USING GRID TECHNOLOGIES TO SUPPORT INTELLIGENCE BY PROVIDING A HIGHER LEVEL OF ACCURACY IN FINANCIAL DECISION MAKING

Nik Bessis

University of Bedfordshire, Department of Computing and Information Systems, Institute for Research in Applicable Computing, Park Square, Luton, LU1 3JU, UK

Tim French

University of Bedfordshire, Department of Computing and Information Systems, Institute for Research in Applicable Computing, Park Square, Luton, LU1 3JU, UK

Zhihua Lai

University of Bedfordshire, Department of Computing and Information Systems, Institute for Research in Applicable Computing, Park Square, Luton, LU1 3JU, UK

ABSTRACT

This paper aims to offer proof of concept in the applicability of grid related technologies for supporting intelligence in financial decision-making. We do this, by discussing Simon's well-known decision-making phases model "intelligence-design-choice" alongside with the concept of "bounded rationality". Next, we present our approach for producing calculations of higher level of accuracy using grid technologies. This will serve as the underlined point of reference throughout this paper. To further the explanation of the concepts and practices associated with using grid technologies to support intelligence in financial decision-making, a mini-case is employed incorporating scenarios. Specifically, Pension Funds as our typical real-world case study are used to demonstrate how they could benefit by using grid technologies when seeking to match their future liabilities against their income streams. Finally, we conclude by discussing the implications of using grid technologies to assist intelligence in the financial decision-making sector.

KEYWORDS

Grid technologies, intelligence, decision-making, high-level accuracy, financial sector

1. INTRODUCTION

This section provides a brief grounding in intelligence informed decision-making technologies, their application and integration within the modern organisations.

Scott-Morton first articulated the concepts of Decision Support Systems (DSS) in the early 70s under the general term of Management Support Systems (MSS). Further works on "bounded rationality" from Simon (1977) and "classification types of DSS" from Keen and Scott-Morton (1978), Alter (1980), Holsapple and Whinston (1996) have led us to understand that DSS is a set of concepts associated with supporting the decision making process via the use of appropriate resources. These resources may include but are not limited to users, data, models, software, hardware, or even instruments such as satellites, seismographers, detectors and PDAs. For example think of an emergency situation caused by an earthquake. The emergency management team will be required to make real-time intelligent decisions and act accordingly by assessing multiple dispersed resources (Asimakopoulou et al., 2005). This particular decision making process will require team working and collaboration from a number of dispersed decision makers whose decisions may be depended on each other's interactions. Resource integration at that level will support decision makers since it

will allow them to view satellite images of the affected area, observe seismic activity, forecast, simulate and run "what if" scenarios, collaborate with experts and the authorities. This will assist decision makers to prioritize and ultimately make decisions, which will be disseminated to available rescue teams who will take then care of the operational tasks. This dissemination may typically involve a server broadcasting decisions to heterogeneous mobile devices such as PDAs.

Computer-based developments over the last four decades have facilitated decision makers with numerous tools to support intelligence in operational, tactical and/or strategic level of enquiries within the environment of an organization. One of the purposes of these technologies is to provide decision makers with a holistic view hence, the ability to analyze data derived from a collection of multiple dispersed and potentially heterogeneous sources (Han, 2000). One of the challenges for such facilitation is the method of data integration, which aims to provide seamless and flexible access to information from multiple autonomous, distributed and heterogeneous data sources through a query interface (Ullman, 1997; Calvanese, et al., 1998; Levy, 2000).

On the other hand, the volume of the datasets is typically measured in terabytes and will soon reach petabytes (Antonioletti, et al., 2005). These datasets are variably geographically distributed and their complexity is ever increasing. That is to say, that the extraction of meaningful knowledge requires more and more computing resources. The communities of users that need to access and analyze this data are often large and geographically distributed. The combination of large dataset size, geographic distribution of users and resources, and computationally intensive analysis results in complex and stringent performance demands that, until recently, have not been satisfied by any existing computational and data management infrastructure.

In tackling these problems, the latest research in relation to networking and resource integration have resulted in the new concept of grid technologies, a term originally coined by Foster in 1995. Grid computing has been described as "the infrastructure and set of protocols that enable the integrated, collaborative use of distributed heterogeneous resources including high-end computers, networks, databases, and scientific instruments owned and managed by multiple organizations, referred to as Virtual Organizations" (Foster, 2002). A Virtual Organization (VO) is formed when different organizations come together to share resources and collaborate in order to achieve a common goal (Foster et al., 2002). In this context, resource sharing is 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 occurs. It is therefore anticipated that grid technologies will facilitate intelligence informed decision making in a way that managers and their teams will be able to carry out tasks of increased complexity more effectively and efficiently in the form of one or many interconnected, separable or inseparable Virtual Organizations (Brezany et al., 2003; Bessis & Wells, 2005; Bessis et al., 2007).

Hence the paper's main goals are twofold. Firstly, to discuss how grid technologies can serve as the vehicle to empower intelligence in decision making. We do this, by discussing Simon's well-known decision-making phases model "intelligence-design-choice" alongside with the concept of "bounded rationality". Secondly, to stimulate thinking towards a better understanding of how grid technologies could be used in the financial sector. To further the explanation of the concepts and practices associated with using grid technologies to support intelligence in financial decision-making, a mini-case is employed incorporating scenarios. We conclude by discussing the implications of using grid technologies to assist intelligence in financial decision-making.

1.1 The Grid Concept and its Commercial Exploitation

The concept of grid computing has emerged as an important research area differentiated from open systems, clusters and distributed computing. That is to say, open systems such as Unix, Windows or Linux servers, remove dependencies on proprietary hardware and operating systems, but in most instances are used in isolation. Each deployed application has its own set of servers purchased for a particular purpose within the enterprise. Multiple applications rarely share common servers, resulting in silos of statically linked applications and servers. This configuration results in poor server utilization. In contrast, "the grid builds upon open source architectures and addresses the removal of silos within a connected enterprise" (Xu et al., 2004).

Unlike conventional distributed systems, which are focused on communication between devices and resources, grid computing takes advantage of computers connected to a network making it possible to

compute and to share data resources. Unlike clusters, which have a single administration and are generally geographically localized, grids have multiple administrators and are usually dispersed over a wide area. But most importantly, clusters have a static architecture; whilst grids are fluid and dynamic with resources entering and leaving. The added value that grid computing provides as compared to conventional distributed systems lies in the inherent ability of the grid to dynamically orchestrate large scale distributed computational resources across Virtual Organizations (VOs), so as to leverage maximal computational power towards the solution of a particular problem. More specifically, the grid can allocate and re-schedule resources dynamically in real-time according to the availability or non-availability of optimal solution paths and computational resources. Should a resource become compromised, untrustworthy or simply prove to be unreliable, then dynamic re-routing and re-scheduling capabilities can be used to ensure that the quality of service is not compromised. These advanced features that are integral to grid computing are rarely to be found in large scale conventional distributed networks, particularly those that need to co-operate and coordinate dynamically across organizational and geographical boundaries. Hence, it is the ability of grid communities to orchestrate their activities at the VO level and the service level dynamically (without the need to consider platform dependant features) that characterizes grid solutions as distinct from large-scale conventional distributed computer networks.

In terms of standards, grids share the same protocols with Web Services (XML, WSDL, SOAP, UDDI). This often serves to confuse as to exactly what the differences between the two actually are. The aim of Web Services is to provide a service-oriented approach to distributed computing issues, whereas grid arises from an object-oriented approach. That is to say, Web Services typically provide stateless, persistent services whereas grids provide state-full, transient instances of objects. The recently emerged Open Grid Services Architecture (OGSA) is an informational specification that aims to define a common, standard and open architecture for grid-based applications. The goal of OGSA is to standardize almost all the services that a grid application may use, for example job and resource management services, communications and security. OGSA specifies a Service-Oriented Architecture (SOA) for the grid that realizes a model of a computing system as a set of distributed computing patterns realized using Web Services as the underlying technology. An important merit of this model is that all components of the environment can be virtualized. It is the virtualization of grid services that underpins the ability to map common service semantic behavior seamlessly on to native platform facilities. These particular characteristics extend the functionality offered by Web Services and other conventional open systems. In turn, the OGSA standard defines service interfaces and identifies the protocols for invoking these services. The potential range of OGSA services are vast and currently include data and information services, resource and service management, and core services such as name resolution and discovery, service domains, security, policy, messaging, queuing, logging, events, metering and accounting. OGSA-DAI (Data, Access and Integration) provides a means for users to gridenable their data resources. OGSA-DAI is a middleware that allows data resources to be accessed via Web Services.

Overall, global grid initiatives initially tended to focus on the needs of the UK scientific community (Fox & Walker, 2003) in "an initiative collectively known as E-Science", but in the future, "the business community is expected to increasingly benefit too: grid computing is expected to become a mainstream business-enterprise topology during the rest of the current decade" (Castrol-Leon, 2005). Typical initial application areas have included E-Science data-grids in which University's share their resources across a grid so as to process vast quantities of data involved in areas such as molecular modelling, climate change modelling, and financial and economic modelling.

1.2 Early Adoption of the Grid by 'blue-chip' Banking Industry

The grid is being utilized internally and externally by business organizations to aid their financial decision making and modelling. A number of major banks in the UK, USA and Europe have been 'early adopters' (1999-) of internal and external grid computing models so as to better utilize underused computational nodes in the context of financial services modelling and decision making. For example, the chairman of the influential Landesbank Baden Wurtenburg (LBBW) has expressed that the grid and financial services industry are a marriage made in heaven: "the banking and finance industries are predestined from grid computing solutions. Our business processes can be parallelized and thus made faster and more efficient than ever before" (Platform, 2005). That is to say, by seeking to use underused resources as part of a grid (where

the VOs are typically comprise different internal departments), these organizations hope to create and run advanced simulations and otherwise distribute increasingly data-intensive computational tasks across their existing computational nodes without the need to purchase additional or dedicated resources. Interested readers are refereed to two reports (Davidson, 2002; Carbonnier, 2005) in which grid projects within JP Morgan and Chase Manhatten Banks respectively are described in some detail and which may be viewed as being fairly typical in illustrating the rationale behind early adoption of grid applications within international banking. The financial imperative for wider commercial use of the grid is now undeniable and has recently been articulated as follows:

"Grid computing is not just about an asset change in enterprise environments; it is about supporting a new business model, since there is no killer application for grids. The key question for Finance Directors and CFOs is how to break out of the cycle of asset acquisition and into a capacity service provision model in order to save money against a new budget system. The benefits of grid computing are about helping to bring CAPEX (capital expenditure – i.e. the cost of the network, infrastructure and terminals) and OPEX (operating expenditure – i.e. the cost of the network running) down to acceptable levels. The grid-based payper-use/utility model is attractive because it can transfer cost from a CAPEX to an OPEX model, but we don't believe it will ever be an 'all or nothing' situation for users." (Fellows, 2005).

However, our main goal here is to extend the current application of grid technologies in the financial sector. That is to say, we aim to proof concept that grid technologies can be used to produce results of a higher level accuracy by distributing computational tasks across existing and underused dispersed computational nodes without the need to purchase additional or dedicated resources.

2. ENABLING INTELLIGENCE IN FINANCIAL DECISION-MAKING USING GRID TECHNOLOGIES

The objective of this section is to discuss and exemplify the potential of how grid technologies within a dynamically changing environment can assist intelligence in financial decision-making.

We do this, by discussing Simon's well known decision making phases "intelligence-design-choice" alongside with the concept of bounded rationality. Next, we are going to proof concept our approach for producing calculations of higher level of accuracy using grid technologies. This will serve as the underlined point of reference throughout this paper. Finally, we are going to describe a typical financial services minicase, which alongside with our approach will clearly demonstrate how grid technologies can assist intelligence in financial decision making.

2.1 Simon's "intelligence-design-choice" Decision Making Phases

Simon's (1977) systematic decision-making process includes the three-phases of "intelligence-design-choice". In the first phase of intelligence, someone must clearly define the problem by identifying symptoms and examining the reality. Specifically, the first phase begins with the identification of the organizational goal and objective that is for example, to provide an accurate service to their customers.

Once these have been defined, the organization must move to the design phase that is the second phase of Simon's systematic decision-making process. This phase involves finding or developing and analyzing possible courses of action towards the identification of possible solutions against the identified problem space. This requires an operation under the process-oriented decision-making thinking (Keen and Scott-Morton, 1978) and Simon's (1977) "bounded rationality" theory. That is to say, the organization must appreciate that despite the attractiveness of optimization as a decision-making strategy, its practical application is problematic. This is due to the fact that it is not feasible to attempt to search for every possible alternative for a given decision. Simon exemplified this by defining the term of "problem space". A problem space represents a boundary of an identified problem and contains all possible solutions to that problem: optimal, excellent, very good, acceptable, bad solutions and so on.

The rational model of decision-making suggests that the decision maker would seek out and test each of the solutions found in the domain of the problem space until all solutions are tested and compared. At that point, the best solution will be known and identified. However, what really happens is that the decision maker actually simplifies reality since reality is too large to be handled by human cognitive limitations. This narrows the problem space and clearly leads decision-maker to attempt to search within the actual problem space that is far smaller than the reality. In other words, the attempted problem space is incomplete and refers to the actual problem search space. Thus, the decision maker will most likely not choose the optimal solution because the narrowed search makes it improbable that the best solution will ever be encountered. The approach will lead the decision maker to settle for a satisfactory solution rather than searching for the best possible solution.

If however, the organization had access to more resources, decision makers will most likely choose a better solution because the extended search of the actual problem search space increases the possibility that a better solution will be encountered. Figure 1 (re-produced over page) illustrates such intelligence and choice decisions.

At this stage, the organization must move to the choice that is the third and last phase of Simon's (1977) systematic decision-making process. In other words, the organization needs to make a decision based on the alternatives derived from the previous phase. The organization has three options to choose from: take the risk and do nothing; purchase additional resources; or enter into a grid partnership. The latter option supports the idea that is to utilize the spare-capacity of available dispersed computer based resources in real-time, on an on-demand basis. Grid partners will then orchestrate the optimal workflow (scalability) required between them, making best use of any spare capacity available, so as to process and analyze what is originally required for the organization.

2.2 Producing Calculations of a Higher Level Accuracy using Grid Technologies

Our approach assumes that we would like to perform a calculation using a workstation that is 32-bit capable only (PC-1). By default, the PC-1 can handle maximum 32-bit unsigned integers. That is to say, the PC-1 search space ranges from 0 to $2^{32} - 1 = 4,294,967,295$. Performing the same calculation using a workstation that is 64-bit capable (PC-2) would enlarge the search space from 0 to $2^{64} - 1 = 18,446,744,073,709,551,615$ leading to more accurate results.

To exemplify it, suppose that we would like to find the approximation of π (3.14159265358979323846...) written in a finite fraction by using an exhaustive search algorithm to find the proximity of A/B. To do this, we would have to calculate the absolute value of $abs(A/B - \pi)$.

Using PC-1, the actual search space for both A and B ranges from 1 to 4,294,967,295 while using PC-2, the actual search space for both A and B is enlarged as it ranges from 1 to 18,446,744,073,709,551,615. Specifically, using the PC-1 resource the absolute values for A and B [$abs(A/B - \pi)$] will be limited to no more than 10⁻⁶ decimal places. That is to say, $abs(A/B - \pi) < 10^{-6}$. On this basis, the proximity of A/B will be limited to the identification of the set of integers (355, 113) leading to the absolute value of $abs(355/113 - \pi) < 0.000001$.

Similarly, if the PC-2 resource were to be used, the absolute values for A and B [$abs(A/B - \pi)$] would be limited to no more than 10⁻¹⁰. That is to say, $abs(A/B - \pi) < 10^{-10}$. In this case, the absolute values for A and B [$abs(A/B - \pi)$] would be much more accurate, as PC-2 resource would search in a larger set of unsigned integers producing additional decimal places. On this basis, the PC-2 resource will identify a set of higher proximity integers for A/B (6,283,185,307/2,000,000,000=3.1415926535). This clearly leads to a more accurate absolute value of $abs(6,283,185,307/2,000,000,000 - \pi) < 0.000000001$.

Overall, if PC-1 and PC-2 where combined in a grid environment, these would produce more accurate results in a more efficient manner. This is because PC-1 could have been assigned the task to handle the range between 1 and 4,294,967,295 (the search space of SS-1) whilst PC-2 could have been assigned the task to handle the range between 4,294,967,295 and 18,446,744,073,709,551,615 (the search space of SS-2). The combined search spaces (SS-1 and SS-2) for both PC-1 and PC-2 providing results of higher-level of accuracy can be seen in figure 1.

There are a number of middleware solutions supporting the co-ordination and allocation of jobs to be done in a dispersed environment including OGSA-DAI, OGSA-DAI DQP, Condor-G, Globus Toolkit and Unicore. With this in mind, we are going next to explain a mini-case in which a higher-level accuracy achieved is best desirable.



Figure 1. Producing higher level accuracy calculations by enlarging the search space using grid technologies

2.3 Financial Services Mini-case

Pension Funds are used as our typical financial services decision making scenario. Pension Funds are constantly seeking to match their future liabilities to an income stream. One of the financial instruments they need to continuously monitor on a real-time worldwide basis are undated (so called perpetual) government or undated corporate bonds. Specifically, there is a need to constantly monitor the gross yield (y%) of such instruments to adjust in real time for currency movements in order to match their future predicted liabilities over 20-40+ year time periods so as to match current and predicted income streams.

To calculate the gross yield (y%) of any particular bond (there are more than 10, 000 such bonds available at any one time globally) great accuracy is necessary since the sums typically involved in buying, "stripping", swaps or selling such bonds often exceed $\pounds 100+$ million for each and every transaction. In addition of course the conversion rate into UK sterling equivalent needs to be taken into account with equal precision.

Suppose for example that Pension Fund "ABC" has calculated that for February 2007 a purchase of £150 million pounds of perpetual bonds is required at an indicative gross y% (spread) of between 6.956546547654% and 6.9666547564%. To calculate likely candidate global bonds requires a real time search to be made every 30 minutes or so of the current sterling equivalent price and hence y% of every available bond (n= 10000+) is simply calculated as i/p (where i = "coupon" with a fixed interest rate and p = current price in UK pounds).

For example, undated "UK War loan" coupon, i = 4.5% and the current price (as at 1000 hrs GMT Friday 23rd Feb 2007), p = 64.02 pounds per 100 pounds issue price of nominally issued stock. On this basis, the proximity of $y\% = i/p \rightarrow y\% = 4.45/64.02 = 0.069509528272....$

Thus it would appear that this particular stock potentially meets the criteria (many stocks will do so), the "best" stock is the one that matches as closely as possible but does not exceed the highest value within the indicative range (i.e. 6.9565465....). For the purposes of the initial selection y%, the notion of "best" stock ignores: costs of purchase and sales as well as accrued interest due i.e. the so-called "dirty" price of the bond. The "clean" (idealized) price and hence y% is used for initial selection. In practice further calculations will be necessary to adjust for interest due or accrued interest refining the search space.

Specifically, the above bond's clean price needs to be adjusted because the annual coupon is paid in four quarterly installments. The next installment (4.45/4) is due on March 1st. Hence the present price is ex dividend which means the buyer won't receive the next installment of interest due on March 1st. Thus the seller needs to pay extra to the buyer to account for the interest lost. This means the "dirty" price of the bond is adjusted "lower" so as to give the buyer a slightly higher real y% to compensate them for the dividend (interest payment) due on March 1st which they will not be entitled to if they buy the bond now during the ex dividend period.

This aforementioned highly simplified scenario clearly demonstrates that if Pension Funds enter into a grid partnership they expected to benefit greatly by achieving greater rich, scalability and accuracy levels when seeking to match their future liabilities against income streams. This is due to the fact that Pension Funds would have an enlarged search space to look at and therefore, this could make possible for a better solution to be encountered. More generally, such grids are leveraged by such institutions so as to more

accurately price a wide range of financial instruments. Typically, these include bonds, derivatives, interest rates swaps and other financial hedging instruments. Thus grids are seen to be a useful measure for risk reduction within an extended search space in their financial decision making.

3. CONCLUSION

This paper has endorsed the logic that the concepts and practices associated with grid related technologies could assist decision managers in making intelligence informed decisions. It is anticipated that the decision to use grid technologies will unfold new opportunities as it will enlarge the actual search space boundaries within the term of "problem space" as described by Simon (1977). By default, a problem space represents the boundary of an identified problem and contains all possible solutions to that problem: optimal, excellent, very good, acceptable, bad solutions and so on. By searching in a narrow space, the decision maker will most likely not choose an optimal solution because the narrowed search of the actual problem search space makes it improbable that the best solution will ever be encountered. It might then still be possible not identify the optimal solution but it is more likely to increase the opportunities for a better solution to be encountered. Overall, it will facilitate methods towards normative thinking as required for a better quality of service.

Clearly the grid potentially vastly increases the size and complexity of the problem spaces that can realistically be addressed not only by financial related businesses but also by other types of organizations. Problems that have hitherto been regarded as being intractable either because of the size of the data-sets needed, their distributed nature or the sheer complexity of the multi-dimensional analysis required can now be re-examined. Within e-Science these problem spaces encompass traditional scientific domains such as nuclear physics but now also typically include areas such as climate change, where vast quantities of data and simulations requiring multi-dimensional analysis are needed.

Within the business community large banks have been amongst the first to exploit the enhanced power of the grid to leverage extra value from vast legacy systems. On this basis, we attempted to produce calculations of higher level of accuracy using grid technologies. We exemplified this by finding the approximation of (3.14159265358979323846...) written in a finite fraction by using an exhaustive search algorithm to find the proximity of A/B. The higher level of accuracy produced served as the underlined point of reference throughout the paper. Finally, we described Pension Funds as a typical financial services mini-case, which alongside with our approach clearly demonstrated how grid technologies could assist intelligence in financial decision-making. As it has been shown through our illustrative case study, financial related organizations are able to address previously intractable problems and to leverage competitive advantage from grid computing. This is only the beginning – decision makers will soon be able to address or re-address complex multi-dimensional problems within their businesses using grid solutions as their standard or normative preferred tool. Thus, the grid should not be seen as being merely a tool of scientists or academicians but rather as a new and powerful business decision support tool, having real cutting edge potential to solve business problems and enhance competitive advantage.

However, for the power of the grid to be fully realized by business decision makers, a risk assessment is required. For example, trust issues remain one of many risk factors that need to be considered before grid computing is adopted. Since the grid by definition involves the creation of virtual partnerships between VOs, like any partnership there are risks as well as rewards. In the future, grid computing will only be seen to serve and support decision makers if these risks are properly assessed and accommodated. Like all enabling technologies, investment needs to be made in properly harnessing the power of the grid without exposing the business to undue risk. This is one of the challenges that still remain to be solved if grid computing is indeed to become a normative tool of the business community.

REFERENCES

Alter, S. L. (1980). Decision Support Systems: Current Practices and Continuing Challenges. Reading, MA: Addison-Wesley.

Antonioletti, M., et al. (2005). The Design and Implementation of Grid Database Services in OGSA-DAI. Concurrency and Computation: Practice and Experience, 7 (2-4), 357-376.

- Asimakopoulou, E., et al. (2005). Studies of Emergency Management Procedures in Greece, Italy and the UK. In, Procs. 3rd International Conference on Construction in the 21st century, Advancing Engineering, Management & Technology, 15-17 September 2005, Athens, Greece.
- Bessis, N. and Wells, J. (2005). Grid Technologies Revive the Basic IT Infrastructure. Information Systems Unplugged: Developing Relevant Research, In, Procs.UKAIS 2005:10th Int. Conference in Information Systems, March 22-24, Newcastle, UK.
- Bessis, N. et al. (2007) "Using grid technology for data sharing to support intelligence in decision making", in Managing Strategic Intelligence: Techniques and Technologies, M. Xu, Ed. Idea Group Publishing Inc., (In Press).
- Brezany, P., et al. (2003). Towards an Open Service Architecture for Data Mining on the Grid. In, Procs. DEXA Workshops 2003, (pp. 524-528).
- Calvanese, et al. (1998). Information Integration: Conceptual Modelling and Reasoning Support. In, Procs. CoopIS'98.
- Castro-Leon, E., and Munter, J. (2005) 'Grid Computing: Looking Forward Part 1: Technology Overview', Intel White Paper, Retrieved April 30, 2006, from http://www.intel.com/cd/ids/developer/asmona /eng/202676.htm
- Carbonnier, J. P. (2005). BNP Paribas Bolsters Grid Project. Retrieved July 30, 2006, from http://www.datasynapse.com/pdf/DWT_BNP_Paribas_Bolsters_Grid_Project_05_3005.pdf
- Davidson, C. (2002). JP Morgan unveils Compute Backbone Project, Waters Financial Intelligence, October 2002. Retrieved April 30, 2006, from http://www.platform.com/Company/Case.Studies /CS.JPMorgan.06.08.2005.htm
- Fellows, W. (2005). Enterprise grid computing to shape IT business models, Finance Director Europe, March 21,2005. Retrieved April 30, 2006, from http://www.financedirectoreurope.com /articles/451group.htm
- Foster, I. (2002). What is the Grid? A Three Point Checklist. Grid Today, July 22, 2002: 1(6) Retrieved April 30, 2006, from http://www.gridtoday.com/ 02/0722/100136.html
- Foster, I. Et al. (2002). The physiology of the grid: an open grid services architecture for distributed systems integration. Retrieved April 21, 2007, from http://www.globus.org/alliance/publications/papers/ogsa.pdf
- Fox, G., and Walker, D. (2003). E-Science gap Analysis, Cardiff University. Retrieved April 30, 2006, from http://grids.ucs.indiana.edu/ptliupages/publications/Gap Analysis30June03v2.pdf
- Han, J. (2000). Data Mining: Concepts and Techniques, Morgan Kaufmann.
- Holsapple, C. W., and Whinston, A. B. (1996). Decision Support Systems: A Knowledge-Based Approach. St. Paul, MN: West Publishing.
- Keen, P. G. W. and Scott-Morton, M. S. (1978). Decision Support Systems: An Organisational Perspective. MA: Addison-Wesley.
- Levy, A. (2000). Logic-based Techniques in Data Integration. In Jack Minker (Ed.) Logic Based Artificial Intelligence. Kluwer Publishers.
- Platform Web site (2005). Retrieved April 30, 2006, from http://www.platform.com/ Company/Case.Studies/
- Simon, H. (1977). The New Science of Management Decision. Englewood Cliffs, NJ: Prentice Hall.
- Ullman, J. (1997). Information Integration using Logical Views. In, Procs.6th ICDT. (pp. 19-40), Delphi, Greece.
- Xu, M. et al. (2004). Service Virtualisation: Infrastructure and Applications. In, The Grid 2, Blueprint for a New Computing Infrastructure (Eds., Foster, I. and Kesselman C.). Elsevier, CA.