RC23005 (W0209-037) September 10, 2002
Other

IBM Research Report

Case Studies: Supply Chain Optimization Models in a Chemical Company

Young M. Lee
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

E. Jack Chen
BASF Corporation
3000 Continental Drive - North
Mount Olive, NJ 07828

Research Division
Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

Case Studies: Supply Chain Optimization Models in a Chemical Company

Young M. Lee* — *IBM T.J. Watson Research Center, Route 134, P.O. Box 218, Yorktown Heights, NY 10598, ymlee@us.ibm.com*

E. Jack Chen — *BASF Corporation, 3000 Continental Drive - North, Mount Olive, New Jersey 07828, chenej@basf.com*

In this paper, we give a short overview of the supply chain management models that have been used in the past few years by one of the largest international chemical companies. These models have made significant impacts in improving strategic, tactical, and operational supply chain processes. Then, we describe three supply chain optimization model case studies in detail: distribution network optimization; capacity requirement planning; and Web-based production planning/scheduling. For each case study, we describe motivation for developing the supply chain optimization model, requirements, modeling methods, deployment and business impact of the model. Using these case studies we intend to share our lessons learned, and address supply chain management issues that are especially relevant to chemical industry. These models utilize mathematical programming, discrete-event simulation, and Web-enabling technologies.

* The corresponding author.

1. Introduction

New business models and efficient management of supply chain are becoming critical success factors in today's highly dynamic and competitive business environment, driven by rapid advances in the information technologies and operations research methods (Geoffrion and Powers 1995). The goal of supply chain management (SCM) is to procure raw materials, manufacture products, and deliver the products to customers at desirable price and service. SCM requires coordination of the flow of products, services, and information among supply-chain entities, such as suppliers, manufacturers, distributors and customers (Keskinocak and Tayur 2001). Many companies are using enterprise resource planning (ERP) tools to improve or optimize their supply chain. However, ERP systems often produce unrealistic production scenarios that result in excess inventories, sub-optimal utilization of resources and ultimately poor customer service (Hsiang 2001). Therefore, it is necessary to model and optimize supply chain even if ERP systems are in place.

Modeling and optimizing supply chain management is much more affordable now due to relatively inexpensive computer hardware and abundant availability of supply chain modeling tools. The popularity of SCM tools is partly due to the advancement of the Internet, which allows easy access to such tools by supply chain decision makers. The Internet also facilitates supply chain coordination and collaboration with the suppliers and customers (H. L. Lee 2002).

The study and work described in this paper are based on supply chain management modeling activities of a large, international chemical manufacturing company. The company has been developing and using a wide range of supply chain management tools to better assess, analyze, and improve their supply chain. As typical supply chain characteristics of chemical industry, the company mainly produces functional products, which are defined as ones that have long product lifecycle and stable demand, with relatively stable manufacturing process. Most manufacturing processes are continuous processes which require high initial capital expenditure to setup, and run as built-to-forecast processes. Production rates have a minimum and maximum range for physical reason as well as economic reason. Profit margin is relatively low; therefore, economies of scale are very important. The life cycles of chemical products are usually long.

For customer demand, the company has been using demand forecasting tools to predict future customer demand that uses historic patterns and anticipated business changes. Chemical companies mostly produce functional products, which tend to have more predictable demand than the innovative products such as high-end computers (Fisher 1997, H.L. Lee 2002). Nevertheless, the customer orders placed for manufacturing can exhibit significant fluctuations due to the "bullwhip effect" (Lee, Padmanabhan & Whang, 1997). Accurate forecasting, especially for individual product and longer time horizon, is very difficult because customer demands depend on many dynamic factors, such as economic, social, behavioral factors, and unexpected events. However, customer demand is often the main input for many strategic, tactical, and operational SCM tools, such as distribution planing, transportation planning, manufacturing planning, and capacity planning. Therefore, it is important to forecast customer demands as accurately as possible. Typically, demand forecast is done in aggregated product level and shorter time horizon to ensure that the forecast is reasonably accurate and meaningful. These demand forecast models are often based on time-series modeling.

The company has been developing and using many distribution network optimization models for finished goods distribution using tools, such as SAILS (Strategic Analysis and Integrated Logistics Systems, Insight, Inc.) and MIMI (Manager for Interactive Modeling Interfaces, AspenTech, Inc.). Some of these models are for packaged goods, which are usually shipped by trucks and stored in warehouses. For packaged goods, since different products can be shipped together and stored together, there are opportunities for consolidations of storage and transportation. Some other models are for bulk liquid products, which are carried by tank rail cars or tank trucks and stored in storage tanks or splitted into smaller volume in transloading (rail to truck) facility. Since bulk liquid products cannot be mixed together for transportation or for storage, bulk liquid network models are mainly used to identify optimal facility locations. Some of the distribution network models focus on global level and covers entire business region such as the NAFTA (North America Free Trade Agreement) region, and some other models focus on individual businesses, describing more detailed distribution process of a business. The output of the network model in one level is often used as input for a model in another level. These distribution network optimization models are supplemented by inventory analysis tools; for example, discrete-event simulation models that analyze dynamic effects of various replenishment policies and outbound (customer) shipment patterns including seasonal effect of sales. These simulation tools generate dynamic inventory profile at distribution facilities, which is critical in deciding the size of storage tanks and warehouse space. Distribution network models were developed using the Mixed-Integer Linear Programming (MILP) methods.

The company has also been developing and using production planning and scheduling optimization models. Production planning models usually focus on a single manufacturing plant or several manufacturing plants that produce common products, compute the optimal production amount of certain products in each production line for each time period, usually in weekly or monthly buckets. The models take into consideration demand, production capacity, storage capacity, and raw material availability. Finite scheduling models focus on daily or hourly sequencing of manufacturing equipment and other resources, and try to minimize the Work-In-Process (WIP) and inventory level. Production planning and scheduling models are especially important when manufacturing plants are running at full or almost full capacity. In chemical industry, many manufacturing processes are continuous; therefore, well-managed product changeover is very important in minimizing the interruption of the manufacturing process, which is time-consuming and costly. Production planning models are usually built using MILP methods, and the finite scheduling models are built using the combinations of MILP and heuristics.

Capacity planning is another area in which the company has been actively developing and using decision support tools. For examples, discrete-event simulation models have been used to determine the size of the railcar fleet that are used in transporting bulk materials from manufacturing plant to storage site and eventually to customers. We considered factors, such as transit time, customer dwell time (the length of time that the railcar stays in customer's premise), loading, unloading, maintenance and other activities, and identified the optimal number of railcars to ship products to customers on time for each business group as well as for the whole corporation. We also developed simulation models to determine optimal production and storage capacities.

In transportation area, we have been using a shipment consolidation model to optimally consolidate less-than-truck-load (LTL) shipments into a truck-load (TL) shipment and to compute optimal routes that minimize the total transportation time.

In the following chapters we describe three case studies of supply chain optimization models mentioned above, and discuss important issues in developing the models, implementing the solutions and the benefits. The three case studies are a finished good distribution network optimization model, a storage capacity requirement planning simulation model, and a Web-based production planning/scheduling optimization model.

2. Case Study 1: Finished Goods Distribution Network Optimization Model

The chemical company has been growing rapidly in the past few years through various acquisitions and divestitures. As seen as a common business practice in this dynamics business environment, the company has been constantly looking into adding businesses that improve its business portfolio, and divesting businesses that are not part of its core businesses. When a business is acquired, its distribution network is also inherited into the corporate network. Similarly, when a business is sold, a slice of the corporate network is removed. Over time the corporate distribution network became inefficient, and consisted of many independent and fragmented networks. The company wanted to assess the current distribution network, identify opportunity for network consolidation and improvement, and to implement the new optimal solution to realize the benefit as soon as possible. The company also wanted to put in place a network analysis tool that can be used periodically to study network as the network evolves with business.

The assessment study of the company's current network clearly showed that the network had several inefficiencies. One of them was operating too many distribution centers (DCs); more than 100 DCs were being used in the NAFTA region at the time of this study. When so many distribution centers are used, the amount of products stored and shipped from each DC is relatively small. Moreover, there were many small shipments that originate from each DC. The relatively small storage and transportation volumes make it difficult to obtain volume-discounted warehousing and transportation rate from service providers. Another inefficiency was customer assignment; some customers were serviced by DCs that are unreasonably far away. In certain cases, a DC located in the west coast of the U.S. was shipping products to customers in the east coast and vise versa. This was partly because individual business unit within the corporation was using only its own distribution network without utilizing other facilities available for the whole corporation. When the customer assignment is not efficient, the outbound shipment (customer-bound) lanes tend to be unreasonably long, thus increasing the transportation time and costs. Moreover, each DC needs to have certain level of safety stock to accommodate uncertain demands and unforeseen production problems (Brown et al. 2002). Long distance transportation to customers also requires relatively high amount of safety stocks in warehouses. Therefore, having both too many DCs and long transportation lanes contributeed to the high level of overall safety stock, and caused relatively large inventory holding costs. Thus, it was evident that the distribution network needed to be optimized.

The primary objectives of the network optimization project were to identify the optimal number and locations of strategic DCs, and to optimally assign customers to appropriate DCs to lower overall distribution costs and improve customer services. By consolidating DCs, the economy of scale can also be realized. With larger amount of products stored and shipped from each DC, we can have leverage for negotiating better storage and handling costs with warehouse service providers. With smaller number of transportation lanes and larger volume in each lane, we can also have leverage for negotiating better transportation costs with transportation service providers. The shorter distance between DCs and customer will also improve the customer services. The overall safety stock requirement will go down, therefore, lowering inventory carrying costs. Also, working with a smaller number of service providers simplifies the whole distribution process. Operating smaller number of DCs also allows more opportunities to consolidate cross business LTL shipments into TL shipments, which is much less expensive; therefore reducing the overall transportation costs.

Distribution network usually degrades over time, similar to entropy in thermodynamics, which naturally moves toward a higher degree of disorder. Especially in today's dynamics business world, it is very likely that a business will change in many unexpected ways altering the supply chain substantially. Distribution networks have to be re-evaluated every few years and be re-optimized. Another objective of the project was to make available a strategic distribution network optimization model that can be re-used periodically with simple changes of data and constraints, and to identify opportunities for network improvement.

2.1 Model

A very large scale MILP model was developed to model and optimize the finished goods distribution network. The model consists of roughly 40 manufacturing sites, 100 candidate DCs, 300 demand regions, 100 aggregated product groups, and million shipment transactions per year.

In the past twenty years, many studied have been done in modeling and solving complex distribution network problems (Geoffrion and Graves 1974, Bradley et al. 1977, Brown and

McBride 1984). However, it is very difficult to develop a comprehensive, global distribution network model that would optimize over all aspects of businesses (Camm et al. 1997) in a large company because the model size would be too large to manage. Therefore, we modeled the distribution network in two levels; global and business level. In the global level, we modeled the entire business, e.g., the NAFTA region, and this model was primarily used for identifying optimal number and locations of strategically important distribution facilities. In this global level model, we used a higher level of product aggregation and simpler transportation rate structure to focus on the corporate-wide network rather than details of each business group. The solution from the global level model was passed onto the business level models. In business group level modeling, we focused on individual businesses, modeling many details of a distribution network of one business at a time. The business level models were primarily used for computing customer assignments and detailed cost calculations, which are needed for implementing the network solution. Individual business network models contained various business specific constraints, such as service level requirements, and used much less product aggregation and more accurate, lane by lane, transportation rate structures.

Transportation rates for replenishment flows (plants to DCs) were assumed to be all TL shipments, and we used a customized tariff for the company. Transportation rates for direct shipments (plants to customers) and outbound shipments (DCs to customers) consisted of TL and LTL shipments, and were computed as weighted average rate using the historic profiles of shipments and TL and LTL tariffs, which are company specific tariffs and are dependent on weight breaks. The historic shipment profile were computed by analyzing one year's transaction data from corporate ERP system.

The distribution network optimization model is a strategic tool that helps decide the optimal network structure by generating aggregated information, such as annual throughput at each facilities, annual transportation costs, and average service levels. However, the model is not adequate to address the dynamic effects, such as shipment size, frequencies of replenishment, outbound flow on the network, and seasonal demand fluctuations (Cheung et al. 2001). The dynamic effects are critical in analyzing dynamic profile of inventory and determining the safety stock, DC capacity requirement and inventory policies. We developed a discrete-event simulation model using eM-Plant (Tecnomatix, Inc) to supplement the network optimization model by analyzing dynamic profile of inventory. The simulation model was used to analyze various replenishment policies and outbound shipment patterns to determine the DC capacity requirements.

2.2 Results

The distribution network optimization models were successfully optimized for many business scenarios. A distribution network model is a rough description of a real distribution process with many assumptions and approximations. In business world, obtaining the global optimization solution and taking it blindly as the optimal business solution is risky. The distribution network model deals with many uncertain data, such as future customer demand, warehousing costs, transportation rates, and efficiency of managing distribution network. Therefore, finding the mathematically optimal solution is not as significant as finding insights and a set of solutions that can be rationalized and implemented easily. We experienced, in certain cases, that it takes many hours of computation just to improve a good solution by a fraction of 1%, which is much smaller than the error of data and is rather insignificant. We treated the distribution network optimization model as a decision support tool, which provided useful information to business decision makers so that a good decision is made and understood. The network solution we presented to the decision makers was a set of optimal network scenarios that clearly explained the tradeoffs among important components of distribution network such as distribution costs and customer services.

Our distribution network study indicated that the optimal number of DCs is between 6 and 10, with more than 95% of products reaching customers within next day delivery service (within 450 miles from a DC in our case). This result was also intuitive because 6 to 10 circles with radius of 450 miles should be sufficient to cover all the customers regions of the continental U.S. and Canada. The savings from the optimized network was several million dollars in transportation and warehousing costs, which was about 10% improvement of the network. The customer service improvement was about 35%, with most of customer serviced in next day delivery. The optimal distribution network was reviewed with all the business groups in the company, and was approved. Another benefit of the modeling was the fact that during the modeling process the organization has obtained much better understanding of the distribution network as well as its business.

There are many challenges in modeling distribution network, especially for large company that consists of many business units with their own business goals and needs. One of the challenges is obtaining the data required for the model. The optimization computation is based on data on model parameters, such as shipment transaction data, transportation costs, DC variable costs (handling and storing), and DC fixed cost. If the data were not accurate, the optimal results would be wrong, too. The company has been in the process of standardizing ERP systems when the model was being built. Therefore, there were more than one data sources; some businesses had their data in one ERP system, and others had data in another. Each ERP system had somewhat different data elements, format, unit of measure and naming convention, and it was time-consuming to unify the data from multiple sources into a consistent form.

Another difficulty of the modeling was to have all the business groups to participate in the modeling activity. It was extremely important that all the business groups provided the necessary data for the modeling, and validated the network solutions. The primary objective of the distribution network optimization was to consolidate distribution activities of all the business groups into a corporate-wide strategic network. The benefit was geared toward the corporate-wide optimization, not individual business optimization. Majority of business groups would benefit by participating in the optimization; however, a few business groups may not save money and even lose money as a consequence of implementation of new network solution. It was very difficult to convince those business groups to sacrifice for the benefit of the whole corporation since business leaders are compensated by the performance of their own businesses not by that of the overall corporation. And, there is always resistance to change. Implementation of the network solution involves changes in the business processes, and it is not a painless task. Moreover, there is often lack of trust on mathematical solution. Business leaders often feel that supply chain management is too complicated to model mathematically.

Modeling distribution network is relatively easier than implementing the network solutions. The implementation includes shutting down some DCs, which involves termination of employments. It also includes opening new DCs, which requires evaluating and selecting DC service providers from several candidates or even building private DCs. Requirements for new DC are computed from the network optimization model, and they include annual throughput, expected turns of products, special storage requirements such as for flammable and refrigerated products, frequencies and sizes of replenishment, and outbound shipments etc. Implementing the network also involves selecting new transportation service providers and canceling existing agreements. The transportation service requirements are transportation lanes, shipment profiles (frequencies and sizes), etc. The requirements for DCs and transportation are communicated to service providers with RFP (request for proposal). Once the proposals are received, they are carefully evaluated, and one that promises the least cost and the best service is selected. The network optimization models are critical not only for designing the network but also for generating the essential information required for the implementation of the optimal network.

3. Case Study 2: Web-Based Production Planning/Scheduling Optimization Model

One of the polymer manufacturing plants in the chemical company was facing problems of production capacity and production planning flexibility. Customer demand has been increasing, and due to the dynamics of economics it was more difficult to predict future customer demand. The manufacturing process is a continuous process, and the process has to be interrupted often to switch over from producing one product to another, and each interruption idled the production for two weeks. There was one production planner at the plant who has to continuously communicate with a product manager at the business headquarter, and it took a few days for the planner to generate a production plan based on the demand input from the product manager. Quite often, a production plan has to be modified to accommodate changing customer demand, and it also took a few days to change the plan. The planner has been using a rather old planning spreadsheet to display the input data and to generate planning report. The company called for a better tool that can improve the quality of production plan, reduce the planning time and has the flexibility of rapidly modifying production on-demand. It was also important that the tool is used both by production planner at the plant and a product manager at the headquarter. We developed a MILP-based production planning optimization model that runs on the Web for this problem.

The plant has multiple production lines and produces hundreds of millions of pounds of a polymer per year. Each production line is a continuous process that operates 24 hours a day, 365 days a year except during maintenance period when the production is interrupted for a few days. The raw materials are brought into the facility, usually by pipeline, and are fed continuously into the process. The plant produces multiple product grades with different physical properties. However, the product changeover cannot be done easily. When the plant switches from one product to another in a production line, the line will produce an off-grade product for a few days to weeks until the process reaches a steady state and produced a product with the desired specifications. Off-grade products can be sold but at a much lower price than products that meet the specifications, thus minimizing the number of product changeovers is important. Furthermore, the plant is operated at full capacity; therefore, it must have a well-planned product-changeover schedule to maintain overall production levels. Each production line can produce only certain product grades, and is constrained by a minimum and maximum production rate. The permanent storage tanks for the finished products have a limited space and temporary storage space is costly to use. Each production line has a minimum length for production campaigns during which product changeovers are not permitted. The goal of planning production is to compute a production plan that minimizes the inventory holding costs and product changeover costs while satisfying all the customer demands for finished goods and other process constraints.

The Internet greatly facilitates the deployment of highly interactive applications. With the Internet, it is now possible to deliver computational services that were once available only to those employees with specialized computer training and access to special computers and tools. The Internet-based computational applications can be accessed from virtually any Web browser on any computer anywhere in the world at any time to perform complex computational tasks. Optimization is one such computational tool that can provide lots of benefits when it is available on the Internet. Optimization has been used widely in industry for solving complex business problems. The company has been using MILP to optimize distribution networks, production planning, and scheduling. However, until recently, the users of such optimization applications had to have powerful computers with special optimizing engines and other data interface utilities. Many of the optimization models used in the company have been standalone applications and have lacked standard interfaces with other enterprise applications, such as data warehouse (DW) or ERP systems. Therefore, it has been difficult to deploy such optimization tools to multiple users throughout the company. Also, communicating optimization results among users have not been easy.

With the optimization tools on the Web, virtually anyone within the allowed community on the Internet or Intranet can access the complex optimization tools without any special hardware or software. It is now easy to make optimization technology available to many people. The optimization engine can reside on only one powerful server with enough computing resources (or in rare cases, a few servers). Moreover, a network of computers can serve as a parallel and distributed processing server environment for solving computationally intensive problems. Communicating the optimization results among the users, especially among business managers, engineers, and production planners, is easy with the Web-based tool, because the optimization results are stored in a centralized server and can be accessed by and presented to the users through a flexible and powerful medium, such as Hyper Text Markup Language (HTML). Furthermore, maintaining a Web-based optimization tool is much easier than maintaining traditional optimization tools installed individually on each user's computer. One can modify or enhance a Web-based optimization tool on one server, and all the users can access the change immediately. Supply Chain Optimization tools are particularly well suited to Internet innovations.

Therefore, we designed, implemented, and deployed an interactive Web-based optimization tool for this production planning optimization problem. The framework we developed is general and modular, and it can be used for developing similar tools for other businesses with the corporation. This tool permits users to change the objective functions and constraints of optimization models using a Web-browser and to run optimization and view the results in HTML pages. The users do not need to use FTP or TELNET protocols. In our framework, the input and output presentations are dynamically generated from a JSP (Java Server Page) that resides on a Web server. The Web is a client-server application; the client is a local computer and the server is a remote host (computer). The input data are taken from the clients and passed to the application (optimization) server, where an optimization model is executed remotely. Typically, the server is a powerful, high-end computer. After the optimization is complete, the results are passed from the server back to the client computers in the form of a standard HTML document, which users can view on the browser. The client computer could be any computer with a Web browser and an Internet access.

3.1 Model

We formulated the production planning problem as a MILP model, and used XPRESS-MP to model the problem and to optimize the model. Lee and Chen (2002) describe the details of the mathematical formulation for this model.

Java technology has revolutionized computer use, and many Web-based applications are being developed in Java. However, most optimization modeling and optimization packages, such as ILOG OPL (Optimization Programming Language) and XPRESS-MP (by Dash Associates), are based on C and other traditional programming languages. Thus, it was impossible to call those optimization subroutines from Java directly until recently. Fortunately, Java provides the Java Native Interface (JNI), which allows Java to interface with other popular programming languages, such as C or C++, Visual Basic, and Fortran, which can be interfaced with most optimization packages. We developed a framework for calling optimization subroutines from Java via JNI with Web browsers.

Figure 1 shows a three-tiered architecture for web-enabled optimization tools. The first tier on the client side processes the input data and presents the output data. The second tier is the Web-server, which manages the server-side processing and communicates with third tier servers such as the database server and the application server. The JSP and Java programs, the database system, and the optimization engine can all run on one server. However, because of the security and performance reasons, it is better to run them in separate servers, for example, on a Web server, a database server and an application (optimization) server. Users can use a Web browser in the client computer to edit input data, invoke execution of the optimization, and receives results via Web pages. The main implementations and processing tasks are carried out on the server side. This framework provides the flexibility of a programming language in a production environment, and developers can customize connections among models, data sources, and user interfaces.

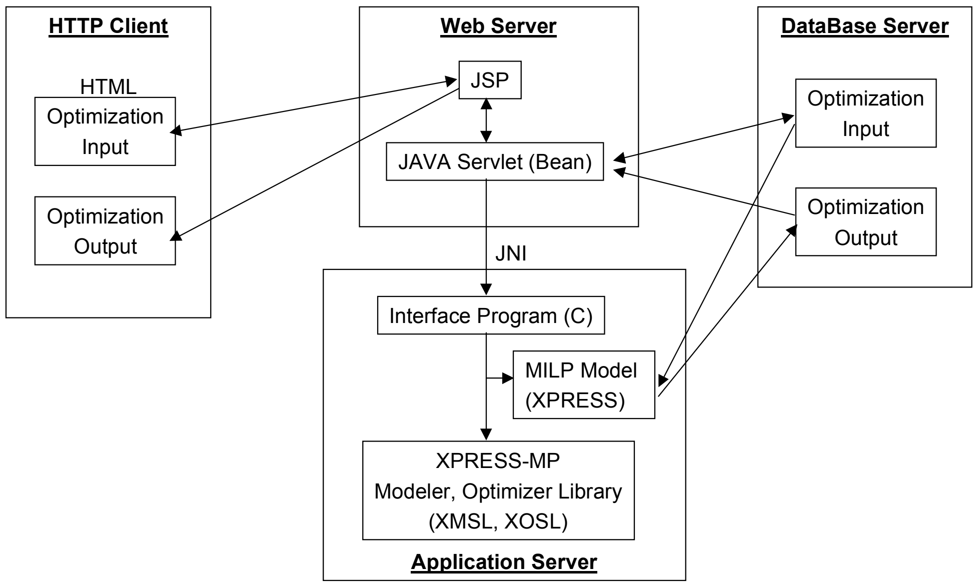

Figure 1: Architecture of the Web-based optimization.

Some optimization-service Web sites deliver Web-based optimization by providing FTP utilities; users upload their local model and data files to the host computers and remotely invoke the execution of the optimization model. Some other optimization-service Web sites provide text boxes that allow users to edit their model and data files in the server. These implementations are fairly easy and straightforward. However, they require users to know quite a bit about the optimization model and the optimization engine to use the tools.

In chemical industry, however, users of optimization tools typically are not people trained in mathematics or optimization. Therefore, to be useful, optimization tools must be easy to use. In our implementation, we provided a user-friendly interface to allow people with little knowledge of mathematical modeling to easily operate the optimization model. Users of the models control their optimization goals, such as minimizing cost or maximizing profits, and constraints by changing the parameters in the user interface screen; however, they often don't need to understand the mathematical models and solver to use the model.

The optimization process starts when the users make an HTTP request for a JSP from a client computer. The JSP technology enables rapid development of dynamic Web applications that are platform independent. JSP separates the user interface from content generation, making it possible to change the overall layout of the Web page without altering the underlying dynamic contents. Pekowsky (2000) describes how JSP works with HTML in detail. The JSP program makes a connection to optimization input files via a Java Servlet (JavaBean) and dynamically displays the names of the optimization input data tables associated with the model.

Using the input-file-selection page, the user selects one or more tables from the files, for example, demand forecast; then a JSP generates and displays an HTML page with those data. The data are displayed in a tabulated format; values that can be updated are displayed in text cells. Users can change the data by typing over the displayed data in the text cell on the browser. The modified data updates the files through the JSP and a servlet. Most solvers provide the option to decouple the model and data files. The high-level algebraic formulations describe optimization models in concise, symbolic formats, and an accompanying data file specifies the model instance to be solved. However, different solvers may require different formats for the model and data files. A servlet will translate data from the browser into solver-specific formats. Moreover, some solvers provide the option to read in data from spreadsheets and databases based on the SQL queries in the optimization model. If solvers do not have the option to interface with spreadsheets and databases directly, one can write a Java servlet to retrieve data from databases, translate data into solver-specific formats, and write them to data files. It is also possible to retrieve or update data from multiple databases using different access methods and protocols as long as they are available on the network.

Once the users update the input data, the JSP program calls a servlet that runs a JNI with a C-based program, which in turn runs the optimization engine with an optimization model. The C-based program will initiate a command instructing the optimization model to read in data from data files specified in the model and to bind the data with predefined variables before it initiates the optimization command. Thus, the model and data are completely independent. When the optimization is completed, the results are updated to the output files. A JSP program then dynamically displays the selected output report tables to the user's browser as HTML pages. The user can select any tables and view the optimization results. When the user selects the output file of interest from the drop-down menu; a JSP generates and displays an HTML page with those data. A downstream application uses the optimization results to generate daily production data. A downstream application uses the optimization results to generate daily production schedules, hourly raw material feed rates, and production reports, such as Gantt charts.

The framework will work with any optimization packages as long as they have C-based optimization library. Lee and Chen (2002) detail the implementation of this architecture.

3.2 Results

The production-planning model we implemented has 1,452 variables (308 binary variables) and 1,113 equations. The optimization usually takes a few seconds on a Sun/Solaris workstation, running Netscape Enterprise Server. The response time depends on other factors as well, such as the state of the application server and the load of network traffic. Before we implemented the model, the production planner took several days to plan a production schedule. The short response time of the integrated model has allowed the production plant to adjust its production schedule quickly to accommodate any sudden market changes. The quality of production planning has also improved. The planning model has helped the production planners to reduce inventory and to improve the utilization of production and storage capacity. In addition, business managers in different locations are now able to view the results and make intelligent business decisions quickly. The maintenance and technical support of the model have become much easier too. We modify and enhance the tool in only one server, and all the users can access the change immediately.

One of the benefits of integrating the optimization engine with the Web is that we can easily implement a parallel and distributed processing capability into the infrastructure. So far, application of parallel processing has been limited. Until recently, parallel computers could be found only in research laboratories or large universities. Furthermore, system software to support large-scale distributed processing remains scarce (Luo et al. 2000). On the other hand, an inherent characteristic of the Web is its distributed processing nature. The Internet can emulate the parallel processing architecture of expensive parallel hardware, while the Internet protocol unifies diverse networking technologies and administrative domains. A network of several computers within the Intranet or Internet can collaborate to solve complex optimization problems, effectively utilizing computers that may be unused otherwise. In solving an MILP problem, for example, a main server can generate sub-optimization problems through the branch-and-bound method, and several computers within the network can optimize the sub problems in parallel while the main server orchestrates the overall optimization strategy. Commercial optimization software, such as XPRESS-MP, can support such parallel-processing architecture. Therefore, the marriage between Web-based optimization and parallel and distributed processing seems natural.

The Web-based optimization infrastructure we developed is a generic framework that can be applied to a variety of optimization applications. The JNI interface between the Java class and the C-Interface program is generic and can be customized easily for other Web-based computational applications.

Once an Operations Research tool becomes available through the Intranet or Internet, it can be further integrated with other enterprise applications. For example, the input parameters of an optimization model can be updated in databases regularly through some other enterprise applications. End users of the optimization tool do not need to modify the input data themselves. Moreover, the optimization results can be stored in databases for other applications.

4. Case Study 3: Storage Capacity Requirement Planning Simulation Model

The company was planning a major production capacity expansion of one of its major product plants. The upstream of the process, i.e. manufacturing, was already scaled up by a group of engineers, and the company wanted to make decisions on the downstream processes, i.e. storage, mixing, packaging and transportation. There were a number of existing silos for mixing and storage, and a packaging equipment for the plant. Adding new equipments, especially silos, is very expensive. However, sufficient silo space is very important. A shortfall of the silo space will interrupt the manufacturing process because the process is continuous and the product coming out from the manufacturing end would not have any place to go. The shortfall also affects the transportation process thus affecting customer service level as well. The company called for an accurate analysis to decide whether, how many and what size of silo and packaging equipment are needed for the plant. We developed and used a discrete-event simulation model to analyze the problem. The simulation model helped us in determining capital equipment requirements and assessing alternative strategies for the logistics operations.

The plant produces three different grades of a dry chemical (denoted as A, B and C) at a specific production rate. These three different grades are produced in a continuous cycle with a fixed quantity for each grade. The product is transferred to a storage tank, from which it is distributed to another facility of further processing and packaging. A larger portion of products is sent to railcar for shipment, and the rest is sent to truck for shipment. Furthermore, the sequence of railcar and truck shipment is random and mixed. The capital outlay of such facilities is tremendous, and the designer needed a credible, valid, detailed model of operation.

There are several large volume silos available for the plant. However, only one silo can be used to receive the production outflow from the plant at any given time. The outflow from the silos cannot take place until the silo has completely filled. This is necessary because a batch number will be assigned to a particular silo load so that the source and quality of the product can be traced. Only one outflow from the silos can take place at any given time. Grade A of the product requires special blending and needs to be kept in the silos for at least twenty-four hours. There is one RailSilo used to load railcars, which has a loading capacity that is a multiplication of railcar load to avoid less-than-full-load railcars. The RailSilo cannot have flow-in and flow-out at the same time. There is one BagSilo used for the bagging process. The BagSilo has a smaller capacity than other silos, however, it can have flow-in and flow-out at the same time. The flow-out rates from all the silos are all fixed.

While bulk railcar shipments do not require special packaging, truck shipments need to be bagged first. The bagging-process will produce a certain volume of bagged products every few minutes. The bagging machine requires a few minutes of maintenance after processing a certain volume of the product. It takes a few minutes to change over between different grades of products. The bagging machine breaks down occasionally and needs to be stopped for repair.

When trucks arrive at the plant, they are weighed at the weigh station, and the process take a few minutes. A fixed fraction of the arriving trucks are here to pick up our bagged product. The remaining fractions are here for other purposes. There is a fixed number of loading docks, and it takes a few minutes to load the truck, which also has a fixed capacity. Once the truck is loaded, it needs to be weighed again before it can leave the premises. Both the inbound and outbound trucks use the same weigh station. If there are more than one truck waiting for the weigh station, the order of trucks go to weigh station will be based on first come first served basis.

The main objective of the simulation study is to ensure the process configuration and capacity can support continuous outflow of the manufacturing plant, and to optimize the number and size of storage silos. There are several standard size silos under consideration. The decision is not only to select more silos with smaller size or fewer silos with larger size but also to optimize the combination of the number and size of silos. Moreover, the activities of bagging process, the activities of railcars and trucks, such as the inter-arrival time of railcars and trucks, are analyzed to ensure continuous material flows are maintained without interruption. Once the average inter-arrival time is determined, the size of railcar fleet can be calculated indirectly. The simulation model can help us not only to verify the feasibility of our configurations but also to search for the optimal configuration among several alternatives.

4.1 Model

Many production activities in chemical industry involve continuous material flows, such as liquid, gas or solid, and it is very costly to interrupt and restart the production process. Simulation is a useful tool to study dynamics in such processes in a simulated environment. Simulation models do not only provide quantitative information that can be used for decision-making but also increase the level of understanding of how the process works. Most models are used to simulate discrete events. Discrete-event simulation concerns the modeling of a system as it evolves over time by a representation in which the state variables change instantaneously at separate points in time (Law and Kelton 2000) and has a commendably long and successful track record in the improvement of manufacturing process (Law and McComas 1997).

Although in the chemical manufacturing plant, materials are mixed and transferred as continuous flow through a maze of tanks and pipes, we did not have to model the continuous components to effectively study throughput issues. We defined the product in a batch that uses various resources for a period of time simply by the amount of fluid being transferred and the rate of transfer. We used discrete-event simulation to model the continuous material flow in this plant. In many real world applications the behaviors of discrete event and continuous process are often interdependent. Note that several simulation packages have the capability to build hybrid discrete/continuous models. Some researchers have developed simulation models to analyze the hybrid nature of chemical manufacturing plant (Watson 1997, Saraph 2001). Some simulation issues in this area are conceptualizing production operations for simulation, discretization of continuous processes and building adequate level of detail in the models (Chen et al. 2002).

The output of the manufacturing plant is continuous at certain metric tons per minute. The transfer of a continuous flow from Silo X to Silo Y via Pipe S was simulated as a delay based on the amount being transferred and a fixed transfer rate. We discretized the continuous material flow to a fixed weight moving unit. The output was then converted as one unit every period. The weight per unit was initialized from a data table in the model. For example, if the output rate is 6 metric tons per hour, and the weight per discretized unit is 2 metric tons, then the output rate becomes 3 units per hour or one unit every 20 minutes. In general, with a smaller discretized unit weight, the simulation model can simulate the continuous material flow more accurately. The model can be regarded as continuous if the discretized unit weight is the weight of a grain of the product. However, it also complicates the simulation model because most discrete-event simulation software uses the next-event time advance mechanism for the simulation clock (Law and Kelton 2000).

We wanted to build a simulation model that allowed us to analyze the logistics system adequately without modeling unnecessary details. We chose two metric tons per unit for our model because it is the smallest incremental weight that the product are bagged and processed in this logistics system. This discretized unit weight allowed us to analyze the system adequately without complicating the modeling of the bagging process. If there are two different packaging sizes, for example 2 and 5 metric tons, the 2 metric tons per unit discretization will complicate the implementation of the simulation model. In this case, we would use one metric ton, which is regarded as the smallest incremental weight per unit.

One of the purposes of the simulation analysis was to find out the minimum required number and size of storage silos; therefore, the outflow control from the plant always searched the available silos from left to right as the downstream station. Thus, excessive silos will not be used by the system. The I/O control between the main silos and RailSilo, BagSilo determined which main silos should have outflow and which downstream silos the material should flow to. The outflow of main silos was based on the "first available" rule. The flow-in time was recorded when the material in the silo was ready to flow out. For example, grade A product may be stored in silo1 before grade B product is stored in silo2. But the flow out of grade A product cannot take place until the material has been processed in the silo for at least 24 hours. Therefore, the I/O control will select silo2 for outflow instead of silo1.

The bagged product was stored in the warehouse until a truck made a request. The warehouse was viewed as a sink of upstream stations, i.e., the warehouse had virtually unlimited storage capacity. However, the warehouse acted as a source for downstream stations. The material was stored in the warehouse until a truck was ready to be loaded. To reduce the warm-up period, we assumed that there is a certain volume of initial inventory in the warehouse.

One of the difficulties in developing this model was to simulate changes of the statuses of the silos. Once a silo was completely filled, there was no further inflow until the silo was completely emptied. The outflow of the silo became available immediately when the silo was filled, except grade A which needed to be kept for at least 24 hours. A complication arises because there is a lag between the outflow from the upstream station to the inflow of the downstream station. It will be too late to switch outflow from the upstream station when the receiving silo is completely filled, because the material in the pipeline will be lost. Thus, it was important to synchronize all the processes in this model. For example, the plant needed to send its outflow to other silos when the material in the pipeline filled the receiving silo completely. We accomplished the synchronization by making the material move instantaneously. As soon as the material left a station, it immediately appeared in its destination. The transfer time between stations was simulated after it reached its destination. The outflow control was embedded in the silo object, which can adjust the flow out rate. For example, if the current material flow is from Silo2 to RailSilo, then one unit will be removed from Silo2 every few minutes according to the flow out rate. The unit was added to the RailSilo as soon as it has been removed from upstream. This was possible because the capacity of the inflow rate of downstream was always greater than the upstream outflow rate.

The arrival of railcars and trucks were modeled as Poisson processes with mean inter-arrival time of a few hours. Previous experience indicates that the stochastic arrival process can be adequately simulated with the Poisson process, i.e. exponential inter-arrival, and the interval between break down and the time required to fix a machine can be simulated with a Weibull distribution (Law and Kelton 2000).

The visualization of the simulation model was very useful for users to validate the model. Visualization was also critical in communicating the outcome of a simulation study to the management. Decision-makers often do not have the technical knowledge to understand the statistical outcome of a simulation run. But through the visualization, the managers was able to see the status of the silos and the flow of material. The process of building the simulation model also gave an opportunity for the plant personnel and upper management to better understand the logistic process.

4.2 Results

Sargent (2000) described several methods to validate simulation models, such as animation, historical data validation, face validity, extreme condition tests, internal validity, and traces. To reduce uncertainty, we used historical data to build and drive the simulation model whenever possible. Face validity refers to asking people who are familiar with the process whether the model and its behavior are reasonable. We used their feedback to determine whether the logic in the conceptual model was correct. We validated the model through several extreme conditions, where the analytic solutions were attainable. The model output was then compared with the verified analytical results. For example, if we set up the simulation model to terminate in one month, we can verify whether all the material adds up. We can trace the material in certain states, such as the quantity shipped by railcar and truck, the quantity stored in different silos, the quantity processed by the bagging machine etc. Accurate statistical analysis is central to the validity of any simulation project (Law and Kelton 2000). Since we were simulating stochastic systems, we could not conclude our results with one simulation run. Internal validity refers to make several independent runs of the model to determine the stochastic variability. A high variability may indicate the system is sensitive to its input parameters, and the appropriateness of the simulation results needs to be investigated more closely.

The users agreed that our model was an accurate representation of the real system. To alleviate any concerns of the robustness of the results due to the random variations inherent in simulation, each scenario was run multiple times with different time horizons, one, two and three years. The modeling approach described above was used to evaluate various alternatives. Many of the alternatives were defined and modified only in the data tables. This flexibility allowed the user to read in data, run a scenario, and get results very quickly. No scenarios required modifications to the model itself. Moreover, when the modifications are necessary, the model can be easily and quickly changed due to the object-oriented design of the model. The results from the simulation provided a clear picture as to a best choice of planning.

Several scenarios with different numbers and different sizes of silos were used in our experiments. The scenario study provided valuable information, because the cost structure of the size of the silos was not linear. The optimal combination of the number and size of silos was determined with simulation of a pre-determined set. After several preliminary experiments, we determined that three mid-size silos are most cost effective and are able to support the continuous operation of the manufacturing plant. The followings are experimental results corresponding to the model. Since we hypothesized that three main silos will be enough to support continuous flow, we set up four silos in the model so that we will be able to verify our hypothesis. Of course, we can model the system with three silos and check whether an overflow occurred, however, we will not have the utilization information of the non-exist silo.

Table 1
Silo utilization statistics

Silo	**Utilization**
Silo1	63.98%
Silo2	63.91%
Silo3	10.93%
Silo4	0.00%

Table 1 lists the silo statistics for one particular replication when the model was simulated with one-year time frame. The report indicates that the fourth silo has not been used, thus, it is possible to remove the fourth silo without causing disruption of the production flow. The low utilization of Silo3 is also very re-assuring. If the fourth silo has been used, it is an indication that three silos are not enough to support continuous flow. Table 2 lists the bagging machine statistics. The report shows the utilization of the bagging operation is quite low at 43.65%, i.e., it is idle 56.35% of the time. In this plant, the designer purposely built a high-capacity bagging machine to accommodate the anticipated future expansion of the production capacity. Furthermore, the bagging machine is relatively inexpensive to build and operate. The report also shows that only 2.85% of the simulation time was used in changeover between different grades of product, and 1.66% of the time were used in maintenance.

Table 2
Bagging machine statistics

Bagging Machine	**Percentage**
Idle	56.35%
ChangeOver	2.85%
Maintenance	1.66%
Grade A	5.98%
Grade B	19.50%
Grade C	13.66%

The changeover time between different products is different. For example, it may take 20 minutes to switch from bagging Grade A to Grade B and take 30 minutes to switch from bagging Grade B to Grade C. The changeover information is stored in data tables, therefore, the bagging process will be able to simulate multi-products without any modification. The simulation results also provided information regarding the number of changeovers and the average time between changeovers.   This information was important in determining the campaign volume.

5. Conclusions

Chemical industry has unique supply chain characteristics such as continuous and stable production processes, large changeover cost, handling of bulk material, high volume per SKU (Stock Keeping Unit), long life cycle of products, and low profit margin. In this paper, we described an overview of practical supply chain management applications that have been used in a chemical company.

We also focused on three case studies and discussed the motivations of developing such tools, values that those tools have added to the company, issues that needed to be dealt with, and lessons learned. For the first case study, we described a large-scale MILP model that was developed to optimize distribution of finished goods. The optimization model consolidated distribution network that consisted of many independent and fragmented networks. The model generated substantial savings in distribution costs and drastically improved customer services. For the second case study, a generic computation framework for web-based optimization was described. The framework was developed using a server-side Java programming, and a practical production planning optimization model was successfully developed and deployed using the framework. The model improved the quality of production plan, flexibility of production plan change and accessibility of the tool. For the third case study, we described how we used a discrete-event simulation to model a logistic process and to determine capacity requirements of the storage and packaging facilities that allow a continuous production outflow and customer shipments. The simulation model reduced a capital expenditure substantially.

Acknowledgement

The authors thank the anonymous referee for valuable comments, which greatly improved the quality of this paper.

Online References

Underlined terms in the paper indicate online references.

INFROMS Resources (www.informs.org/Resources/Computer_Programs)

Java (developer.java.sun.com/developer/onlineTraining)

Neos Server for Optimization (www.mcs.anl.gov/neos)

PaperSuite 2.0 (www.optamaze.com)

Remote Interactive Optimization Testbed (riot.ieor.Berkeley.edu/riot)

SAS Institute (www.sas.com/solutions/supplychain/demos/index.html)

Servlet (java.sun.com/products/servlet/2.1/)

Statit tool set (www.statware.com)

References

Bradley, G.H., G.G. Brown, G. Graves. (1977). "Design and Implementation of Large Scale Primal Transshipment Algorithms", Management Science, 24:1, pp.1-34.

Brown, G, J Keegan, B. Vigus, K. Wood. (2002). "The Kellogg Company Optimizes Production, Inventory, and Distribution", Interfaces, 31:6, pp.1-15.

Brown, G., R. McBride, R. (1984). "Solving Generalized Networks", Management Science, 30:12, pp.1497-1523.

Chen, E. J., Y. M. Lee, P. L. Selikson. (2002). "A simulation study of logistics activities in a chemical plant," *Journal of Simulation Practice and Theory.* 10 (3-4): 235-245.

Cheung, W., L. Leung, Y.M. Wong. (2001). "Strategic Service Network Design for DHL Hong Kong", *Interfaces*, 31:4, 1-14.

Fisher, M. (1997). "What Is the Right Supply Chain for Your Product", *Harvard Business Review*, 75/2: 105-116.

Geoffrion, A. M., G.W. Graves. (1974). "Multicommodity Distribution System Design By Bender Decomposition", Management Science, 20:5, pp.822-844

Geoffrion, A. M., R. Power. (1995). "Twenty Years of Strategic Distribution System Design: An Evolutionary Perspective", *Interfaces*, 25:5, pp.105-127.

Hsiang, T. (2001). "The Illusion of Power", OR/MS Today, Vol.28, No.1, pp.34-36.

Keskinocak, P., S. Tayur. (2001). "Quantitative Analysis for Internet-Enabled Supply Chains," *Interfaces*, 31:2, 70-89.

Law, A. M., W. D. Kelton. (2000). *Simulation Modeling and Analysis*, 3rd Edition. McGraw-Hill, New York.

Law, A. M., M. G. McComas. (1997). "Simulation of manufacturing systems*," Proceedings of the 1997 Winter Simulation Conference," Atlanta, December 1997,* IEEE Press, New York. 711-717.

Lee, H. L. (2002). "Aligning Supply Chain Strategies with Prodiuct Uncertainty", California Management Review, 44:3.

Lee, H. L., V. Padmanabhan and S. Whang, (1997). "The Bullwhip Effect in Supply Chains", Sloan Management Review, 38/3: 93-102.

Lee, Y. M., E. J. Chen. (2002). "BASF Uses a Framework for Developing Web-Based Production-Planning-Optimization Tools," *Interfaces*, 32(6): 15-24.

Luo, Y.C., C.H. Chen, E. Yücesan, I. Lee. (2000). "Distributed Web-Based Simulation Optimization. *Proceedings of the 2000 Winter Simulation Conference," Orlando, Florida, December 2000*, IEEE Press, New York. 1785-1793.

Pekowsky, L. (2000). *Java Server Pages*. Addison-Wesley, New York, NewYork.

Saraph, P. V. (2001). "Simulating biotech manufacturing operations: issues and complexities," *Proceedings of the 2001 Winter Simulation Conference, Washington, D.C., December 2001*, IEEE Press, New York. 530-534.

Sargent, R. G. (2000). "Verification, validation, and accreditation of simulation Models," *Proceedings of the 2000 Winter Simulation Conference, Orlando, Florida, December 2000*, IEEE Press, New York. 50-59.

Sweeney, D.J., Dill, F.A., Wegryn, G.W., Evans, J.R., Camm, J.D., and Chorman, T.E. (1997). Blending OR/MS, Judgement, and GIS: Restructuring P&G's supply chain. *Interfaces,* 27(1):128-142.

Syrjakow, M., J. Berdux, H. Szczerbicka. (2000). "Interactive Web-Based Animations for Teaching and Learning," *Proceedings of the 2000 Winter Simulation Conference, Orlando, Florida, December 2000*, IEEE Press, New York. 1651-1659.

Watson, E. F. (1997). "An Application of Discrete-Event Simulation for Batch-Process Chemical-Plant Design," *Interfaces*, 27(6): 35-50.