Grid Computing for Real Time Distributed Collaborative Geoprocessing

Yi Shi[1], Ashton Shortridge[2] and Jon Bartholic[3]

Institute of Water Research, 115 Manly Miles Building, Michigan State University, East Lansing, MI, 48823 USA, Phone: 1-517-432-5585, shiyi1@msu.edu[1]
Department of Geography, 314 Natural Science Building, Michigan State University, East Lansing, MI, 48824 USA[2]
Institute of Water Research, 115 Manly Miles Building, Michigan State University, East Lansing, MI, 48823 USA[3]

Abstract

Grid computing has emerged as an important new field in the distributed computing arena. It focuses on intensive resource sharing, innovative applications, and, in some cases, high-performance orientation. In this paper we describe Grid computing concepts and technologies and establish compelling reasons for why they may prove useful for important geospatial research and application issues. Then we review high performance parallel geoprocessing in general. Finally, we propose an example GIS application scenario that will employ underlying Grid computing technology intensively. We report results on a prototype employing a universal soil loss equation implemented on an array of supercomputer processors.
Keywords: grid computing, parallel processing, virtual organisation, GIS, watershed modelling

1 Introduction

The Internet's exponential growth has greatly increased the significance of distributed computing technologies for information systems generally. This development did not go unnoticed by the geographic information science (GIScience) community, which early focused on identified distributed, Internet-based computing as a key research topic (UCGIS, 1998). Indeed, GIScience has long been engaged with distributed computing. Such engagement is natural:

- Computationally intensive spatial applications often tax the capacity of single machines.

- The GIS user base may itself be spatially distributed, with a need to share data, models, and results.
- Spatial operations and decisions are frequently best made in the field, requiring flexible access to a remote database.

An early representative application is Sequoia 2000, a project that linked global climate change researchers throughout the University of California system via an integrated computer environment (Stonebreaker, 1995). The project involved dozens of earth scientists and computer scientists, as well as industry representatives, working together to create a system capable of managing very large environmental databases and serving data and visualisations rapidly across a dedicated wide area network.

Experience from Sequoia 2000 and similar projects indicated the presence of both opportunities for and impediments to the employment of distributed computing. Significantly, impediments are both technological and non-technological in nature. Technological challenges include:

- integrating data from different sources, stored in different databases
- transmitting very large quantities of information over the network
- developing flexible interfaces for different users

Organisational challenges include:

- combining members of very different research communities - computer science and earth / spatial science
- communicating research objectives and requirements clearly
- orchestrating linked but independent environmental models developed in different laboratories

A variety of GIScience initiatives are intended to improve the utility of distributed computing for geospatial applications (see e.g. OGIS, 2002; Hecht, 2001; UCGIS, 1998). Generally they have focused on technical aspects, although the organisational challenges of large projects have long been acknowledged.

This paper reports on two related developments that may remove many of these impediments and greatly facilitate distributed geoprocessing. The first is the concept of the Virtual Organisation (VO). This powerful model for accomplishing group objectives in a dynamic manner independent of the geographic location of group members has gained currency in business management since the mid-1990's (Palmer & Speier, 1997). The second is the technological realisation of the Grid, a system of powerful processors and massive storage that is seamlessly knit together via the Internet by protocols and low-level programming interfaces. Grid-based technologies enable extensive Internet-based infrastructure for the VO. Section 2 describes the virtual organisation concept and Grid-based technologies that can support it. Section 3 indicates primary functions of large GIS applications and geoprocessing applications that appear to be amenable to the VO paradigm. Section 4 describes a prototypical geospatial virtual organisation and reports on the utility of Grid-based infrastructure to execute a distributed geospatial application upon it. We conclude in Section 5 with a discussion on the prospects for Grid-based geocomputation.

2 Virtual Organisation Paradigm and the Grid

A virtual organisation is a network linking geographically dispersed individuals or agencies with partially overlapping objectives in order to pool complementary core competencies (Jägers et al., 1998). The VO is a framework for accomplishing shared objectives in a flexible and efficient manner. A typology proposed by Shao et al. (1998) identifies four characteristic features of a virtual organisation:

- Connectivity, which establishes linkages where none previously existed.
- Purpose, which creates a common incentive for agents in the virtual organisation.
- Technology, which enables connectivity.
- Boundary, which separates those who are part of the VO from those who are not.

The role of trust is critical, since decisions must be made without the same degree of face-to-face interaction in more traditional settings, and because data, strategic decisions, and objectives must be shared with semi-independent participants in the VO. Because of this greater autonomy, VO participants tend to be more equal, with greater independence than is typical in a hierarchical organisation.

The VO paradigm has been highly useful for certain business activities, particularly those that are technologically oriented and require flexibility and rapid adaptation (Jägers et al., 1998). It has also been adapted in academic settings at large and small scales. At large scales, it is common for multidisciplinary, research teams with members in several academic institutions to collaborate on projects. At small scales, multiple-authored papers like this one are typically developed in concert, with portions of work done by different authors in physically different places. In either case, electronic transmission of documents and data, as well as ideas, keeps the work knit together. The network enables the rapid flow of communication, which is critical for the virtual organisation.

The Grid computing concept is intended to enable co-ordinated resource sharing and problem solving in dynamic, multi-organisational virtual organisations. The basic idea is to provide computation power to everyone who can access it, just like electric power Grid providing electricity. The sharing in this computational Grid refers to not only file exchanges but also direct access to a spectrum of computers, software, data, and other resources, as is required by a range of collaborative problem-solving and resource-brokering strategies emerging in industry, science, and engineering. This sharing is necessarily highly controlled, with resource providers and consumers defining clearly and carefully just what is shared, who is allowed to share, and the conditions under which sharing occur. A set of individual and/or institutions defined by such sharing rules form a Grid computing based virtual organisation (Foster et al., 2001).

The Grid computing based VO represents a new approach to collaborative computing and problem solving in computation- and data-rich environments. For example, the application service providers, storage service providers, computer cycle providers, university researchers and consultants engaged by a government agency might form such a VO to perform real time decision support for environmental management. Common technical concerns and requirements of different VOs include (Foster et al., 2001):

- Highly flexible sharing relationships, ranging from client-server to peer-to-peer and brokered.
- Complex and high levels of control over how shared resources are used, including fine-grained access control, delegation, and application of local and global policies.
- Sharing of varied resources, ranging from programs, files, and data to computers, sensors, and networks.
- Diverse usage modes, ranging from single user to multi-user and from performance sensitive to cost-sensitive and hence embracing issues of quality of service, scheduling, co-allocation, and accounting.

Current distributed computing technologies do not address these concerns and requirements. For example, current Internet technologies address communication and information exchange among computers but not the co-ordinated use of resources at multiple sites for computation. Business-to-business exchanges focus on information sharing (often via centralised servers). Enterprise distributed computing technologies - such as CORBA and Enterprise Java - focus on enabling resource sharing within a single organisation. Emerging “Internet computing” companies seek to harness idle computers on an international scale but, to date, support only highly centralised access to those resources. In summary, current technology either does not accommodate the range of resource types or does not provide the flexibility and control on sharing relationships needed to establish VOs (Foster et al., 2001).

Grid computing technologies were produced over the past five years to meet these concerns and requirements. They include protocols, services, and tools that address precisely the technological needs for scalable VOs. These technologies cover security solutions that support management of credentials and policies when computations span multiple organisations; resource management protocols and services that support secure remote access to computing and data resources and the co-allocation of multiple resources; information query protocols and services that provide configuration and status information about resources, organisations, and services; and data management services that locate and transport data sets between storage systems and applications (Foster & Kesselman, 1997).

We believe that the virtual organisation paradigm is broadly applicable to many large spatial management and modelling projects. This paradigm, coupled with Grid-enabling technology, concerns numerous aspects of distributed computing in such projects and changes the way that computers may be used to solve GIS problems. Perhaps more profoundly, it offers a way to reframe multi-organisational GIS and spatial modelling efforts as virtual organisations, enhancing the collaborative process. The following section considers the distributed computing applications of large GIS applications in detail.

3 Geoprocessing and the Grid

GIS is concerned with three sorts of distributed computing objectives: distribution of data, of software, and of processing. Each of these important subcategories is discussed below.

First, geographic information systems employed by public and private organisations collect, store, and maintain very large amounts of data about diverse themes (Longley et al., 2001). It is not unusual for different offices to be responsible for maintaining and updating individual layers in a regional GIS database. This division of labour has clear institutional advantages, but a significant challenge has always been maintaining the wide availability of the most current version of the entire database (Sugarbaker, 1999). While the sharing and distribution of spatial data has typically occurred only within the organisation, new Internet GIS technologies have enabled the serving of spatial information via the Internet to client organisations and the public (Coleman, 1999). A related issue is the on-line availability of massive quantities of spatial data by agencies whose purpose is distribution. Much United States base data, including digital elevation and land cover, may be downloaded at the time it is needed. Key Internet nodes serve as spatial data repositories, which can process queries and retrieval requests from users (Smith & Rhind, 1999). A key aspect of the Grid-based system for scientific research is the ability of software agents to identify and retrieve relevant data sources for processing. Significantly, the user needs not to be aware of the physical location of the data in the network. Instead, processing occurs at a (possibly remote) server that independently dispatches a search for necessary data (Foster, 2002). As a geoprocessing illustration, consider a hydrologist interested in modelling suspended sediment load in a river system. As input for the primary research, the hydrologist needs to execute an overland flow model for the watershed. The model is executed remotely, and the process itself searches for and obtains relevant elevation and landcover data, freeing the scientist from this potentially onerous search, as well as the download, pre-processing, and upload of large data quantities.

Second, a complex GIS operation may involve many autonomous steps. Traditionally, monolithic software systems have attempted to handle all possible routines. Drawbacks with the monolithic model for the user include overly complicated software design, large system elements that are never used, and inefficient processing due to the general-purpose design. An alternative approach is a distributed component GIS, in which individual spatial processing components are obtained from different sources from the network. These components are linked together to build the desired multi-step GIS procedure. Research on such software distribution falls generally under the rubric of distributed computing and interoperability, which is at least equally, concerned with open data structures. A current trend in GIS software industry reflecting this approach is to build Internet based distributed component GIS. Basic component functions are defined by OpenGIS consortium standards formed by different vendors. Actually, Grid technologies can be used as an infrastructure upon which truly distributed component GIS can be constructed. Grid concepts are designed to facilitate large data analysis tasks, with processing occurring on many different machines, and to encourage the development of scientific software portals, which provide an interface to specialised problem-solving applications on remote servers.

Third, many geoprocessing applications are extremely resource-intensive due to large volumes of data and intensive algorithms. For such applications, parallel processing enables the division of the work onto multiple machines, with potentially large reductions in time to complete the task. This goal has attracted the attention of researchers involved in diverse, high-load geoprocessing activities. Hunsaker et al. (1996) developed a parallel implementation for spatial data error simulation. The task required the production of many equi-probable realisations of a land cover data set using a complex spatial statistical model. Terrain modelling involves operations to generate terrain meshes from scattered points, detecting local surface properties, and calculating view sheds. Executing these operations on multiple processors is especially useful when grids are large or dense. Magillo & Puppo (1998) summarised progress towards implementing parallel terrain modelling. Spatial environmental modelling of snow conditions across the United States involves intensive use of resources. Research by El Haddi et al. (1996) sought to improve model performance by parallelizing the application. Spatial data mining that involves knowledge extraction from vast stores of spatial data could also benefit from parallelization. In all the cases described here, parallelization involves some combination of algorithm development and data model architecture design.

An interesting and relevant technical issue for intensive scientific computation concerns the different rates of technological change for computing power, storage, and network speed, typically measured in doubling time. These rates are roughly 18 months, 12 months, and 9 months, respectively (Foster, 2002). Foster notes that petabyte-magnitude data archives are in the planning stages based on the assumption that this trend will continue. However, processing capacity is not keeping pace. The implication is that local computer power will not be up to the task of processing this volume of data. Network speed, on the other hand, is increasing even more rapidly than storage, as is exemplified by the 40 Gbs/second rate for the Terragrid network (Benner, 2001). As a result, communication-intensive tasks will become relatively more efficient than processor-intensive tasks, encouraging network-distributed solutions for major, processor-intensive tasks.

In general, two separate problems exist for distributing geoprocessing. The first is the problem of parallelizing the relevant algorithms so that they may be employed on multiple machines. The second is concerned with identifying an optimal spatial partition of the database so that the load for each processor is roughly equivalent. Shekhar et al. (1996) and others (e.g. Hoel & Samet (1994)) developed parallel algorithms and spatial declustering techniques. Spatial operations that may be amenable to parallelization possess several qualities. The algorithms may have regularities for parallel processing so they can efficiently use idle cycles on different computers within a VO. Second, the problem may be one that can be spatially partitioned into smaller geographic units and distributed to different processors. Third, although spatial analysis may be intensive due to large data volume rather than complex processing, the results of the analysis may be much smaller than the data volume used for it. For intensive spatial analysis with parallel regularities in the algorithms, the data domain can be divided spatially. Then subsets of data are transmitted along with the (sub)algorithm to different machines. Once these smaller, more tractable subtasks are done, the outcomes are collected, reassembled spatially, and presented or stored.

A substantial challenge for the parallelization of geoprocessing applications has been implementing the low-level processing necessary for enabling multiple processors to work in concert. The Grid facilitates parallel processing by handling much of this. Toolkits have been developed that supply protocols for data handling and communications, fault detection, and cross-platform portability (Foster, 2002). This frees the programmer to concentrate on higher-level issues, such as specific parallelization implementations, improved algorithm design, or speedier heuristics.

The problems and applications that GISs address appear well suited to take advantage of the virtual organisation concept and the underlying Grid computing technology that supports it. Because the implementation is so recent, however, there have been very few publicised efforts to develop Grid-based geoprocessing applications (though some are underway (NPACI, 2002)). The next section describes our testbed for implementing a representative application, and reports on some preliminary results.

4 A Watershed Management Testbed for Grid Computing

In this section, we propose a watershed management scenario that will use underlying Grid computing technology intensively. Watershed management is a holistic activity that involves co-ordination among different organisations and collective decision-making based upon information from different sources. The ultimate goal is to optimise the use of land, water and vegetation in an area to alleviate drought, moderate floods, prevent soil erosion, and improve water quality, land use and agricultural production on a sustained basis. There is also an urgent need to educate the general public on watershed problems and to develop and implement agendas to meet these problems.

Due to its organisational characteristics and modelling goals, watershed management is an excellent theme upon which to develop a VO/Grid computing testbed. This is a problem-oriented framework that contains data and methods for facilitating decision-making in a particular geographic region of a watershed.

Fig. 1 presents a hypothetical example indicating how the Internet can be used, via Grid technologies, to develop a watershed management virtual organisation. In this illustration, a university research institute develops several computational hydrological models for a government agency on a contract basis. It sets up a GIS-based model server to provide remote modelling services to that agency. These

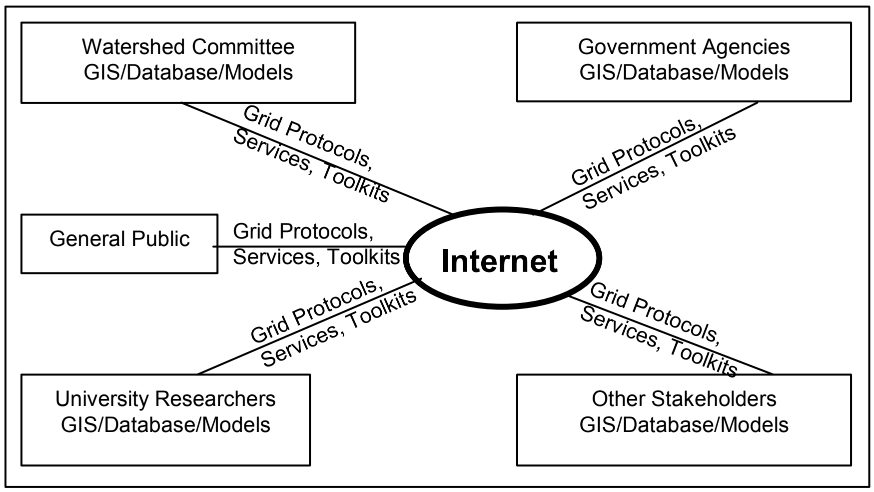

models access real time data collected from field devices maintained by many actors in the VO. In addition, for particular computation- and data-rich models, the

Fig. 1. General Internet-based integration framework for watershed management. Grid protocols handle the flow of data, processing, reports, and other information through the system

model server delegates the processing tasks to one or more of the heavy-duty supercomputers or a cluster of PCs within the VO to meet the real time requirements for the model. Each agent may selectively publish its data and geoprocessing services for the general public to access. Also a watershed management committee can set up a system to integrate data and geoprocessing services from different stakeholders and implement its own real time decision support system. The committee can also selectively publish its data, reports and geoprocessing services to the general public. All of these systems are linked through the Internet so that they can access each other's data and geoprocessing services based on certain protocols defined by Grid technologies.

This research concerns just a few aspects of the full watershed VO just described. In particular, we are interested in testing the data distribution and parallel processing capability of Grid-based technologies. We intend to test the capability of the system to distribute, process, and present results for a basic hydrologic model running on a large geographic region. Since one goal of a watershed VO might be to process results rapidly enough to enable quick scenario checking and short time scale decision making, real time geoprocessing is of particular concern. To test the potential of Grid computing technologies to meet this real time geoprocessing requirement, we developed a parallelized soil loss model using the Message Passing Interface (MPI) standard library that underlies the Grid. The model was run on both a CrayT3E supercomputer located in the University of Texas and local PC clusters we built at the Michigan State University using Grid packages. The architecture of the system is illustrated in Fig. 2.

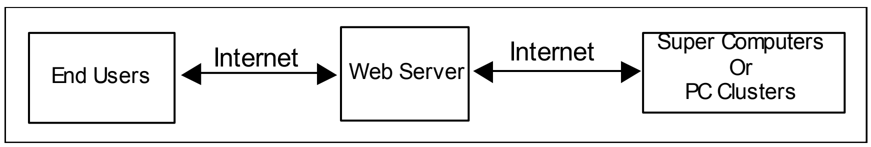

Fig. 2. Prototype system architecture for parallel processing application

The soil loss model has the following form (from Mitasova & Mitas, 1999):

$$E = R * K * LS * C * P \tag{1}$$

where E is the average soil loss, R is the rainfall intensity factor, K is the soil factor, LS is the topographic factor, C is the cover factor, and P is the prevention practices factor. Factors for a 1000x1000 cell raster grid were developed for the calculation of E. This model was run on a single processor on the Cray T3E; processing required 6.525791 seconds. The parallel model was then executed on 10 processors, which required 0.827482 seconds. Performance gain is significant, although this improvement must be weighed against network transfer time, as the following paragraphs indicate.

This spatial application employs purely local map algebra operations upon one or more input raster layers; that is, the resulting value in cell [i,j] in the output raster is solely a function of values at [i,j] in the input raster(s). In general, for applications involving local operations, processing time may be calculated for a single processor:

$$T_s = nRows \times nCols \times T_c \tag{2}$$

where T_s is the time to complete the task on a single processor, *nRows* and *nCols* are the dimensions of the raster, and T_c is the time to perform the local function on a single cell. If the operation is parallelized, processing speed is:

$$T_p = 2 \times nLayers \times \frac{(nRows \times nCols) \times bytes_c}{NetworkSpeed(bytes/s)} + \frac{nRows}{nProcs} \times nCols \times T_c \quad (3)$$

where *nLayers* is the number of input raster layers, $bytes_c$ is the number of bytes per cell value that need to be transmitted, and *nProcs* is the number of distributed processors handling spatial subsets of the operation. Processor speed is held constant for both single and parallel implementations, which is highly conservative (typically, one would take advantage of very powerful distributed processors). Setting Eqs. (2) and (3) equal to one another in order to find the break even between single and parallel processing times (T_s=T_p) and simplifying the result in the following expression:

$$T_c = \frac{2 \times nLayers \times bytes_c}{NetworkSpeed(bytes/s)} + \frac{T_c}{nProcs}. \quad (4)$$

The right-hand side of Eq. (4) consists of two elements: the added cost of parallelization, which is the time required to transfer the data, and the benefit, which is the reduction in processing time. As *nProcs* increases, the processing time drops. For a given application then, one can identify the number of processors required to offset the network speed bottleneck.

5 Discussion and Conclusions

This paper has described Grid computing as an important recent development in the general area of distributed computing. This is due to its ability to link multiple agencies with a network of shared data, software, and processors. The innovation is the degree of integration that a VO employing such technologies can achieve.

We identify three major driving factors for the implementation of Grid-based geoprocessing. As network speeds continue to increase relative to single processing power, parallel implementations become relatively efficient. Second spatial partitioning for local operators is simple and very general. Partitioning for neighbourhood operations ranges from simple (slope) to complex (flow accumulation). Complex partitioning problems restrict the utility of parallel implementations. Finally, algorithm complexity is a factor favouring parallelization, since users can take advantage of more powerful processors than are available on the local network.

The trend towards intensive multi-agency geospatial projects involving the collaboration of dozens of scientists, policy-makers, and members of the public makes the virtual organisation paradigm attractive. This research identifies the utility of real-time distributed model processing for such organisations. We maintain that Grid technologies show great promise for enhancing a variety of geospatial applications, particularly those with intensive computing requirements and a multi-organisational structure.

References

Benner J (2001) A Grid of supercomputers [online]. Wired News, August 9, 2001, Available from: http://www.wired.com/news/technology/0,1282,45977,00.html (Last checked 3/10/02)

Coleman DJ (1999) Geographical information systems in networked environments. In: Longley PA, Goodchild MF, Maguire DJ, and Rhind DW (eds) Geographic Information Systems, New York, Wiley, pp 317-329

El Haddi AA, Shekhar S, Hills R, McManamon A (1996) Mirror-image round robin spatial data partitioning: a case study with parallel SEUS. In: Third International Conference/Workshop on Integrating GIS and Environmental Modeling. January 21-25, 1996, Santa Fe, New Mexico, USA

Foster I (2002) The Grid: a new infrastructure for 21st Century Science. Physics Today February 2002, 42-47

Foster I, Kesselman C, Tuecke S (2001) The anatomy of the Grid: enabling scalable virtual organisation. Intl. J. High Performance Computing Applications 15(3): 200-222

Foster I, Kesselman C (1997) Globus: A Metacomputing Infrastructure Toolkit. Intl. J. Supercomputer Applications 11(2): 115-128

Hecht L (2001) OGC aims to enable information communities [online]. GEO World, August 2001. Available from: http://www.geoplace.com/gw/2001/0801/0801opng.asp (Last checked 3/10/02)

Hunsaker C, Ehlschlaeger CR, Davis FW, Goodchild MF (1996) Estimating spatial uncertainty as a function of scale: implications for landscape ecology. In: Third International Conference/Workshop on Integrating GIS and Environmental Modeling. January 21-25, 1996, Santa Fe, New Mexico, USA

Hoel EG, Samet H (1994) Performance of data-parallel spatial operations. In: Proceedings 20th Intl. Conf. On Very Large Data Bases. September 1994, Santiago, Chile, pp 156-167

Jägers H, Jansen W, Steenbakkers W. (1998) Characteristics of Virtual Organizations [online]. In: Organizational Virtualness: Proceedings of the VoNet Workshop, April 27-28, 1998, Bern, Switzerland pp 65-76. Available from: http://www.virtual-organization.net/files/articles/nl2-1.pdf, last checked 3/10/02.

Longley PA, Goodchild MF, Maguire DJ, Rhind, DW (2001) Geographic Information Systems and Science, New York: Wiley

Magillo P, Puppo P (1998) Algorithms for parallel terrain modeling and visualization. In: Healey R, Dowers S, Gittings B, and Mineter M (eds) Parallel processing algorithms for GIS, London: Taylor & Francis, pp 351-386

Mitasova H, Mitas L (1999) Modeling soil detachment with RUSLE 3d using GIS [online]. Available from: http://www2.gis.uiuc.edu:2280/modviz/erosion/usle.html. (Last checked 3/15/02)

NPACI (2002) SDSC Receives Six ITR Awards. EnVision 17(4):20

OGIS (2002) OpenGIS Consortium Inc. Homepage [online]. Available from: http://www.opengis.org (Last checked 3/10/02)

Palmer JW, Speier C (1997) A typology of virtual organizations: an empirical study [online]. In: Proceedings of the Association for Information Systems 1997 Americas conference. August 15-17, 1997, Indianapolis, USA Available from: http://hsb.baylor.edu/ramsower/ais.ac.97/papers/palm_spe.htm (Last checked 3/10/02)

Shao YP, Liao SY, Wang HQ (1998) A model of virtual organizations. Journal of Information Science 24(5):305-312

Shekhar S, Ravada S, Kumar V, Chubb D, Turner G (1996) Parallelizing a GIS on a Shared-Address Space Architecture. IEEE Computer 29(1):42-48

Smith NS, Rhind D W (1999) Characteristics and sources of framework data. In: Longley PA, Goodchild MF, Maguire DJ, Rhind DW (eds) Geographic Information Systems. New York: Wiley, 655-666

Stonebreaker M (1995) An overview of the Sequoia 2000 Project [online]. Digital Technical Journal 7(3) 39-48. Available from: http://www.research.compaq.com/wrl/DECarchives/DTJ/DTJJ03/DTJJ03PF.PDF (Last checked 3/10/02)

Sugarbaker LJ (1999) Managing an operational GIS, In Longley PA, Goodchild MF, Maguire DJ, and Rhind DW (eds) Geographic Information Systems, New York: Wiley, pp 611-620

UCGIS (1998) Distributed Computing, Revised UCGIS White Papers [online]. Available from: http://www.ucgis.org/research_white/distrib.html (Last checked 3/10/02)