22

Geographical information systems in networked environments

D J COLEMAN

This chapter examines basic approaches to distributed architectures in GIS, with special emphasis on exploitation of local- and wide-area networks and, more recently, the Internet and wireless communications. The first section reviews the design models inherent in host-based systems, distributed networks, and field-based systems. The second section describes selected applications of these design models, with particular emphasis placed on Internet-based geographical information retrieval and computer-supported cooperative work. The final section deals with system performance considerations in networked environments and describes new developments affecting the performance of future systems.

1 INTRODUCTION

Technical design models supporting mainstream GIS technology have evolved from early host-based efforts on mainframe and mini-computers, through stand-alone systems operating on personal computers and workstations, and on to today's distributed computing environments across local- and wide-area networks. Each advance has extended overall flexibility in terms of the relative location of users, processing capabilities, and data storage units (see also Batty, Chapter 21). Recent advances in broadband and wireless communications technologies – as well as the dramatic increase in Internet usage and extension of Internet browsing technology – promise to extend further the reach and range of GIS users working in offices or laboratories, in the field or at home.

This chapter examines the basic approaches to distributed architectures in GIS, with special emphasis on the exploitation of local- and wide-area networks and, more recently, the Internet and wireless communications. The first section of the chapter reviews the design models inherent in host-based systems, distributed networks, and field-based systems. The second section describes specific applications of the latest design models, with particular emphasis placed on emerging issues associated with Internet-based geographical information retrieval and computer-supported cooperative work. The final section of the chapter deals with the issue of system performance in networked environments under a variety of conditions. After a brief summary of considerations involved in objectively assessing system performance, the author describes new developments affecting the performance of future systems.

2 HOST-BASED SYSTEMS AND EARLY PERSONAL COMPUTING ENVIRONMENTS

2.1 Host-based systems

Early centralised computing environments were characterised by small numbers of large-scale mini- or mainframe computers, with shared storage devices attached via hardware input/output channels and multiple users connected via terminals possessing varying levels of on-board 'intelligence' or processing power (Katz 1991). Such environments defined the architecture for most major data processing and information systems applications until the mid 1980s, including the GIS installations found in major forestry organisations, utilities, municipalities, and land records management programmes.

This architecture implied greater control over data integrity and system security; the database was managed centrally, with responsibility for system and data administration entrusted to experienced data processing specialists. However, performance of such systems would often degrade in unpredictable ways when growing numbers of users demanded computing resources and database access. Conflicts with system administrators over development and maintenance priorities also often resulted in dissatisfaction among endusers in many large organisations.

2.2 Stand-alone, PC-based systems

By 1986, PC-based GIS software packages had begun moving geoprocessing out of the hands of information system managers (Miller 1990). Besides their low cost, these systems offered more predictable response times since the user was the only one on the system. However, the proliferation of stand-alone PC-based systems meant it was much more difficult to share data among several different people in the organisation. Also, the enduser often had to become his or her own system and database administrator. While PC-based systems undoubtedly accounted for the dramatic growth in GIS usage through the late 1980s, they also put greater onus on managers in large organisations to keep effective track of the data being collected and processed by an increasing number of endusers with little experience in routine data management procedures.

3 COMPUTER NETWORKING AND DISTRIBUTED SYSTEMS

Beginning in the mid 1980s, higher-performance workstations connected through local and wide area networks (LANs and WANs) became a viable alternative to host-based and stand-alone configurations. Connectionless LAN and 'LAN-interconnect' services began displacing earlier connection-oriented services which required a dedicated link between the enduser and the host computer (Figure 1).

Based on packet-switching transmission protocols which did not require dedicated connections between user and host, networking permitted users to share access to scarce equipment resources (e.g. printers, plotters, databases etc.), made possible inter-site communication applications like electronic

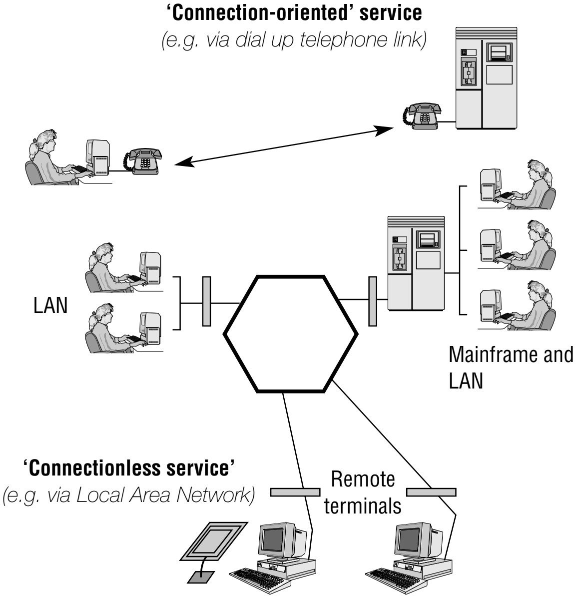

Fig 1. Examples of alternative wide-area networking services: 'connection-oriented' vs 'connectionless' models.

mail and file transfer, and enabled distributed processing. It also allowed users to expand their facilities with a degree of vendor independence and to incorporate special purpose processors, storage units, or input/output devices as required. The term *distributed computing* was coined to describe a situation where processing tasks and data are distributed among separate hardware components connected by a network, with all these various components capable of being accessed in a relatively transparent manner (Champine et al 1980).

Distributed systems today generally offer a combination of: (1) greater access to GIS data stored and managed on a central server; (2) faster response time due to local computing; (3) tighter system security; (4) less complexity; and (5) in many cases, lower-cost computing solutions than more traditional mainframe or minicomputer solutions. At a higher level, the notion of distributed computing often provides a better fit to the complex structures and often multidisciplinary nature of modern organisations and offers greater user involvement in information management activities.

Early investigators of formal network-based approaches to distributed GIS database management systems (DBMS: e.g. Ezigbalike et al 1987; Webster 1988) laid out excellent foundations for future research, but were limited by both the distributed DBMS technology then commercially available at reasonable cost, and the relatively limited use of LANs or WANs by the GIS community at that time. The subject has been revisited more recently by Ingoldsby (1991) and Laurini (1994), among others.

3.1 The client–server model

Most distributed computing today is based on a client–server architecture (Katz 1991). In this model, a collection of workstations (or *clients*) relies on one or more *servers* residing elsewhere on the network for access to data files, application software and, in certain cases, more powerful computing resources. Such servers are really high-volume storage devices with processors which have been optimised to provide high-speed retrieval of large disk-based data or database files. In such an environment, the data retrieval aspects of a database query can be carried out largely independently of the data processing and display tasks.

The time required to execute such a data retrieval operation depends on the individual performance of the storage, processing, and communications components involved in the client–server system. While improvements in all three of these technologies have made key contributions, it has been the introduction of remote file management services which finally provided the transparency required for distributed computing. Remote file system implementations like Sun Microsystems' Network File System (or NFS) on UNIX workstations, PC-NFS and Novell's Netware on PC-compatibles, and Appletalk on Macintoshes are now present in many organisations.

3.2 Network architecture

A computer network architecture can be considered as a set of functions, interfaces, and protocols which enables devices to communicate with one another on-line. The architecture is composed of a layered collection of communication, networking, and application functions implemented such that – while each layer is designed to operate independently – higher-level operations are built on functions provided by the lower layers (Chorafas 1980).

Network protocols are the formal sets of rules or specifications for coding messages exchanged between two communication processes on a network (Voelcker 1986). Protocols govern data control and format across a network, and a variety of protocols exist to ensure that these communications are conducted effectively.

Two different stacks or suites of layered approaches are in common use today:

- OSI: the seven-layer Open System Interconnect suite of protocols developed by the International Standards Organisation;
- TCP/IP: the four-layer Transmission Control Protocol/Internet Protocol suite originally developed for the ARPAnet research network in the USA.

Originally developed by computer user groups and European telephone companies, the OSI model helped unify world telephony and provided a clear framework and explanation of the functions required for computer communications. However, it was regarded by some as being too cumbersome for high-speed networks (Wittie 1991).

By comparison, the TCP/IP suite of protocols became a *de facto* standard by the early 1980s as a result of its early use in the implementation of the US Defense Department-funded Internet. Its longer-term popularity was secured through subsequent bundling with the 1983 release of the Berkeley UNIX 4.2 operating system. While TCP/IP protocols do not precisely fit into the more general OSI model, the functions performed by each OSI layer correspond to the functions of each part of the TCP/IP protocol suite and provide a good framework for visualising the respective relationships between the various protocols.

3.3 Local area networks

LANs physically or logically connect together multiple workstations, terminals, and peripheral devices via a single cable or shared medium (Pretty 1992). Through the 1980s alone, over 100 000 LANs were set up in offices and laboratories around the world to link workstations to printers, share files and send electronic mail (Wittie 1991). LAN usage in the general computing community grew at a rate of 80 per cent per year between 1985 and 1991 (Pretty 1992), with networks extending into schools, libraries, laboratories, and offices around the developed world using telephone lines, optical fibres, and satellite links.

Several accepted and standardised types of LAN technology now share the market, including IEEE 802.3 (CSMA/CD or Ethernet), IEEE 802.4 (Token Bus), IEEE 802.5 (Token Ring), and ANSI FDDI (fibre distributed data interface). To date, the Ethernet (Figure 2) and Token Ring technologies have dominated the market.

3.4 Metropolitan and wide area networks

Until the early 1990s, most wide-area networking and LAN-interconnect services did not deliver the performance required either to move large GIS data files quickly or to offer real-time access to remote users over very long distances (Craig et al 1991). Where data delivery was an issue, magnetic tapes and diskettes were viewed as the media of choice for the distribution or exchange of data among GIS data producers and users (Newton et al 1992).

Modern metropolitan area networks (MANs) and WANs employ different protocols and technology in order to offer speeds comparable to LANs while operating over greater distances. By 1992, fibre-based packet switching services like SMDS in North America and FASTPAC in Australia began providing high-speed (34–45 Mbit/sec) links between LANs across and between major cities. Today, dedicated FDDI networks connect users within limited areas at rates up to 100 Mbit/sec, and higher-speed asynchronous transfer mode (ATM) services promise to support a wide range of real-time multimedia applications across long distances.

Results of the performance testing experiments described by Coleman (1994) indicated that future users of these broadband services would be able to access remote disks, processors, and output devices at near-LAN levels of performance. Today, such services have already influenced the nature of data delivery and software support practices to a small but growing group of users in the GIS community (e.g. Annitto and Patterson 1995; Streb 1995).

3.5 The Internet

Especially since 1993, it has been the connectionless services built atop the Internet which have come to define the current paradigms for mass market network usage. The Internet is a collection of interconnected campus, local, provincial, national, and corporate networks in more than 150 countries. Originally designed to connect researchers in university, government, and industrial defence establishments, the Internet is now estimated to

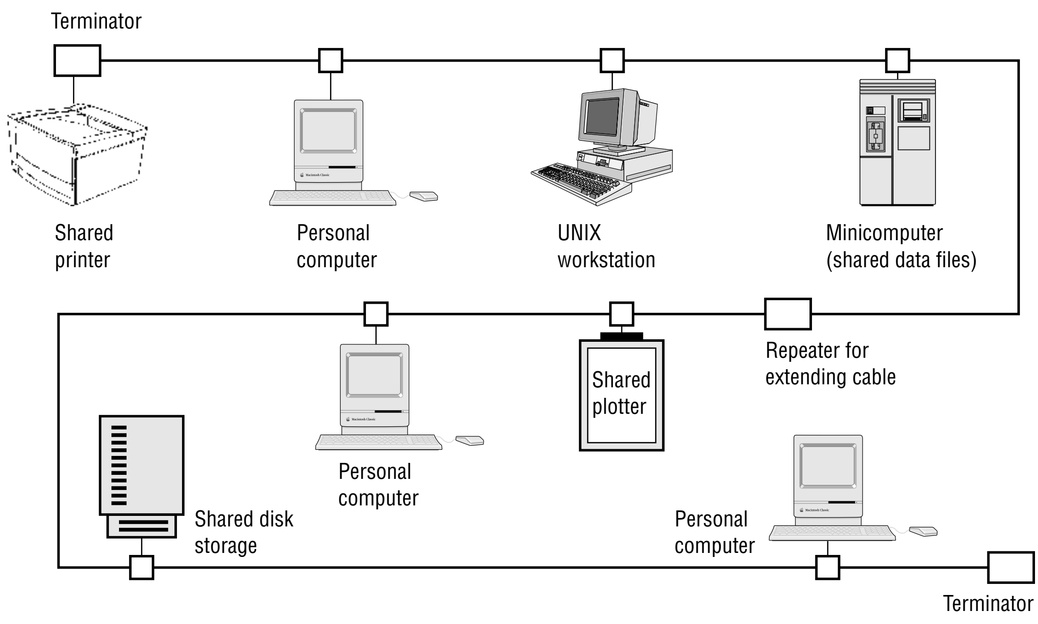

Fig 2. Simplified Ethernet local area network.

service more than 20 million users through a loose collection of more than 5000 registered networks (Reinhardt 1994; Thoen 1995). While electronic mail, file transfer, and remote login services originally accounted for the majority of Internet traffic, it has been ease-of-use and the multimedia capabilities of the World Wide Web (WWW) that have attracted most of the mass-market attention and resultant increases in usage. Over 145 000 Web sites have appeared since its introduction in 1993, and this number is continuing to rise monthly (Webcrawler Survey 1996).

3.6 Wireless, field-based systems

Recent advances in wireless communications, notebook computing, and the integration of GIS and global positioning system (GPS) technology have extended the application of traditional GIS databases into field operations (Lange and Gilbert, Chapter 33). Using normal telephone lines and, more recently, cellular telephone modems, field users tying into enterprise-wide networks may now access information resources previously available only from within the branch office or headquarters. New implementations are extending functionality beyond traditional records management and modelling roles into such areas as property inspection, field updating, facilities maintenance, customer service, and emergency response (Elliot 1994).

At least one industry source predicts the mobile computer/mobile field automation market will grow from $30 billion in 1995 to $80 billion by the year 2001 (FieldWorks Inc. 1996). FieldWorks estimates the 1995 total market for 'ruggedised' portable computers at $500 million with an annual growth rate of 35 per cent.

This predicted growth in the ruggedised portable market is based on current trends within large service firms and government organisations to automate the data collection and communication capabilities of field personnel through greater use of reliable mobile computers. As growth in notebook over desktop sales suggests, purchasers want systems which combine the functionality of desktops with the ability to fulfil specific applications beyond word processing and spreadsheet manipulation. The integration and use of GIS software and GPS technology within mobile computers for field data collection and database updating is one example of such extended functionality.

4 APPLYING THESE NEW MODELS

4.1 Enterprise computing implementations

Large utilities and municipalities were among the first to identify the operational requirements to integrate smaller GIS- and facilities management-related databases on stand-alone systems with larger corporate databases residing on mainframes (Popko 1988). As LANs became more widespread, enterprise computer systems evolved from host-based (or 'single-tier') to two-tier environments with networked PCs and workstations in a client-server architecture replacing connected terminals. This modular approach is attractive to many organisations since it allows them to take quicker advantage of price/performance improvements on modular hardware components and the other advantages of network computing described earlier (Mimno 1996).

While advantageous in many respects, users have found that these two-tier environments may not necessarily retain some of the advantages found in centralised host-based implementations, notably availability, expandability, and reliability of service (Strand 1995). As a result, newer three-tier environments have been developed to facilitate the placement of applications and data in locations in order to optimise these three factors (see Figure 3).

Existing mainframes at Tier 3 – possibly connected over long distances via MANs and WANs – are used to support legacy applications and provide access to large databases. To implement a client/server computing strategy, additional layers of hardware and software are added as a front end to the host computer. These layers consist of shared servers at Tier 2 (interconnected by high-speed LANs), and LAN-based PCs and laptop computers at Tier 1 (Mimno 1996). GIS and desktop mapping applications are included at this first tier, although the data may reside on either the Tier 2 servers or even the Tier 3 hosts.

This enterprise information architecture is highly flexible and can be modified easily to accommodate changes in business requirements. As more large customers adopt this approach, commercial GIS software firms and mainstream DBMS vendors alike are modifying their offerings in response to customer demands. In particular, spatial data management models and processes are becoming a more intrinsic component of an organisation's larger information management architecture. ESRI's Spatial Database Engine and MapInfo's 1996

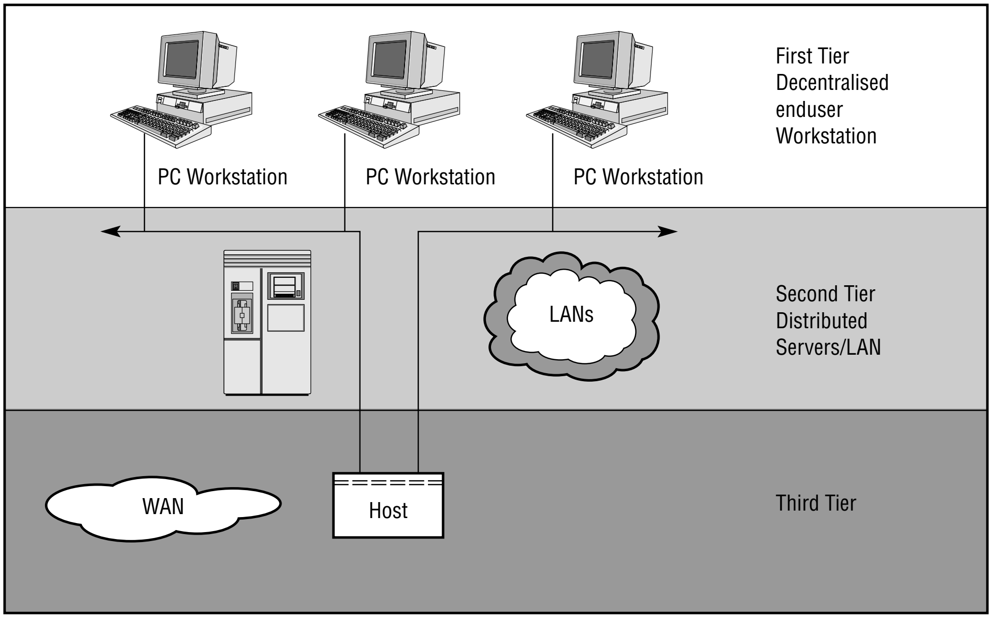

Fig 3. Three-tier enterprise computing environments.
Source: Strand 1995

acquisition of the Spatialware server technology from Unisys Corporation are two examples of how GIS vendors are providing even tighter high-performance links into mainstream relational DBMS packages like ORACLE and others, while based on *de facto* industry standards for object linking and embedding (OLE), structured query language (SQL), third- and fourth-generation languages, and ODBC. Table 1 contains a summary of recent industry offerings in this regard.

At the same time, the international database standards community (through ISO/IEC/JTC1/SC21 – WG3) is extending the proposed SQL MM ('MultiMedia') to accommodate spatial relationships, indexing arrangements, and operators typically found in GIS packages. Mainstream commercial database vendors are already incorporating some such enhancements into their own RDBMS offerings, with ORACLE's SDO and Informix/Illustra's Spatial DataBlades offerings being two such examples. These developments, combined with the recent formal attention being paid to interoperability between software packages (Sondheim et al, Chapter 24), suggest that future enterprise implementations may rely less on 'full-solution' packages maintained by a single vendor or consortium of vendors. Rather, large institutional users may opt for a series of smaller, lower-cost software components or applications possessing limited functionality on their own, but which seamlessly interact with other database, word processing or advanced modelling applications which may be present in those organisations (see Elshaw Thrall and Thrall, Chapter 23).

Table 1 Tighter linkages between GIS and RDBMS: 'Spatial Middleware' (from Costello 1996).

Vendor	Client Software product	Middleware Product
Vision	AutoCAD	Vision Server
Intergraph Corporation	MicroStation	Jupiter
ESRI	ArcView MapObjects	Spatial Database Engine (SDE)
MapInfo Corporation	MapInfo	SpatialWare

To make optimum use of this flexibility, application developers must have the ability to allocate data and processes anywhere in the multi-tier enterprise information architecture – and also reallocate them as operational requirements change. At the time of writing, concerns have been expressed that first-generation client/server application development tools like Visual Basic, PowerBuilder, and others may not have the capabilities and performance required to support more complex, enterprise-wide applications adequately (Mimno 1996). While it is likely that newer versions of these and similar tools will address many of these concerns, they must still evolve as the needs and practices of different user groups in these large organisations change over time.

4.2 Data access and delivery

Dozens of major land information and resource inventory programs around the world have demonstrated on-line organisation and electronic distribution of government imagery, mapping, and related land information across proprietary networks. Early examples included Land Information Alberta (McKay 1994) and the Manitoba Land Information Utility (Oswald 1994) projects in Canada, the early commercial ImageNet service in the United States, and the Land Ownership and Tenure System (LOTS) in South Australia (Sedunary 1988). Numerous other early examples may also be found elsewhere in these countries and in Europe.

Originally, such projects offered connection-oriented access to secure databases utilising either direct LAN-interconnect services or high-speed modems. Especially since late 1993, however, the Internet-based WWW has emerged as an alternative means of accessing, viewing, and distributing spatial information. Used in combination with modern Web 'browsing' software packages like Netscape and Microsoft's Internet Explorer, the WWW is emerging as a mainstream tool for

- the distribution of public-domain spatial information;
- online ordering of commercial datasets;
- indexing and cataloguing of related spatial datasets available off-line (Dawe 1996).

An analysis of more than 25 home pages classified WWW usage in the GIS community into four overlapping categories (Coleman and McLaughlin 1997). Setting aside those sites focusing on product or program advertising, the remaining three categories included:

1. *Data distribution.* At such sites, users may search for specific spatial information features or datasets based on either keyword searches or Boolean queries of existing database fields (e.g. Crossley 1994; Nebert 1994; Pleuwe 1994). In some cases, the user may obtain only pointers to datasets stored off-line (e.g. Marmie 1995). In more sophisticated implementations, the user may retrieve image, map graphics, or attribute information covering a given area of interest (e.g. Conquest and Speer 1996).
2. *Custom map creation and display.* These sites represent the results of special-purpose development projects aimed at the composition, display, and downloading of user-defined custom map products. They may or may not involve the use of a GIS package running in the background. Currently limited in terms of the extent and variety of data coverage, they nevertheless represent an exciting development in providing 'just-in-time' mapping to the general public (Pleuwe 1997).
3. *GIS/WWW integration*. These sites represent the results of integrating front-end query capabilities (supported through standard Internet and WWW interfaces and protocols) with the capabilities of commercial DBMS and GIS software packages residing in the background. User-defined queries are translated into corresponding SQL commands and passed to the 'back-end' GIS database for handling. The resulting response is passed back through the gateway to the user and, in the case of map-based responses, the resulting map is translated into a graphics or bit-map format suitable for fast transmission and viewing across the Internet. User-developed examples of such GIS/WWW integration are documented by Conquest and Speer (1996), Nebert (1996) and Liederkerke et al (1995), among others. Software vendors offering WWW interfaces to spatial data servers or map-based spatial data browsers include Autodesk, ESRI, Genasys, MapInfo, and Universal Systems.

The lines distinguishing these categories are fading quickly. Various electronic telephone and business directories offered in Australia *(http://www.whitepages.com.au)*, Chile *(http://www.chilnet.cl/index.htm)*, and the USA *(http://www.bigbook.com)* now allow users both to obtain contact information and to find the location of selected individuals or businesses on custom-generated electronic street maps (see Elshaw Thrall and Thrall, Chapter 23). Other creative early examples of how these systems may appeal to general users include:

1. The 'Tripquest' service (Figure 4) *(http://www.mapquest.com)*, which enables users to specify a starting point and destination of a

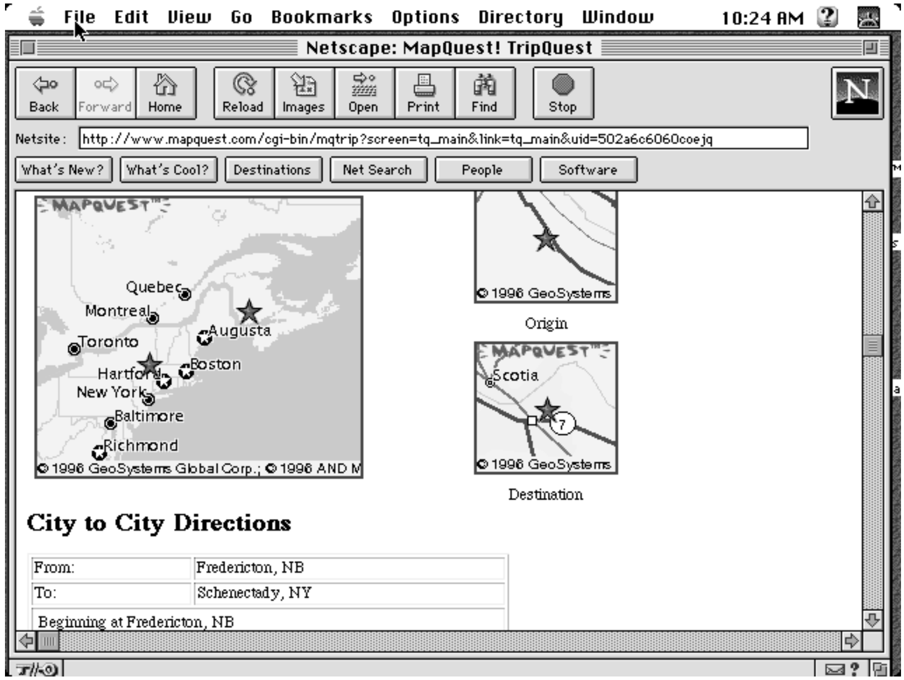

Fig 4. 'Tripquest' service.

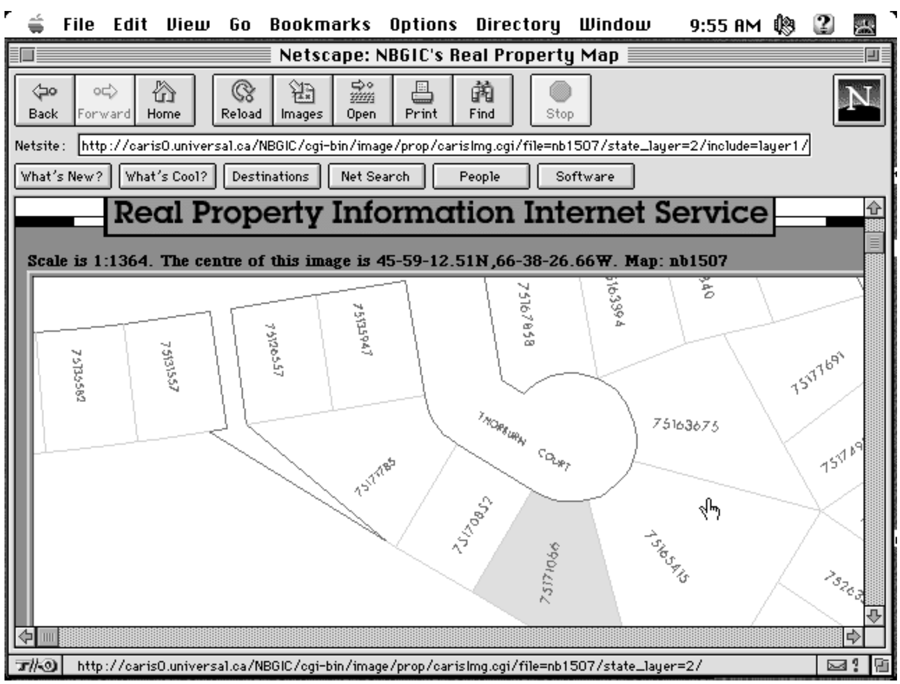

Fig 5. New Brunswick Geographic Information Corporation Property Information Service.

road trip, then receive detailed driving instructions based on a road network database contained at the server site.

2 The LandData BC site *(http://www.landdata.gov.bc.ca)*, which offers on-line access to a wide variety of mapping and land information products from different provincial government departments in British Columbia, Canada.

3 The New Brunswick Geographic Information Corporation Property Information Server (Figure 5), which allows NBGIC subscribers to access and download property ownership, registration, assessment, and parcel-centred mapping information covering any property in New Brunswick, Canada via keyword search, database query or map-based index (Arseneau et al 1996).

The advent of Java – a portable, object-oriented Internet language developed by Sun Microsystems Inc. – promises to remove many of the constraints inherent in early WWW protocols and further extend the capabilities of Web-based data browsing systems (Strand 1995). By moving much of the requisite display, processing, and analysis functionality to the client end of the Internet connection, performance delays due to server load and Internet bandwidth limitations may be greatly reduced (Figures 6 and 7).

Early examples of Java-based applications in GIS have tended to focus on limited tasks, including improved client-based spatial data browsing (Mapguide 1996) and network analysis engines (Fetterer et al 1996). More recent research efforts (Choo 1996) are aimed at using Java to provide an even wider suite of GIS display and modelling capabilities across the Internet. With a number of vendors and standards organisations (e.g. the Open GIS Consortium) now including it as an important component of an overall network-based GIS architecture, Java-based technology may eventually be one of the keys to developing more open systems of distributed processing of geographical data.

At the time of writing, more extensive and well-maintained lists of WWW sites falling into all these categories may be found at the 'Metadata and WWW Mapping' home page maintained by Katz *(http://www.blm.gov/gis/nsdi.html)* and the 'GeoWeb' site *(http://wings.buffalo.edu/geoweb)* maintained by Pleuwe.

4.3 Collaborative production and group decision-making

Hardware, software, and procedures to support computer-supported cooperative work (CSCW) have been discussed and compared for more than 30 years (Englebart 1963). In the corporate world, shared access to corporate resources and innovative new approaches to collaborative production using groupware tools like Lotus Notes (Coleman 1995) and, more recently, the WWW (Pilon et al 1996) are now being investigated.

When considered in combination with emerging groupware products, the developments in GIS and broadband communications mentioned earlier may also offer a new approach to both collaborative group decision-making and digital map production – enabling some processes to be conducted concurrently rather than sequentially (for examples see Churcher and Churcher 1996; Coleman and Brooks 1995; Faber et al 1995; Karacapilidis et al 1995).

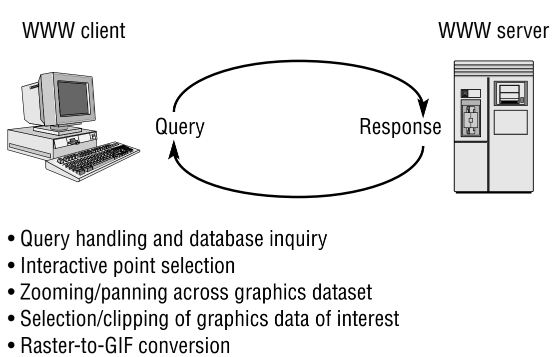

Fig 6. Internet spatial data request handling using HTTP.

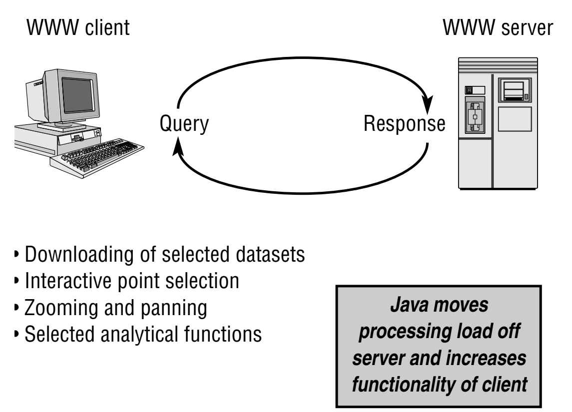

Fig 7. Internet spatial data request handling using Java language.

Network-based collaborative applications depend on more than just electronic mail. Specifically, they are predicated on the presence of a shared 'database' or collection of files, and rules which permit the definition of group members' roles, task status reporting and tracking, and gateways to electronic mail and other sources of data. Such systems should permit the organisation of correspondence, comments, reports etc. associated with a project or product and should support the management of multiple versions of objects (e.g. images, vector-based charts, video, and sound). Finally, a limited number of applications may require two or more remote users to be able to view the same file simultaneously, modify or add comments to specific entities in the same file where necessary, and communicate via voice, video, or e-mail while making these changes.

A number of organisations are already prototyping the WWW as the medium for collaborative decision-making (Gordon et al 1996), enhanced design, and concurrent engineering applications. Moreover, recent extensions to WWW viewers like Netscape are now introducing much of the communication and threaded-discussion, file transfer, data management and simultaneous data viewing functionality required in these applications (Forrester Research Inc. 1996; Ziegler 1995).

However, while some major corporations have made significant productivity gains through adoption of early groupware packages. The costs and cultural changes involved in such adoptions are likely to blunt the near-term impact of CSCW in the marketplace. Technically, vigorous adoption of collaborative technology will depend on:

- acceptance of the Internet or corporate intranets as the preferred medium for communications and information delivery;
- acceptance of common standards for threaded discussions and calendar management;
- subsequent development of a wider variety of inexpensive and compatible components or 'plug-in' tools capable of operating atop the standards extended from or similar to Netscape's ONE and Microsoft's COM (Forrester Research Inc. 1996).

Culturally, successful groupware implementations are likely to be shaped more by prevailing organisational cultures and constraints than by the technology itself. Since organisations work in different ways, it is best to let them customise their own groupware applications using modular toolkits rather than expecting any single shrink-wrapped package or technology to serve their needs adequately (Schrage 1995).

5 SYSTEM PERFORMANCE IN NETWORKED ENVIRONMENTS

Forthcoming generations of system architecture and broadband telecommunication technology promise to change fundamentally the way many organisations manage, transfer, and utilise their spatial data. However, before adopting such technology, both suppliers and potential customers require a clear understanding of potential network usage and the performance, capacity, and cost trade-offs involved in existing and emerging applications. A study carried out for the Government of Canada (IDON 1990) suggested that designing and building cost-effective GIS networks within and between federal government organisations would require defensible answers to six important groups of questions. These questions concerned the kind of information to be moved, the data volumes involved, response-time requirements of the users, the distance the data were being transferred, the frequency or regularity of data transmissions and, perhaps most important, the available funds which could be devoted to such operations.

Given the ability of today's WWW servers to track the volume and nature of user transactions at and between given sites, it is now easier to address such questions than it was seven years ago. Further, given the ubiquity of the Internet today, the justification of such a network may appear a moot exercise. However, the continuing requirement for 'hard numbers' reinforces the need to determine defensibly the performance of a particular application or group of applications across a network under known conditions, in order to predict the performance of the same application(s) under other conditions.

Numerous authors have suggested the fundamental importance of performance determination within the overall framework of the system life cycle and structured system and database design processes (e.g. Ferrari et al 1983; Jain 1991). However, while performance analysis has formed an important component of many GIS selection processes since the early 1980s, test procedures and results have largely remained unavailable because of the constraints of commercial confidentiality and competitive pressures.

Researchers have been developing more systematic and rigorous approaches to the determination of GIS performance on stand-alone configurations since the mid 1980s (e.g. Goodchild and Rizzo 1986; Hawke 1991; Marble and Sen 1986; Wagner 1991). Investigations at the University of Edinburgh (Gittings et al 1993) and the University of Tasmania (Coleman 1994) have extended these studies into client/server environments across LANs and WANs. The latter two references in particular provide extensive reviews of the literature related to GIS performance testing.

A key limitation of GIS performance testing research to date has been its lack of extendibility. Absolute performance figures are tied closely to the hardware configurations employed. Further, even under controlled network conditions, it is extremely problematic to model end-to-end GIS performance and to predict how response times will be affected by changes in server load and network data traffic.

As an increasing number of applications take advantage of the Internet – which, in itself, can send data packets across randomly changing combinations of high- and low-speed network connections with ever growing traffic loads – it will be increasingly difficult to make any claims (apart from criticisms) concerning application performance in a given setting. Further, if events over the past 15 years are any indication, the storage, memory, and processing requirements of application software and operating systems have a tendency to expand to surpass eventually the capabilities of a given generation of computers and networks alike.

That being said, the following advances in technology and services will – along with many others – help improve GIS application performance over the next five years:

- the increasing optimisation of high-performance disk arrays and parallel processors for use in networked GIS environments (Sloan 1996);
- the proliferation of fibre-based broadband communication services in both the workplace and the home;
- the introduction of higher-speed satellite and cellular data communication links into areas not served by normal services;
- the introduction of higher-speed Internet services (e.g. the proposed 'Internet II') and value-added service offerings which improve performance by 'detouring' subscriber traffic around high-volume links and nodes;
- the accelerating introduction of special plug-ins, Java 'applets' and other modular software components which balance client and server processor loads and further optimise data traffic between workstations on a network;
- (perhaps most importantly) improvements in the collections of underlying algorithms and procedures defining the application software packages themselves, removing many existing bottlenecks and optimising older code to run on newer systems.

References

Annitto R, Patterson B 1994 A new paradigm for GIS data communications. *Journal of the Urban and Regional Information Systems Association* 7: 64–7

Arseneau B, Kearney A, Quek S, Coleman D 1997 Internet access to real property information. *Proceedings GIS 97 Conference, Vancouver.* GIS World Inc.

Champine G A, Coop R D, Heinselman R C 1980 *Distributed computer systems*. Amsterdam, Elsevier Science

Choo Y K 1996 'Interactive distributed geographical information systems (IDGIS)'. Unpublished Master's Research Project, Department of Rangeland Ecology and Management, Texas A&M University. *http://starr-www.tamu.edu/choo/idgis/intro.html*

Chorafas D 1980 *Computer networks for distributed information systems*. New York, Petrocelli Books

Churcher M, Churcher C 1996 GROUPARC: a collaborative approach to GIS. In Pascoe R T (ed.) *Proceedings, Eighth Annual Colloquium of the Spatial Information Research Centre*. Dunedin, University of Otago: 156–63

Coleman D J 1994 'Geographic information systems performance in a broadband communications environment'. Unpublished PhD dissertation, Department of Surveying and Spatial Information Science, University of Tasmania

Coleman D J 1995 An overview of groupware. In Coleman D J, Khanna R (eds) *Groupware technology and applications.* Englewood Cliffs, Prentice-Hall

Coleman D J, Brooks R 1995 Applying collaborative production approaches to GIS data collection and electronic chart production. Santa Barbara, NCGIA and see *http://www.ncgia.ucsb. edu/research/i17/i17papers.html*

Coleman D J, McLaughlin J D 1997a Information access and usage in a spatial information marketplace. *Journal of Urban and Regional Information System*s 9(1)

Conquest J, Speer E 1996 Disseminating ARC/INFO dataset documentation in a distributed computing environment. *Proceedings 1996 ESRI User Conference.* Redlands, ESRI. *http://www.esri.com/resources/userconf/proc96/TO200/PAP165/P165.HTM*

Costello B 1996 Trends in client-server architecture in GIS environments. Slide from invited presentation at 1996 Geomatics Atlantic Conference, Fredericton, New Brunswick, Canada

Craig W J, Tessar P, Ali Khan N 1991 Sharing graphic data files in an open system environment. *Journal of the Urban and Regional Information Systems Association* 3: 20–32

Crossley D 1994 WAIS through the Web: discovering environmental information. Paper presented Second International WWW Conference Mosaic and the Web, Chicago, 17–20 October

Dawe P 1996 'An investigation of the Internet for spatial data distribution'. Unpublished Master of Engineering report, Department of Geodesy and Geomatics Engineering, University of New Brunswick, Fredericton, Canada

Elliot W 1994 Moving AM/FM to the field: technology, trends and success stories. *Proceedings of 1994 AM/FM International Conference*: 845–57

Englebart D 1963 A conceptual framework for the augmentation of man's intellect. In Howerton P (ed.) *Vistas in information handling*. Washington DC, Spartan Books 1: 1–29

Ezigbalike I, Cooper R, McLaughlin J 1987 A query management strategy for land information network implementation. *Proceedings of 1987 Annual Meeting of the Urban and Regional Information Systems Association (URISA)* 3: 105–12

Faber B G, Wallace W, Cuthbertson J 1995 Advances in collaborative GIS for land-resource negotiation. *Proceedings of GIS 95 Ninth Annual Symposium on Geographic Information Systems*: 183–9

Ferrari D, Serazzi G, Zeigner A 1983 *Measurement and tuning of computer systems*. Englewood Cliffs, Prentice-Hall

Fetterer A, Goyal B, Agarwal N 1996 Background: Java interactive routing for Minneapolis: explanation of WebRoute project. Spatial Database Laboratory, Department of Computer Science, University of Minnesota. *http://www.ggrweb.com/geojava/*

FieldWorks Inc. 1996 Eden Prairie. *http://www.canaska.com/fw.htm*

Forrester Research Inc 1996 Teams on the Internet. *The Forrester report* 7(6): 2–8. Cambridge (USA), Forrester Research Inc. *http://www.forrester.com*

Gittings B M, Sloan T M, Healey R G, Dowers S, Waugh T C 1993 Meeting expectations: a view of GIS performance issues. In Mather P M (ed) *Geographical information handling*. Chichester and New York, John Wiley & Sons

Goodchild M F, Rizzo B 1986 Performance evaluation and workload estimation for geographic information systems. *Proceedings, Second International Symposium on Spatial Data Handling*: 497–509

Gordon T, Karacapilidis N, Voss H 1996 Zeno: a mediation system for spatial planning. *Proceedings of ERCIM workshop on CSCW and the Web, Sankt Augustin, Germany, 7–9 February. http://orgwis.gmd.de/W4G/proceedings/zeno.html*

Hawke D 1991 Characterising the performance of geographic information systems. *Proceedings of Third Annual Colloquium of the Spatial Information Research Centre, Dunedin*. University of Otago

IDON Corporation 1990 *Final report: a federal geographic information systems (GIS) network requirements study*. Contract report prepared for the Data Communications and Networking Subcommittee, Interagency Committee on Geomatics, Government of Canada

Ingoldsby T R 1991 Transparent access to geographically related data from heterogeneous networked systems. *Proceedings of 1991 Annual Conference of the Urban and Regional Information Systems Association* 3: 14–24

Jain R 1991 *The art of computer systems performance analysis: techniques for experimental design, measurement, simulation, and modeling*. New York, John Wiley & Sons Inc.

Karacapilidis N, Papadias D, Egenhofer M 1995 Collaborative spatial decision making with qualitative constraints. *Proceedings of Third International ACM Workshop on Advances in Geographic Information Systems*: 53–9

Katz R 1991 *High performance network and channel-based storage*. Sequoia 2000 Technical Report 91/2. Computer Science Division, Department of Electrical Engineering and Computer Sciences, University of California, Berkeley

Laurini R 1994 Sharing geographic information in distributed databases. *Proceedings of 1994 Annual Conference of the Urban and Regional Information Systems Association (URISA)*: 441–54

Liederkerke M van, Jones A, Graziani G 1996 The European Tracer Experiment system: where GIS and WWW meet. *Proceedings of 1996 ESRI User Conference.* Redlands, ESRI

MapGuide (1996) Mill Valley, Autodesk Inc. *http://www.mapguide.com*

Marble D F, Sen L 1986 The development of standardized benchmarks for spatial database systems. *Proceedings, Second International Symposium on Spatial Data Handling*: 488–96

Marmie A 1995 Promoting and providing GIS data via the Internet. *Proceedings, 1995 ESRI User Conference*. Redlands, ESRI. *http://www.esri.com/resources/userconf/proc95/to250/p2101.html*

McKay L 1994 Data brokering in land information – Land Information Alberta. *Proceedings, 22nd Annual Conference of AURISA* 1: 171–82

Miller A 1990 GIS in a networking environment. *Proceedings of GIS 90, International Symposium on Geographic Information Systems*: 255–62

Mimno P 1996 *Building enterprise-class client/server applications*. Lisle, Dynasty Technologies, Inc. *http://www.dynasty.com/product/mimno_wp.htm*

Nebert D D 1994 Serving digital information through the WWW and Wide Area Information Server technology. Paper presented at Second International WWW Conference Mosaic and the Web, Chicago

Nebert D D 1996 Supporting search for spatial data on the Internet: what it means to be a clearinghouse node. *Proceedings of 1996 ESRI User Conference, Palm Springs* Redlands, ESRI. *http://www.esri.com/resources/userconf/proc96/TO100/PAP096/P96.htm*

Newton P W, Zwart P R, Cavill M 1992 Inhibitors and facilitators in high-speed networking of spatial information systems. In Newton P W, Zwart P R, Cavill M E (eds) *Networking spatial information systems*. London, Belhaven

Oswald R 1994 Manitoba Land Related Information System land information utility. *Proceedings, 1994 Conference of the Urban and Regional Information Systems Association (URISA)*: 179–92

Pilon D, Whalen T, Palmer C 1996 *Communications toolkit*. Technical Paper. Network Services and Interface Design Laboratory, Communication Research Centre, Industry Canada. *http://debra.dgbt.doc.ca/~daniel/cscw.text.html*

Pleuwe B 1994 The GeoWeb project: using WAIS and the World Wide Web to aid location of distributed datasets. Paper presented at WWW 94 Conference, Chicago, 17–20 October. *http://wings.buffaloedu/~plewe/paperwww.html*

Pleuwe B 1997 *GIS online: information retrieval, mapping, and the Internet*. Santa Fe, Onward Press

Popko E 1988 *An enterprise implementation of AM/FM*. Internal paper. Geo-Facilities Information Systems Application Center, IBM Corporation, Houston, Texas

Pretty R W 1992 LANs, MANs, and WANs: introductory tutorial on computer networking. In Newton P W, Zwart P R, Cavill M E (eds) *Networking spatial information systems.* London, Belhaven Press

Reinhardt A 1994 Building the data highway. *Byte Magazine* 19: 46–74

Schrage M 1995 Groupware requires much more than bandwidth. *Business Communications Review*: 35–8

Sedunary M E 1984 LOTS and the nodal approach to a total land information system. In Hamilton A C, McLaughlin J D (eds) *Proceedings of 'The Decision Maker and Land Information Systems' – FIG Commission III International Symposium*

Sloan T M 1996 'The impact of parallel computing on the performance of geographical information systems'. Unpublished MPhil dissertation, Department of Geography, University of Edinburgh

Strand E J 1995 GIS thrives in three-tier enterprise environments. *GIS World* 8: 38–40

Streb D 1995 Using cable television to distribute GIS access. *Proceedings, 1995 ESRI User Conference*. Redlands, ESRI

Thoen B 1995 Web GIS: toy or tool? *GIS World* (October) *http://www.csn.net/~bthoen/webgis.html*

Voelcker J 1986 Helping computers communicate. *IEEE Spectrum* (March): 61–79

Wagner D F 1991 'Development and proof-of-concept of a comprehensive performance evaluation methodology for geographic information systems'. Unpublished PhD dissertation, Department of Geography, Ohio State University, Columbus

Webcrawler Survey 1996 *World Wide Web size and growth*. Global Network Navigator, Inc. *http://webcrawler.com/WebCrawler/Facts/size.html*

Webster C J 1988 Disaggregated GIS architecture. Lessons from recent developments in multi-site database management systems. *International Journal of Geographical Information Systems* 2: 67–80

Wittie L D 1991 Computer networks and distributed systems. *IEEE Computer*: 67–76

Ziegler B 1995 Internet software poses threat to Notes. *The Wall Street Journal* (7 November)