

Optimization with the SAS® System: What It Is, What's New, and Why You Should Be Using It

Edward P. Hughes and Trevor D. Kearney
SAS Institute Inc., Cary, NC

ABSTRACT

This paper presents an introduction to optimization and an overview of SAS capabilities in optimization, with a focus on mathematical optimization with SAS/OR® software. We review the current suite of optimization modules and techniques available in SAS/OR software, highlighting recent additions. We discuss the value that optimization adds to business intelligence, and we describe several recent cases in which optimization has played a central role in the creation of SAS decision support solutions.

INTRODUCTION

SAS software includes a diverse and well-established array of optimization capabilities designed to analyze and solve many different types of optimization problems. The greatest concentration of modeling, analysis, and problem-solving techniques occurs in SAS/OR software, but some complementary techniques also are present in SAS/STAT®, SAS/IML®, and Enterprise Miner®.

Users at a wide range of major U.S. and international corporations have found these capabilities to be indispensable in building decision support and decision-guiding solutions aimed at identifying the best and most profitable decision strategies and action plans for their enterprises. Meanwhile, more and more enterprise solutions developed and marketed by SAS are adding value for their users by leveraging the decision-guiding power and insight added by SAS optimization.

Still, among SAS users at large SAS optimization is often not very well known or well understood. This paper is intended to shine a new light on this "hidden jewel" of the SAS System and encourage its broader use.

In this paper we describe the nature and purpose of optimization, and consider the various types of optimization problems that can be built and solved with SAS/OR software. We explore the value that optimization adds to business intelligence. Finally, we examine several recent uses of SAS optimization in enterprise decision solutions.

WHAT IS OPTIMIZATION?

Simply put, optimization is the process of choosing the actions that result in the best outcome. With this definition, it's clear that every living thing—and every enterprise—on the planet is, in some sense, involved in optimization. Whether the goal is survival of the species through propagation, maximization of profits through production decisions, or minimization of risk with purchasing strategies, we're all striving to optimize something—or, more likely, several things.

That's not to say that we all succeed at optimization. We experience varying degrees of success, influenced by a number of factors. We might not have the freedom to take the actions that really produce the best results. We might have incomplete or incorrect information on how various actions interact with each other, how they are limited by our circumstances, or how they influence the outcomes. There may be so many choices of possible actions that it's impossible to evaluate them all. We might not even have a reliable way to determine if the outcome we produce is remotely close to being the best outcome.

OPTIMIZATION: KEY ELEMENTS AND KEY QUESTIONS

A well defined approach to optimization, such as SAS offers, can help to reduce or eliminate many of these difficulties. It aids in describing the three key elements of any optimization problem:

- **Decision variables**, representing decisions or choices that are available to you. Examples might include production levels, route selections, and resource allocations.
- An **objective** that you're trying to achieve. Examples include maximizing profit, minimizing distance traveled, and minimizing unused raw materials.
- **Constraints**—requirements or rules that place limits on how you can pursue your objective by placing limits on the values of the decision variables. Some examples are factory processing capacities, customer demands, materials available, and budgets.

The purpose of an optimization is to achieve the objective within the limits established by the constraints. Moreover, solving an optimization problem can help you answer other critical questions about your current and future actions and decisions and the outcomes that they produce. These questions can include:

- Is what we are trying to accomplish possible?
- How are we doing currently?
- Can we do better? If so, how?
- What's the best that we can do?
- What happens when conditions change?

These are fundamental, compelling questions whose relevance remains intact across departmental, industrial, and geographical boundaries. Answering them gives you valuable information not only about what decisions to make now, but also about how to change your decisions as conditions change.

OPTIMIZATION: ANSWERING THE KEY QUESTIONS

The outlines of the process for tackling optimization problems are straightforward. The first step is to build a model of your situation, which means that you need to identify the important factors (such as the amount of each product to make, the associated costs, required machinery, etc.) and define how they interact with each other. In this model, all interactions and restrictions are defined by using mathematical equations and inequalities.

After building the model, the next step is to solve the model and identify a set of decisions that optimizes your objective. SAS/OR software provides a number of excellent, scalable techniques for solving a variety of optimization models. Each type of optimization model is defined by the types of mathematical expressions and constraints that are used in building it, and for each type of model the solution methods are distinct.

- **Linear programs (LPs)** use only linear expressions in constraints and the objective (for example, $4x + 3.5y - 2z \leq 7$), and the decision variables can take on any value in specified allowable ranges.
- **Mixed-integer programs (MIPs)** are linear programs in which some of the decision variables must only take on integer values.
- **Integer programs (IPs)** are linear programs in which all of the decision variables may only take on integer values.
- **Nonlinear programs** use nonlinear expressions (such as x^2 ,

cos (y), 1/x, and others) in the constraints and/or the objective.

Once you've found a solution, you need to implement it in the situation that you're modeling—and it may happen that you find that the solution can't really be implemented. Perhaps you're deciding how many tires to make and the model suggests a production level of 1200.35 for one model. In this case, you'd go back to the model and add a restriction that it must consider only integer values for number of tires to make. This kind of give-and-take between the model and what it's representing is natural. It's how models get improved. Good modeling is an art.

Finally, no model is good for long if it can't respond to change. Better models can give you advance information about how potential changes will affect the solution you've found. SAS/OR software is well equipped to do this.

What do you get in return for building, refining, and solving your optimization model? Plenty. Just by considering all of the available options—something that an organized model-building approach gives you—you can discover how to squeeze a tremendous amount of additional performance out of the system you're modeling. Moreover, in the process of building and solving models you must describe in detail how the system functions. Often, this forces you to think in detail about aspects that you've "glossed over" previously, leading you to a deeper understanding of the system. This can lead in turn to better ideas for improving the system—ideas that may lie outside modeling, in the realm of innovation.

OPTIMIZATION AND BUSINESS INTELLIGENCE

"There is always a well-known solution to every human problem—neat, plausible, and wrong."—H.L. Mencken

Business intelligence leverages diverse and possibly far-flung sources of enterprise data, creating a central repository of information. Business intelligence uncovers key relationships, surfaces performance metrics, predicts future performance, identifies targets for improvement, and much more. The process of amassing excellent business intelligence provides unquestionable value to the enterprise by providing the raw material for tactical and strategic decision-making.

Some, though, confuse supplying the necessary information for making decisions with the process of identifying and justifying optimal decisions. It's important to understand the distinction. Effective business intelligence gathering can create a huge variety of views into the structure, operations, and performance of an enterprise, but it cannot tell you how to use this information to decide what to do. It cannot help you sort through the myriad of decisions and decision strategies that are possible to identify the few—or the one—producing the best results.

Faced with such an array of choices, most people will fall back on "rules of thumb," time-worn decision rules, or "gut feel," choosing the decisions that appear safe, customary, or intuitively best. These approaches might indeed produce optimal results or, as H.L. Mencken claims in the quote above, they might lead to disaster. They might simply produce mediocre results. Without a structured and robust decision-support and decision-guidance system, though, you'd never even know just how much better your enterprise's performance could have been.

At the enterprise level, structured decision support, through optimization, forecasting, and other advanced analytical techniques, is the indispensable follow-on to business intelligence.

OPTIMIZATION WITH SAS

SAS/OR software can be used to solve a wide variety of optimization problems. As described previously, the basic optimization problem is that of minimizing or maximizing an objective function subject to constraints imposed on the variables of that function. The objective function and constraints can be linear or non-linear and the constraints can be bound constraints, equality or inequality constraints, or integer constraints.

Traditionally, optimization problems are divided into Linear Programming (LP; all functions are linear) and Nonlinear Programming (NLP). Some variations of LP problems are assignment problems, network optimization problems, and transportation problems. Nonlinear Regression (fitting a nonlinear model to a set of data and the subsequent statistical analysis of the results) is a special NLP problem.

SAS/OR software has separate procedures (or facilities within procedures) for solving each type of problem. Data describing models are supplied in a form suited for the particular type of problem, and the respective optimization algorithms are also tailored to take advantage of the special structures and characteristics of the problems that they solve, thus reducing solution times considerably.

SAS/OR software includes six optimization-oriented procedures.

- **PROC LP** solves LPs (using a general simplex algorithm) and mixed integer programming problems.
- **PROC NETFLOW** solves network optimization problems, which are LPs with a dominant network (node and arc) structure. Both a specialized simplex method and an interior-point algorithm are available.
- **PROC INTPOINT** is an additional procedure focused solely on the use of the interior-point algorithm in solving linear programs and network optimization problems.
- **PROC NLP** solves general nonlinear programming problems, and also solves quadratic programming problems and least-squares problems (special forms of NLPs).
- **PROC TRANS** solves transportation problems, in which items must be moved from a set of "supply" locations to a set of "demand" locations at lowest cost.
- **PROC ASSIGN** solves assignment problems, in which one set of items must be assigned to another (example: assigning subcontractors to specific jobs) at lowest total cost.

PROC LP

The LP procedure in SAS/OR software solves linear programming problems by using the simplex method, a well-established technique for solving LPs. Integer and mixed-integer programming problems are solved via the branch-and-bound technique (based on the simplex method), which systematically explores, evaluates, and selectively discards the possible integer-valued solutions. For all algorithms PROC LP includes a great number of options that you can use to customize and adjust the solution process. PROC LP can be run interactively, enabling you to alter the solution process mid-stream if necessary and to store intermediate results for later use and analysis.

PROC LP includes extensive post-optimality analysis (also known as sensitivity analysis) capabilities, enabling you to examine (without re-solving the problem) how the optimal solution to a problem changes in response to alterations in the constraints, the objective, and other factors.

PROC NETFLOW

Network models describe a wide variety of real-world applications ranging from production, inventory, and distribution problems to financial applications. These models typically represent a process that changes material as it flows through a network. The oil and gas industry pioneered the application of mathematical

programming techniques for modeling refinery planning. Many similar models can be applied to other process manufacturing, such as mills and grain, breweries and beer, or refineries and sugar.

A network model consists of:

- **nodes** that represent supply, demand, or transshipment points, and model components such as warehouses, tank farms, factories, refineries and customers. At a supply node, material enters the network; at a demand node, material leaves the network; and at a transshipment node, the total flow of material into the network equals the total flow of supply out of the node. This last restriction is known as conservation of flow.
- **arcs** carrying material between nodes. Each arc originates at one node and terminates at another. Arcs have a per unit traversal cost (that may be zero or negative if the cost is really a profit), a capacity (the largest amount of flow the arc can convey), and a minimum flow bound (useful to force some minimum quantity through an arc).

Since networks can be depicted pictorially, they clarify the model structure and are useful for explaining what the model does and how solutions should be interpreted. Excellent models can be built using nodes and arcs; however, it is often necessary to specify additional constraints on the flow of material through a network. These additional constraints, called **side constraints**, can add tremendously to the realism and usability of a network model. Side constraints are needed, for example, to model situations in which the amount of finished products resulting from a manufacturing process is a specific proportion of the amount of raw material fed into the process.

PROC NETFLOW solves a network model by determining the minimum cost flow through the network from the supply nodes to the demand nodes so that

- the conservation of flow is maintained at nodes
- the side constraints are satisfied
- the flow through each arc is on or between the arc's lower flow bound and capacity

With PROC NETFLOW you can solve network optimization problems with either the interior point algorithm or the network simplex algorithm. The interior point algorithm is often faster when problems have many side constraints, and this advantage grows as the size of the network and the number of side constraints increase.

PROC NLP

PROC NLP offers a set of optimization techniques for minimizing or maximizing a continuous nonlinear function of n decision variables with lower and upper bound, linear and nonlinear, equality and inequality constraints. Problems of this type are found in many settings in which linear functions and linear equations cannot faithfully describe the situation being modeled.

The NLP procedure provides a number of solution methods for this type of problem that take advantage of special structure in the objective function, the constraints, or both. Available solution techniques include:

- Quadratic active set technique
- Trust-region method
- Newton-Raphson method with line search
- Newton-Raphson method with ridging
- Quasi-Newton methods
- Double-dogleg method
- Conjugate gradient methods
- Nelder-Mead simplex method

- Levenberg-Marquardt method
- Hybrid quasi-Newton methods

SAS OPTIMIZATION: WHAT'S NEW

In Version 8 of SAS, SAS/OR software added a number of enhancements to its optimization capabilities. Specifically:

- PROC LP added optional pre-processing capabilities to reduce the size of integer and mixed-integer programming problems, thus reducing that amount of time needed to solve them.
- PROC NETFLOW added an interior-point method to its existing network simplex solution method.
- PROC INTPOINT was added, using an interior-point method to solve LPs and network optimization problems.
- All SAS/OR procedures became compatible with the SAS Output Delivery System (ODS), enabling the results of optimizations to be published in a variety of formats including HTML.

SAS OPTIMIZATION IN ACTION

In this section we discuss briefly three very different uses of SAS optimization to solve business problems. Two are new SAS solutions in development, and the third was developed in concert with a SAS/OR customer and is currently in use. The common theme running through all three examples is the ability of the SAS model-building and optimization techniques to adapt to fit the problem at hand.

SUPPLIER RELATIONSHIP MANAGEMENT

SAS optimization has been utilized to augment the SAS Solution for Supplier Relationship Management (SRM) in two important areas: supplier ranking and supplier portfolio optimization. Both extend the SAS SRM solution beyond reporting on the past and current purchasing patterns, to shaping and guiding future purchasing decisions.

SRM RANKING

SRM Ranking, a current component of the SAS SRM Solution, is designed to rate and rank suppliers when considering multiple performance criteria simultaneously. It provides an automatic method for scoring and comparing suppliers, eliminates the need to determine weights for calculating the scores, and can account for business rules in the process.

BUSINESS PROBLEM

In many businesses, data on supplier performance is collected and used for evaluating suppliers. Unfortunately, grading and comparing suppliers based on this data often is not straightforward, due to the presence of numerous and possibly conflicting evaluation criteria. For example, if one supplier outperforms all others according to one performance criterion but fails to achieve satisfactory levels on other criteria, it becomes unclear how to proceed with the comparison.

The traditional solution is to assign a fixed weight to each criterion to form an aggregated, weighted score for each supplier. Usually, weights are chosen to support specific business rules (such as weighting quality measures more heavily than financial measures to reflect their greater importance). Several problems can arise from this approach, including the following:

- Weights are subjective, difficult to agree on, and have a tremendous effect on the final scoring.
- It is difficult to balance relatively strong and relatively weak performances for different criteria.
- Differences in units of measurement for the criteria can distort the influence of the weights used in the scoring.

The business problem can be stated most simply as **how to best rate suppliers on the basis of multiple, possibly conflicting, performance measures and account for business rules.**

SOLUTION

SRM Ranking implements an innovative solution to this business problem. It solves the problem of differing measurement units among criteria by normalizing the performance data. This eliminates units from the measures, removes distortions associated with differences in units, and provides for more balanced comparisons.

SRM Ranking solves the difficult problem of determining weights by evaluating each supplier in isolation and optimizing the supplier's performance relative to all other suppliers (based on Data Envelopment Analysis). Automatically calculating optimal weights for each supplier's performance criteria avoids the problems resulting from assigning fixed weights to all suppliers.

SRM Ranking uses the optimally calculated weights to compile the relative scores used to compare and rank the suppliers. SRP captures business rules by enabling you to place limits and other restrictions on the weights used for the various performance criteria. This establishes rules on the relative importance of the criteria, and can also be used to account for imprecise performance information.

The SRM Ranking methodology is driven by the performance data. After normalizing each performance criterion, SRM Ranking uses linear programming to calculate weights on each supplier's performance criteria that optimize the supplier's overall performance rating. This rating is the weighted sum of the supplier's individual performance scores, with the weights being determined individually (and optimally) for each supplier. When this process is complete a ranking of all suppliers is possible.

POSSIBLE USES

SRM Ranking has a broad range of possible uses, and the underlying techniques used have almost unlimited applicability. Within the confines of supplier management, the most immediate use of SRM Ranking is in the periodic review of supplier performance. In this role, SRM Ranking can easily provide its ratings, rankings, and tier assignments as input to the SRM Portfolio Optimizer, discussed in the following sections of this paper. Another opportunity to use SRM Ranking occurs in the Request For Proposal (RFP) or Request For Information (RFI) process, during which the purchasing company may need to shorten the list of bidding suppliers based on their past performance. In this scenario SRM Ranking's performance ratings could easily drive the reduction in the pool of bidding suppliers.

SRM PORTFOLIO OPTIMIZER

The SRM Portfolio Optimizer assists in structuring and restructuring portfolios of suppliers, with a goal of maximizing the buyer's benefit while meeting specific requirements on portfolio makeup. SPO relies on rationalized supplier data describing suppliers and the goods that they provide, available from a number of sources such as Dun & Bradstreet. SRM Portfolio Optimizer provides guidance for answering strategic questions such as:

- Who should we buy from?
- Should we spend more or less with a supplier?
- What should our expected risk be?
- What should we buy from a given supplier?
- How should we alter our buying practices?

BUSINESS PROBLEM

Often, businesses don't have enough information on their suppliers of parts, raw materials, and other critical items. They cannot readily and easily determine how much they are spending,

what they are purchasing, who their top suppliers are, or the answers to many other important questions. Surfacing such information effectively is one key to improving the buyer's position and to negotiating better relationships with suppliers.

Another key to managing supplier relationships is an organized approach to moving from your current supplier portfolio to your desired supplier portfolio. In improving a supplier portfolio, one of the most often-mentioned goals is better negotiating leverage. Typically, though, buyers face restrictions as they pursue this goal. These restrictions or requirements may originate from internal business rules on supplier selection, from regulations specific to the business's industry, or from other sources. Some examples of these restrictions include:

- "At least 5% of our purchases should be made with small businesses."
- "We should have 5 to 10 suppliers of paper goods."
- "To ensure quality, buy at least 25% of all supplies from ISO compliant businesses."
- "The average financial stress score (FSS) of our supplier portfolio should not exceed 1.7 and the average supplier evaluation risk (SER) should not go beyond 2."

What's needed is a method for reshaping the supplier portfolio, focusing on the goal of maximizing negotiating power while adhering to whatever restrictions may apply.

SOLUTION

SRM Portfolio Optimizer's solution technique is based on mixed-integer programming and is aimed at making procurement choices that maximize negotiating leverage within the restrictions created by the business rules governing the makeup of the supplier portfolio. As is true with SRM Ranking, the user interface provided by SRM Portfolio Optimizer does not require users to understand the fine points of mathematical optimization, but only to grasp the concepts motivating the demographic, budgetary, geographic, and other business rules on the supplier portfolio.

The user selects from a menu of possible business rules to apply, and customizes each rule used by specifying parameter values such as minimum percentages, budgetary upper and lower bounds, and regional distribution targets. SRM Portfolio Optimizer then finds the combination of purchases and vendors that meets or exceeds the specified requirements and maximizes purchasing leverage. SRM Portfolio Optimizer achieves this by consolidating purchases to buy mainly from suppliers for whom the company is a major client, while simultaneously maintaining a balanced and diversified portfolio as required by the business rules.

POSSIBLE USES

The usage possibilities for SRM Portfolio Optimizer are quite similar to those for SRM Ranking. First, SRM Portfolio Optimizer is well suited for use in the periodic review of a supplier portfolio. This may be especially important if, for example, the purchasing firm completes a merger or acquisition and must deal with an augmented supplier portfolio. Additionally, it can (like SRM Ranking) assist in the Request For Proposal (RFP) or Request For Information (RFI) process. In this role, SRM Portfolio Optimizer can be refocused to minimize overall spend while satisfying the requirements outlined in the RFP or RFI, identifying a lowest-cost set of bidders.

OFFER OPTIMIZATION IN BANKING AND FINANCIAL SERVICES

Offer Optimization, a solution under development in the SAS Supply Chain Center, addresses a common problem in database marketing and customer relationship management in banking and financial services: how to market the right product to the right

customer via the most appropriate channel. This solution, as described by Storey and Cohen (2002), is an improvement over earlier methods that chose customers on a product-by-product basis and did not take a global view of the marketing effort.

BUSINESS PROBLEM

The banking industry regularly stages promotional campaigns to improve customer value by offering new products to existing customers. In such campaigns the overriding goal is to offer the right product to the right customer at the right time, and via the right channel. The measure of “right” is the return on the investment of time and resources in the campaign; that is, the product, channel, and time are “right” if the customer who accepts the offer contributes to a maximum return on the marketing campaign.

In practice this problem can be enormously complex and difficult to solve. The financial institutions involved often wish to market several products simultaneously, and in most cases they operate under a complex set of business rules and constraints. Balancing maximization of the return on the marketing investment against meeting all of the business constraints creates a daunting problem. Since multiple products may be involved in the same campaign, the decision is not only “which customers to offer a product” but also “which product(s) to offer each customer.” Finally, the scope of the problem is typically huge, since it’s not unusual to see a campaign aimed initially at a base of several million customers.

ASSUMPTIONS AND BUSINESS CONSTRAINTS

The most critical assumption for Offer Optimization concerns the ability to estimate the response of any individual in the customer base to any product offer included in the campaign. This is expressed in the form of the expected profit from this pairing, and relies in turn on the assumption that prior campaigns have been extensively analyzed, resulting in response models that estimate the likelihood of acceptance for any combination of product and customer.

Several business constraints are also assumed to exist. A sample of these constraints includes:

- Minimum/maximum number of product offers to be made in the campaign.
- Minimum expected profit from offers of each product.
- Channel capacity and funding limits.
- ROI minimum level for the overall campaign.

SOLUTION

The most natural formulation for this problem is as an integer program, with one binary decision variable for each possible combination of customer, product, and channel. If the variable has the value 1, the corresponding product is to be offered to the corresponding customer via the corresponding channel; if it has the value 0 then the offer is not to be made.

Unfortunately, when dealing with millions of customers such a problem is simply too large to solve effectively. Instead, we shift to a multi-stage approach in order to create and solve a more workable problem.

1. Use statistical clustering techniques to cluster the customer base according to the cost of making offers and the expected profit from offers.
2. Using the centroid of each cluster to approximate the behavior of the cluster, construct a linear programming problem that determines what proportion of each cluster to offer each combination of product and channel.
3. Use a sorting technique (for single products) or solve a smaller assignment problem (for multiple products) to determine precisely which customers in each cluster should be offered the channel/product combinations selected for the cluster. Allocations in this step respect the proportions determined in the previous stage.

This approach is being undertaken in a joint development

environment at one of the leading North American and international financial institutions. Implementation of this solution resulted in an increase in ROI from 165% to 258% and an increase in expected profits from \$2.65 million to \$3.58 million.

PRODUCTION, DISTRIBUTION, AND INVENTORY PLANNING IN THE SUGAR INDUSTRY

In 1998 a corporation that refines and distributes sugar partnered with SAS to explore methods for modeling and managing its supply chain. This corporation sells approximately 4.6 billion pounds of sugar to its customers, accounting for 25% of the U.S. sugar market at a value of \$1.5 billion.

During this joint effort the current supply-chain practices at the corporation were analyzed and a network model was designed to identify the minimum cost solution for production, inventory and distribution planning for over 80 plants and warehouses, 250 products and 2,000 customers.

GOALS

With a long history of on-time customer delivery, the sugar corporation views maintaining customer satisfaction as a top priority. With this in mind, they realize that efficient supply chain management requires a reliable packaging plan, efficient inventory management, and a timely and cost conscience distribution schedule. These costs include production costs, inventory-holding costs, handling costs, packaging costs, and freight costs.

SUGAR PRODUCTS

Sugar products are classified into four groups: bulk, liquid, consumer packages, and industrial packages. Production packaging is performed at eight different locations using multiple lines and shifts. The packaged products range in size from 1 to 10 pounds for consumer products and from 25 to 100 pounds for industrial products. Products are packaged under many different labels, over 250 different sugar products.

While screening bulk sugar to a particular granularity produces most sugar products, others are produced by mixing sugar with other components (for example, adding cornstarch to produce powdered sugar). Packaging is completely driven by customer demand (and safety stock levels) except when bulk sugar must be packaged for storage in public facilities. In this situation, maximum inventory limits are required to avoid unnecessary production of lesser-demanded products (for example, private-label consumer packages).

DISTRIBUTION

Freight cost is the primary cost associated with delivering a product to a customer. With more than 70 distribution locations and over 2,000 customers to serve, efficient inventory management is essential. The major cost drivers for warehousing products include the freight cost to the warehouse, the product handling charges, and the monthly charge for storing the product. Also, if a product moves to a transfer point (a location where rail shipments are transferred to trucks), there may be additional handling costs (for converting bulk sugar to liquid sugar, for example). Sugar is transported via truck and rail, although some storage and distribution facilities cannot handle both modes of transportation.

MODEL FORMULATION

Based on the business issues and the structure of the organization, a minimum cost network with side constraints was formulated to model packaging, distribution and inventory. The solution determines the flow of bulk sugar from production and packaging facilities, conversion to packaged products, movement to other refineries, storage facilities and distribution centers, and delivery to customers. The physical layout of this network is shown in Figure 1.



Figure 1. The Layout of the Network

The packaging, inventory, and distribution processes were modeled as a minimum cost network flow problem with linear side constraints. Nodes represent facilities (plants, warehouses and transfer points) as well as customer orders, while arcs represent directions of potential product flow (for example, packaging, storage and delivery). The nodes were replicated for each product (SKU) and each period (month) of the planning horizon.

PLANT CAPACITY SIDE CONSTRAINTS

Linear side constraints are needed to represent a variety of plant capacities as well as to insure desired product inventory levels. Since a plant's capacities are based on product type (bulk, liquid, consumer package, or industrial pack-age), sums of product flows (arcs) across several SKUs are necessary to model these constraints. These types of constraints include facility storage capacity and facility inbound, outbound and throughput handling capacity. Safety stock constraints and maximum stocking levels are designated for each SKU and are determined by summing the inventory arcs for each SKU across all plants. Rather than having two separate side constraints, these min and max constraints are modeled more efficiently by creating a nonarc variable with lower and upper bounds that indicate the minimum and maximum inventory levels. These nonarc variables are simply linear programming variables that are expressed as a sum of the inventory arc flows.

PACKAGING PRODUCTS

At packaging plants, some sugar products may be converted (packaged) to create other sugar products. Parts of the model are concerned with sugar passing through packaging line machinery. One or more packaging lines may exist at each of the packaging plants. Some packaging lines will accept more than one input sugar product for packaging. Also, some packaged products require the addition of some non-sugar ingredient (for instance, corn-starch for powdered sugar). Side constraints were used to specify these recipes, and ensure that output capacities for the various output products are not exceeded. Shortfall or excess nodes and arcs are used to supply non-sugar ingredients into the system, or to drain away losses or volume reductions due to density changes.

ADDITIONAL PROBLEM CONSTRAINTS

Additional side constraints were added to the model formulation to control interplant product shipments. These constraints limit the amount of product shipped into a plant in a given period to be less than or equal to the amount of product to fill customer orders from that plant or to be consumed in packaging at that plant within the next n months (where n is an integer and is the same for all products). In other words, products should not be moved from one plant to another unless they will be used to source a customer order from that plant or will be used in the packaging of another product at that plant in the next n months.

Additionally, in order that appropriate "on-going" inventory levels are left at the end of the planning horizon, the last period of the model can be represented by one or more periods' worth of data. By combining supply and demand for future periods into the last period, the inventory levels at the end of the next to the last period

will reflect the impact of future demand. Thus, solving a problem with a 13-month planning horizon (with the last period equal to 5 or 6 months) will provide a better solution for the 12-month problem.

MODEL GENERATION

Business data were extracted from the corporations ERP system based on an SAP AG software system and its freight management system. Other business data were managed with SAS software. SAS provides extensive tools for this extraction that enable users to directly access any required data.

Planning over a 13-month time horizon with approximately 80 plants (packaging, storage, distribution and transportation transfer points), 250 different sugar products (SKUs) and over 2,000 customers, the resulting network model contained roughly 220,000 nodes, 1 million arcs, 3,000 nonarc variables and 26,000 linear side constraints. The problem was solved using the primal-dual predictor-corrector Interior point algorithm of the NETFLOW procedure on a PC running Windows NT in about 2.5 hours.

OPTIMAL SOLUTION

The optimal solution, contained in a single SAS data set created by PROC NETFLOW, specifies the minimum cost schedule indicating the use of packaging, distribution and inventory assets that satisfies customer demand. These results are used in the current planning cycle, which involves decisions regarding the budgeting of all freight and warehousing expenses.

Primary users include the production planning, logistics, and distribution managers. Presently, the results are being used to make strategic level decisions, but future plans call for the uploading of information to the ERP system for support of operational level decisions. The solution provides essential information for these areas of planning:

- production packaging
- labor and resource requirements
- warehousing expenses
- freight expenses
- inventory management
- order distribution

Analysis of the results has provided additional insight into several business processes. By not limiting the capacity of some storage facilities, the model identified the optimal storage levels. In some cases, this showed that a facility was being underused, while in others it suggested that additional public storage should be acquired. Given the flexibility of data input and problem generation, it is an easy matter to perform what-if analysis with various scenarios. Through multiple problem runs and resolution of numerous data in-feasibilities, several potential problem areas were identified. These included

- how to manage insufficient sugar supply
- how to handle customer sourcing requirements
- how to reconcile inventory inconsistencies
- how to improve inventory turns
- how to incorporate fixed production planning

In some cases, the model formulation was modified to incorporate additional information in the final solution, while in others, problem-specific reports were generated to provide summary results for informed decision making.

CONCLUSIONS

Optimization enables enterprises to discover the true value in their business intelligence by leveraging it in the creation and solution of models that shape and guide tactical and strategic decision-making. Exploration of these models highlights innovative decision strategies, adds consistency and support to the decision-making process, and enables quick and effective

adaptation to changing market conditions.

SAS/OR software incorporates a robust, versatile suite of optimization modeling and solution capabilities. With the scalability to solve enterprise problems and the flexibility to model real-world business processes in detail, SAS optimization can add insight, innovation, and credibility to the decision-making process in virtually any enterprise.

REFERENCES

Kearney, Trevor D. (2001), "Optimization: Practice, Theory, Future Trends," Internal Document, Cary, NC: SAS Institute Inc.

SAS Institute Inc. (2000), "Supplier Management: Optimally Managing the Supplier-Buyer Relationship," Internal Document, Cary, NC: SAS Institute Inc.

Storey, Andrew, and Cohen, Marc-david, "Exploiting Response Models – Cross-Sell and Up-Sell Opportunities in Banking," paper submitted to KDD 2002 conference.

AUTHORS' CONTACT INFORMATION

Your comments and questions are valued and encouraged.
Contact the authors at:

Ed Hughes
SAS Institute Inc.
Office R-4109, SAS Campus Drive
Cary, NC 27513
Phone: (919) 531-6916
Fax: (919) 677-4444
Email: Ed.Hughes@sas.com

Trevor Kearney
SAS Institute Inc.
Office R-4206, SAS Campus Drive
Cary, NC 27513
Phone: (919) 531-7726
Fax: (919) 677-4444
Email: Trevor.Kearney@sas.com

SCSUG 2003 PRESENTOR

Mary G. Crissey
SAS Institute Inc.
San Antonio, TX 78247
Phone: (210) 637-5855
Email: mary.crissey@sas.com