Business Intelligence

Vol. 5, No. 1

CUTTER CONSORTIUM

The Future of Connective Technology:

Greater Integration Through Semantic Modeling

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

Semantic modeling has great potential for dealing with the vast streams of new data that organizations will encounter in the future. Making semantic modeling a reality requires the 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 *Executive Report* discusses semantic modeling and prototype applications of M in ERP systems, retail operations, and agriculture.

Executive Report

Cutter Business Technology Council

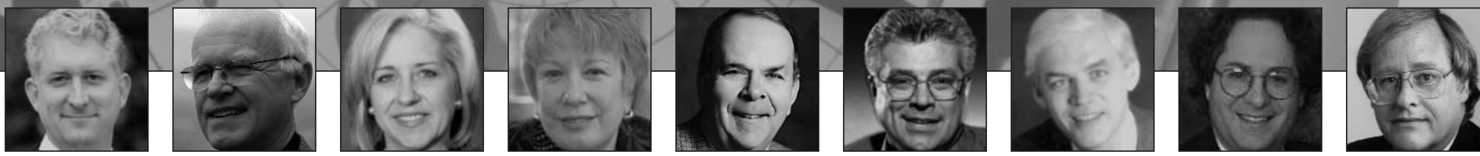

Rob Austin | Tom DeMarco | Christine Davis | Lynne Ellyn | Jim Highsmith | Tim Lister | Ken Orr | Lou Mazzucchelli | Ed Yourdon

Access to the Experts

About Cutter Consortium

Cutter Consortium is a truly unique IT advisory firm, comprising a group of more than 100 internationally recognized experts who have come together to offer content, consulting, and training to our clients. These experts are committed to delivering top-level, critical, and objective advice. They have done, and are doing, groundbreaking work in organizations worldwide, helping companies deal with issues in the core areas of software development and agile project management, enterprise architecture, business technology trends and strategies, enterprise risk management, metrics, and sourcing.

Cutter offers a different value proposition than other IT research firms: We give you Access to the Experts. You get practitioners' points of view, derived from hands-on experience with the same critical issues you are facing, not the perspective of a desk-bound analyst who can only make predictions and observations on what's happening in the marketplace. With Cutter Consortium, you get the best practices and lessons learned from the world's leading experts; experts who are implementing these techniques at companies like yours right now.

Cutter's clients are able to tap into its expertise in a variety of formats including content via online advisory services and journals, mentoring, workshops, training, and consulting. And by customizing our information products and training/consulting services, you get the solutions you need, while staying within your budget.

Cutter Consortium's philosophy is that there is no single right solution for all enterprises, or all departments within one enterprise, or even all projects within a department. Cutter believes that the complexity of the business-technology issues confronting corporations today demands multiple detailed perspectives from which a company can view its opportunities and risks in order to make the right strategic and tactical decisions. The simplistic pronouncements other analyst firms make do not take into account the unique situation of each organization. This is another reason to present the several sides to each issue: to enable clients to determine the course of action that best fits their unique situation.

For more information, contact Cutter Consortium at +1 781 648 8700 or sales@cutter.com.

The Future of Connective Technology:

Greater Integration Through Semantic Modeling

BUSINESS INTELLIGENCE
ADVISORY SERVICE

Executive Report, Vol. 5, No. 1

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

We live in a world filled with data. The success of business depends on the underlying flow of data and information for effective management. Since the 1960s, the advent of low-cost data collection methods, such as bar codes, along with advances in database technology have drastically improved the amount, quality, and timeliness of data in all organizations [45]. This long-term trend has contributed to significant improvements in productivity, especially in the areas of logistics, supply chain management, quality assurance, marketing science, and the financial management of complex organizations.

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 by as much as 40%-60% for many organizations. EMC, a leading manufacturer of data storage devices, recently noted that "companies are struggling to figure out how to turn all those bits and bytes from a liability into a competitive advantage" [34].

Dealing with increasing volumes of data will require innovative standards and information architectures to improve integration and communication between hardware, software, and business entities. However, the bigger question is this: how are we going to analyze and make sense of these large volumes of data?

A new research initiative, the MIT Data Center, addresses this 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 [9, 43]. This greater connectivity will spur new waves of productivity as managers learn to take advantage of the models and 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. David Brock — coauthor of this report, principal research scientist at MIT, and founder of the MIT Data Center — is credited with the idea of M. Designed as 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 will occur through an Intelligent Modeling Network that spans organizations. The conceptual design of M is such that network growth, in terms of adding more 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.

This *Executive Report* discusses the framework, details, and background of proposed standards for a language and protocol — M — that will enable computers to describe and share models and to assemble new models automatically from a general repository [6, 7]. This will substantially increase the "clockspeed" [17] of modeling and the computational efficiency of applying models to perform the functions of "sense," "understand," and "do" that compose the underpinning of creating smart objects within supply chains, in addition to other business activities of importance in achieving competitive advantage. The new computer language infrastructure includes open standards with two specific purposes: (1) communication of models between computers to create interoperability; and (2) the ability to run distributed models across the Internet.

In a sense, this effort is a step beyond linking the physical world, the underlying concept that has made Auto-ID technology successful. Networks, of physical objects or of abstractions like models, share the premise that leaps in productivity arise from the free flow of information. Creating an Intelligent Modeling Network will accelerate the flow of information to the advantage of many businesses and will form the backbone of a new type of Internet. Simply put, forging stronger links between models and data will result in productivity gains for the business.

It is important for IT managers to understand the direction of various types of connective technology research, including Semantic Modeling and M, as a means of planning for future computing systems. Some element of this planning becomes inevitable if firms desire to get the greatest benefit from the explosive growth in data available within businesses and entire supply chains. Computer languages and architectures currently exist that could enable immediate intra-organizational implementation of interoperable systems on a limited scale. Understanding these technologies is an important first step in organizing computing functions to accommodate the increasing amounts of data expected during the next several years. This report forms a solid base for IT professionals to gain insight into the emerging field of Semantic Modeling.

The sections that follow describe initial research on designing a network for abstract objects like models, including the underpinnings of Semantic Modeling and an overview of M.

The final part of this report describes three prototypes of Semantic Modeling currently under development at the Data Center. The prototypes deal with enterprise resource planning (ERP) systems, retail operations (lot sizing for short lifecycle products), and agricultural modeling (harvest risk) [42].

The ***Business Intelligence Advisory Service Executive Report*** is published by Cutter Consortium, 37 Broadway, Suite 1, Arlington, MA 02474-5552, USA. Client Services: Tel: +1 781 641 9876 or, within North America, +1 800 492 1650; Fax: +1 781 648 1950 or, within North America, +1 800 888 1816; E-mail: service@cutter.com; Web site: www.cutter.com. Group Publisher: Christine Generali, E-mail: cgenerali@cutter.com. Production Editor: Linda M. Dias, E-mail: ldias@cutter.com. ISSN: 1540-7403.

THE MODERN CONTEXT OF MODELING

There is no question that recent developments such as Auto-ID technology [8, 12, 35] will further increase the amount of data available for business decision making by using computing systems that sense and interact with the physical world. In the field of logistics management alone, these computing systems open new opportunities in terms of track and trace [25, 39], theft detection [26], improved service parts inventory management [24], and the control of production and logistics within military [14] and civilian supply chains. However, analyzing the large volume of raw data produced by Auto-ID technology in an orderly way requires the use of new mathematical models to provide representations and understanding.

Managers from all business disciplines frequently comment that the process of building mathematical models lacks productivity. Implementing mathematical models is complex, time-consuming, and requires advanced technical capabilities and infrastructure. Although there is a strong history of applying models to help managers make decisions about complex systems, specialists often develop these comprehensive models internally within business organizations or academia. This is commonly an application-specific job, and the same model-building technique must be reinvented for each new situation. Though internal development can lead to significant breakthroughs, this approach depends on trial and error, mathematical intuition, and an extensive knowledge of technical publications.

In the 1980s, software companies began embedding models into software packages installed on network servers, enabling organizational-wide modeling ability. This approach improved the productivity of modeling but limited users to a relatively small set of proprietary methods for problem solving. In all cases, internal development, or packaged software, models have become highly structured with few opportunities for creative applications. Proprietary systems also reduce the possibility of sharing models between business applications that exist outside the computing environment under which the original model implementation took place.

Part of the problem can be traced to traditional thinking about information theory. Today computers are faster, memory is cheaper, and bandwidths are plentiful, yet the tasks performed on these machines, such as e-mail, documentation, and data storage, are nearly the same as they were 10 years ago. Computers primarily store, manipulate, and transmit data to people. Without direct human interaction, computers essentially do nothing.

Yet computers have far greater unrealized capability. With current technology, it is possible to design large-scale Internet systems that might allow computers to store and analyze vast quantities of information and to share these results automatically with other computers throughout the world. Networks of computers have the potential to operate independently or collectively without human interaction.

The failure to take full advantage of the computer's potential lies not in the hardware or communications technologies but in the lack of languages and standards that allow systems to share data and interface models across multiple applications and domains. The consensus is that this lack of integration is a barrier to increased productivity for a wide range of situations.

Semantic Modeling challenges the long-standing philosophy that emphasizes individual effort in formulation and implementation of mathematical models. The ultimate goal is to build an integrated modeling structure to accelerate the development of new applications.

RECENT DEVELOPMENTS THAT SHOW THE FUTURE

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

Amazon.com has recently announced A9, a tool that can complete searches of information located on HTML Web pages as well as in the text of thousands of books [22]. Eventually, A9 hopes to incorporate the ability to perform even more specialized searches by accessing other proprietary databases. The chief executive of A9 has 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 his or her own view of information gathered by Internet searches. All of A9's activities point toward the Semantic Modeling goals of greater integration, improved search capabilities, and a relative approach to organizing information.

Outside of the Internet, other developments also point toward greater integration. For example, in the US, there are billions of embedded microcontrollers in cars, traffic lights, and air conditioners, which provide specialized instructions for control based on sensing specific aspects of the environment. These microcontrollers all act in total isolation. Ember, a company headquartered in Boston, Massachusetts, has developed a "mesh network" that holds the potential of allowing all of these microcontrollers to communicate with one another [11]. 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 rewiring. There are almost endless opportunities to establish communication connections for a wide variety of microcontrollers.

Just as Internet searches cannot gather 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 due to several factors, including an inability to apply models to data quickly. Overcoming these limitations is a complex task. One option to meet this challenge involves building networks based on semantics. The next section explores this idea in greater depth.

A SEMANTIC-BASED INTERNET SEARCH

Existing Internet standards do not provide any semantics to describe models precisely or to interoperate models in a distributed fashion. For the most part, the Internet is a "static repository of unstructured data" that is accessible only though extensive use of search engines [16, p. 377]. Though these means of finding data have improved since the inception of the Internet, human interaction is still required, and there are substantial problems concerning semantics. In general, "HTML does not provide a means for presenting rich syntax and semantics of data" [16, p. 7].

For example, one of the authors of this report recently did a search for "harvest table, oak" hoping to find suppliers of home furniture. Instead, the search yielded a number of references to forestry and the optimal time to harvest oak trees. Locating the URLs relating to furniture required an extensive review of a number of different Web sites. This time-consuming process of filtering can only be accomplished though human involvement.

With inaccurate means of performing specific searches based on one semantic interpretation of data, information, or models, it is nearly impossible for the Internet to advance as a productive tool for modeling.

Types of Webs

The problem of semantics stems from the fact that keywords are the means used to describe the content of Web pages. Each keyword may have multiple meanings, creating a situation of great difficulty when attempting to accomplish an exact search. The difficulty increases by an order of magnitude when attempting to do phrase-based searches. Without exact search capability, it is impossible to create any sort of machine-understandable language for the current "Web of Information."

Though the search engine issue has not been resolved, industry forces are pushing for a new type of Internet, characterized as the "Web of Things." Driven by developments in Auto-ID technology and ubiquitous computing, the Web of Things aims to link physical objects to the Internet using radio frequency identification (RFID) tags as real-time communication devices and to "shift from dedicated computing machinery (that requires user's attention, e.g., PC's) to pervasive computing capabilities embedded in our everyday environments" [46].

Aiding this effort is EPCglobal, Inc.,[1] an international standards organization formed by EAN International (to be renamed GS1) and the Uniform Code Council (UCC; to be renamed GS1 US). EPCglobal administers the EPC numbering system, which provides the capability to identify an object uniquely. With serial identification for physical objects, searches accomplished through Internet search engines or proprietary IT infrastructures will be much more effective in finding an exact match. This provides the ability to do track and trace across entire supply chains as well as other computerized functions important to logisticians. Linking the physical world, using Auto-ID technology and ubiquitous computing, will form the basis for a revolution in commerce by providing real-time information and enabling smart objects [37, 41, 44].

Impressive as the effort to create the Web of Things is, it still does not address the question of semantics in describing objects beyond the use of a simple serial number. There exist a large number of abstractions, such as mathematical models, that cannot be characterized by a unique serial number no matter how sophisticated the syntax. Without the ability to provide unique identification of an abstraction, the Internet will serve little useful purpose in linking mathematical models together in a way similar to the manner that the Web of Things will eventually link the physical world.

In the future, the definition of a model and the sharing of models through a network will become as important as the model itself. To accomplish this higher goal, the Internet must become a "Web of Abstractions," in addition to a Web of Information and a Web of Things.

Creating a Web of Abstractions requires a semantic definition of models that is precise and can be machine-understandable. Given this capability, models can be searched, organized, categorized, and executed — sequentially and in parallel — creating multiple, large-scale synthetic environments. These synthetic modeling environments will exist only in virtual reality and offer the potential for creating a dynamic metastructure for specific classes of models.

Through a Web of Abstractions, models can be matched much more quickly to practical problems, along with the available data, and shared beyond single end-user applications. This capability is of great value both to practitioners and researchers alike who are interested in gaining the maximum value in modeling logistics for practical decision making.

The Representation of Model Schema

Previous research in computer science consistently states that the missing structure necessary to create a Web of Abstractions is an ontology. Simply stated, "An ontology specifies what concepts to represent and how they are interrelated" [16, p. 34]. This structure provides order when conducting searches and serves the important purpose of creating a crude form of intelligent behavior. For example, one group of researchers involved in the early aspects of using computers to create artificial intelligence concluded that "the clue to intelligent behavior whether of men or machines, is

[1]See www.epcglobalinc.com.

highly selective search, the drastic pruning of the tree of possibilities explored" [15, p. 6]. Properly constructed, an ontology reduces search time for abstractions, creating a free flow across a network. With the thousands of models that do not find widespread application in practice, the capability to conduct a quick and accurate search improves the chances that more applications will occur.

There are two important aspects to consider when using an ontology to organize abstractions like mathematical models for machine-understandable searches. First, the ontology assumes that a semantically precise definition of an abstraction (model) exists. Absence of this in the current schema presents a problem because the classification of mathematical models depends on keywords that might have different meanings under different contexts (e.g., planning and scheduling).

Second, because meaning arises by the way one model is connected or related to other models, an ontology also serves an indirect definitional function. This is important in visualizing the big picture of the relationships between different types of models. It also drastically decreases search time by reducing the number of possibilities in reaching an exact semantic match. However, there are significant drawbacks concerning the establishment of an ontology that is robust enough to include all mathematical models in existence.

The Limitations of Representing Models Using Ontologies

By definition, ontologies are rigid and inflexible and assume that one absolute definition exists for each knowledge element. The idea is to establish a set structure of definitions and relationships between different abstractions (models) that are canonical and eternal. This means that the usefulness of an ontology for modeling depends on intensive study and rigorous examination of the canon put forth. It is unrealistic to believe that any independent body of academics or practitioners could formulate an all-inclusive canon that would stand the test of time. The ontology approach is a throwback to the philosophy of scholasticism that dominated Western thought during the high middle ages. History has proven that canonical structures, meant to organize and communicate knowledge, often have the unintended outcome of restricting the adoption of further innovations that exist outside the bounds of the canon.

In addition, rigid ontological structures lack the ability to adapt based on inductive reasoning. There is no ability to learn automatically from specific examples that occur through time and to generalize to form a new element of knowledge contained in the ontology. This was the major limitation of expert system architectures and a leading reason for the decline in the application of expert systems in practice.

A final major drawback involves the difficulty in merging separate, distinct ontologies into a whole. For all the advantages of a rigid structure in organizing abstractions (models) and reducing search time, there is no easy translation or interface to integrate two different classes of models. We believe that advances will take place only through the free exchange between widely disparate fields of modeling. Without this ability, efforts in establishing computer languages to share and interoperate models will be difficult.

A Relative Approach to Model Representation

To overcome the disadvantages of traditional ontologies in computer science, we advocate the abandonment of a single, unified structure to represent abstractions (models).[2] The reality is that the representation of objects and their interrelation is almost entirely dependent on a person's viewpoint. In other words, as opposed to a single ontological representation for models, we propose a more flexible means of description, so that others may construct

[2]Edmund Schuster is working as a research affiliate at MIT to establish the Data Center and is active in the development of M. As mentioned earlier, David Brock, principal research scientist at MIT, is credited with the idea of M as a modeling language.

their own particular representations and unique ways for connecting them together.

Furthermore, our approach provides the means for building dynamic, on-the-fly model taxonomies; that is hierarchical organizations of models that are generated as a function of an individual's point of view. In our system, there is no one classification scheme (ontology) but rather, multiple ones. Simply put, several ontologies can exist simultaneously with no contradictions.

With this approach, a model is an atomic element that may subscribe to one or more classification hierarchies. These taxonomies may be mutually agreed upon industry standards — essentially, commercial data dictionaries, proprietary schemes, or dynamically generated groupings for particular applications. In all cases, the representations, relations, and organization of models will be dynamic and configurable to the task. Later in this report, we provide an example of model representation that is integral to our view of the schema necessary to create the Web of Abstractions.

In the next two sections, we discuss the practical and theoretical aspects of combining advances in computer science with the existing body of mathematical models that have been developed by logistics researchers over a period of many years. The prospect of doing Semantic Modeling for large-scale business applications draws upon the intersection of computer science and the practice of modeling. We anticipate other disciplines such as linguistics, graph theory, and discrete mathematics to be important in the development of Semantic Modeling.

SEMANTIC MODELING

Most would agree that modeling is a craft industry analogous to the production of automobiles prior to the advent of the assembly line. Although models are ubiquitous management tools, they are, for the most part, isolated from one another. In other words, a model from one domain, such as weather forecasting, does not interact with another, such as logistical systems.

The reason for this is obvious. Until very recently, humans were the only ones who built, used, and shared models. Our limited cognitive ability naturally restricts the number and diversity of models we can accommodate. Computers, on the other hand, have the ability to execute and communicate models with vast numbers of other computers. With ever-increasing processing power, data storage, and networking bandwidth, the computing grid is poised to revolutionize our ability to understand and manage the physical world. With its standards and languages, the Internet provides the backbone for communication, but it does not provide the mechanism for describing and integrating diverse models. The future is a form of modeling on demand similar to other efforts in establishing a computer grid that resembles electric power distribution [28].

Our goal is to turn modeling into a mass production system based on standardization, scale, and interoperability. In summary, this means that a Semantic Modeling language capable of achieving this functionality must include the following [16, p. 8]:

1. "A formal syntax and formal semantics to enable automated processing of their content"
2. "A standardized vocabulary referring to real-world semantics enabling automatic and human agents to share information and knowledge"

Achieving this goal will mean that practitioners can produce models in a timely manner with greater productivity and relevance. This anticipates a new era for computers in terms of insight and awareness, and it implies the ability to organize data and define the inputs and outputs of models in a semantically precise way.

The mechanism we put forth to mass produce models and create interoperability draws inspiration from current efforts to improve the search capabilities for the Web of Information. The World Wide Web Consortium (W3C) is responsible for initiating select efforts to improve overall Web searching capabilities.[3] Some of the initial work conducted by W3C forms a reference base for our research in developing and implementing a Web of Abstractions.

Each abstraction (model) has unique elements that can be defined just as a language has a specific syntax and grammar. Defining these elements alone will be of no benefit unless there is a protocol, or computer language, to communicate and execute the elements of models across a large network like the Internet. Our efforts in establishing Semantic Modeling are grounded in the idea of having data and models defined and linked in a way that can be used by machines not just for display purposes but also for automation, integration, and reuse across various applications. Accelerating the reuse of model elements across vast networks of users will lead to the mass production of models, which will in turn be of great benefit to practitioners. In addition, distributed modeling, a set of geographically separated model elements working simultaneously in parallel, adds additional prospects for large-scale parallel computing.[4] This capability will improve the utilization of desktop computers and provide grids of almost unlimited modeling power.

[3]See www.w3c.org/2001/sw.

[4]Software Agents for Distributed Modeling and Simulation (www.informatik.uni-rostock.de/~lin/AnnounceIEEE/node2.html).

Though W3C provides something called a Resource Description Framework (RDF) that defines the basics of representing machine-processable semantics [16, p. 9], no formal computer language has been put forth that enables the sharing of models or doing large-scale modeling in parallel. The following section gives an overview of our vision for a computer language and protocols that achieve Semantic Modeling.

SYSTEM ARCHITECTURE

The fundamental idea is to design a family of standards that enables the creation of models that integrate automatically into an executing synthetic environment. In this way, developers can formulate models within their particular areas of expertise and know that the resulting models will interoperate in a shared environment. We believe it is possible, with sufficient care in the definition, to create such a language that is both precise and expressive in its description yet constrained in its breadth to ensure compatibility.

The goal is to create synthetic environments that receive data from the physical world (e.g., through Auto-ID technology) and then produce inferences, interpretations, and predictions about the current and future states of the environment.

This interpolated or extrapolated state data is essential for any automated decision system. In other words, the estimated environmental states support networks of decision making algorithms so that they can make informed decisions and deliberate plans (that feed back to the physical world.) This type of modeling is essentially the underlying basis for automated control, monitoring, management, and planning.

M is currently in the initial stages of research and development at the Data Center. Comprising several important elements, M's purpose is to serve as the fundamental language to link models and data together.

In essence, 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 preassigned, precise semantic definitions through M — provide a mapping of connections between models and other models, data and models, and 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.

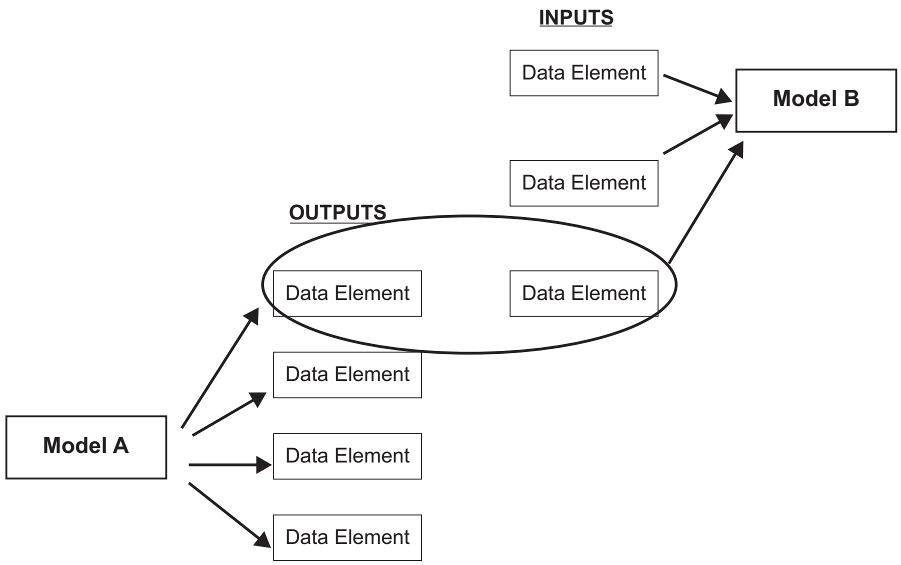

Figure 1 — Connecting models.

To achieve these connections, the structure of M must be made up 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, 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 ACP, which locates the appropriate connections.
- **Automated Control Protocol (ACP)** helps model outputs and inputs locate one another within a network, even if the individual models 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.
- The **dictionary** is 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. In cases where a word has multiple meanings, only one definition will be used. 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. We anticipate the use of graph theory, linguistics, and discrete mathematics to refine the conceptual framework for M and Semantic Modeling.

The basic premise is that models and data are similar to building blocks where a precise semantic definition aids in making connections. 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 RDF Schema put forth by the Semantic Web, which includes a syntactical convention and a "schema, which defines basic ontological modeling primitives on top of RDF" [16, p. 9].

In summary, Figure 2 shows the interaction of the major components of M.

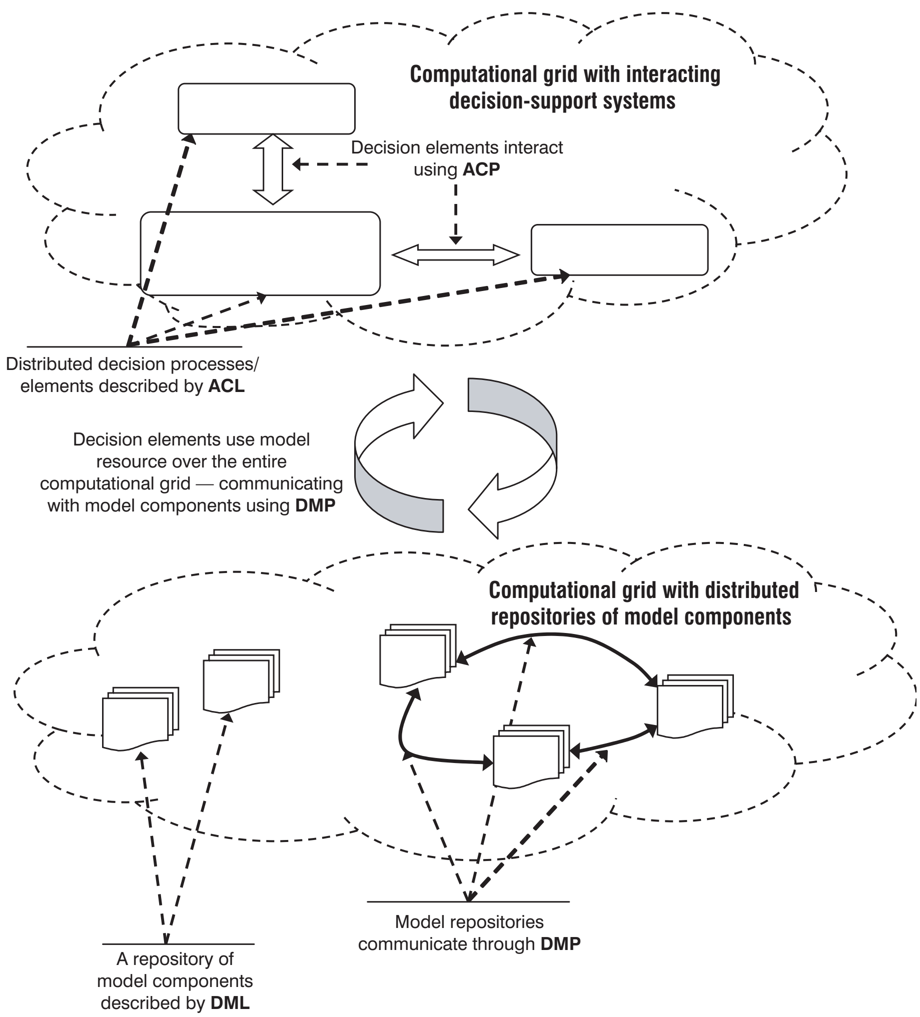

Figure 2 — Proposed distributed system using DML, DMP, ACL, and ACP.

AN EXAMPLE FROM LOGISTICS

Researchers at the *2001 Logistics Educators Conference* presented an interesting article about the implication of advanced planning and scheduling systems (APS) on supply chain performance [10]. The article also contained an appraisal of the changes needed in academic curriculums to ensure that students receive proper education about the role of APS in supply chain management. Based on these comments, we decided to investigate the literature of finite capacity scheduling (FCS), an important subsegment of APS, to find an initial example for demonstrating the aspects of Semantic Modeling.

In general, there are many solution methods for FCS. A nonexhaustive list includes mathematical programming, simulation, heuristics, genetic algorithms, neural networks, theory of constraints, and expert systems. Of this list, the first three are frequently found in practice, with the most common being heuristics. About 80% of commercial scheduling packages use heuristic solution approaches [31].

A detailed analysis reveals that each model for FCS exhibits primal properties based on the solution method or algorithms employed [36]. Table 1 summarizes the capabilities of each model in its pure application without modification.

Understanding that each model class for FCS listed in Table 1 — math programming, simulation, and heuristics — does not fully address all attributes commonly found in commercial FCS problems is important in supporting the belief that future advances will come from combining existing models in new ways to address a wider range of attributes.

A recent article provides substantial background about FCS from the perspective of practical implementation, including several references to a group of models that provide different FCS capabilities [40]. Essentially the entire group deals with the same scheduling problem. This body of research provides insight for a simple example that highlights how elements from different models can combine to produce new models with better performance, thus demonstrating the importance to practitioners and researchers of developing a computer language and protocols to

Table 1 — Comparison of Different Scheduling Approaches

Attribute	Math Programming	Simulation	Heuristic
Hold time		X	X
Queue time		X	X
Customer service		X	
Forecast bias		X	
Set-up cost	X		X
Holding cost	X		X
Overtime cost	X		X
Capacity	X		X
Production lot size	X		X
Production sequence	X		X
Customer due date	X	X	X
Family structure	X		

X = Functional

facilitate this process with some degree of automation.

The example set forth below deals with various types of models used to schedule production for manufacturing lines common to the consumer-goods industry. With high demands for customer service, it is important for consumer-goods companies to schedule the production of end items with proper consideration given to the risk of being out of stock as well as to the capacity constraints that might limit production in times of peak demand. Based on statements made in the literature, all of these models were implemented at the same consumer-goods company during a span of 15 years. The following provides a description of each model:

- **Model A — Deterministic Simulation** [38]. With bias-adjusted safety stocks that use customer-service levels as an input, production planning occurs for each item independently. All items run on a production line are summed to give a total capacity load. This model initially assumes that infinite capacity is available for production and does not consider set-up or inventory-carrying cost. However, the model does provide a method for safety-stock planning that considers dynamic forecasts and the impact of forecast bias in planning safety-stock levels.

- **Model B — Mathematical Programming** [1]. Exploiting the fact that consumer goods have a family structure defined by package size, production can be planned using a two-tier hierarchichal structure in which product families are sequenced, with disaggregation taking place to form end-item schedules. This approach provides optimal solutions based on cost and utilizes an innovative mathematical formulation that yields near instantaneous solutions to mixed-integer math programming problems.

- **Model C — The MODS Heuristic, Sequence Independent** [3]. An approach to scheduling using the Modified Dixon Silver (MODS) method to calculate near-optimum production schedules based on inventory and set-up costs and inventory set-up time.

- **Model D — The MODS Heuristic, Sequence Dependent** [13]. Building on the MODS method, this approach utilizes the nearest neighbor variable origin (NNVO) heuristic as a second step to sequence production based on a "from-to" table of changeover costs between items.

Relationship to Proposed System Architecture

By looking at working models as an aggregation of interchangeable elements, the possibilities for identifying new combinations are numerous. Using our system definitions, the DML describes various elements of models, such as the bias-adjusted safety-stock method used in Model A, that are modular and interoperable. The ACP provides a mechanism for various model elements to locate each other across a network like the Internet. Analyzing the examples of Models A, B, C, and D, it appears that the developers located model elements as a function of many years of study in the FCS area combined with mathematical intuition.

In the situation where distributed modeling takes place, the DMP allows for communication between active models located on separate computing platforms. For example, bias-adjusted safety stock (Model A) might be calculated on one computing platform with the results being transferred to another platform that contains the MODS heuristic (Model C). In this case, the DMP establishes the order to run the models and the timing of data transmissions. The final part of our system architecture is the ACL, which establishes the formal connections based on the DML descriptors of model inputs, outputs, and data. The ACL is needed because the decisions from one model (outputs) might become data (inputs) for another model. This is the case for Model A, which can provide safety stocks (output) as an input to Models B, C, and D. The ACL matches the outputs of one model to the appropriate inputs for another model.

Establishing Semantics for Logistics Models

The starting point for the goal of building an interoperable system based on DML, ACP, DMP, and ACL is a semantically precise definition of a model. Given that most model descriptions depend on keywords, which might have a number of different meanings, we propose an alternative approach to define a model. The intent of DML is to label models semantically in such a way that common elements can be machine understandable and interoperable.

Our approach to the semantic-labeling problem involves forgoing attempts to describe the various algorithms employed in each model. Rather, we focus on the data (inputs) and the decision variables (outputs) required for each model as a unique base for machine understanding and the grouping together of common models. This assumes that a special, unique relationship exists between a model and its data.

As a practical matter, we believe that a definition of a model in terms of data inputs will provide a more precise semantic than a definition derived from classifying the algorithm used for each modular component (model). Keyword definitions for the complex algorithms that comprise models are notorious for having different semantic meanings. In addition, the keyword descriptions often have no meaning at all to business practitioners that do not have extensive formal training in logistics or management science.

Table 2 illustrates how data inputs can become a tool for establishing semantic meaning.

From Table 2 we note that Models A, B, C, and D all share the data inputs D1, D2, D3, D4, and D7. This provides a natural way to categorize Models A, B, C, and D into the same group. This also implies that models using the same data will deal with the same initial problem (in this case, scheduling of production lines for the consumer-goods industry) and that all four models are interoperable with respect to the data. Any of the four models could be applied to the same data set to gain the result of a production schedule. By defining a model in terms of its data inputs, the outcome is a precise semantic that allows assignment of the model to a common group.

Further, the use of input data as a means of establishing semantics also aids in distinguishing differences between models in a group. Likely, the data inputs for a group of models will not be identical if different solution methods (algorithms) are used. From Table 2, we notice that none of the four models share all the same data inputs yet all of these models are capable of producing a schedule (output) for a manufacturing process characteristic of the consumer-goods industry. This offers a way to identify differences between models within the same group as categorized by data. This also provides an indirect indication of the solution methods (algorithms) employed.

For example, Models B, C, and D share the commonality of requiring a capacity limit (D11), inferring that these models belong to a class of FCS systems and perhaps are interoperable. In another case, Table 2 shows that Models A, B, C, and D all have service level as a parameter (D7), implying that this class of models includes some aspect of safety stock. Other safety-stock models, not mentioned in this example, might offer alternative ways to calculate safety stocks using the same data requirements. Because all of these models share the same set of data inputs, they are interoperable with Models A, B, C, and D.

The reader must keep in mind that we view models in an atomic, elemental way. Taking an example from chemistry, a single element like calcium (Ca) can become part of many different molecules, such as calcium hydroxide ($Ca(OH)_2$) or calcium chloride ($CaCl_2$), through chemical reactions. In a similar way, a single model, for example bias-adjusted safety stock (Model A), can be combined with Models B, C, and D to create entirely new model forms. Data inputs, as part of DML, hold the key for developing an open architecture for models to combine automatically such as in chemical reactions.

Table 2 — Data Inputs to Models A, B, C, and D

Data Input	Model A	Model B	Model C	Model D
D1. Beginning inventory	X	X	X	X
D2. Forecast demand (by week)	X	X	X	X
D3. Historical shipments (by week)	X	X	X	X
D4. Historical forecast (by week)	X	X	X	X
D5. Hold time (days)	X			
D6. Queue time (days)	X			
D7. Service level (percentage in stock)	X	X	X	X
D8. Set-up cost ($/changeover)		X	X	X
D9. Set-up time (hrs/setup)			X	X
D10. Holding cost ($/week)		X	X	X
D11. Capacity limit (hrs/day)		X	X	X
D12. Family structure (end items per group)		X		
D13. Overtime cost ($/hr)			X	X
D14. Sequence dependent set-up cost (from-to table of changeover costs)				X

To summarize, the descriptors we put forth as the basis for DML include data inputs as the primary semantic for grouping models and the initial basis for machine understanding. Model outputs are only important in providing (1) general guidance concerning the objective of the modeling effort and (2) some definitions of model outputs that may in turn become model inputs in other situations.

We do not believe that semantic description of algorithms based on keywords will play a significant role in the design of DML. One important means of classification that we have not mentioned involves the assumptions of the model. The use of assumptions as a precise semantic of a model provides an interesting area for future research.

An Example of Multiple Ontologies

As an illustration of the fact that multiple ontologies exist with respect to the definition of a model and its relationship to other models, we now examine a final example involving Models A, B, C, and D.

Depending on viewpoint, the library of models could be used in two different ways:

- **From a production planner standpoint —** given a specific beginning inventory, end-item demand forecast, and target safety-stock levels, the models, could provide a computer-generated schedule of the timing and amount of production needed at a manufacturing plant
- **From a supply chain manager standpoint —** given a specific beginning inventory, end-item demand forecast, and target safety-stock levels, the models could provide an accurate projection of inventory levels in plant warehouses. This information could be used to determine the overall size of the warehouse.

There is evidence in the literature that this group of models has in fact been used in both of these ways. This brief example shows that the same library of models has different meanings different relationships depending on the viewpoint of end users. This aspect of relative relationships makes the establishment of rigid ontologies difficult to achieve in practice. Though we have an idea how to handle this obstacle in producing machine-understandable semantics, there certainly needs to be more research conducted in this area before totally abandoning single ontology architecture.

It appears that the key to building multiple ontologies depends on the relationships between models. When faced with systems characterized by intricate relationships, engineers sometimes employ graph theory to provide representations for complexity. Using this approach, we believe the edges of the graph hold the answer to establishing different ontologies for the same group of models.

THE FIRST BUSINESS APPLICATIONS

Choosing a set of prototype business 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 ERP systems, applications in medicine, the automotive industry, agriculture, the entertainment industry (video games), environmental science, retail, financial services, manufacturing planning and control systems, legal services, and engineering [9]. Applications in the automotive industry alone, including driver information systems, make up an entire discipline. The following presents an overview of three chosen from this initial group.

ERP Systems

Simply stated, an ERP system identifies and plans the "resources needed to take, make, ship, and account for customer orders" [4]. To achieve these important tasks, ERP incorporates 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 fit exactly, 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 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 possible only 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.

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 day-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 catalog and online retailer Lillian Vernon Corporation of Rye, New York. The company was established in 1951 and markets gift, houseware, gardening, seasonal, and childrens' 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. More than 1,700 new products are introduced each year, and the total product line averages over 6,000 items [27].

With such a large assortment of items, many with relatively short lifecycles and seasonal sales, inventory management is a complex issue. Given the uncertain demand for such items, determining the proper lot size of merchandise to order is a common problem. To illustrate the breadth of the problem, Figure 3 shows examples of four typical demand patterns for seasonal and ongoing merchandise.

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. An estimate of the cost of ordering the lot.
7. Weighted average cost of capital (inventory-carrying cost).
8. Cost of lost sales.
9. Price breaks on lot size.
10. Transportation method and 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 fast way to incorporate it into existing modeling systems. The application of M offers the opportunity to

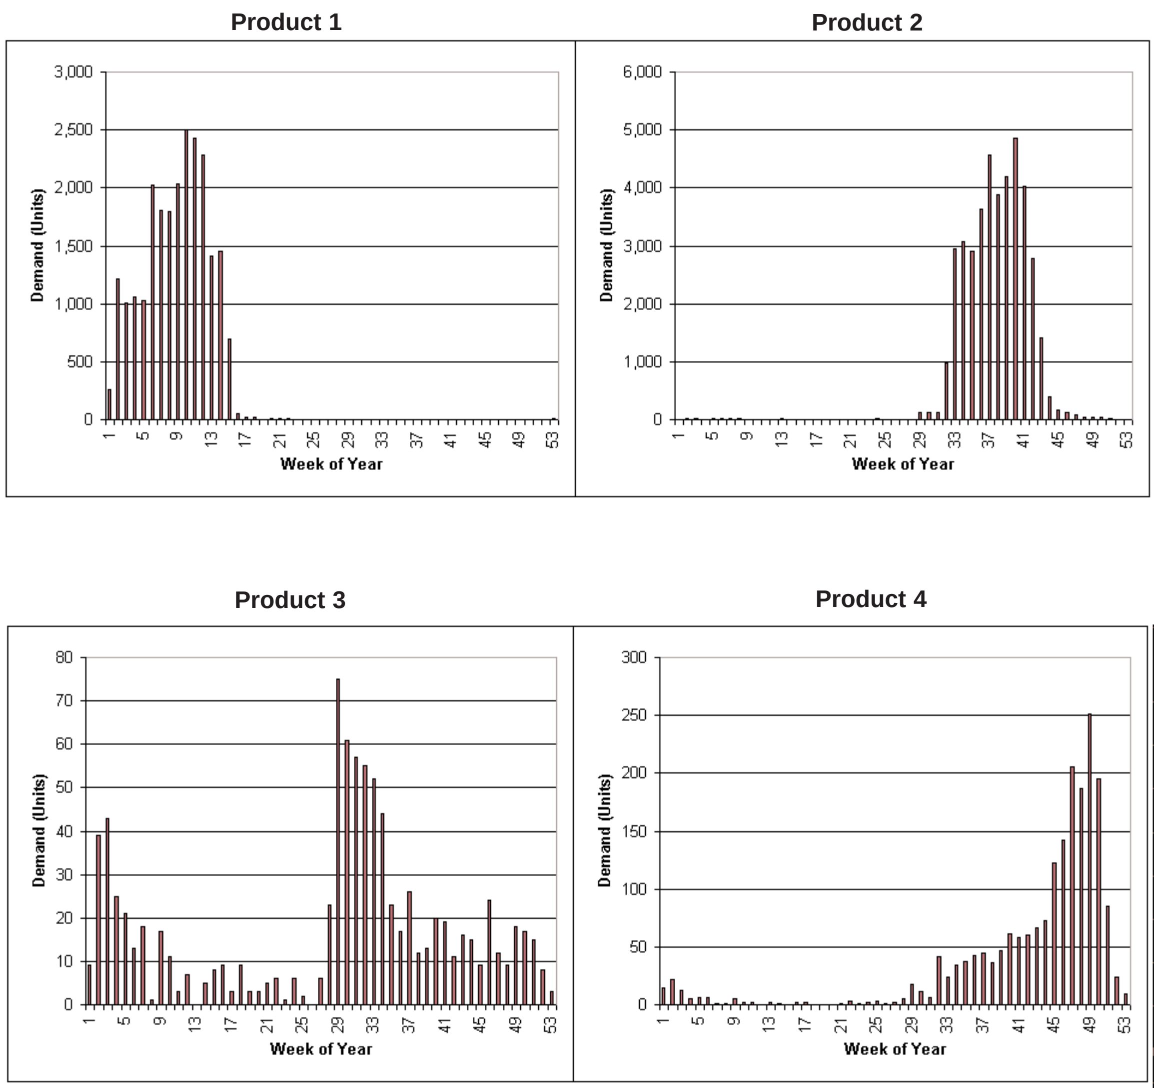

Figure 3 — Weekly demand for various items sold by Lillian Vernon.

make full utilization of data and to match the appropriate model for analysis.

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 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 optimize 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. In situations where uncertain weather conditions present a risk of crop loss, the rate to harvest crops and the corresponding capital investment are critical strategic decisions.

This common problem in agriculture requires the application of mathematical models to calculate risk. The authors recently presented a case study of the Concord grape harvest and the development of a mathematical model to control harvest risk by finding the optimal harvest and processing rate [2].

Mostly grown in the northern US, 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 the grapes before a fall frost terminates operations.

Since it is impossible to predict exactly when a frost will occur, it is important to employ a risk model 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 such as fashion goods. With fashion merchandise, there is a risk of ordering too much or too little; either can result in significant financial loss.

Likewise, the grape harvest represents a one-time event where harvesting too rapidly requires too large an investment in equipment, yet 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 with 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 in which 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 Geographic 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 Survey 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 due to its location 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.

PRACTICAL CHALLENGES

The history of modeling includes a tradition of individual or small team efforts to formulate a single comprehensive model that provides a robust solution for a particular problem. Elements of other models are rarely incorporated into such efforts beyond conducting the standard literature review. To introduce the system we propose in this article will require a culture shift originating in academic institutions that serve as the training centers for the modelers of the future. Developing DML, DMP, ACL, and ACP as a formal set of languages and protocols ia a step forward in changing the culture of model building. Once practitioners experience the power of automatically sharing models between computers, we believe there will be acceptance in adopting our system. As more model builders begin to use the languages and protocols, the power of the network will increase, resulting in productivity gains.

For both we are in the process of developing a search engine interface that resembles an Internet browser to locate model elements residing on a network. The browser uses data inputs as the semantic for conducting the search. Once the appropriate models are located, another computer interface provides a workspace for visualization that shows how various model elements might fit together to form a practical solution. The key to the visualization is to show in two or three dimensions the various combinations of specified models that might be possible. With this type of interface, the proper matching of a model to data and the interoperability of models becomes clear to the user. Ultimately, this will accelerate implementation in practice, resulting in the mass production of models.

To begin the process of development, we are establishing an online community to define the data types used by M as a means for semantic searches. This is a tedious process; however, there is no other way to establish a precise semantic for models. Previous work conducted by industry organizations such as the International Organization for Standardization (ISO) and various US government agencies such as the National Institute of Standards and Technology (NIST) will aid this effort. The online community will also communicate various aspects of Semantic Modeling and the state of development of M.

Given that a prototype of M is achievable within the next year, the question remains, what incentives will exist for model builders and practitioners to use Semantic Modeling? Our approach focuses on future model building and the establishment of a repository for models. However, the hundreds of logistical models currently in use present a problem in that these will need to be coded in the proper language and protocols of M. Since many models are run using proprietary systems, the task of coding will be significant unless new methods of interface and translation are developed. This has to be part of our efforts in developing M.

One idea to provide an incentive for model builders to use M involves a new Internet payment technology [23]. With this scenario, developers could form a representation of their models using M and post to the Internet in machine-understandable format. Those (either humans or machines) seeking to find models would do a search to locate the best model for their application. When the user downloads a specific model found by semantic search, the developer would receive a payment determined in advance or by market forces. In the case of simpler models, a smaller "micropayment" might be more appropriate given the volume of downloads. This would provide financial incentive for developers to select older models for coding that have been long forgotten by practitioners.

We envision a new industry forming where specialized firms constantly review old software or journal articles for signs of models having commercial value when coded into M and distributed using the Internet. In the long term, existing large companies in the business of selling packaged software might yield to a new generation of firms that specialize in producing a repository of models using M. With this scenario, practitioners benefit because model applications would more closely match the problem at hand, unlike the current situation where many firms must radically redesign organizational processes to meet the demands of commercially packaged software. If nothing else, Semantic Modeling offers the possibility of assessing the true value of a model through the free exchange across a network.

A final hurdle for the implementation of M involves the adherence to standards. With every standards-setting opportunity, there is always the chance that adopters will bend standards to meet their own objectives. This was the case in the development of electronic data interchange (EDI) standards as well as others. Good design of the standards along with active industry associations to monitor adherence are the means needed to maintain integrity.

THE UNDERPINNINGS OF SEMANTIC MODELING

As we conclude this overview of Semantic Modeling, it is important to note that the idea of defining elements of models for the purpose of reuse is not new. Previous work has concentrated on the use of structured modeling to define elements for management science techniques [19, 20] and also for building a system for "metamodeling" [32]. 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. [21]

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 models [33, 47]. 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 perform benchmarking [29]. As part of this effort, the researchers 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 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 [5]. 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 1960s, 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 [18]. 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 [30].

Finally, several other groups of researchers have developed languages meant to do functions similar to Semantic Modeling. These include Simple HTML Ontology Extensions (SHOE), DARPA Agent Markup Language — Ontology (DAML-ONT), and Unified Problem-Solving Method Development Language (UPML) [16]. 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.

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 report, 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 the early stages of developing M and the practice of Semantic Modeling, there appears to be great potential to fulfill a need in the industry to improve the integration of models and data.

The prospect of sharing, through standard languages and protocols, the collective efforts of modelers throughout the world is beyond enticing. It has the potential to revolutionize nearly every aspect of human endeavor, as well as provide unprecedented benefit and savings across industry and commerce. Yet the challenges and difficulties are extraordinary, from theoretic achievability to practical implementation. Still the rewards make the journey well worth pursuing, which may lead to a true Intelligent Modeling Network.

REFERENCES

1. Allen, Stuart J., and Edmund W. Schuster. "Practical Production Scheduling with Capacity Constraints and Dynamic Demand: Family Planning and Disaggregation." *Production and Inventory Management Journal,* Vol. 35, No. 4, 1994.

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

3. Allen, Stuart J., Jack L. Martin, and Edmund W. Schuster. "A Simple Method for the Multi-Item, Single-Level, Capacitated Scheduling Problem with Set-Up Times and Costs." *Production and Inventory Management Journal,* Vol. 38, No. 4, 1997.

4. *APICS Dictionary.* Published by the Association for Operations Management, Alexandria, Virginia, USA, 2004.

5. Bond Graphs (www.bondgraphs.com/about.html).

6. Brock, David L. "Beyond the EPC — Making Sense of the Data Proposals for Engineering the Next Generation Intelligent Data Network." The MIT Auto-ID Center (now EPCglobal, Inc. and Auto-ID Laboratories), June 2003.

7. Brock, David L. "The Data Project — Technologies, Infrastructure and Standards for Distributed Interoperable Modeling and Simulation." MIT Data Project Workshop, September 2003.

8. Brock, David L. "Intelligent Infrastructure — A Method for Networking Physical Objects." *MIT Smart World Conference,* April 2000.

9. Brock, David L., Edmund W. Schuster, Stuart J. Allen, and Pinaki Kar. “An Introduction to Semantic Modeling for Logistical Systems.” *Journal of Business Logistics,* forthcoming 2005.

10. Closs, David, and Anand Nair. “Supply Chain Management with Advanced Planning and Scheduling: Implications for Academia.” *Logistics Educators Conference,* 2001.

11. Corcoran, Elizabeth. “Giving Voice to a Billion Things.” *Forbes,* 6 September 2004.

12. Dinning, Mark, and Edmund W. Schuster. “Fighting Friction.” *APICS — The Performance Advantage,* Vol. 13, No. 2, 2003.

13. D’Itri, Michael P., Stuart J. Allen, and Edmund W. Schuster. “Capacitated Scheduling of Multiple Products on a Single Processor with Sequence Dependencies.” *Production and Inventory Management Journal,* Vol. 40, No. 5, 1998.

14. Engels, Daniel W., Robin Koh, Elaine Lai, and Edmund W. Schuster. “Improving Visibility in the DOD Supply Chain.” *Army Logistician,* Vol. 36, No. 3, 2004.

15. Feigenbaum, Edward A., and Julian Feldman. *Computers and Thought.* AAAI Press, 1995.

16. Fensel, Dieter, James Hendler, Henry Lieberman, and Wolfgang Wahlster. *Spinning the Semantic Web.* MIT Press, 2002.

17. Fine, Charles H. *Clockspeed.* Perseus Publishing, 1999.

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

19. Geoffrion, Arthur M. “The Formal Aspects of Structural Modeling,” *Operations Research,* Vol. 37, No. 1, 1989.

20. Geoffrion, Arthur M. “An Introduction to Structured Modeling.” *Management Science,* Vol. 33, No. 5, 1987.

21. Geoffrion, Arthur M. “Structured Modeling: Survey and Future Research Directions, Section 1.” 1 June 1999 (www.anderson.ucla.edu/faculty/art.geoffrion/home/csts/csts1.htm).

22. Hof, Robert D. “Amazon Joins the Search Party.” *BusinessWeek,* 27 September 2004.

23. Huang, Gregory T. “The Web’s New Currency.” *Technology Review,* January 2004.

24. Kar, Pinaki, Ming Li, and Edmund W. Schuster. “A Case Study of Computer Service Parts Inventory Management,” *Service Parts Resource Guide* (1st Edition). American Production and Inventory Control Society, 2003.

25. Koh, Robin, Edmund W. Schuster, Indy Chackrabarti, and Attilio Bellman. “Securing the Pharmaceutical Supply Chain.” The MIT Auto-ID Center (now EPCglobal, Inc. and Auto-ID Laboratories), 2003.

26. Koh, Robin, Edmund W. Schuster, Nhat-So Lam, and Mark Dinning. “Prediction, Detection and Proof: An Integrated Auto-ID Solution to Retail Theft.” The MIT Auto-ID Center (now EPCglobal, Inc. and Auto-ID Laboratories), 2003.

27. Lillian Vernon Corporation (www.lillianvernon.com).

28. London, Simon. “Big Blue Reinvents Itself Again.” *The Financial Times,* 9 October 2003.

29. 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. “Tools for Inventing Organizations: Toward a Handbook of Organizational Processes.” *Management Science,* Vol. 45, No. 3, 1999.

30. Math-Net (www.math-net.de).

31. Melnyk, Steven A. “Finite Capacity Scheduling.” *APICS — The Performance Advantage.* Vol. 8, No. 8, 1998.

32. Muhanna, Waleed A., and Roger A. Pick. “Meta-Modeling Concepts and Tools for Model Management: A Systems Approach.” *Management Science,* Vol. 40, No. 9, 1994.

33. Murphy, Frederic, H., Edward A. Stohr, and Ajay Asthana. "A Representation of Schemes for Linear Programming Models." *Management Science,* Vol. 38, No. 7, 1992.

34. Park, Andrew. "Can EMC Find Growth Beyond Hardware?" *BusinessWeek,* 22 October 2004.

35. Sarma, Sanjay, David L. Brock, and Kevin Ashton. "The Networked Physical World." MIT Auto-ID Center (now EPCglobal, Inc. and Auto-ID Laboratories), 2000.

36. Schuster, Edmund W., and Stuart J. Allen. "Visualization and Finite Capacity Planning in the Process Industries." *APICS International Conference Proceedings,* Nashville, Tennessee, USA, 1998.

37. Schuster, Edmund W., and David L. Brock. "Creating an Intelligent Infrastructure for ERP Systems; The Role of RFID Technology." American Production and Inventory Control Society. White paper published on homepage for APICS, 2004.

38. Schuster, Edmund W., and Byron J. Finch. "A Deterministic Spreadsheet Simulation Model for Production Scheduling in a Lumpy Demand Environment." *Production and Inventory Management Journal,* Vol. 31, No. 1, 1990.

39. Schuster, Edmund W., and Robin Koh. "To Track and Trace." *APICS — The Performance Advantage,* Vol. 14, No. 2, 2004.

40. Schuster, Edmund W., Stuart J. Allen, and Michael P. D'Itri. "Capacitated Materials Requirements Planning and Its Application in the Process Industries." *Journal of Business Logistics,* Vol. 21, No. 1, 2000.

41. Schuster, Edmund W., Daniel W. Engels, and Stuart J. Allen. *Raising the Bar (Code): The Value of Auto-ID Technology.* South-Western Educational Publishing, forthcoming 2005.

42. Schuster, Edmund W., David L. Brock, Stuart J. Allen, and Pinaki Kar. "Prototype Applications for Semantic Modeling." The Data Center, 2004.

43. Schuster, Edmund W., David L. Brock, Stuart J. Allen, and Pinaki Kar. "Semantic Modeling: A New Direction for Knowledge Management." *Cutter IT Journal,* Vol. 17, No. 12, 2004.

44. Schuster, Edmund, W., David L. Brock, Stuart J. Allen, Pinaki Kar, and Mark Dinning. "Enabling ERP through Auto-ID Technology." Strategic ERP Extension and Use, edited by E. Bendoly and F.R. Jacobs. Stanford University Press, forthcoming 2005.

45. Schuster, Edmund W., Tom Scharfeld, Pinaki Kar, David L. Brock, and Stuart J. Allen. "The Next Frontier: How Auto-ID Could Improve ERP Data Quality." *Cutter IT Journal,* Vol. 17, No. 9, 2004.

46. Weiser, Mark. "The Computer for the Twenty-First Century." *Scientific American,* Vol. 265, No. 3, 1991.

47. Welch, James. S. "PAM — A Practitioners' Approach to Modeling." *Management Science,* Vol. 33, No. 5, 1987.

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 helping to organize a new research effort involving the large-scale analysis of data. His interests are the application of models to logistical and planning problems experienced in industry. He has a BS in food technology from Ohio State University and a master's degree in public administration from Gannon University with an emphasis in management science. Schuster also attended the executive development program for physical distribution managers at the University of Tennessee and holds several professional certifications.

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 BS 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. Allen began his research career in the field of non-Newtonian fluid mechanics and has published more than 50 journal articles in engineering and management science. He has also owned and operated three businesses in Wisconsin and New York state.

David L. Brock is Principal Research Scientist at MIT and cofounder of 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. Dr. Brock 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 bachelor's degrees in theoretical mathematics and mechanical engineering, as well as a master's degree and a Ph.D. from MIT.

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's degree in mechanical engineering from the Indian Institute of Technology, Kanpur, and a master's degree in logistics from MIT.

Index of published issues

> Business Intelligence Advisory Service

Upcoming Topics

- ***Creating the Integration Framework for a BI/BPM Implementation*** **by Al Moreno and Greg Mancuso**
- ***Data Quality*** **by Michael Brackett**

This index includes Business Intelligence *Executive Reports* and *Executive Updates* that have been recently published, plus upcoming *Executive Report* topics. Reports that have already been published are available electronically in the Online Resource Center. The Resource Center includes the entire Business Intelligence Advisory Service archives plus additional articles authored by Cutter Consortium Senior Consultants on the topic of business intelligence.

For information

on beginning a subscription or upgrading your current subscription to include access to the Online Resource Center, contact your account representative directly or call +1 781 648 8700 or send e-mail to sales@cutter.com.

Executive Reports

Executive Updates

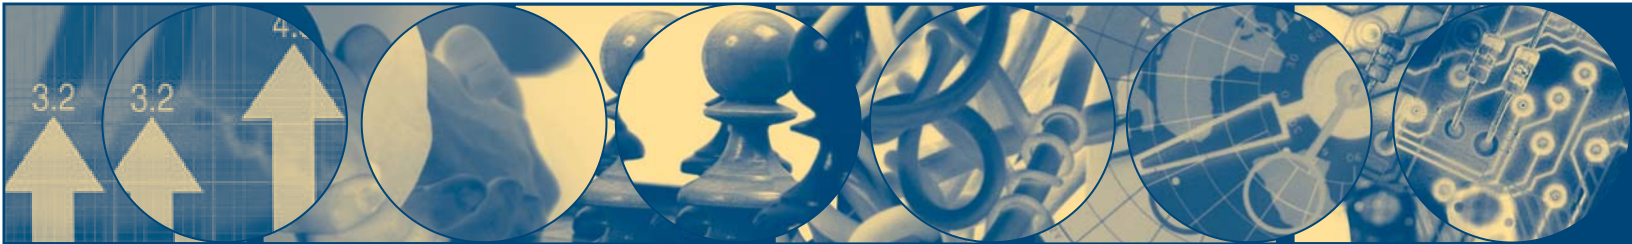

Business Intelligence Practice

The strategies and technologies of business intelligence and knowledge management are critical issues enterprises must embrace if they are to remain competitive in the e-business economy. It's more important than ever to make the right strategic decisions the first time.

Cutter Consortium's Business Intelligence Practice helps companies take all their enterprise data, augment it if appropriate, and turn it into a powerful strategic weapon that enables them to make better business decisions. The practice is unique in that it provides clients with the full picture: technology discussions, product reviews, insight into organizational and cultural issues, and strategic advice across the full spectrum of business intelligence. Clients get the background they need to manage technical issues like data cleansing as well as management issues such as how to encourage employees to participate in knowledge sharing and knowledge management initiatives. From tactics that will help transform your company to a culture that accepts and embraces the value of information, to surveys of the tools available to implement business intelligence initiatives, the Business Intelligence Practice helps clients leverage data into revenue-generating information.

Through Cutter's subscription-based service and consulting, mentoring, and training, clients are ensured opinionated analyses of the latest data warehousing, data mining, knowledge management, CRM, and business intelligence strategies and products. You'll discover the benefits of implementing these solutions, as well as the pitfalls companies must consider when embracing these technologies.

Products and Services Available from the Business Intelligence Practice

- The Business Intelligence Advisory Service
- Consulting
- Inhouse Workshops
- Mentoring
- Research Reports

Other Cutter Consortium Practices

Cutter Consortium aligns its products and services into the nine practice areas below. Each of these practices includes a subscription-based periodical service, plus consulting and training services.

- Agile Software Development and Project Management
- Business Intelligence
- Business-IT Strategies
- Business Technology Trends and Impacts
- Enterprise Architecture
- IT Management
- Measurement and Benchmarking Strategies
- Enterprise Risk Management and Governance
- Sourcing and Vendor Relationships

Senior Consultant Team

The Senior Consultants on Cutter's Business Intelligence team are thought leaders in the many disciplines that make up business intelligence. Like all Cutter Consortium Senior Consultants, each has gained a stellar reputation as a trailblazer in his or her field. They have written groundbreaking papers and books, developed methodologies that have been implemented by leading organizations, and continue to study the impact that business intelligence strategies and tactics are having on enterprises worldwide. The team includes:

- Verna Allee
- Stowe Boyd
- Ken Collier
- Clive Finkelstein
- Jonathan Geiger
- David Gleason
- Curt Hall
- Claudia Imhoff
- André LeClerc
- Lisa Loftis
- David Loshin
- David Marco
- Larissa T. Moss
- Ken Orr
- Raymond Pettit
- Ram Reddy
- Thomas C. Redman
- Michael Schmitz
- Karl M. Wiig