system is 'core', 'central' memory or RAM (Random Access Memory), which is the fast access memory used by the central processor (see Maguire, 1989, for an excellent general review of the architecture of computers in the context of geographical enquiry). As recently as the late 1970s, the high cost of RAM was reflected in the small amount available even in large mainframe computers. MS-DOS, still the dominant operating system of the IBM PC, is limited to 640 Kb (1 Kb $\approx 10^3$ bytes). But more recent multitasking operating systems such as OS/2 and Unix have been able to take advantage of improvements in manufacturing technology that have reduced the cost of RAM in 1990 to roughly $100/Mb. Cheap RAM means that computers can perform many tasks simultaneously, and analyse large data sets more efficiently.

Manipulation and analysis

Until the late 1970s the dominant configuration for GIS consisted of a mainframe computer serving a number of connected users, each operating a 'terminal'. All information, graphic or text, passed through a communication line and was constrained by the line's capacity, which might be no more than 30 characters/second. Coding systems, such as Hewlett Packard's HPGL (Hewlett Packard Graphic Language), or Tektronix's 4010 code, were used to express graphic information in the form of sequences of text characters. A terminal capable of displaying colour graphic information might cost over $10 000.

By 1990 the functions of mainframe computing and graphic display had been combined into powerful personal 'workstations', and the previously clear lines of distinction between personal computing and mainframes had become hopelessly blurred. The concept of a workstation emerged in the late 1980s as a specialized desktop system for scientific and engineering applications, offering a characteristic computing environment with little resemblance to that of the PC or Apple Macintosh, and dominated by the Unix operating system. The typical workstation of the early 1990s, costing no more than the graphic display terminal of 1980, has the computing power, storage capacity and display capabilities to support a powerful GIS in a single, integrated desktop 'platform'. It may have a central processor speed of over 10 MIPS (Million Instructions Per Second), 10 Mb of RAM, and 1 Gb of hard disk. Workstations are commonly networked, and can share distributed databases (the 'diskless' workstation has no local storage, obtaining all of its data over a network). The workstation's graphic display resolution is normally expressed in terms of rows and columns of picture elements (pixels), and the number of colours simultaneously displayable on the screen. For GIS applications, minimal requirements are typified by 480 rows by 640 columns, 16 colours displayable (the EGA standard of the IBM PC); high quality is typified by 1024 rows by 1280 columns, 256 colours displayable (10^6 pixels = 1 'megapel').

Although the contents of GIS databases continue to be dominated by structured geographical information, including well-defined objects with coded attributes, other forms are increasingly important in some applications. In digitizing, for example, it may be useful to display an unstructured image of a map or an air photograph on the screen, allowing the user to input the locations of features directly. This is a commonly used approach in digitizing utility installations such as underground pipes. In emergency management applications it may be useful to allow pictures of buildings to be retrieved and displayed on the screen. Sound may also be important. The GIS community is increasingly recognizing that structured geographical data are only one form of spatial information, and that other forms can be exploited very effectively using appropriately designed platforms. (For a discussion of multimedia databases see Shepherd 1991 in this volume.) The power of the eye (and ear) to interpret and recognize pattern is so great that the high degree of interpretation present in a structured GIS database may not always be necessary.

Because of its reliance on map input, GIS has been seen as a two-dimensional technology.

Surfaces, such as that of the earth, can be handled only if they are single-valued functions of the two spatial dimensions (see Weibel and Heller 1991 in this volume), and the earth's curvature must be dealt with by use of map projections (see Maling 1991 in this volume). Raper and Kelk (1991 in this volume) and Raper (1989) deal with current developments in 3-D GIS, which are occurring in a rapidly changing technological setting. In the context of display, Tektronix introduced a device in the mid 1980s that allows objects on the screen to be seen in stereo vision using polarizing filters. The screen displays alternately the left and right images using twice the normal refresh rate. These are then polarized in opposite directions by a filter installed in the front of the screen. The user wears eyeglasses containing polarizing filters, but otherwise interfering very little with normal vision.

Several of the 1990 generation of workstations, typified by the IBM RS/6000, include graphics processors (or 'adaptors') with 3-D display functions, capable of computing, displaying and texturing 3-D solid objects on the screen at high rates. With specifications of 10^5 3-D vectors per second, these systems are able to offer real-time display and manipulation of scenes containing illuminated solid objects in a desktop computing environment for investments of the order of $10 000. For GIS, this means the ability to simulate oblique views of terrain with superimposed geographical distributions; to replace abstract cartographic symbols and classes with 'artists' impressions' of real physical appearance; to animate displays; to work with geographical distributions over the curved surface of the globe; and to build and manipulate true 3-D databases in applications such as oceanography, geology and subsurface hydrology.

The workstation has replaced the old centralized view of computing as a mainframe with connected users, as modified by the independent personal computer, with a network of distributed computing power and distributed databases, communicating through common standards and protocols. Unix has emerged as the standard operating system of choice in this environment, replacing both the proprietary operating systems of the mainframes, and the simple environments of the PCs. Unix has been enormously successful at offering a uniform computing environment over a variety of vendor products, and at integrating systems from the simplest to the most powerful. At the same time, the conventional form of user interaction has moved strongly away from the old, text command format to the graphic user interface (GUI) pioneered by Xerox and Apple, among others, in which the user communicates largely by pointing at objects on the screen. Frank and Mark (1991 in this volume) discuss the implications of the move to GUI and visual metaphors for GIS and spatial databases. As yet, none of the competing GUI standards has emerged as dominant in the GIS field.

Output

Numerous devices can be used to create an output image in the form of a map. In the early days of automated cartography, the line printer was even used, by overprinting characters to create various shades of grey, despite the limitations of a picture element measuring 1/6 inch by 1/10. Early plotters offered enormous improvement by allowing the user to direct the movement of a pen, but required a substantially different software design to create an image by random pen movements ('vector') rather than by sequentially printed picture elements ('raster'). As computer processors have become cheaper, it has been possible to offer devices with substantial internal computing power for conversion between raster and vector, blurring the distinction between plotter and printer. Thus the electrostatic plotter, probably the most popular GIS map output device of 1990, generates images by assembling uniform-size picture elements sequentially (with typical resolutions of hundreds of dots per inch) but can be driven as a plotter by sending information in the format of pen movements. The distinction between 'plotter' and 'printer' has become largely meaningless as the relatively unreliable technology of moving pens has given way to various forms of electrostatic printing.

THE GIS MARKETPLACE

GIS is a complex and diverse field, more a loose consortium of interests than a mature industry. Its products reflect this, and it is possible to speculate endlessly about whether the future holds greater uniformity or greater diversity. Numerous surveys of GIS products have been published; those in the *AGI Yearbook* (an annual publication assembled by the Association for Geographic Information in the United Kingdom and published by Taylor and Francis) and the *GIS Sourcebook* (published annually by GISWorld, Fort Collins, Colorado) are comparatively accessible and updated regularly. Estimates of the total number of GIS products on offer range up to 1000. The purpose of this section is to provide an organizing framework in terms of several key issues, as a means of describing the technological setting of this volume.

Arguably the most fundamental means of distinguishing between GIS products is through an understanding of each system's data model. Since geographical reality is infinitely complex, it is necessary to generalize, approximate or abstract in order to create a representation in the finite, discrete space of a digital store. A data model is defined here in the sense of Peuquet (1984) as 'an abstraction of the real world which incorporates only those properties thought to be relevant to the application or applications at hand, usually a human conceptualization of reality'. Tsichritzis and Lochovsky (1977: 21) also offer a useful definition: 'a data model is a set of guidelines for the representation of the logical organization of the data in a data base. It is a pattern according to which the data, and the relationships among the data, can be logically organized. It consists of named logical units of data and the relationships among them.' (See also Guptill, 1991 in this volume, and Maguire and Dangermond 1991 in this volume.)

Although numerous data models have been devised, no one system offers all of them. Moreover some data models are better able to represent a given geographical phenomenon than others, so the choice of data model(s) offered by a given system is driven by, and in turn constrains, that system's applications.

The second way of differentiating current systems is through the supporting platform, and its associated operating system, database management system, and so on. While some vendors offer similar products over the full range of platforms from personal computers through workstations to mainframes, others are limited to one or perhaps two operating systems, and, therefore, to the platforms that support them.

Finally, it is possible to differentiate current systems according to their areas of application. Although some vendors attempt to market their products across the entire range of GIS applications, others have specialized into particular market niches.

The following discussion does not attempt to provide a comprehensive list of products, or detailed information on their availability. For these, the reader is referred to current lists and reviews, such as those cited earlier.

Data models

Perhaps the strongest cleavage in the marketplace is between systems dominated by raster and vector data models, although the distinction is rather blurred (see below). A raster database models each layer or theme as an assemblage of cells or pixels, specified in some standard order. No coordinates are stored, as they are implied by the position of each cell in the sequence. Raster systems differ in the constraints imposed on the number and types of values allowed in each cell of a layer. Arbitrarily shaped features on the earth's surface, such as lakes or roads, must be represented by giving appropriate values to cells. For example, a lake may be represented by giving the value 'L' to cells whose centre points lie in the lake, and a 'Lake' layer can be constructed by giving '0' to all cells not classified as 'L'. In this form the existence of a named lake, for example Lake Superior, is implicit, and its dimensions can only be found by searching for a contiguous block of 'L' cells. The various forms of the MAP package originally developed by Tomlin (1983) are examples of this approach. Some raster systems, for example TYDAC's SPANS, support the notion of a feature more explicitly by coding all of the Lake Superior cells with a unique feature number, and associating these with a table identifying the characteristics of each feature.

Vector systems, on the other hand, employ a dominant data model that describes features as points, lines or areas through sets of coordinate pairs. ESRI's ARC/INFO, for example, allows a point, line or area coverage to contain the respective class of feature, with an associated table of attributes. Rasters can be handled in a number of ways in several parts of the system, just as vectors can be used as a convenient method of input in raster systems such as IDRISI, but in both cases the data model must be transformed to take full advantage of the system's capabilities.

The functions offered by each system depend closely on the data model. Transformation of coordinate systems, and use of map projections are common features of coordinate-based vector systems, but less well represented in raster systems. Most raster systems include the ability to generate an oblique view of an elevation data set, but this is rare in vector systems. Although raster systems offer comparatively consistent functionality, there is an enormous range of variability both of data models and functionality within the vector domain. A utility company installing IBM's GFIS, for example, would take an entirely different approach from one using ESRI's ARC/INFO. In part this is because a collection of features, for example areas, can have a much wider range of meanings than a layer of raster cells. Some systems, such as ESRI's ARC/INFO, require that the areas in a coverage be 'planar enforced', meaning that they may not overlap, and must exhaust the plane, so that every place lies inside exactly one area. Other systems, such as Intergraph's TIGRIS, allow coverages to contain any mixture of points, lines and areas with or without overlap. Planar enforced systems often 'store topology', meaning that they store adjacency relationships between areas. Others such as TIGRIS may have 'integrated topology', meaning that all intersections, containments and adjacencies between the arbitrary collection of features in a coverage are stored. Some systems, such as Prime's SYSTEM/9, allow the user to define complex objects as collections of simple objects, and to give them their own graphic symbols on output. Data models are an important key to understanding the design of a GIS product, and its functionality, and are much less predictable in the vector domain. (See Maguire and Dangermond, 1991 in this volume, for discussion of the link between functionality and GIS.)

Platforms

From a historical perspective, it is possible to identify three classes of computing: mainframes, consisting of large multiuser processors connected to interactive terminals; personal computers, which are small independent computing systems; and workstations – larger and more powerful than personal computers, with multitasking operating systems and networking capabilities. The distinction between the three classes is not precise. A fourth class, the 'minicomputer', was important in the 1970s and early 1980s but is now largely absorbed by the workstation and mainframe categories. Each class has had its dominant operating systems: proprietary systems such as DEC VMS, IBM VM/CMS or PRIME Primos for mainframes, Microsoft MS-DOS for PCs and various forms of Unix for workstations. Each has also had its associated market, with PCs dominant in small business and home applications and workstations oriented to science and engineering.

A few GIS vendors have attempted to offer similar products across the entire range of platforms, and ESRI's ARC/INFO, which started as a mainframe product, is the clearest example. Others, such as Intergraph, have historically provided software for customized hardware environments. TYDAC is an example of a vendor that began in the PC marketplace under MS-DOS, and has migrated to OS/2 on the same platform, and to workstation Unix. As the networking advantages of Unix become clearer, and this operating system is increasingly available on mainframes and PCs, it will likely come to dominate the GIS field as a platform-independent environment. Nevertheless MS-DOS has a powerful hold on the small computer marketplace, and the Macintosh GUI environment is also popular.

Applications

In this last section the differentiation of the GIS marketplace by application is examined. The chapters in Section III of this volume provide additional comments on GIS products in various application contexts. In municipal and utility applications, where GIS is used as a means of gaining geographically based access to large inventories of facilities, parcels and records, the sheer size of the application dictates the use of a large mainframe. Small workstations and PCs may be integrated as platforms for editing or intensive analysis, but the GIS of choice is likely to be one that runs under a mainframe operating system, or increasingly Unix (e.g. ARC/INFO, Genasys and GeoVision).

In market research, planning and emergency applications data volumes are much smaller, and the need for personal input and management of data much higher. Vector-based systems such as Strategic Mapping's Atlas GIS, Generation 5's Geo/SQL, MapInfo and Caliper's GIS Plus have been able to penetrate this market, using small MS-DOS platforms, and several Macintosh products are significant also. In resource management applications, where remote sensing input is significant and the concept of continuous geographical variation is more acceptable, raster-based PC systems, such as MAP, IDRISI and EPPL7 have been successful. Finally, TYDAC's SPANS has been able to provide a limited degree of integration of raster and vector models on the PC platform.

Although many GIS products are moving to the workstation platform, among the earliest were systems designed for scientific and modelling applications, such as GRASS. Because of their intermediate status between PCs and mainframes, and the integrating power of Unix and networking, workstations are now the platform of choice for many GIS, and this trend can only accelerate as the distinction between the three types of platforms becomes less distinct.

CONCLUDING REMARKS

The technological setting of GIS is changing rapidly, as workstation computing power increases by a factor of roughly two every year, as hardware of all types continues to become cheaper, and as networking protocols and standards become more widespread. Among the very few constants in this dynamic field are those involving humans – the high cost of digitizing, and the constant demand for faster computing, larger and more accurate databases, and better and more friendly user interfaces. There seems to be no limit to the potential complexity of software systems, as the technology has already gone far beyond the point where any one individual can understand every detail of the design of an operating system or a GIS. Storage and computing power will continue to fall in price over the next decade, 3-D and animated graphics will become cheaper, more widespread and easier to use. Unix will strengthen its grip as the universal operating system for the new distributed style of computing. GIS vendors will also continue to offer a wide variety of products adapted to the needs of particular applications.

In this environment of continuous change it is difficult sometimes to see the more fundamental longer term changes. For a long time the controlling factor in GIS applications was hardware cost, but now software cost is at least as important. Communication speed was critical in the early 1970s when GIS relied on terminal–mainframe interaction. Now that these challenges have diminished, new ones have become important – for example, there is no reasonable way of processing the enormous volumes of data about to come from Eos, and current GIS offer a very limited view of the rich world of geographical data modelling. Databases are still largely map based, and filled with structured geographical information.

In the coming decade the technological problems which plagued earlier generations of GIS products will be far less important than the human ones – lack of trained staff, the high personnel costs of digitizing, poor planning and management, resistance to technological change within institutions, and so on. At the same time expectations will continue to rise in step with the technology, and to demand ever more precise data, in ever larger volumes. Although many of the simpler technical problems have been solved, and costs have fallen dramatically, the technological environment of GIS will continue to constrain the expectations of its users.

REFERENCES

Buchmann A, Günther O, Smith T R, Wang Y-F (1989) *Design and Implementation of Large Spatial Databases. Lecture Notes in Computer Science 409.* Springer-Verlag, New York

Coppock T, Rhind D W (1991) The history of GIS. In: Maguire D J, Goodchild M F, Rhind D W (eds.) *Geographical Information Systems: principles and applications*. Longman, London, pp. 21–43, Vol 1

Crain I K (1990) Extremely large spatial information systems: a quantitative perspective. *Proceedings of the 4th International Symposium on Spatial Data Handling* International Geographical Union, Columbus Ohio, pp. 632–41

Davis F W, Simonett D S (1991) GIS and remote sensing. In Maguire D J, Goodchild M F, Rhind D W (eds.) *Geographical Information Systems: principles and applications*. Longman, London, pp. 191–213, Vol 1

Dickinson H J, Calkins H W (1988) The economic evaluation of implementing a GIS. *International Journal of Geographical Information Systems* **2** (4): 307–28

Ehlers M, Edwards G, Bedard Y (1989) Integration of remote sensing with GIS: a necessary evolution. *Photogrammetric Engineering and Remote Sensing* **55** (11): 1619–27

Fisher P F (1991) Spatial data sources and data problems. In: Maguire D J, Goodchild M F, Rhind D W (eds.) *Geographical Information Systems: principles and applications*. Longman, London, pp. 175–89, Vol 1

Frank A U, Mark D M (1991) Language issues for GIS. In: Maguire D J, Goodchild M F, Rhind D W (eds.) *Geographical Information Systems: principles and applications*. Longman, London, pp. 147–63, Vol 1

Goodchild M F (1987) Application of a GIS benchmarking and workload estimation model. *Papers and Proceedings of Applied Geography Conferences* **10**: 1–6

Goodchild M F, Rizzo B R (1987) Performance evaluation and workload estimation for geographic information systems. *International Journal of Geographical Information Systems* **1**: 67–76

Guptill S C (1991) Spatial data exchange and standardization. In: Maguire D J, Goodchild M F, Rhind D W (eds.) *Geographical Information Systems: principles and applications*. Longman, London, pp. 515–30, Vol 1

Hutchinson M F (1988) Calculation of hydrologically sound digital elevation models. *Proceedings of the 3rd International Symposium on Spatial Data Handling*. International Geographical Union, Columbus Ohio, pp. 117–33

Imhoff E (1982) *Cartographic Relief Presentation*. de Gruyter, Berlin

Jackson M J, Woodsford P A (1991) GIS data capture hardware and software. In: Maguire D J, Goodchild M F, Rhind D W (eds.) *Geographical Information Systems: principles and applications*. Longman, London, pp. 239–49

Maguire D J (1989) *Computers in Geography*. Longman, London

Maguire D J, Dangermond J (1991) The functionality of GIS. In: Maguire D J, Goodchild M F, Rhind D W (eds.) *Geographical Information Systems: principles and applications*. Longman, London, pp. 319–35, Vol 1

Maling D H (1991) Coordinate systems and projections for GIS. In: Maguire D J, Goodchild M F, Rhind D W (eds.) *Geographical Information Systems: principles and applications*. Longman, London, pp. 135–46, Vol 1

Peuquet D J (1984) A conceptual framework and comparison of spatial data models. *Cartographica* **21** (4): 66–113

Raper J F (1989) *Three Dimensional Applications in GIS*. Taylor & Francis, London

Raper J F, Kelk B (1991) Three-dimensional GIS. In: Maguire D J, Goodchild M F, Rhind D W (eds.) *Geographical Information Systems: principles and applications*. Longman, London, pp. 299–319, Vol 1

Robinson A H (1953) *Elements of Cartography*. Wiley, New York

Samet H (1989) *The Design and Analysis of Spatial Data Structures*. Addison-Wesley, Reading Massachusetts

Shepherd I D H (1991) Information integration and GIS. In: Maguire D J, Goodchild M F, Rhind D W (eds.) *Geographical Information Systems: principles and applications*. Longman, London, pp. 337–60, Vol 1

Shiryaev E E (1987) *Computers and the Representation of Geographical Data*. Wiley, New York

Tomlin C D (1983) *Digital cartographic modelling techniques in environmental planning*. Unpublished PhD Thesis, Yale University.

Tomlinson R F, Calkins H W, Marble D F (1976) *Computer Handling of Geographical Data*. UNESCO, Paris

Tsichritzis D C, Lochovsky F H (1977) *Data Base Management Systems*. Academic Press, New York

Weibel R, Heller M (1991) Digital terrain modelling. In: Maguire D J, Goodchild M F, Rhind D W (eds.) *Geographical Information Systems: principles and applications*. Longman, London, pp. 269–97, Vol 1