Encyclopedia of Networked and Virtual Organizations

Goran D. Putnik
University of Minho, Portugal

Maria Manuela Cunha
Polytechnic Institute of Cávado and Ave, Portugal

Volume II
G–Pr

Information Science REFERENCE

INFORMATION SCIENCE REFERENCE
Hershey • New York

Acquisitions Editor: Kristin Klinger
Development Editor: Kristin Roth
Senior Managing Editor: Jennifer Neidig
Managing Editor: Sara Reed
Copy Editors: Amanda Appicello, Alana Bubnis, Ashley Fails, April Schmidt, Katie Smalley, and Larissa Vinci
Typesetters: Jeff Ash and Jamie Snavely
Cover Design: Lisa Tosheff
Printed at: Yurchak Printing Inc.

Published in the United States of America by
Information Science Reference (an imprint of IGI Global)
701 E. Chocolate Avenue, Suite 200
Hershey PA 17033
Tel: 717-533-8845
Fax: 717-533-8661
E-mail: cust@igi-pub.com
Web site: http://www.igi-pub.com/reference

and in the United Kingdom by
Information Science Reference (an imprint of IGI Global)
3 Henrietta Street
Covent Garden
London WC2E 8LU
Tel: 44 20 7240 0856
Fax: 44 20 7379 0609
Web site: http://www.eurospanonline.com

Library of Congress Cataloging-in-Publication Data

Encyclopedia of mobile computing and commerce / David Taniar, editor.
p. cm.
Summary: "Nowadays, mobile communication, mobile devices, and mobile computing are widely available. The availability of mobile communication networks has made a huge impact to various applications, including commerce. Consequently, there is a strong relationship between mobile computing and commerce. This book brings to readers articles covering a wide range of mobile technologies and their applications"--Provided by publisher.
Includes bibliographical references and index.
ISBN 1-59904-002-6 (hardcover) -- ISBN 1-59904-003-4 (ebook)
1. Mobile computing--Encyclopedias. 2. Mobile communication systems--Encyclopedias. 3. Mobile commerce--Encyclopedias. I. Taniar, David.
QA76.59.E47 2007
004.16503--dc22

2006039745

British Cataloguing in Publication Data
A Cataloguing in Publication record for this book is available from the British Library.

Grid Computing for Social Science

Kenneth J. Turner
University of Stirling, UK

Paul Lambert
University of Stirling, UK

K. L. Tan
University of Stirling, UK

Vernon Gayle
University of Stirling, UK

Richard O. Sinnott
University of Glasgow, UK

Ken Prandy
University of Cardiff, UK

Erik Bihagen
Swedish Institute for Social Research, Sweden

Marco H. D. van Leeuwen
International Institute for Social History, The Netherlands

INTRODUCTION

Grid computing is named by analogy with the electrical power grid. Power stations are linked into a universal supply that delivers electricity on demand to consumers. Similarly, computational resources can be linked into a grid that delivers computing or data on demand to the user's desktop. The origins of grid computing lie in networked computing, distributed computing, and parallel computing. Grid computing coordinates distributed resources that are not subject to central control, using standard protocols and interfaces to meet the required levels of service (Foster, 2002).

International standards have been defined for open and vendor-independent networking. However, this simply allows exchange of information among computers that operate largely independently. Distributed computing was developed to use networks for effective sharing of resources: information, processing, and devices. This allows multiple computers to cooperate in manipulating information, no matter where it is stored or processed.

Parallel computing evolved separately as a means of allowing closely coupled processors to work on the same tasks. However, it is often more cost-effective to harness the power of many processors through distributed computing than to build specialized parallel processing systems.

Grid computing continues these trends by offering a standard architecture for large-scale distributed and parallel computing. Grid computing has proven particularly attractive where substantial computation is needed to analyze the rapidly growing amounts of data used by many collaborators across a wide range of disciplines. For example, it is envisaged that the Large Hadron Collider will regularly produce many petabytes of data (over 10^{15} bytes per second from small particle collisions).

In the context of the grid, a virtual organization is a set of individuals or institutions that agree to share distributed computing facilities. Such an organization is virtual because it crosses conventional organizational boundaries. A virtual organization is typically closed (the membership is specified) and may be dynamic (membership may change over its lifetime).

Grid computing provides facilities to support virtual organizations, notably authentication and security. This chapter describes a case study that demonstrates the value of virtual organizations for collaboration among social scientists.

BACKGROUND

Distributed Computing

A wide range of standards was developed for networked computing under the banner of open systems interconnection (ISO/IEC, 1994). Although this architecture remains an important framework for networking, both it and the proprietary solutions it superseded have largely faded from practical use. Instead, Internet-related standards have become the predominant means of linking networked systems.

Remote procedure calls (Birrell, 1984; ISO/IEC, 1988) were developed as a simple but effective way of linking computation in different systems. They allow procedures (methods) to be invoked in the same way, irrespective of whether they are on local or networked processors. The same approach is supported by modern programming languages such as C# or Java (Downing, 1998).

Initially, proprietary standards were developed for distributed computing, for example Microsoft distributed component object model (Abernethy,Morin, Chahin, & Morin, 1999). The need for vendor-independent standards led to the common object request broker architecture (Bolton, 2001). International standards were also defined for open distributed processing (ISO/IEC, 1995).

Distributed architectures have been defined in specialized areas such as telephony. For example the architecture of the intelligent network (ITU, 2000) supports distributed call processing, while the telecommunications intelligent network architecture (Dupuy, Nilsson, & Inoue, 1995) focuses on telephony services.

Grid Computing

Grid computing has grown out of web computing. The growth of the web as a means of providing information led to strong industrial interest in using a similar approach for distributed computing. Web services are supported by many commercial solutions such as BEA Weblogic (Mountjoy & Chugh, 2004), IBM Websphere (Francis, Herness, Knutson et al., 2002), Microsoft .NET (Platt, 2003), and Sun Open Net Environment (Mogha & Bhargava, 2002).

Web services allow networked applications to cooperate. Extensible markup language (XML) is a common foundation in web service standards for representation of structured data. The interface to a web service is defined by the web services description language (World Wide Web Consortium, 2001). Communication between web services is supported by the simple object access protocol (World Wide Web Consortium, 2003). Repositories of web services are created and interrogated through universal description, discovery and integration (OASIS, 2002).

Web services were adopted as the basis for grid computing, allowing grid applications to support distributed resource sharing and computation. Grid applications often support long-lived services. This requires stateful services that preserve state across different invocations. To meet this requirement, specialized solutions were initially developed for the grid. However, web services can also benefit from being stateful. The web Services Resource Framework (Graham, Marmakar, Mischinsky et al., 2006) was therefore developed as a collection of standards that allow Web services to expose the resources they use. This has allowed a much closer convergence between web services and grid services. However, the grid also has specialized standards in areas such as data access, file transfer, resource management and security. A grid service is said to run in a *container* that provides the infrastructure required.

Berman, Fox, and Hey (2003) provide a high-level overview of what the grid aims to achieve. Grid computing offers a number of distinctive advantages that draw on features for:

- *Virtual organizations* that transcend conventional boundaries, and may come together for only a particular task or to share resources
- *Grid portals* that provide simplified, user-oriented access to grid-enabled resources
- *Single sign-on* that requires authenticating once to use multiple resources such as data repositories or computational servers
- *Virtualized resources* that allow platform-independent and location-independent access to heterogeneous computing and data resources, or to specialized equipment

- *Security* that controls access, but allows third parties to act on behalf of the user
- *Distributed* and *parallel computing*

Standards for the grid are being defined by the Global Grid Forum (GGF). Work is progressing towards standards for an open grid services architecture (Foster, Kesselman, Nick, & Tuecke, 2002). Although web and grid services now share a common basis, grid computing make use of many additional features such as those listed above. Globus Toolkit 4 (Sotomayor & Childers, 2006) and the Commodity Grid Kits are widely used to implement grid applications.

Occupational Data Analysis

Grid computing is a very broad area. To illustrate what can be achieved, this chapter focuses on an important application in social science: the means whereby social survey researchers collect and process data on the occupations of survey respondents. Such data is routinely collected because it is widely believed that occupational position reveals a great deal about an individual's life circumstances. However, the data that is originally collected—typically a textual description of a job title—is not in itself easily summarized or analyzed. Usually, social scientists wish to develop a summary indicator of occupational circumstances. This is typically assigned to survey respondents by means of an external occupational information resource, which connects the original occupational record with some other summary data about the occupational position. As an example, sociologists often use occupational information resources in order to assign occupational positions to locations within a social classification system, such as a 'social class' scheme.

Social scientists thus make widespread use of occupational information databases that contain summary data on particular occupational positions. Many alternative resources are available to the social science research community. Most often they are distributed freely, as small electronic files downloadable from voluntarily maintained Internet sites; for example, see Ganzeboom, 2006, and Lambert & Prandy, 2006. However the way in which such information resources are distributed and exploited is rather uneven in social science practice. Most often, occupational information resources are published with informal instructions only; the datasets are seldom annotated with consistent metadata about their structure and provenance. Indeed they are not usually available via repositories, inhibiting their use by a wider community.

Further difficulties arise when social scientists wish to connect occupational information resources to the records of occupational positions stored in social survey datasets. Survey data is often subject to strict controls on access. The data frequently contains several different pieces of information on occupational positions stored in different locations, which the analyst may wish to relate to each other. Although occupational records are an important measure in many social science analyses, social scientists rarely obtain the benefits of effective resource sharing and efficient connections between survey data records and occupational information resources.

An example of occupational information resources can be found in the occupationally-based social classifications supplied by the Camsis project (Cambridge Social Interaction and Stratification Scale, 2007). This project calculates measures that indicate the average levels of social advantage or disadvantage (in, for instance, economic and environmental terms) that are typically experienced by those in different occupational positions. The calculation is based on statistically modeling patterns of social interaction exhibited by the incumbents of different occupations (Lambert & Prandy, 2003). The approach generates Camsis scale scores (measures that may be assigned to different occupational positions). Databases of these scores are examples of occupational information resources that are distributed and coordinated through the Geode.

Collaboration is desirable among social scientists using occupational data. Those who originate occupational information often wish to cooperate directly. However, there are many other users who simply wish to link their social survey data with an appropriate occupational information resource. Moreover, much social research transcends national boundaries, so many occupational information databases refer to multiple countries. Some research has tried to standardize occupational data collection and analysis across countries, for example Elias and Birch (2006). However, the prevailing model remains pluralistic, whereby numerous different occupational information resources are available across different countries and from different time periods.

There is a strong need to create virtual organizations, often ad hoc and short-lived, that can cooperate on occupational data issues. This includes work on developing new occupational information databases, and on international standards for occupational classifications. Grid computing provides a secure and distributed framework for such collaborations.

GRID COMPUTING CASE STUDY

Occupational Data on the Grid

As a case study in grid computing, this section describes the work undertaken by the Geode project (grid enabled occupational data environment, 2007). The aim of this project was to allow social scientists using occupational data to benefit from grid computing. The Geode project ran from October 2005 to March 2007. The project involved social scientists in different institutes, and brought social scientists and computer scientists into close collaboration.

A high-level view of the Geode architecture is shown in Figure 1. The user's interface is provided by an occupational data portal. A system administrator is responsible for maintaining the overall system, including the users, virtual organizations and data repositories Subject to suitable permissions, an end user can add a local (i.e., locally curated) or external (i.e., remote) datasets. End users can also search, access and link to datasets in the repositories.

On user request, the portal passes details of a dataset to be added to the local or external data repository. An indexing service registers a list of repositories, allowing searches to be made. Datasets are accessed via a data service that supports format-independent access. A data linking service allows survey data to be linked to these datasets.

Portal Interface

Portals are commonly employed for user-oriented access to *web* services. A Java community standard (JSR 168) has been developed as a standard framework for the components of a portal, called portlets. Since portals are ideal for less technical users, a *grid* portal has been developed by Geode for use by social scientists working on occupational data.

The Geode portal has been developed using GridSphere (2007). The approach clearly separates the presentation (user interface) from the logic (underlying grid services). GridSphere provides direct support for the grid, notably in the area of security. Authentication can be performed by several means, ranging from a username/password to digital credentials. The portal interface for Geode allows an administrator to create users, authorize use of data, add and maintain occupational datasets, and manage metadata for these. For ordinary users, the portal provides facilities to provide, link and search occupational datasets.

Data Access

Data abstraction is supported by OGSA-DAI (database access and integration, 2007). This provides a unified and networked interface to datasets held in a variety of formats, including relational databases and comma-separated flat files. OGSA-DAI acts as the middleware for data management in a grid (a data service as shown in Figure 1).

Occupational information resources such Camsis scale scores are currently supplied by social scientists on a voluntary basis in a variety of different electronic formats. The proprietary packages SPSS (2007) and Stata (2007) are particularly common. However relational databases, spreadsheets and even comma-separated value lists are also employed. The Geode project has developed uniform support for these formats, allowing social scientists freedom in how their data may be supplied or accessed.

The suppliers of occupational information resources rarely provide metadata to describe the technical format of a dataset or its provenance. The Data Documentation Initiative (2007) supports XML schemas for annotating social science datasets with metadata. This approach has been specialized by Geode for databases containing occupational information. This gives greater consistency, making it easier to search and link occupational data online. The metadata is registered with an indexing service when the corresponding datasets are added via the Geode portal.

Some datasets are under the control of Geode, while others are contributed by collaborating partners. For this reason, a distinction is made in Figure 1 between the local data repository (aggregate occupational data classifications managed by Geode) and external data repositories (managed by others). The architecture is distributed and scalable, making it easy for other social scientists to provide and use occupational information. Collaborators may supply their data in native format.

Figure 1. Architecture for occupational data grid services

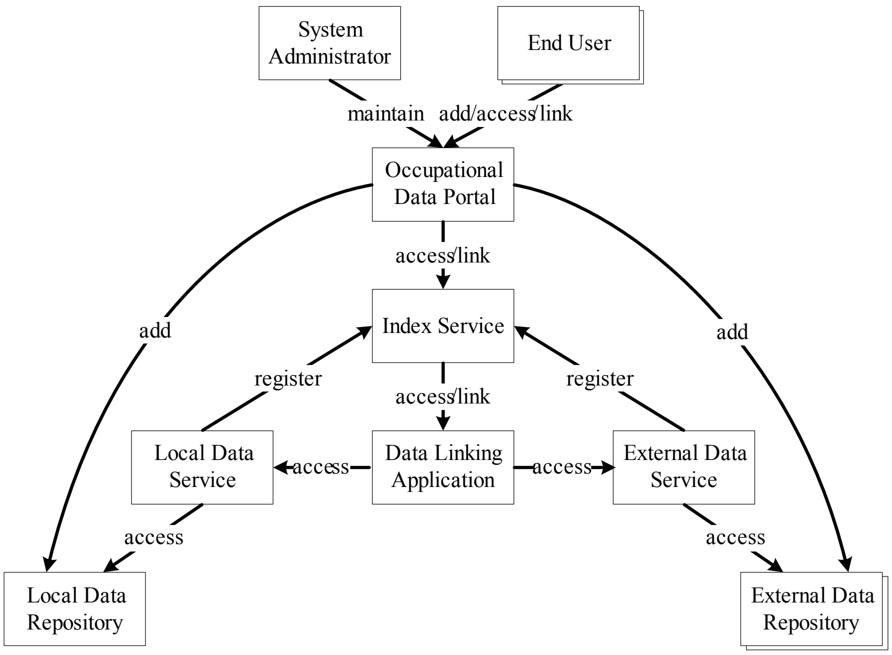

Alternatively if they can run their own OGSA-DAI data service, this may be linked into the Geode framework. A data linking application service is provided to allow micro-social survey information to be linked to occupational classifications.

Datasets and occupational applications are registered with the index service provided by the Globus Toolkit. This creates a searchable repository of data services and application services provided by Geode. The design allows one or multiple instances of the index service to be created, thus permitting a centralized or distributed implementation of the facilities for analyzing occupational data.

OGSA-DAI supports format metadata (technical aspects such as the database schema) as well as custom metadata (a specialization of the Data Documentation Initiative for occupational data). Rather than using data directly, users employ the information via application services provided through the portal.

Security

Security is an important issue, partly because some datasets require strict controls on access. The goal of open access for social scientists must also be tempered with the need for proper authorization.

The Globus Toolkit implements the grid security infrastructure, that in turn derives from web services security. This supports authentication, authorization and credential delegation. It also permits single sign-on—a user need authenticate only once in order to access multiple resources. Logging onto the Geode portal allows users to delegate their credentials (i.e., use proxy certificates) so that the underlying services can perform data manipulation on their behalf. OGSA-DAI is configured to use this credential information when data access is authorized.

Geode can store credentials for retrieval by services that need them. This means that credentials do not have to be moved around among the services. An alternative choice would be MyProxy (2007), but this is currently limited to Unix implementations. Geode is intended to be platform-independent, including Microsoft Windows (which is much more widely used than Unix by social scientists). The preferred solution is the Delegation Service supported by the Globus Toolkit. As this currently supports delegation only to services within the same container, it may have to be deployed in multiple containers.

Status and Evaluation

A grid portal has been developed for Geode and made available to occupational data researchers. The portal acts as a gateway to occupational data resources. Users of the portal can discover and use any occupational information resources that have been exposed to the Geode service. Users are also able to add their own occupational information resources to the service. Using existing occupational data with Geode has minimal requirements, involving a very limited selection of identifying records (such as details of the county, time period and occupational units to which the information is relevant). Additional details about the occupational information may later be added by authorized users within the Geode virtual organization.

A specialized grid service is supported for linking survey data to occupational information resources. The linking mechanism offers a highly attractive opportunity for social scientists who have previously been required to undertake an extended series of data manipulation tasks on their own computers in order to achieve the same linkage.

Encouraging the use of metadata has improved the accessibility and consistency of occupational datasets. Searching occupational data is now more rigorous. Virtualization of data resources has made a wider range of datasets available to the community, without requiring users to have specialist technical knowledge. In addition, collaboration among social scientists has been aided through the creation of virtual organizations. Geode has provided valuable facilities for linking surveys and occupational classifications. The gain to researchers is greater productivity in their analyses.

Although the Geode work has focused on occupational data analysis, the design allows for other disciplines that require linking of statistical data. For example, the authors also envisage the same approach being useful for census data.

FUTURE TRENDS

Grid computing was initially driven by the needs of hard sciences such as astronomy, biology and physics. Such applications are characterized by large volumes of data, highly complex calculations, and distributed collaborations. As grid computing matures, its use is being extended to other communities. For example, this article has illustrated how grid computing can benefit social scientists. Less technical subjects, such as business and economics, are also likely to gain from the grid. Large-scale computing may not be required in such cases. However, considerable advantages can still be obtained in other ways such as the ease of creating virtual organizations, single sign-on, and secure access.

Although grid computing shows considerable promise, it is still far from being a 'push button' technology. End users need grid solutions that can be deployed without extensive IT support. As the underlying computer science is developed, easier to use packages will undoubtedly be made available.

Currently, many grid applications are available only for Unix variants. As Microsoft Windows is the dominant platform in computing, this situation is very likely to change as grid computing moves away from specialized scientific applications and into mainstream usage.

CONCLUSION

The origins of grid computing lie in networked computing, distributed computing and parallel computing. The basic features of grid computing have been introduced. The grid offers important advantages such as virtual organizations, portals, single sign-on, resource abstraction and virtualization, security and authorization, distributed and parallel computing.

As a concrete example, the use of occupational data in social science has illustrated how grid computing can be deployed in practice. It has been seen how this allows uniform and secure use of occupational data by a distributed community of social scientists. Appropriate use of metadata now supports consistent access to data in different formats.

REFERENCES

Abernethy, R., Morin, R., Chahin, J., & Morin, R. C. (Eds.). (1999). *COM/DCOM unleashed*. Indianapolis, IN: Sams Publishing.

Berman, F., Fox, G. C., & Hey, A. J. G. (Eds.). (2003). *Grid computing: Making the global infrastructure a reality*. Chichester, UK: John Wiley and Sons.

Birrell, A. D. (1984, September). *Secure communication using remote procedure calls* (Tech. Rep. No. CSL-84-2). California: Xerox Corporation.

Bolton, F. (2001). *Pure CORBA*. Indianapolis, IN: Sams Publishing.

Cambridge Social Interaction and Stratification Scale (2007). Retrieved November 2007, from http://www.camsis.stir.ac.uk

CERN (2007). *Large hadron collider*. Retrieved November 2007, from http://www.cern.ch/lhc

Commodity Grid Kits (2007). Retrieved November 2007, from http://www.globus.org/cog

Data Documentation Initiative (2007). Retrieved November 2007, from http://www.icpsr.umich.edu/DDI

Downing, T. B. (1998). *The Java RMI: Remote method invocation*. Chichester, UK: John Wiley and Sons.

Dupuy, F., Nilsson, G., & Inoue, Y. (1995, November). The TINA consortium: Towards networking telecommunications information services. *IEEE Communications Magazine*, 78–83.

Elias, P., & Birch, M. (2006). The Review of ISCO88: A European Perspective. Paper presented at *European Socio-Economic Validation Conference*. Lisbon, Portugal.

Foster, I. (2002, July). What is the grid? A three point checklist. *Grid Today*, *1*(6).

Foster, I., Kesselman, C., Nick, J. M., & Tuecke, S. (2002). Grid services for distributed system integration. *Supercomputer Applications*, *35*(6).

Francis, T., Herness, E., Jr., R. H., Knutson, J., Rochat, K., & Vignola, C. (2002). *Professional IBM Websphere application server*. Chichester, UK: John Wiley and Sons.

Ganzeboom H. B. G. (2006). Tools for deriving status measures from ISK0-88 and ISCO-68. Retrieved October 2006, from http://home.fsw.vu.nl/~ganzeboom/PISA

Graham, S., Marmakar, A., Mischinsky, J., Robinson, I. & Sedukhin, I. (Eds.). (2006). *Web services resource*. Billerica, MA: Organization for the Advancement of Structured Information Standards.

Grid Enabled Occupational Data Environment (2007). Retrieved November 2007, from http://www.geode.stir.ac.uk

GridSphere (2007). Retrieved November 2007, from http://www.gridsphere.org

Global Grid Forum (2007). Retrieved November 2007, from http://www.ggf.org

ISO/IEC. (1988). *Information processing systems—Basic remote procedure call (RPC) using OSI remote operations*. Geneva, Switzerland: International Organization for Standardization.

ISO/IEC. (1994). *Information processing systems—Open systems interconnection—Basic reference model*. Geneva, Switzerland: International Organization for Standardization.

ISO/IEC. (1995). *Information processing systems—Open distributed processing—Basic reference model*. Geneva, Switzerland: International Organization for Standardization.

ITU. (2000). *Intelligent network—Q.120x series intelligent network recommendation structure*. Geneva, Switzerland: International Telecommunications Union.

Lambert, P. S., & Prandy, K. (2006) *Camsis project*. Retrieved October 2006, from http://www.camsis.stir.ac.uk

Mogha, R., & Bhargava, R. (Eds.). (2002). *Sun ONE studio programming*. Chichester, UK: John Wiley and Sons.

Mountjoy, J., & Chugh, A. (2004). *Weblogic: The definitive guide*. Sebastopol, CA: O'Reilly and Associates.

MyProxy (2007). Retrieved November 2007, from http://grid.ncsa.uiuc.edu/myproxy

OASIS. (2002). *Universal description, discovery and integration*. Billerica, MA: Organization for The Advancement of Structured Information Standards.

OGSA-DAI (2007). *Open grid services architecture-database access and integration*. Retrieved November 2007, from http://www.ogsadai.org.uk

Platt, D. S. (Ed.). (2003). *Introducing Microsoft. NET*. Redmond, WA: Microsoft Press.

Prandy, K., & Lambert, P. S. (2003). Marriage, social distance and the social space: An alternative derivation and validation of the Cambridge scale. *Sociology, 37*(3), 397–411.

Sotomayor, B. & Childers, L. (2006). *Globus toolkit 4: Programming Java services*. San Francisco: Morgan Kaufmann.

SPSS (2007). *Statistical package for the social sciences*. Retrieved November 2007, from http://www.spss.com

Stata (2007). Retrieved November 2007, from http://www.stata.com

World Wide Web Consortium. (2001). *Web services description language (WSDL)*. Geneva, Switzerland: World Wide Web Consortium.

World Wide Web Consortium. (2003). *Simple object access protocol (SOAP)*. Geneva, Switzerland: World Wide Web Consortium.

KEY TERMS

Authentication: Validating the identity of a user requiring access to a resource, often achieved in the grid community through the public key infrastructure.

Authorization: Validating and enforcing the privileges of a user requiring access to a resource, often achieved through access control lists or security policies.

Credential: A digital certificate that identifies a user.

Credential Delegation: The means whereby users can pass their credentials to proxies.

Data Abstraction: A technique for providing a database-independent view on data.

Data Documentation Initiative (DDI): A set of XML schemas for annotating social science data.

Distributed Computing: Cooperative information processing by networked computers.

Globus Toolkit (GT): A widely used collection of fundamental tools for grid computing, developed by the Global Grid Alliance.

Grid Computing: A software technology that coordinates distributed resources not subject to central control, using standard protocols and interfaces to meet required levels of service provision.

Grid Service: A web service that conforms to grid computing standards.

Micro-Social Survey: A social science dataset that covers detailed information about numerous individual respondents, often dealing with occupation as an important aspect.

Networked Computing: Information exchange among networked computers.

Occupational Data Analysis: A technique used in social science to link survey data to occupational information.

Occupational Information: Data about different occupational positions, typically obtained through summary (aggregation) analysis of individuals**.**

Occupational Position: Data that describes the particular occupation of an individual or subject, such as a job title.

Open Grid Services Architecture (OGSA): A high-level architecture for grid computing.

Open Distributed Processing (ODP): A family of standards for distributed computing.

Open Systems Interconnection (OSI): A family of standards for networked computing.

Parallel Computing: Data processing by a cluster of closely coupled processors.

Portal: A web-based interface that provides a uniform interface to a range of networked services.

Proxy: A third party or service that acts on behalf of the user with their authority and identity.

Proxy Certificate: The credentials of a user passed to a proxy.

Public Key Infrastructure (PKI): A secure means of sending data over a public network, encrypting it on transmission with the recipient's public key and decrypting it on arrival with the recipient's private key.

Resource: A networked facility that provides access to data or computation.

Single Sign-On: The means whereby a single authentication procedure allows access to networked resources.

Simple Object Access Protocol (SOAP): A standard for conveying structured information between web services.

Universal Description, Discovery and Integration (UDDI): A standard for registering and searching web services in a repository.

Virtual Organization (VO): An organization that transcends institutional boundaries, coming together to collaborate for a specified purpose on a short-lived or long-lived basis.

Web Service: A networked service that allows cooperative information processing with other similar services.

Web Services Description Language (WSDL): A standard for defining the interface to a web service.

Web Services Resource Framework (WSRF): A family of standards for exposing long-lived resources used by web services.