The Anatomy and Physiology of the GeoGrid

Keith C. Clarke[1], Qingfeng Guan[1,] and Tong Zhang[1], [2]

1. Department of Geography, University of California, Santa Barbara
Santa Barbara, CA 93106-4060 {kclarke@geog.ucsb.edu}
2. Department of Geography, San Diego State University

Abstract: In Foster et al.'s description of the Grid, the “grid problem” is defined as how to create “flexible, secure, coordinated resource sharing among dynamic collections of individuals, institutions and resources—what we refer to as *virtual organizations*.” Their solution is to create a grid architecture (protocols, services, applications programming interfaces, and software development kits). This architecture, in addition to an Open Grid Services Architecture that defines uniform sematics, includes ways of creating, naming and discovering service instances, is transparent and interoperative, and is portable across platforms. The technological impact of these systems has had considerable influence on federal research funding in the US, on new fields such as the semantic web, and on data mining and knowledge discovery. Yet what has been the impact of these cyberinfrastructure concepts on GeoComputation, GIScience and on the Autocarto community? We discuss what aspects of the cyberinfrastructure are particularly challenging to geospatial applications. What are the differences between the Grid and the GeoGrid? Is research in GIScience leading to developments in interoperability and geoservices? What are the research opportunities and unanswered questions that inhibit GIScience as a field from contributing toward a GeoGrid? Lastly, does GeoGrid fundamentally change the vision of GeoComputation that emerged from the work of Openshaw and others, such that the vision is in major need of an update? What would a *geospatial* virtual organization look like? This conceptual paper hopes to detail the advances to date, examine work outside of GIScience that is relevant, and outline a possible future research agenda that will encourage the next generation of GIScientists to take a broader and more service-oriented view of GeoComputation.

Keywords: grid, cyberinfrastructure, semantic web, geoservices

1.0 Introduction

Autocarto dates back to a meeting in September 1974, in Reston, Virginia following the previous year's International Symposium on Computer-Assisted Cartography. The automated cartography theme reflected the first stage of the digital transition in cartography, that is converting the map production process to the computer. Influential in the early years was the integration of new source data in the form of remotely sensed imagery. With further advances in computing, the "data manipulation" stage of processing emerged as a new opportunity to apply scientific methods to geographical problem solving, leading to geographic information systems as tools to facilitate spatial analysis, and a new variant of cartography, named Analytical Cartography (Tobler, 1976).

As GIS matured to yield a new approach to problem solving, that Michael Goodchild termed Geographic Information Science (Goodchild, 1997), we similarly saw a closer bond between advanced computing methods and GIScience labeled GeoComputation (Openshaw and Abrahart, 2000). Perceptively, Chrisman (1997) expanded the definition of GIS to include people and organizations, especially those impacted by GIS and a technology. Other important components in this abbreviated intellectual history include digital map libraries and the rise of the Internet and web cartography. As we move toward an era of geobrowsers and massive on-line geodatabases on which large proportions of citizens now rely for everyday services, it is probably time to ask: who (or what) is the community of contemporary Autocarto? We contend that it has moved on beyond a small community of academic, government and industry craft practitioners, and now encompasses society at large. The effects of this change in audience will be profound, and are probably best confronted, and exploited, early.

Long before the public view of the internet became a vast collection of flat and largely text-oriented "pages" of HTML, there was a more ambitious vision of what the web could bring, primarily that of its inventors. Key to understanding this vision is the concept of a *virtual organization*, that is a set of individuals and/or institutions defined by a suite of sharing rules for jointly using information resources (Foster et al., 2001). In the revised view of network-based computing now termed either the "Grid" or "cyberinfrastructure", this communal sharing is concerned not with computer file exchange but with "direct access to computers, software, data, and other resources as is required by a range of collaborative problem-solving and resource brokering strategies" (Foster et al., 2001, p. 2). One measure of how this collaborative model is emerging to dominate computing is the fact that by 2002, 14% of web requests to Microsoft's popular Terraserver were direct (software) calls to the Terraserver WebServer, and an additional 5% were call to the TerraService Map Servers (Barclay et al., 2002). Clearly the internet is increasing seeing visits to web services not by people using browsers, but by browsers and services serving people and communities. Indeed, contemporary browsers barely resemble any more the tools of the isolated user seeking information, but clusters of agents swarming around their users specific demands, anticipating and enriching whatever queries are given and using the internet in a very different way.

So what are the Autocarto community's virtual organizations? What tasks do they perform and how well are those tasks supported by the information that is delivered? Given the emerging sets of protocols, systems architectures and distributed services now collectively called the Grid, is there a distinct community that might be called the GeoGrid, with its own demands that are being or are not being met by cyberinfrastructure research and development? And what does the community stand to lose if it is not involved in the research and development of the GeoGrid? As unanswerable as these questions currently are, in this paper I speculate about the next generation of methods for computer mapping, which includes the creation of new virtual communities.

In this paper we first examine the components--the anatomy and physiology--of the grid (Foster et al., 2001; Foster et al., 2002). Next, we narrow the discussion to the GeoGrid, i.e. the components of the grid specific to handling and using geospatial data. Lastly, we discuss the nature of the virtual organizations that exist or could emerge as the Autocarto community transitions into a virtual organization and expands to encompass far more everyday activities.

2.0 The Anatomy and Physiology of the Grid

The Grid has a set of component parts, an anatomy, and a set of functions and behaviors, a physiology. The anatomy of the Grid consists of many components that already exist: networks, computers, peripheral devices, sensors and sensor webs (including GPS), displays, and even supercomputers. Other necessary elements are the formal standards and specifications that establish linkages and manage the transactions among components, such as protocols and transfer standards. A stated goal of grid computing is to remove from the user the details of exactly how (i.e. in terms of hardware and software) a computing task is performed. This implies that where the user now performs hardware-specific tasks (e.g. establish a connection, log-in, search and manipulate files, transfer information) they can be “hidden” so that they function from the user perspective as a transparent service. For example, in a data clearinghouse model, the user needs to search a portal such as the Geospatial One-stop, then link to a data site, then query to see whether data are online, in what format, at what scale, and by what tiling system. In a grid model, the user simply sets up a search by theme and extent, and data are assembled, transformed, reprojected and brought into user software without further user interaction.

Similarly, and also heavily based on the standards, lies the physiology. This consists mostly of software tools and utilities, based on the standards, that facilitate interoperability and support the service delivery model. Wheras the client-server model was a major development in computing, the Grid therefore offers the next transition of computing, and the prospect of a cyberinfrastructure as envisioned in the NSF plans (Atkins, 2003). From the anatomy point of view, the original grid layers (fabric, connectivity, resources,collectives, and applications) need to be modified or extended to reflect the special characteristics of geospatial data and the GeoComputing paradigm. This will in turn be reflected in the extension of Grid tools such as the Globus Toolkit (Globus Alliance, 2006). This task is usually called building the "geospatial grid middleware" (Armstrong et al., 2005).

Foster et al. (1999) defined the Grid as “a proposed distributed computing infrastructure for advanced science and engineering.” The grid must assure flexible, secure, coordinated resource sharing among dynamic collections of individuals, institutions and resources. This collection of people is what was termed a “virtual organization.” Foster et al. (2001) gave examples of: a company seeking to place a new factory; an industrial consortium developing a feasibility study for next generation supersonic aircraft; a crisis management team responding to a chemical spill; and thousands of scientists collaborating around the superconducting semi-collider at CERN in Switzerland. The solution is to create a grid architecture (protocols, services, applications programming interfaces, and software development kits), which, in addition to an Open Grid Services Architecture defines uniform sematics, includes ways of creating, naming and discovering service instances, is transparent and interoperative, and is portable across platforms. This architecture allows the applications community to function as a virtual organization, regardless of location or status.

The architecture of the grid has received an extraordinary amount of attention in the last few years, and even forms the basis of a whole special office at NSF. First and foremost, grid computing rests on protocols. These are mostly standards-based, such as XML, TCP/IP, HTML, etc. These enable web-based service provision, a simple geographic example being the ability to send a server a place name such as “Vancouver”, and receive geographic coordinates associated with the place (Leidner, 2004: Metacarta, 2006). Components of the grid that are special to GIScience, and hence the GeoGrid include those protocols surrounding geospatial data, where a substantial amount has already been done based on early work for the US National

Spatial Data Infrastructure under the auspices of the Federal Geographic Data Committee, and more recently by the Open GeoSpatial Consortium (OGC, www.opengeospatial.org). With substantial support and active interest from government and industry in particular, the OGC has spearheaded many geospatial data standards. A few more have been added via the World Wide Web consortium (www.w3.org), such as XML and the Geospatially-enabled Virtual Reality Modeling Language (GeoVRML).

A recent representative study that integrated open source solutions to solve the problem of disseminating environmental data and simple analysis results in Mexico (Anderson and Moreno-Sanchez, 2003) used the following set of protocols and standards in the process (Table 1).

Table 1: Grid Open Source Tools Used by Anderson and Moreno-Sanchez

Component	**Role**	**Source**
XML	Extensible Markup Language, supports browser and web interoperability and data self-description	www.w3.org
GML	Geography Marup Language. XML subset for geospatial data	www.w3.org
SVG	Scalable Vector Graphics. XML subset for drawing and scaling vector graphics on web browsers	Eisenberg, J. D. (2002)
GRASS	Grass open Source GIS	grass.baylor.edu
XLST	Extensible Stylesheet Language: Transformations, part of XLS	www.w3.org/Style/XSL
SAX	Simple API for XML. A parser for XML code.	www.w3.org
Java2D API	Java Development toolkit (JDK) subset for 2D object handling	Java.sun.com/products/java-media/2D
PHP	Personal Hope Page tools. Scripting language for HTML control.	www.php.net
PostgreSQL	Object-relational database management system supporting SQL	www.postgreSQL.org
PostGIS	PostgreSQL-based RDMS supporting geographic objects. Uses OGC's Sinple features Specification for SQL	www.postgis.org
MapServer	Development environment for building web-based mapping applications	mapserver.gis.umn.edu
Linux	Operating system. Open source UNIX variant	www.redhat.com
Apache Web Server	HTTP compliant web server	www.apache.org

Clearly evident here are the critical elements of the grid. GML, a subset of XML targeted at geospatial data, has the advantage of supporting full geographical referencing, and so has an advantage over other tools such as Flash. The close fit of

XML and OGC is highly advantageous, meaning that tools for web-delivery of high quality graphics (e.g. SVG) and for serving geospatial data across the web (e.g. PostGIS) can interoperate. The compatibility of existing GIS (e.g. GRASS) and programming support (e.g. PHP, Java) are indeed powerful. As a result, the functionality of Web-linked GIS, such as ESRI's Internet Map Server ,can now be subsumed into open source tools such as MapServer. Andseron and Moreno-Sanchez (2003) note that "Isolated, standalone systems are being replaced by integrated components, and large applications are being replaced by smaller, more versatile applications that work together transparently across networks" (Anderson and Moreno-Sanchez, 2003, p 448). While many of these "mashups" deliver data, the trend is toward delivering services, and analytical GIS services at that. PostgreSQL-postGIS, for example, is now capable of performing over 60 spatial operations, including projected distance computations, selection of features left and right of another features, and intersection, with plans to add topology. We are on the verge of being able to deliver all known GIS functionality and data via the web, through open source tools, and ubiquitously across environments.

The advantages of the mixed solution are also ennumerated by Anderson and Moreno-Sanchez (2003) and include:

- Effectively zero software costs
- Easy learning curve for anyone familiar with IT and UNIX, coding and databases
- Small software footprints
- No need to commit to proprietary solutions
- Compatibilty with existing IT infrastructure
- Flexibility to integrate new GIS capabilities
- Principles are straightforward and accessible to a broad audience of GIScientsis and developers
- Interoperability with other systems and applications based on the same open source appliations

These are exciting developments in the field of GIS, where feature-creep and increasingly complex and complicated data models and user interfaces have been the norm for the last few years. Clearly there are remaining problems, especially computational efficiency and grid availability. Nevertheless, the GeoGrid could be said to already exist in terms of its anatomy and physiology. What comes next is how the body develops behavior, learns and increases its own capacity and capability. For this, we need to examine the concept of the virtual organization in a GIS context. Before this, however, we consider the macro-level problems set forward in the introduction.

3.0 Research Issues

Much attention has been given recently to agenda setting in GIScience. While topics like distributed and mobile GIS, data mining and location-based services are present in the agenda for UCGIS (McMaster and Usery, 2005), rarely are they brought out as an emerging overriding theme. Judging by the numbers of papers, geography and GIScience's contributions to the GeoGrid have so far been minimal. Yet there are many critical problems for Grid computing and web services that are directly the consequences of what is already known about spatial data, such as the performance derived from nested data structures and hash indices (Yang at al., 2005). there are major and fundamental differences in how the GeoGrid functions compared to regular web services delivering non-spatial information. This was obvious in the first generation of Autocarto, as the first generation data base management systems were found inadequate for spatial data structures and management. As yet, the GeoGrid forms only a minor part of the research agenda for GISciences, with a possible exception at the National geospatial Intelligence Agency (NRC, 2006). In this report, some of the impediments and promising solutions are discussed and the value of the VO is recognized: “the overarching promise of the 'virtual organization' has high value to NGA in its GEOINT goals.”

In the existing GIScience research, closest to realizing the potential of the GeoGrid is the GeoComputation tradition, best represented by the series of coenferences and on-line web site (www.geocomputation.org). Yet even in these books and conferences, the role of cyberinfrastructure and the vision of virtual organizations has been underrepresented, in spite of the encouraging statements and discussions by Openshaw and others (Opernshaw and Abrahart, 2000, esp. Chapter 11). Perhaps this is because geocomputation initially saw parallel and high performance computing as simply more of the same, i.e. faster and bigger, rather than a change in paradigm. This perception is now changing, and the merit of fully distributed computing seems more focused (Clarke, 2003). In short, technology now no longer seems to be a major problem for the GeoGrid. It is for this reason that a focus on the Virtual Organization becomes of increased importance.

4.0 Autocarto as a Virtual Organization

In order to build the GeoGrid and a geographical VO, we need to address the problems from both the perspective of grid techniques and that of GIScience. On the one hand, to enable the current grid framework and software to support GIServices and applications, there are many technical concerns which we examine elsewhere (Zhang at al., 2006). On the other hand, within GIScience, what can we do to transform the GISystems or GIServices to make them fit into the grid world?

<u>4.1 GIServices as Web Services</u> From the standpoint of current activity, the Service-Oriented Architecture strategy involving Web Service compatible GIServices seems to be the solution. Actually the concept of grid services is to implement grid applications in Web Services, as evident in the trend from WSRF of Globus Toolkit 4.0. The WS-Resource Framework (WSRF) is a set of six Web services specifications that define what is termed the WS-Resource approach to modeling and managing state in a Web services context. The Globus Toolkit is an open source software toolkit used for building Grid applications under development by the Globus Alliance and others. Globus Toolkit is being used to construct Grid infrastructure and to develop Grid applications.

From the geography side, pioneering work on developing new grid-based tools to solve geocomputation problems has been conducted, largely at the University of Iowa under the direction of Shaowen Wang (Wang et al., 2005). The main strategy there is to create a shared Grid-based campus and state cyberinfrastructure that interoperates with the evolving national and global Grid-based cyberinfrastructure. The development is based on currently available Grid tools (Globus Toolkit). The perspective is that given geographic computational tasks, high performance computing (HPC) and grid computing techniques might be of value to the Grid. Thus the convergence of GIScience and Grid will indeed be realized through Web Services.

However, geospatial problems do require particular consideration when applying these emerging techniques. An obvious solution is to develop middleware to facilitate geospatial grid applications based on current grid tools. Geographers know what problems are suitable for the GeoGrid and what geocomputation/geovisualization methods can be employed to solve geospatial problems. However, given the tools and algorithms, many geographers find it is too hard to make them geography-oriented or even work properly in a geography context. This is particularly true for HPC applications for geospatial problems. Grid techniques are available, though making them fit into current GIS models is difficult, and the awareness of their utility within GIScience is limited. Early work has centered on spatial statistics and spatial grids. However, there are many spatial analysis methods which geographers have not tested if they seek better performance within grid environment or what can achieve better performance.

4.2 GIS as a Grid Application Geoffrey Fox (Fox et al., 2006) and his group at the University of Indiana (http://www.crisisgrid.org) approach the solution from a different perspective. They are computer scientists involved in the development of Grid techniques and the Globus Toolkit. For them, GIS is a good applications area of the grid. When they looked into the GIS community, they found the OGC specifications. Consequently, most of their resarch is centered around implementing OGC compatible GIServices in a Grid environment. This team continues to explore further into GIS and spatial analysis, researching geographic scale and cartography so as to design better tools for geospatial problems.

It is likely that these two strategies will eventually merge. Web Services and OGC standards are very important to connect the Grid and GIS. It is critical to let computer scientists know what special concerns GIServices and GISciences have. Thus they can develop powerful middleware to make Grid-based GIServices interoperable and seamlessly powered by the Grid. Thus there is great scope for interdisciplinary collaborative work. Undoubtably, the Web Service paradigm can be extended to reflect the particularities of GIS. Hopefully, OGC will be releasing new specifications for Geocomputation/analytical services as well as for HPC geospatial standards. The increasing collaboration between computer scientists and GIScientists will contribute to realizing the geospatial cyberinfrastructure. This can be more rapidly supported by building an Autocarto Virtual Organization, or AC community.

An Autocarto virtual community will emerge to support this need for collaboration. The internet provides tools for cross collaboration such as discussion forums and Wikis, and there is increasing evidence that these can contribute much. Of great value has been the open source movement, including the Free Software Foundation and the hosts of groups supporting GIS-related mashups and collaboratives. A mashup is a website or web application that seamlessly combines content from more than one source into an integrated experience (Wikipedia, 2006). In a mashup, one site or service relies and depends upon the existence and capability of another. Such an information combination or fusion requires geospatial data and tools, but also requires geographical knowledge, IT experience and educational assistance (Erle et al., 2005). The web, and the emerging Web GIS (Peterson 2003), coupled with the myriad of applications that use the web, from MapQuest to GoogleEarth and the National Map Viewer, are the first wave of GeoGrid applications from which GIScience can benefit. The Internet's extraordinary ability to link scholars and practitioners together across international, age and cultural boundaries is also an important part of the virtual community. Similarly, the ever-growing body of knowledge surrounding GIS (for example, the NCGIA Core Curricula, the Geographers Craft Project, UCGIS's Body of Knowledge project, and CSISS) are essential to build the user community capable of exploiting these powerful and versatile tools.

Taken together, this virtual organization is currently simply the sum of its parts. The GeoGrid offers the potential for the seamless integration of GIS data, tools, knowledge and applications. Given what has already been accomplished by GIS in its brief history, and the vastly increased power offered by the visions of the cyberinfrastructure and Grid, the future appears bright indeed for the Autocarto Virtual Community. Perhaps an AC Wiki is a good starting point. In fact, the contribution and participation of the GIScience community is essential. No element of the Geogrid is incompatible with the methods, theory, principles and scholarship of analytical cartography, geocomputation, or geographic information science. Indeed, from this mashup will emerge the next generation of Autocarto.

5.0 References

Atkins et al. (2003) Revolutionizing Science and Engineering Through Cyberinfrastructure: report of the national Science Foundation Blue-Ribbon Panel on Cyberinfrastructure. http://www.communitytechnology.org/nsf_ci_report/

Armstrong, M. P., Cowles, M. K., and Wang, S. (2005) Using a Computational Grid for Geographic Information Analysis: a Reconnaissance, The Professional Geographer, 57(3), pp. 365-375.

Anderson, G. and Moreno-Sanchez, R. (2003) Building Web-Based Spatial Onformation Solutions around Open Specifications and Open Source Software. Transactions in GIS, 7 (4): 447-466.

Barclay, T. Gray, J. Strand, E. Ekblad, S. amd Richter, J. (2002) *TerraService.NET: An Introduction to Web Services*. Microsoft research Advanced Technology Division. Technical report MS-TR-2002-53

Chrisman N. R. Exploring geographic information systems. New York: Wiley; 1997.

Clarke, K. C. (2003) "Geocomputation's future at the extremes: high performance computing and nanoclients", Parallel Computing, 29, 10, 1281-1295.

Eiseberg, J. D. (2002) SVG Essentials. Sebastopol, CA: O'Reilly and Associates.

Erle, S. Gibson, R and Walsh, J. (2005) Mapping Hacks : Tips & Tools for Electronic Cartography. Sevastopol, CA: O'Reilly and Associates.
Fox, G., Pallickara, S. and Pierce, M. (2006) Building a Grid of Grids: Messaging Substrates and Information Management in M.P. Bekakos, G.A. Gravvanis and H.R. Arabnia. (Eds.) *Grid Computational Methods* http://grids.ucs.indiana.edu/ptliupages/publications/GridOfGrids.pdf
Foster, I. And Kesselman, C. (eds.) (1999) The Grid: Blueprint for a New Computing Infrastructure. Morgan Kaufman.
Foster, I., Kesselman, C. and Tuecke, S. (2001) "The anatomy of the Grid: Enabling Scalable Virtual Organizations", Lecture Notes in Computer Science, Berlin, Springer, Vol. 2150.
Foster, I., Kesselman, C., Nick, J. M. and Tuecke, S. (2002) The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration, Open Grid Service Infrastructure WG, Global Grid Forum.
Globus Alliance. (2006). http://www.globus.org. Accessed 4-19-06.
Goodchild, M. F. (1997) What is Geographic Information Science?, NCGIA Core Curriculum in GIScience, http://www.ncgia.ucsb.edu/giscc/units/u002/u002.html.
Leidner, J. L. (2004) Towards a Reference Corpus for Automatic Toponym Resolution Evaluation SIGIR 2004 Sheffield, UK. (Geographic IR Workshop) http://www.iccs.informatics.ed.ac.uk/~s0239229/documents/leidner-2004-topomark.pdf
McMaster, R. B. and Usery, E. L. (2005) A Research Agenda for Geographic Information Science. Taylor and Francis; CRC Press.
Metacarta (2006) Metacarta web site. Www.metacarta.com. Referenced 3/28/2006.
National Research Council (2006) Priorities for GEOINT Research at the National Geospatial-Intelligence Agency. National Academies Press; Washington, D.C.
Openshaw, S. and Abrahart, R. J. (2000) GeoComputation. London: Taylor and Francis.
Peterson, M. P. (2003). Maps and the Internet. New York: Elsevier.
Tobler, W. R. (1976), Analytical Cartography, The American Cartographer, 3: 21-31
Wang, S. Xiao, S. and Ni, J. A Grid-Based Bridging Domain Multiple-Scale Method for Computational Nanotechnology. International Conference on Computational Science 2005, Atlanta, GA, May 22-25, Lecture Notes in Computer Science, Vol. 3516, pp. 326-333, 2005.
Wikipedia (2006) Mashup (web application hybrid) http://en.wikipedia.org/wiki/Mashup_(web_application_hybrid)
Yang, C., Wong, D. W. Yang, R., Kafatos, M., and Li, Q. (2006) Performance-improving Techniques for Web-Based GIS, International Journal of Geographical Information Science, 19 (3), 319-342.
Zhang, T., Clarke, K. C., and Guan, Q. (2006) The GIS Web Portal: Beyond Data Services", Proceedings, AutoCarto 2006, Vancouver, WA. June 2006. www. autocarto.org