DATA CENTER

Making Sense of your Data

Prototype Applications for Semantic Modeling

Edmund W. Schuster, David L. Brock, Stuart J. Allen, and Pinaki Kar

The Data Center, Massachusetts Institute of Technology, Building 35, Room 234, Cambridge, MA 02139-4307, USA

ABSTRACT

The application of Semantic Modeling has great potential as a means of dealing with the vast streams of new data organizations will encounter in the future. Making Semantic Modeling a reality requires development of a new set of computer languages and protocols, termed M, to connect models to other models, data to models, and data to data. This article discusses initial prototype applications of M within industry. Specific examples include managing models and data for ERP systems, retail operations, and for a new area of agricultural modeling called harvest risk. The article also includes an overview of M and examples of the importance in connecting models and data in an open network. The authors are in the early stages of developing prototype examples of Semantic Modeling.

ABOUT THE AUTHORS

Edmund W. Schuster has held the appointment of Director, Affiliates Program in Logistics at the MIT Center for Transportation and Logistics and is currently working at The Data Center. His interests are the application of models to logistical and planning problems experienced in industry. He has a bachelor of science from The Ohio State University and a master in public administration from Gannon University with an emphasis in management science. Ed also attended the executive development program for physical distribution managers at the University of Tennessee and holds several professional certifications. Ed can be reached at Schuster@ed-w.info

David L. Brock is Principal Research Scientist at the Massachusetts Institute of Technology, and co-founder and a Director at the Auto-ID Center (now EPCGlobal, Inc. and Auto-ID Laboratories). The Center was an international research consortium formed as a partnership among more than 100 global companies and five leading research universities. David is also Assistant Research Professor of Surgery at Tufts University Medical School and Founder and Chief Technology Officer of endoVia Medical, Inc., a manufacturer of computer controlled medical devices. Dr. Brock holds bachelors' degrees in theoretical mathematics and mechanical engineering, as well as master and Ph.D. Degrees, from MIT. Dave can be reached at dlb@mit.edu

Stuart J. Allen is professor emeritus, Penn State – Erie, the Behrend College. He works on design of decision aids for application in manufacturing environments. His educational background includes a bachelor of science in mechanical engineering from the University of Wisconsin, a master of mechanical engineering from Seattle University, and a Ph.D. in engineering mechanics from the University of Minnesota. Stuart began his research career in the field of non-Newtonian fluid mechanics and has published over 50 journal articles in engineering and management science. He has also owned and operated three businesses in Wisconsin and New York State. Stuart can be reached at stuart99allen@yahoo.com

Pinaki Kar is currently an independent consultant working in the pharmaceutical industry on analysis and modeling to support strategic planning, business development, and marketing. He is interested in the application of operations research and statistical techniques for planning and decision support across a wide range of business issues. His experience spans multiple industries that include pharmaceutical, chemical, high-tech, and insurance. Pinaki's educational background includes a bachelor degree in mechanical engineering from the Indian Institute of Technology, Kanpur and a master's degree in logistics from MIT. Pinaki can be reached at pkar@alum.mit.edu

1.0 INTRODUCTION

We live in a world filled with data. Rapidly emerging technologies such as Auto-ID and the Electronic Product Code (EPC) combined with interactive sensor networks will create even larger data streams of greater complexity. By some estimates, the amount of data generated each year is growing as much as 40% to 60% for many organizations (Park 2004). All indications are that the pace of data generation is accelerating. EMC, a leading manufacturer of data storage devices recently noted "…companies are struggling to figure out how to turn all those bits and bytes from a liability into a competitive advantage (Park 2004)."

Dealing with increasing volumes of data will require new standards and information architectures to improve integration and communication between hardware, software, and business entities. However, the bigger question remains "How are we going to analyze and make sense of large volumes of data?"

A new research initiative at MIT called *The Data Center* addresses the important issue of generating value from data. The mission of The Data Center is to create innovative ways of making sense of data through new computer languages and protocols. *Semantic Modeling* provides a general description of these new technologies that will eventually connect data and various mathematical models together for improved analysis, business decision-making, and better day-to-day operations within large and small systems (Brock et al. 2005). Greater connectivity will spur new waves of productivity as managers learn to take advantage of the data within and outside of their organizations. This development represents the next logical step for the Internet.

The specific activities of The Data Center involve the research and development of a new computer language, called "M" that will achieve Semantic Modeling in practice. Designed as an open source code, M serves as the base system capable of linking models to other models, data to models, and data-to-data. All of these activities occur through an *Intelligent Modeling Network* that spans organizations. The conceptual design of M is such that network growth, in terms of adding additional models and data, occurs at minimal cost to end-users. This lowers the marginal cost of expansion, thus creating an incentive for active participation. A large intelligent modeling network will offer great value to industry.

In this article, we examine the first prototypes of Semantic Modeling currently under development at The Data Center. The first prototypes deal with Enterprise Resource Planning Systems (ERP), retail operations (lot sizing for short life cycle products), and agricultural modeling (harvest risk). Before discussing these prototypes in detail, we examine recent developments relating to Semantic Modeling, the underpinnings of the technology, and an overview of M, the new computer language designed to create an integrated modeling environment.

2.0 RECENT DEVELOPMENTS THAT SHOW THE FUTURE

Some important premises of Semantic Modeling already exhibit signs of practical implementation. These include 1) greater integration of data and information, 2) improved search capabilities, and 3) a relative approach to information and data organization.

Amazon has recently announced "A9" a new tool that can accomplish searches of information located on HTML web pages in addition to the text of thousands of books (Hof 2004). Eventually, A9 hopes to incorporate the ability to do even more specialized searches by accessing other proprietary databases. The Chief Executive of A9 has also commented that he wants to help curb information overload by allowing people to organize the web in a more personal way. With A9, each user can have their own view of information gathered by Internet searches.

All of the activities of A9 point toward greater integration, improved search capabilities and a relative approach to organizing information. However, other developments, not confined to Internet searches for information, also point toward greater integration.

In the US economy, there are billions of embedded microcontrollers in cars, traffic lights, and air conditioners that give specialized instructions for control based on sensing specific aspects of the environment. All of these microcontrollers act in total isolation from one another. Ember, a company located in Cambridge, MA, has developed a "mesh network" that holds the potential of allowing all of these microcontrollers to communicate with each other (Corcorian 2004). One practical application of mesh network technology involves the integration of home electrical systems without the need for hardwiring. Ember markets a device that allows a homeowner to turn off all electric lights through a single switch that does not require re-wiring. There are almost endless opportunities to establish communication connections for a wide variety of microcontrollers.

Just as Internet searches cannot reach all potentially useful information, and microcontrollers lack integrated communication within a network, the science and application of mathematical modeling often occurs in isolation with only occasional reporting at conferences and in academic journals. Often these means of sharing ideas are somewhat closed with little information reaching the business world. With the explosion of data streams, models provide a useful means to make sense of data. In the past, the lack of widespread use of models has been dependent on several factors including an inability to apply models to data quickly.

We see a new direction for sharing mathematical models through intelligent networks. Given the need in industry to analyze ever-increasing quantities of data, the path to building such a network will evolve during the next five years.

The general direction of the Internet appears to be following three distinct stages. First, the *web of information* currently in place today allows access to large amounts of information written in HTML through search engines such as Google. However, general issues involving semantics limit the ability to search this information in an efficient manner. The current popularity of search engine companies is evidence that there is a great deal of economic value in conducting a meaningful and quick Internet search.

Second, the *web of things* is just now becoming a tool for business. Through Auto-ID Technology, Radio Frequency Identification (RFID) tags can link objects to the Internet, or a private Intranet. This provides a foundation for a networked physical world that will change the way supply chains operate. The cornerstone of this new system is the Electronic Product Code (EPC) that allows for unique identification and mass serialization. The strength of the EPC is in identifying specific objects using a detailed serial number. In this regard, the EPC is well suited for identification, automation applications involving specific instructions, or the gathering of sense data like temperature, humidity, or vibration, for an object located within a manufacturing plant or a supply chain. However, the EPC is not useful when identification involves a semantic interpretation. There are many objects, such as models, for which a serial number is a poor descriptor. Because of this, the EPC is not a good tool for searching or linking models.

Finally, the eventual system needed to link models together is the *web of abstractions*. Semantic Modeling holds the key in developing new ways of describing abstract objects without using a serial numbering system like the EPC. If a semantically precise definition in machine understandable language exists for abstractions like models, then it is possible to build intelligent modeling networks to exchange and re-use model elements. With precise semantics, it becomes possible to link models together.

Likewise, it is also possible to establish precise semantics to describe data, allowing for the possibility of linking sets of data together and the matching of models to data automatically. New computer languages and protocols that deal with semantics will aid the free flow of models within an intelligent network.
The goal of semantic modeling is to create an interoperable computing environment that is organized in a distributed fashion across different computing platforms. Ultimately, this type of network will potentially link together models developed in different academic and industrial domains. All of this is possible with a precise semantic definition for abstractions like models. In the future, the definition of a model and the sharing of models though an intelligent network will become as important as the model itself.

3.0 THE UNDERPINNINGS OF SEMANTIC MODELING

The idea of defining elements of models is not new. Previous work has concentrated on the use of *Structured Modeling* to define elements for management science techniques (Geoffrion 1987; Geoffrion 1989). The following provides a brief description.

> "The theoretical foundation of structured modeling is formalized in Geoffrion, which presents a rigorous semantic framework that deliberately avoids committing to a representational formalism. The framework is 'semantic' because it casts every model as a system of definitions styled to capture semantic content. Ordinary mathematics, in contrast, typically leaves more of the meaning implicit. Twenty-eight definitions and eight propositions establish the notion of model structure at three levels of detail (so-called *elemental*, *generic*, and *modular* structure), the essential distinction between model *class* and model *instance*, certain related concepts and

> constructs, and basic theoretical properties. This framework has points in common with certain ideas found in the computer science literature on knowledge representation, programming language design, and semantic data modeling, but is designed specifically for modeling as practiced in MS/OR [management science/operations research] and related fields (Geoffrion 2004)."

This approach hints at the possibility of automatically combining models by using a Structured Modeling Language (SML). Others also employ various representation techniques to aid in the formulation of linear programming (LP) models (Welch 1987; Murphy et al. 1992). These efforts became part of proprietary software intended to ease the difficulty of formulating Linear Programming models. In all of these cases, the research occurred prior to the widespread use of the Internet and the existence of ample bandwidth. M takes advantage of these relatively new developments in computer science.

Other academic disciplines have also experimented with variants of Semantic Modeling in areas such as business process design. In one case, academic researchers have developed a large library of business processes in an attempt to build new organizations and to do benchmarking (Malone et al.). As part of this effort, the researchers also developed a definitional language for organizational processes and used a schema similar to an ontology as an aid in searching the library.

For many years, engineers have used something called a Bond Graphs to represent power flow (mechanical, electrical, hydraulic, thermal, chemical and magnetic) as a means of capturing the common energy structure of systems and to increase insight into engineering system behavior (Bond Graphs 2004). This method of linking different energy systems together with a common representation is similar to our efforts in Semantic Modeling. In addition, an interdisciplinary movement, initiated by the engineering community beginning in the 1960's, sought to establish General Systems where models from various academic disciplines, including the social sciences, could be shared with the goal of achieving new applications (General Systems 2004). More recently, the establishment of Math-Net, a global Internet-based information and communication system for mathematics, establishes many knowledge management structures that are similar to Semantic Modeling (Math-Net 2004).

Finally, several other groups of researchers have developed languages meant to do functions similar to Semantic Modeling. These include Simple HTML Ontology Language (SHOE), DARPA Agent Markup Language – Ontology (DAML-ONT), and Unified Problem-Solving Method (UPML) (Fensel et al. 2003). However, in no case did we find any evidence of initiatives to link models together or to establish improved semantics for models in a similar fashion to M.

4.0 "M" AND MODELING

Currently in the initial stages of research and development at The Data Center, Dave Brock is credited with the idea of creating M. Comprised of several important elements, the purpose of M is to serve as the fundamental language to link models and data together. Before discussing the prototype applications in detail, it is best to review the conceptual basics of M and the work completed to date concerning its development as a modeling language.

Fundamentally, M resembles peer-to-peer networking. In this type of architecture, computers running M can communicate and share models and data as equals. There are no servers. The important element in achieving peer-to-peer sharing is a new vision of how to attach a semantically precise definition to a model or data element, along with a series of computer languages and protocols to group, sort, interconnect, and match semantic definitions in a machine understandable way. With this approach, the relationships between a large group of models and data, all pre-assigned precise semantic definitions through M, provide a mapping of connections between models and other models, data and models, and the connections of data-to-data, all within a network. Deeper meaning arises through the visualization of these connections, either individually or group-to-group. FIGURE 1 provides a simple representation of model connections where the output of one model can become the input of another model.

FIGURE 1 – Connecting Models

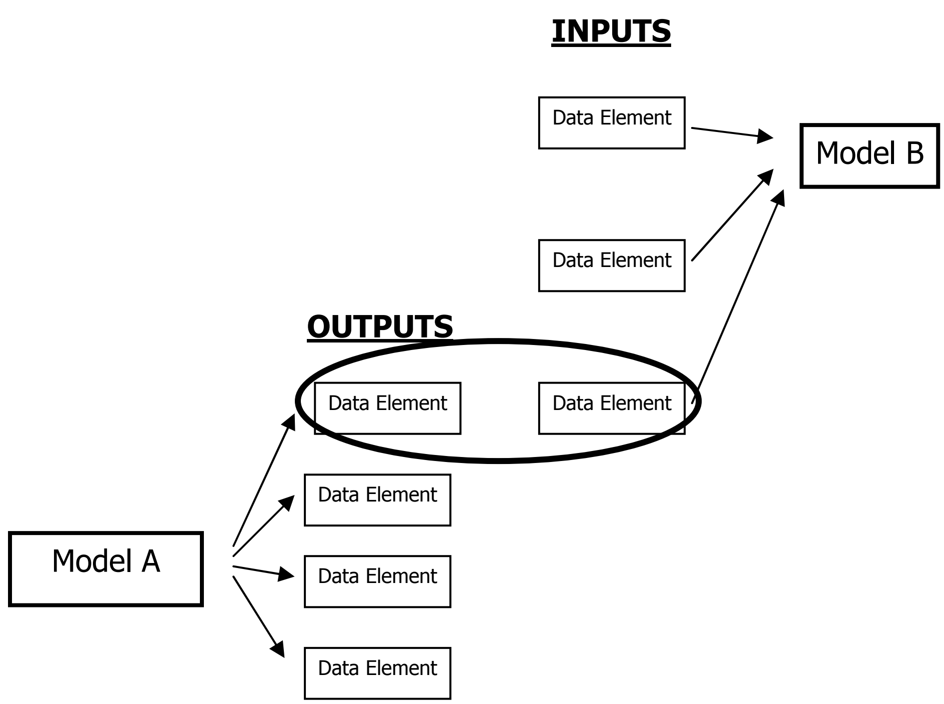

To achieve these connections, the structure of M must be comprised of two languages and two protocols. A comprehensive dictionary of words and various meanings is also included. The following provides brief definitions for each element of M.

Data Modeling Language (DML) is a semantic for describing modular, interoperable model components in terms of individual outputs, inputs and data elements.

Data Modeling Protocol (DMP), once a connection between models and data is established, the DMP coordinates the communication sequence between the computing machines that host models in terms of outputs and inputs.

Automated Control Language (ACL) establishes the connection between models and data based on DML (descriptor of inputs, outputs and data) and the ACP, which locates the appropriate connections.

Automated Control Protocol (ACP) helps model outputs and inputs locate one another within a network, even though the individual models may exist in different host systems and organizations. The ACP identifies potential connections and takes priority over the DMP, which is a coordinating activity after achieving connections through the ACL.

Dictionary a common resource containing words with multiple meanings. The dictionary will utilize established sources such as the Oxford English Dictionary, WordNet, and various specialty dictionaries from the medical field, operations, logistics and other disciplines.

With M, model inputs, outputs, and data elements are described through DML by using words from the dictionary to express a precise semantic. Because multiple words, akin to a phrase or simple sentence, best provide accurate descriptions of outputs and inputs for models and data elements, we envision the use of graphs to express syntax thus giving a precise semantic meaning.

The graphs produced through M to represent outputs, inputs and data elements will need to be of the form that operations, such as sorting, can be applied using computer code. The ACP helps to locate graphs with commonalities that are resident in a network. These commonalities might include 1) similar structure 2) an output of one model that might match the input of another model, 3) a connection between a data element and the inputs for a particular model, or 4) a connection between two or more data elements contained within the network.

Upon enumeration of appropriate matches, the ACL makes a connection and the DMP coordinates operation in parallel across the separate computing platforms. This represents a slight modification from our earlier work (Brock et al. 2004). We anticipate the use of graph theory, linguistics, and discrete mathematics to refine the conceptual framework for M and Semantic Modeling.

However, the basic premise that models and data are similar to building blocks where a precise semantic definition aids in making connections will remain unchanged. As a practical matter, we are currently examining the use of models and data contained in computer spreadsheets as a means of demonstrating the initial feasibility of M and

Semantic Modeling. After prototype testing, M will become a standard set of languages and protocols.

It is important to note that M substantially differs from the Semantic Web. The goal of M is to build an interoperable environment specifically for models and data that depends on a common dictionary to define words used for semantic definitions, but not complete ontologies that attempt to categorize knowledge elements. The relative, distributed approach of M is in contrast to the Resource Definition Format (RDF) put forth by the Semantic Web, which includes a syntactical convention and a "schema, which defines basic ontological modeling primitives on top of RDF (Fensel et al. 2003, p. 9)."

In the final section of this article, we examine several prototypes where M and Semantic Modeling might provide a great deal of value to industry. We selected the prototypes based on several criteria including 1) our previous experience in modeling, 2) data quantity and the expectation of large amounts of new data generation in the future, and 3) the prospect for improvements in the application of modeling in practice.

5.0 THE FIRST PROTOTYPES

Choosing a set of prototype applications for M and semantic modeling is a difficult task because the computer language and concept can apply to a wide range of industries. A number of early prototypes have been identified, including applications in medicine, the automotive industry, agriculture, the entertainment industry (video games), environmental science, retailing, financial services, manufacturing planning and control systems, legal services, and engineering (Brock et al. 2005). Applications in the automotive industry alone, including driver information systems, comprise an entire discipline. The following gives an overview of three chosen from this initial group.

5.1 Enterprise Resource Planning Systems (ERP)

Simply stated, an enterprise resource planning (ERP) system identifies and plans "the…resources needed to take, make, ship and account for customer orders (APICS 2004)." To achieve these important tasks, ERP uses a variety of models and data to plan and control all the resources in a manufacturing or service-oriented company.

With the established success of ERP packages in practice it is realistic to think about what changes in technology might happen that will further enhance ERP. Currently, most organizations implement packaged ERP software that contains a single model for a specific business process. If the model does not exactly fit, substantial modifications are required. Managers often complain that this process of adaptation reduces overall organizational productivity.

One of the first prototypes of M deals with building a network of ERP models that could automatically match to data within organizations. These models include forecasting, production planning and scheduling, lot sizing, logistical, and financials. The ultimate goal is an intelligent modeling network that would partially replace packaged ERP software, providing a more flexible modeling environment for decision-making in business.

Building an in intelligent modeling network as a replacement for ERP makes sense because ERP is at its essence a data management tool. Therefore, it is reasonable that any advancement in the way that data is organized, and matched to models, will have a significant impact on the structure of ERP software.

Such a system is only possible through development of open standards and protocols for collection, sharing, and matching data to models. Without a system based on open standards, interoperability will not be possible and the economics of building suitable interfaces will overwhelm the economic value of the new infrastructure.

5.2 Retail Operations

Direct marketing offers an interesting case for the application of M because large quantities of data exist and there are many opportunities to apply models from management science to determine proper inventory levels. In general, direct marketing companies have impressive data management systems to support data-to-day decision-making. Retailing is a data rich environment. However, so many different models could potentially apply to retail data that a need exists for a flexible modeling system like M.

One of the first experiments in prototyping M involves the national catalogue and online retailer Lillian Vernon Corporation of Rye, New York. The company was established in 1951 and markets gift, house ware items, gardening, seasonal, and children's products. Well known for offering unique merchandise with especially good values, Lillian Vernon shipped more than 3.8 million packages in 2003, employing 3,500 people during the peak holiday season. Over 1,700 new products are introduced each year, and the total product line averages over 6,000 items (Lillian Vernon, Inc. 2004)

With such a large assortment of items, many with relatively short life cycles and seasonal sales, inventory management becomes a complex issue. A common problem is the determination of the proper lot size of merchandise to order given uncertain demand. To illustrate the breath of the problem, GRAPH 1 shows four examples of typical demand patterns for seasonal and ongoing merchandise.

GRAPH 1 – Demand Patterns

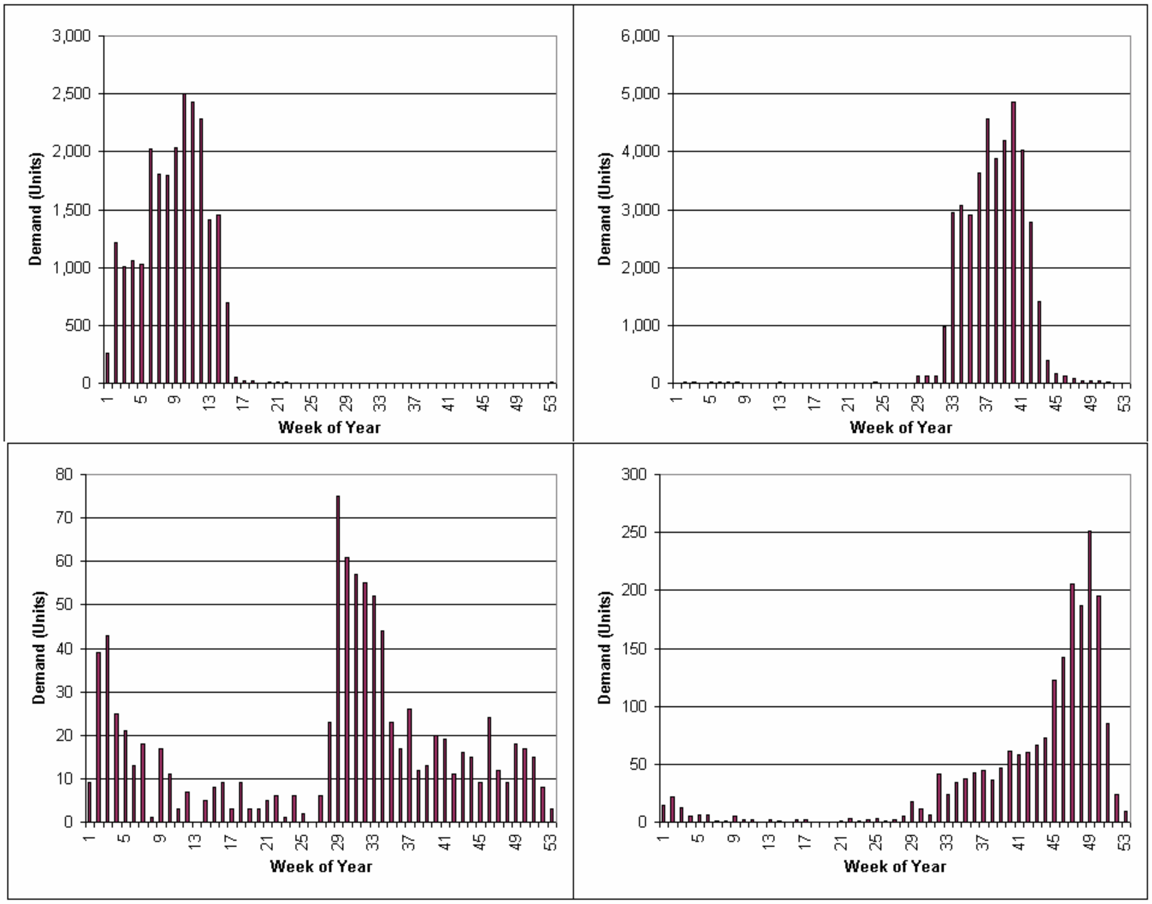

With thousands of different demand patterns, the goal of optimizing risk in terms of customer service and excess inventory becomes a complex challenge in matching the right model to the right data. The operations management literature offers a number of different solution methods to optimize risk for retailers. Most of these require the following common data:

1. Historical actual sales per item, per week.

2. Historical sales forecast per week.

3. Forecast at time the lot sizing decision was made

4. Customer Service level (actual sales compared the lot-size)

5. Salvage (amount remaining, if any, after conclusion of the event)

6. Some estimate of the cost of ordering the lot

7. Weighted Average Cost of Capital (Inventory Carrying Cost)

8. Cost of a Lost Sales

9. Price Breaks on Lot-Size

10. Transportation method/cost

Given a potentially large set of data and demand patterns, we hope to apply the DML to label inputs and outputs of models, along with data elements, to match models to data rapidly using M. In the case of Lillian Vernon, probably all models would operate on a single computing platform, so the DMP and ACP reduce to a simpler situation where model operation and identification of connections between models and data, all occur internally. Likewise, the ACL will make connections to models only inside a closed network.

If we can get simple applications of M, as described in the Lillian Vernon case, to work in a closed system with a subset of data and models, then the next step is to apply M to an open system. For example, there are a number of public sources containing important data on demographics and spatial income distribution. All of this is potentially useful in predicting sales. Much of this data goes unused because there is no rapid way to incorporate it into existing modeling systems. The application of M offers the opportunity to make full utilization of data, and to match the appropriate model for analysis.

5.3 Agriculture

Overall, there is a general lack of practical model use within agriculture. Yet there have been a great number of agricultural models developed at Land Grant Universities that could potentially help growers and agribusiness do a better job of logistics, planning, and resource optimization. Connecting these various models together could lead to the next wave in agricultural productivity.

One particular area of agriculture, harvest risk, offers the potential of introducing models traditionally used in business to the problem of optimizing harvest operations. The result, better utilization of harvest assets, fewer crop losses, and improved crop quality.

Gathering the harvest represents a complex managerial problem for agricultural cooperatives involved in harvesting and processing operations: balancing the risk of overinvestment with the risk of underproduction. The rate to harvest crops and the corresponding capital investment are critical strategic decisions in situations where uncertain weather conditions present a risk of crop loss.

This common problem in agriculture requires the application of mathematical models to calculate risk. Allen and Schuster (2004) present a case study of the Concord grape harvest and the development a mathematical model to control Harvest Risk by finding the optimal harvest and processing rate.

Mostly grown in the Northern United States, Concord grapes are a hardy variety known for exceptional flavor. However, like all agricultural crops, grapes are susceptible to frost damage during fall harvesting operations. Therefore, the goal is to harvest all of the grapes before a fall frost terminates operations.

Since it is impossible to predict in advance exactly when a frost will occur, it becomes important to employ various risk models to determine the best rate to process grapes. The model involves differentiation of a joint probability distribution that represents risks associated with the length of the harvest season and the size of the crop. This approach is becoming popular as a means of dealing with complex problems involving operational and supply chain risk.

The case study notes that Harvest Risk is under researched in agriculture. During the course of model formulation, the authors conducted an extensive literature review and found that there were no similar models for calculating Harvest Risk. This prompted a search for risk models used outside of agriculture to address the problem of a one-time event such as determining the correct lot size for perishable items like newspapers. In many ways, the Harvest Risk problem is similar to making purchases of highly seasonable items like fashion goods. With fashion merchandise, there are risks of ordering too much or too little. Either case can result in significant financial loss.

Likewise, the grape harvest represents a one-time event where harvesting too rapidly implies too much investment in equipment. Harvesting too slowly means an increased probability of losing crop because of a frost. These types of tradeoffs are very important for a variety of business and agricultural problems.

Looking outside a discipline to find mathematical models that might have relevant application is a time consuming task. The authors have noted that their line of research for the Harvest Risk problem dates over eight years. Most development and application of mathematical models occurs in highly specialized domains where researchers and managers have large amounts of specific knowledge but very little general knowledge about other disciplines. It takes years to accomplish meaningful research will realistic application.

The concept of Semantic Modeling helps to solve this problem because it allows for rapid application of models to data regardless of the domain where the model was originally developed. In essence, Semantic Modeling and M allow for the free flow of models over a network in much the same way that the Internet facilitates the free flow of information through interconnected web pages. Simply stated, Semantic Modeling is an advanced form of connective technology. Using this technology, modelers can quickly search for models from other disciplines that might solve the problem at hand.

In addition, Semantic Modeling aids in integrating various data sets. For example, the Harvest Risk model relies on a point estimate of temperatures for a specific grape growing region. Differences in elevation and other physical and environmental factors can result in significant temperature variation within a small area. When a frost hits a growing region, it is seldom evenly distributed.

Semantic Modeling, like Geographical Information Systems (GIS), has the capability of integrating various data sets to get a detailed view of the temperature characteristics for a region. For example, data from the US Geological Service could be integrated into the Harvest Risk model to account for differences in elevation for a specific growing area. This would give a much more accurate picture of what proportion of the Concord crop is susceptible to frost because of being located in lower elevations where cool air tends to accumulate. Sometimes a few feet in elevation can make a big difference in frost damage. Other data from the National Oceanic and Atmospheric Administration (NOAA) could also provide details on surface temperature variation within a growing region. Combining these data sets creates a more robust model that provides an accurate representation of Harvest Risk on a spatial basis.

6.0 CONCLUSION

Semantic Modeling will play an important role in linking models from a wide number of different disciplines to an array of different problems in business. Beyond the current discussion in this article, opportunities exist to link other abstract objects that require a precise semantic meaning, such as engineering designs, elements of financial reporting in a conglomerate, or important aspects of news feeds that might qualify as an object. Though the authors are in early stages of developing M and the practice of Semantic Modeling, there appears to be great potential to fulfill a need in industry to improve the integration of models and data. Our next article will focus on a review of computer languages relating to Semantic Modeling, pointing out the differences as compared to the conceptual design of M.

REFERENCES

Allen, Stuart J. and Edmund W. Schuster (2004), "Controlling the Risk for an Agricultural Harvest," *Manufacturing & Service Operations Management,* Vol. 6, No. 3.

APICS Dictionary (2004), Published by the American Production and Inventory Control Society, Alexandria, VA.

Bond Graphs (2004),
http://www.bondgraphs.com/about.html

Brock, David, L., Edmund W. Schuster, Stuart J. Allen, and Pinaki Kar (2005), "An introduction to Semantic Modeling for Logistical Systems," *Journal of Business Logistics*, forthcoming.

Corcoran, Elizabeth (2004), "Giving Voice to a Billion Things," *Forbes,* September 5.

Fensel, Dieter, James Hendler, Henry Lieberman and Wolfgang Wahlster (2003), *Spinning the Semantic Web*, Cambridge, MA: MIT Press.

General Systems (2004),
http://pespmc1.vub.ac.be/SYSTHEOR.html

Geoffrion, Arthur, M. (1987), "An Introduction to Structured Modeling," *Management Science* 33:5.

Geoffrion, Arthur, M. (1989), "The Formal Aspects of Structural Modeling," *Operations Research* 37:1.

Geoffrion, Arthur, M. (2004),
http://www.anderson.ucla.edu/faculty/art.geoffrion/home/csts/csts1.htm

Hof, Robert, D. (2004), "Amazon Joins the Search Party," *BusinessWeek*, September 27.

Lillian Vernon Corporation (2004),
http://www.lillianvernon.com/

Malone, Thomas W., Kevin Crowston, Jintae Lee, Brian Pentland, Chrysanthos Dellarocas, George Wyner, John Quimby, Charles S. Osborn, Abraham Bernstein, George Herman, Mark Klein, and Elissa O'Donnell (1999), "Tools for Inventing Organizations: Toward a Handbook of Organizational Processes," *Management Science* 45:3.

MURPHY, Frederic, H., Edward A. Stohr, and Ajay Asthana (1992), "A Representation of Schemes for Linear Programming Models," *Management Science* 38:7.

MATH-NET (2004),
http://www.math-net.de/

Park, Andrew (2004), "Can EMC Find Growth Beyond Hardware?," *BusinessWeek*, November 1.

Welch, James. S. (1987), "PAM – A Practitioners' Approach to Modeling," *Management Science* 33:5.

NOTES

