Proceedings of the 2003 Winter Simulation Conference
S. Chick, P. J. Sánchez, D. Ferrin, and D. J. Morrice, eds.

BENEFITS OF USING A SUPPLY CHAIN SIMULATION TOOL TO STUDY INVENTORY ALLOCATION

Andrew J. Siprelle
David J. Parsons
Robin J. Clark

Simulation Dynamics
416 High Street
Maryville, TN 37804, U.S.A.

ABSTRACT

Using Simulation Dynamics' Supply Chain Builder, a supply chain tool has been developed to study alternative inventory allocations in a network. The model represents the supply chain of a nationwide food production and distribution network. The company produces thousands of product lines and is studying the best strategies for allocating inventory to distribution centers to avoid cross-shipments of product. This paper describes the data needs and logic of this simulation tool along with its effectiveness in comparing strategies for locating SKU inventory in a supply chain.

1 INTRODUCTION

Simulation Dynamics was asked to build and use a simulation model to help a major food manufacturer and distributor. A primary objective was to maintain or improve service levels, while reducing inventories and costs.

A project team used simulation to quantify how well alternative optimized networks would function through variation in demand and supply. At a high level, optimization was used to generate optimum networks, and simulation was used to evaluate whether the company could live with those networks.

The company also turned to simulation for answers to the following questions:

- What is the relationship between inventory policies and the resulting inventory levels, customer service levels, and redeployment of stock?
- Does the location of inventory storage for different classes of product have an effect on total inventory levels and redeployment of stock?
- Would better forecasting methods reduce the amount of inventory in the system and the redeployment of stock?

The purpose of this paper is to describe the model used as it relates to material allocation and redeployment, and to document the effectiveness of this model in evaluating the impact of different product deployment strategies on inventory levels.

2 MODEL DESCRIPTION

The simulation tool is a discrete event model, developed using Supply Chain Builder from Simulation Dynamics on Imagine That's Extend™ simulation platform. The model was created using standard Supply Chain Builder blocks (Phelps, Parsons and Siprelle, 2000) and a database built specifically for simulation. Some blocks were customized for use with this particular model.

2.1 Model Elements

Supply Chain Builder is a system of elements and processes (Parsons and Siprelle, 2000) that together create a dynamic system. The model consists of the basic Supply Chain Builder elements representing all the locations, items (materials and resources), inventories, and shipments in the network. These elements and the relationships among them are stored in the model database described more fully in the next section.

2.1.1 Locations

The model simulates a network of production facilities (plants) and non-production facilities (distribution centers) that respond to consumer demand for finished goods (SKUs). Over one hundred locations are modeled, most of them small distribution centers.

2.1.2 Materials and Inventories

"Brand" is a general term we use for a material type that is associated with a particular processing system. Each plant produces a range of finished goods – SKUs – that are produced from a single brand. Inventories at plants consist of SKUs and the brands used to produce the SKUs.

Equipment used to produce the brand needed to fill an order is collectively considered a processing system. Each plant has a set of parallel processing systems. A system is usually capable of processing only certain brands. The model database assigns, or *allocates*, the brands that may be produced on a given processing system.

Raw materials are not specifically modeled in this system. Although raw material supplies could easily have been modeled, they are not included because they were never a true constraint to the manufacturing system.

2.1.3 Transportation Methodology

Several approaches to modeling shipments were considered, but the approach decided upon uses a delay time associated with moving material from one location to another (dock to dock). This is a valid assumption because transportation is not a constraint in this supply chain. A table in the model database describes the delay time from each origin to every possible destination in the model.

Early in the project a difference from actual company accounting was noted with this approach: the company often had truckloads of inventory sitting at a distribution center that was considered "in transit" inventory, while our model, on the other hand, considered such truckloads as "available" inventory for the distribution center. The project team agreed to model this inventory as "available" due to the fact that the company would immediately unload and use the material on a truck if it was needed.

2.2 Data

A database adds depth to a model by providing the opportunity to represent multiple production facilities and huge volumes of inventory, locations, suppliers, and customers (Phelps, Parsons and Siprelle, 2001). In this model we have included two separate databases, a Reference Database and the model database, which is the framework used to drive the simulation. Two database tools, an automated legacy data feed and the Database Builder, were created by SDI to populate and pre-validate the databases.

2.2.1 Model Data and Data Tools

The model requires a significant amount of historic data used as demand data during the model run. Often historic data from the company's database must be restated in terms appropriate for a forward-looking model. We created a customized program that can read-in the company's legacy database and convert it into a format the model can use. This reformatted data is then uploaded to a Reference Database.

The Reference Database is a superset of actual historic data that includes all the data needed to simulate not only the existing supply chain's network of locations, products, and distribution flows, but also has the potential for crafting a network that never before existed. The historic data can be "re-wired" to simulate not only a copy of the existing network, but any variation on the network we desire to model.

The Database Builder program takes corrected and validated data from the Reference Database and creates a network configuration to construct a simulation-ready database framework for the model run. This program can generate completely different distribution networks for thousands of SKUs in as little time as 15 minutes. As it is generating the new model database, the Database Builder applies error detection rules that pre-validate the data before the model experiments run.

The Database Builder can create completely new arrangements of locations, flows, and products. The resulting model database can restate history in terms of a network that never existed, but that would be interesting to explore.

2.2.2 Output Reports

In addition to the rolled up reports from the model itself, detailed reports are created for verification and analysis. Supply Chain Builder can write detailed data to text files while the model runs. The data is typically written to text files during the model run simply because of their sheer size (up to 600 MB per file). Collected text files include daily records (such as inventory for each location for each SKU for each day) and weekly records (i.e., weekly consumption by location and SKU, Forecast Daily Demand (FDD) by location and SKU, Order Point by location and SKU, and Percent Demand Met by Location and SKU).

Other text file reports created for detailed analysis include:

- Missed Consumption Report – Lists every missed consumption event along with the amount missed.
- Systems Runs Report – Lists every production run. The data describes the plant, system, and brand along with the start date/time, changeover completion date/time, finish date/time, and the pounds produced.
- Shipment – Lists every shipment record.

From the detailed model output above, reports are generated using MS Access to present the information in a manner consistent with how the company keeps their actual data. This allows validation to be performed more effectively. The data from the simulation was summarized in a number of different ways to provide information in the following areas:

- Customer service
- Demand
- Deployment
- Inventory
- Production

- Redeployment
- Transportation

2.2.3 Data Validation

Once the database has been built and is simulation-ready, initial model outputs are meticulously analyzed to compare the model to the actual network. The database tools used in this model interface readily with widely available tools, such as Microsoft Excel and Access, to assist in scrutinizing the data.

The model generates summary reports inside the model database while the simulation is running. The reports are used as a first cut at verification and validation by describing such things as:

- Production at a location.
- Average SKU inventory at a location.
- Total consumption by inventories.
- Average percent demand met overall (weighted).
- Average percent demand met for each customer facing inventory location.
- Total redeployed from origin to destination.
- Total shipped from origin to destination.

After a model run, the above output data is collected in an Excel spreadsheet and summarized into subgroups of related data, organized by production and capacity. A system of summary reports are used to identify errors or unusual findings. The summaries flow from broader to more finely detailed to allow identification of possible problem areas and "digging down" deeper into the details of those areas.

Typical summaries include: total capacity vs. total demand at a plant; surplus capacity for a period; percent of demand met for an inventory; total days of production of a brand on a system; and SKU production details.

2.3 Inventory Parameters

Several variables are used to model inventory replenishment when demand is met. Among the most important are Safety Level (Days), Order Point, and Target Level (Days). Of critical importance to these variables are Average Daily Demand, Forecast Daily Demand (FDD), and Net Current Inventory figures. This section describes these parameters and how they are used in the model.

2.3.1 Average Daily Demand

The project team was supplied with over a year's worth of actual daily demand data for each customer-facing inventory in the network. A *demand model* (we refer to it as a model, because it is a sub-model of the larger model environment) was used to describe daily and weekly variations in demand from all locations. The demand model can also use multiple demand scenarios.

To reduce the amount of data in the model and to take into account seasonal and promotional variations, the historic daily demand data was averaged for each week to create a weekly average daily demand figure for each inventory. This figure was recorded as "Average Daily Demand" in the *Daily Demand* table in the database used by the model. To account for daily variations in demand, a random factor, derived from a statistical study of the daily variations around the weekly average, is applied to the weekly average daily demand.

2.3.2 Forecast Daily Demand

Future demand for products must be estimated in order to manage inventory and to know when orders for stock resupply should be placed. Any discrepancy between the estimate and the actual demand is termed forecast error. The model similarly forecasts demand. Forecast daily demand (FDD) is a variable based on the average daily demand entered in the Daily Demand table. FDD is a moving average, calculated daily. FDD is computed as the average of a user specified number of days forward in the Daily Demand table. A consumption block in the model has two parameters, "Days to look back" and "Days to look forward," which are used in the calculation to compute a results field, "Forecast Daily Demand," in the Inventories table of the database.

The model provides a multiplier that may also be used to incorporate greater forecast dispersion so that the error is accurately described. A study of the company's actual forecast and demand values was performed that produced a Mean Absolute Percent Error (MAPE) by inventory class.

The formulas used to calculate MAPE are:

$$APE(for \cdot all \cdot locations, weeks, and \cdot SKU's) = \frac{Absolute(Forecast - Demand)}{Forecast}$$

$$MAPE(for \cdot all \cdot locations \cdot and \cdot SKU's) = \frac{\sum_{1}^{52} APE}{52}$$

This is incorporated into our forecast to gave similar forecast error as that of the company.

2.3.3 Safety Level (Days) and Order Point

FDD is used in many calculations in the model, including calculating when a location should place an order for a SKU. The *order point* (in units) is the point at or below which an order will be placed.

A parameter called "Safety Level (Days)" is also used in this calculation. This parameter is included in the *Inventories* table for the model. Safety Level (Days) is expressed as days of material on hand, rather than an actual amount in units, to respond to changing levels of demand over time. The model multiplies the FDD by the days of safety to determine the units it needs to keep in stock each day as safety stock. In this model, the company defined the safety level as the amount of inventory required to achieve in excess of 98% of customer demand.

A *base inventory* may be used in the formula to allow the user to specify a fixed number of units as a base for any inventory. In a system where demand for an inventory could at times be zero, a base inventory will prevent inventory depletion. Base inventory is comprised of a user defined parameter multiplied by the average daily demand (average over the entire year) for the SKU.

The order point, then, is a continuously calculated number based on parameters (Base Inventory and Safety Level (Days)), variables (FDD), and time to fill the order, which can be variable or fixed. A block in the model (SC Order block) allows the user to specify a time the order policy will use as its time to fill orders. The user may choose to enter that information in a field in an inventory table, instead of in the block. The formula for order point is:

*Order Point = base inventory + FDD * (safety (days) + time to fill)*

2.3.4 Net Current Inventory

When the *net current inventory* of a SKU at any non-production location in the supply chain falls below the calculated order point, that location orders material from the inventory of its suppliers (sources).

Net current inventory is a number calculated throughout the simulation and refers to in-stock inventory plus any inventory currently on order minus the amount needed to fill existing orders minus a *reserve* amount (if any). Reserve is discussed more fully in section 2.5.2.

Net Current Inventory = material on hand + amount on order – orders to fill – $reserve_{TODAY}$.

When the net current inventory of the supplier of this location falls below the supplier's order point, orders continue to be placed further upstream. The process continues until enough material has been received at the location to fill all current orders.

2.3.5 Target (Days)

How much to order is based on FDD and a parameter called "Target (days)." Target (days) specifies a desired upper level of inventory and like safety (days) is multiplied by the FDD to determine a number of units. The number of target units is the desired inventory level reached when material on order is delivered.

The amount to order, then, is the amount required to bring the net current inventory up to the target (days). The formula is:

*Order amount = (target (days) * FDD + base inventory) – net current inventory*

Figure 1 illustrates a simple example in which the average daily demand for SKU 1 at distribution center A is 25 units. The inventory is initially stocked with 200 units and these units are consumed each day until the inventory reaches the order point of 50 units. The target for this inventory is 4 days. An order amount of 50 is computed to bring the inventory up to the target level of 100 units (4 days times the average daily demand of 25). The order cycle (days between orders) is typically the target days minus the safety days (in this case safety days are set to 2). If demand increases, the order point, target, and safety increase.

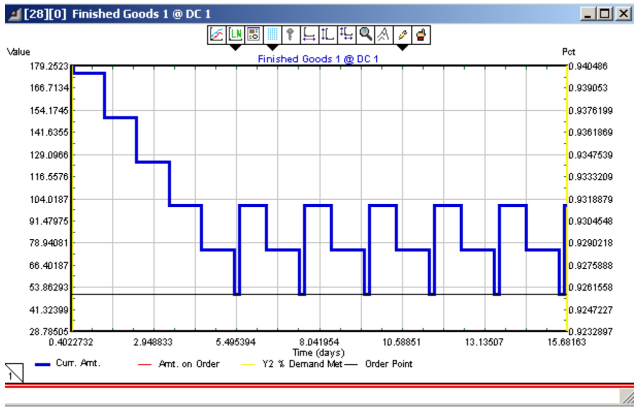

Figure 1: Example Illustrating Order Cycle

In simplest terms, the order point should be above the safety level, so that inventory never falls too low, and the order placed should be an amount to bring the inventory back up to target.

2.4 Inventory Classifications

The company uses an inventory classification system based on cost factors. The cost factor is comprised of total manufacturing cost and volume sold. The resulting classification breaks the SKU's down into A, B, and C classes. We use this classification to determine where to place the cycle stock in the supply chain.

After running the model, we noticed there were several A and B class SKU's that were promotional SKU's. Their manufacturing cost and volume put them in the A and B classification, even though one would normally think of them as C class SKU's. Inventory problems developed, i.e., getting the inventory to the DC too late for promotions due to the immediate spike in demand; and overproducing because the model thought the SKU had more of a continuous demand at the level of the spike. For the purpose of correctly modeling the ordering policies of promotional SKUs, another classification system was developed for ordering policy.

Table 1: Rules for Ordering Policy Classifications

Policy	Max days w/o demand	Min days with demand	Peak demand
Continuous	<= 14	>= 30	< 20 times average
Intermittent	> 14	>= 30	< 20 times average
Promotion	> 14	< 30	> 20 times average

Without demand is defined as any day with demand at 20% or less than the average. In order to qualify as *continuous*, a SKU cannot have a period of no demand longer than 14 days. To qualify as *promotional*, a SKU must have at least one period of demand smaller than 30 days, or have a peak that is 20 times greater than average demand.

Inventories without substantial periods of no demand can generally be managed with relatively low safety levels. Inventories with substantial periods of no demand require greater forecast anticipation in order to build inventories when demand restarts. Promotional demand requires ordering policies that anticipate the period of demand, order the required amount, and then do not order additional product during the period of demand.

2.5 Constructs and Logic

To model a realistic consumer demand pattern and its effect on a supply chain, several constructs are employed by Supply Chain Builder. The model must react to production policies that result in excess product and adapt to periods in which demand is greater than capacity just as the network it is simulating does. The model must "know" and apply alternative strategies for supplying locations when their inventory goes low. This section describes some of the constructs pertinent to inventory allocation.

2.5.1 Cycle Inventories and Push

It is generally not profitable for a plant to manufacture a small amount of a product when the production systems are constrained. The plant must produce enough product during a production run for equipment changeover and use to be practical. Plants, therefore, usually have minimum order or run sizes. Medium to low volume brands may have a production cycle as long as a month or more. In many cases, the minimum run of a product will be more than enough to fill orders and restock downstream inventories to maintain safety stock. This results in excess production, also termed *cycle inventories or cycle stock*. Where does the plant put the excess product?

The solution for excess brand production is for the plant to push the material forward into packaged finished goods (SKUs). The model simulates this push to finished goods in proportion to the FDD for each SKU while considering net inventory. Once cycle stock is packaged, it can be pushed to storage at downstream facilities. But which downstream facilities will receive the cycle stock?

The model determines the distribution of product on the basis of FDD and the current net inventory at the downstream locations. In other words, if two distribution centers have the same demand and net inventories, the same amount of product would be pushed to them. This will usually bring their inventory to a much higher level than their target for pull ordering. Cycle stock, therefore, is defined as inventory in excess of the maximum or target days of supply that would be reached by normal pull ordering.

2.5.2 Reserves (Anticipation Inventories)

Supply Chain Builder has the capability to anticipate future periods of demand over capacity and build a stockpile of SKU production. The model can determine the need for a reserve by looking forward in the demand projection to determine periods when global demand is greater than global capacity. The model then computes when and how much to stockpile.

A reserve amount protects inventory by "hiding" it from the reordering process. When an order point is computed for an inventory, any reserve amount is subtracted from the actual inventory. It thus appears, for the purposes of ordering, that there is less material on the floor than there really is. As a result, orders are placed earlier than would otherwise have been the case. However, the entire inventory, including the part protected by the reserve, is available for meeting downstream orders. Reserves can also be protected when pushing cycle stock downstream. The need to push stock downstream into the network results from excess production. Reserves can be pushed downstream or kept upstream.

The model computes reserves on a periodic basis, usually weekly. Each time reserves are computed, capacity and demand are compared for future periods extending out to a planning horizon set by the model user. If there is a future period where demand is over capacity, a determination is made as to whether there is sufficient capacity surplus between now and then to pre-build the computed shortfall. If there is sufficient capacity, then no reserves are needed now. If there is not sufficient future excess capacity, then reserves are required now. As one gets closer to the start of the period of demand over capacity, the amount of reserves required will increase. It will reach a maximum at the start of the period of demand over capacity, and be worked off to meet the excess demand. The effect is that all customer orders are met, even during peak demand periods.

The model uses parameters to vary the sensitivity of the reserves calculation, for example, there is a parameter that can make the model produce reserves either sooner or later. The model also allows a varying range to look ahead. If the system is a capacity constrained system, a longer look ahead is needed, as it takes longer to build up to the demand over capacity periods.

2.5.3 Reallocation (Production Transfer)

Although reserves can meet the challenge presented by periods of demand over capacity, they have the disadvantage of requiring substantial storage space and the costs associated with carrying large inventories. There is also the compound effect that larger reserves must be built over a longer period of time. An alternative solution to periods of demand over plant capacity is to shift production to other plants that have capacity over demand.

In simple terms, if a plant can't produce enough (due to a capacity constraint) to fill an order, the order can be *reallocated* to another plant if that plant has enough capacity.

Determination of transfers is done in a three step process:

1. Determine the total demand over capacity.
2. Determine which items can be transferred to any plant with excess capacity.
3. Go through the items one at a time and transfer up to its share of the total demand over capacity. While this is done, the capacity of alternate production sites must be tracked, since it is being used up as each item is transferred.

2.5.4 Redeployment

In addition to production transfer, orders may be filled from existing inventories at non-production facilities, such as other distribution centers in the supply chain. This process is termed *redeployment*. Redeployment in the model is defined as resupply of inventory from other than the inventory's primary supplier, except for resupply from alternate plants (as with production transfer).

The function of redeployment is to balance inventories in the network that are out of proportion due to errors in forecasting demand. As an inventory runs low, either an order to make more product is placed or stock already in the system must be redeployed to the low inventory. When stock is redeployed, it must leave the supplier with enough material to still have at least a given number of days of supply on hand. This rule is discussed more fully in section 2.5.5.

Stock redeployment, while it does solve the problem of balancing inventories, is the movement of stock laterally within the network. A lateral move implies an extra move and therefore extra expense.

2.5.5 Sourcing and Rules for Resupply

In order for production transfers and redeployment to work in the model, each inventory item uses a sourcing pattern (list of all possible locations that can supply an item to a customer location) established in the database. The model uses the pattern to apply a set of sourcing rules that determine when and where production transfer and/or redeployment will occur.

The location listed first in the sourcing pattern table is referred to as the primary supplier. All other locations are referred to as alternate suppliers. There can be an unlimited number of alternate suppliers.

An order for resupply will always be filled by the primary supplier unless that source is not able to fill the order. In that case, a choice of alternate supplier is basically performed in two steps:

1. The inventories of alternate plant suppliers are checked to see if there is one that can fill the order completely. If none are found, consider redeployment.
2. Before redeployment can occur, a *days on hand test* is applied. Stock at the ordering inventory must be at or below a user specified number of days of supply. If the inventory of the item at the customer location is at or below the days on hand test, the inventories of the alternate non-plant suppliers are checked to see if there is one that has an amount on hand equal to or greater than its days on hand test after filling the order. If a non-plant supplier is not found, the customer must wait until the supplier determined in the first step can fill the order.

The purpose of these rules is to prevent a problem that can arise from redeployment, that is, non-plant locations could alternate back and forth, sending the same material to each other. For example, location A meets the days on hand test and orders from B when B can fill the order, but if location B also meets the days on hand test, it could simultaneously order from A, since A still has stock on hand. This problem is solved by only allowing B to send stock to A if it has more than the triggering days on hand.

3 EXPERIMENTATION

Experiments were performed using Supply Chain Builder's Experiment Manager block. By using Extend's cloning feature (Krahl, 2002), the Experiment Manager block allows the user to collect into one block the key inputs for a model run from block dialog parameters and database tables. With all the parameters of interest together in one place, the model user can easily experiment with different settings and values. A series of scenarios can be run and tested with replication. The key input parameters used for setting up the experiments for this model were:

- Safety (Days) – Maximum days of inventory on hand before receiving an order (listed by location type and ordering policy).

- Target (Days) – Maximum target days of inventory to have on hand once the order arrives (listed by location type and ordering policy).
- Plant Storage – Amount of storage capacity (in pallets) available at plant. If inventory is at or over this amount of pallets, push orders are generated.
- Forecast Error – Random factor applied to the FDD so model forecast error will approximate actual forecast error (listed by inventory class).
- Reserves Location – Location where the reserves should be kept.

One experimental design was to test the effect of plant storage capacity and amount of forecast error on the overall inventory in the system and on the amount of stock redeployment. The storage options we studied were: no plant storage, the current level of plant storage, and unlimited plant storage.

In the current network some plant warehouse storage exists. Some plant storage space is allocated to "A" class SKUs that have high demand and where direct shipping is possible. Not all plants have storage capacity, so the effect of storing all product inventories at downstream locations was also of interest. The "unlimited storage" option allows all SKU classes to be stored at the plant without capacity limit.

For each storage option, we experimented with the amount of forecast error to see if forecasting was as important to inventory levels as storage strategies were. The forecast analysis of actual company demand revealed an additional bias with regard to very low demand products (C products). Our first scenarios were run with this high forecast for all storage options.

The random factor applied to the forecast in the model was decreased to reflect forecasting without the bias observed in the company's forecasts and another set of experiment were run, then they were run with a forecast error further decreased.

MAPE values resulting from the experiments are presented in the following table.

Table 2: Forecast Error for Inventory Classifications

Forecast Method	MAPE (A)	MAPE (B)	MAPE (C)	Bias (C)
Fcst with Bias	31.7%	37.2%	53.6%	+25%
Fcst w/o Bias	32.2%	37.1%	52.0%	0%
Reduced Fcst w/o Bias	23.9%	27.3%	38.7%	0%

4 RESULTS AND ANALYSIS

For each experiment, we studied the effect of safety (days) and base inventory for each inventory class on customer service. Customer service is measured by % demand met, with 98% being our benchmark level. We attempted to find the scenario that would provide the least amount of inventory in the system while still maintaining the goal service level.

When the raw data from the nine experiment scenarios were analyzed, we found that using even the best possible forecast did not impact redeployment of stock as dramatically as location of stock reserves and cycle inventories.

Using the current level of storage and forecast (Current System with Fcst Error with Bias) as our baseline, the table below presents the results of the nine experiments as a percentage of the baseline.

Table 3: Comparison of Plant Storage Solutions with Forecast Error Variation

	Fcst Error with Bias	Fcst Error with no Bias	Reduced Fcst Error (no Bias)
Current System			
Total Inventory	100%	96%	81%
Redeployment	100%	103%	44%
No Plant Storage			
Total Inventory	107%	102%	85%
Redeployment	126%	129%	60%
Unlimited Storage at Plant			
Total Inventory	95%	92%	76%
Redeployment	49%	41%	20%

5 CONCLUSION

The experiments performed to date on this model plainly indicate that location of excess inventory production and reserves has a significant effect on the amount of stock that must be redeployed and therefore on costs associated with holding and moving inventory. The model showed us that SKU storage at upstream plant sources reduced total inventory and, especially, reduced the amount of stock redistributed in the system. This effect was demonstrated when we modeled unlimited upstream plant storage vs. either the currently existing level of plant storage or storage at downstream locations (no plant storage) such as distribution centers.

Significantly, we found the same effect even when we experimented with reducing the amount of forecast error and eliminating bias to over-forecast.

Additionally, we find that locating inventories upstream at production facilities reduces the impact of forecast error.

Our results allow us to conclude that the next step for the company to consider should be to perform a cost benefit analysis on increasing warehouse size and retaining more stock upstream.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the contributions of Michael Denzien to this document and especially to the portion of the document devoted to the SDI database tools.

REFERENCES

Krahl, D. 2002. The Extend simulation environment. In *Proceedings of the 2002 Winter Simulation Conference*, ed. E. Yucesan, C.-H. Chen, J. L. Snowdon, and J. M. Charnes, 217-225. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers. Available online via <http://www.informs-cs.org/wsc02papers/026.PDF> [accessed July 3, 2003].

Parsons, D. J. and A. J. Siprelle. 2000. A supply chain case study of a food manufacturing merger. In *Proceedings of the 2000 Winter Simulation Conference*, ed. J. A. Joines, R. R. Barton, K. Kang, and P. A. Fishwick, 208-214. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers. Available online via <http://www.informs-cs.org/wsc00papers/144.PDF> [accessed July 3, 2003].

Phelps, R. A., D. J. Parsons and A. J. Siprelle. 2000. The SDI Industry product suite: Simulation from the production line to the supply chain. In *Proceedings of the 2000 Winter Simulation Conference*, ed. J. A. Joines, R. R. Barton, K. Kang, and P. A. Fishwick, 208-214. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers. Available online via <http://www.informs-cs.org/wsc00papers/031.PDF> [accessed July 3, 2003].

Phelps, R. A., D. J. Parsons and A. J. Siprelle. 2001. SDI Supply Chain Builder: Simulation from atoms to the Enterprise. In *Proceedings of the 2001 Winter Simulation Conference*, ed. B. A. Peters, J. S. Smith, D. J. Medeiros, and M. W. Rohrer, 246-249. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers. Available online via <http://www.informs-cs.org/wsc01papers/029.PDF> [accessed July 3, 2003].

AUTHOR BIOGRAPHIES

ANDREW J. SIPRELLE is President and founder of Simulation Dynamics, a firm that provides simulation consulting, training, and custom models. Mr. Siprelle's industry experience includes creation and analysis of models for strategic and capacity planning, market analysis, and the application of industrial statistics. Applications of simulation have been in the areas of business, manufacturing, government, and service operations. Before starting Simulation Dynamics, he worked at the Aluminum Company of America. He received his B. S. in Industrial Engineering and Operations Research from Virginia Polytechnic Institute. His e-mail and web addresses are: <Siprelle@SimulationDynamics.com> and <www.SimulationDynamics.com>.

DAVID J. PARSONS is a principal of Simulation Dynamics. He is responsible for development of SDI's Supply Chain Builder, a supply chain simulation and analysis tool. His experience with simulation began in 1965 with experiments in the use of natural selection algorithms to evolve architectural designs. During the 1980's, he designed, built, and operated several dairy-processing plants using simulation of key systems as an integral tool for design, value engineering, and trouble-shooting. Mr. Parsons received a B.A. from Harvard College and a Master of Architecture degree from the Harvard School of Design. His e-mail and web addresses are: <Parsons@SimulationDynamics.com> and <www.SimulationDynamics.com>.

ROBIN J. CLARK is a Consulting Engineer with Simulation Dynamics, Inc. He creates custom models and provides customer support and training. He has a B.S. degree in Physics from Tennessee Technical University and M.S. degree in Management Science from the University of Tennessee. His background experience before coming to SDI was in operations research and optimization models. His e-mail and web addresses are: <Clark@SimulationDynamics.com> and <www.Simula tionDynamics.com>.