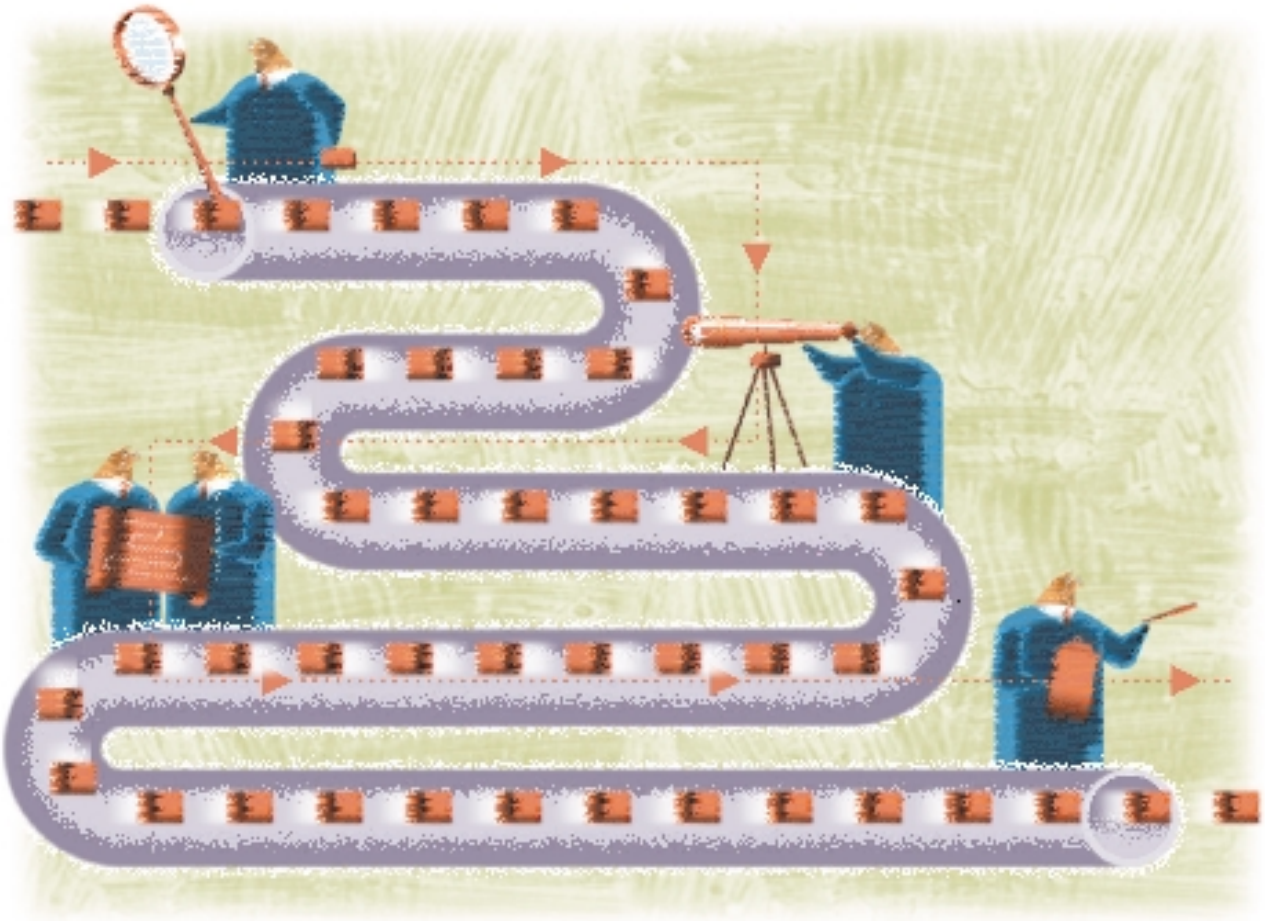

Daniel Guidera

CAPACITY PLANNING

The Antidote to Supply Chain Constraints

By André J. Martin

Companies involved in the retail supply chain face a recurring problem: how to handle constraints—those shortages of capital, people, equipment, and space that keep product from flowing smoothly through the supply chain. Optimization solutions can help, but they're expensive. A more cost-effective way to manage constraints is to anticipate them. One proven technique is time-phased capacity planning.

In the past year, the retail world has been assaulted by a new catch phrase: "Optimize your supply chain by eliminating constraints." A constraint occurs when a supply chain does not have enough capital, people, equipment, or space to acquire, transport, manufacture, and/or sell product. Many software companies are promoting optimization algorithms as a way to solve or eliminate retail supply chain constraints. But is the management of constraints the real issue? And if so, do we need optimization algorithms to solve retail supply chain constraints? Based on personal experience, I would say that the answer 90 percent of the time is no.

There is no doubt that constraints do exist and cause major problems across the retail supply chain—which not only includes the retailers themselves but also the suppliers, carriers, and other providers that serve them. In short, the supply chain implications of constraints in the retail sector are far-reaching. Retailers struggle every day with questions such as: *Do we have enough capacity to receive and put away suppliers' deliveries (receiving constraints)? Do we have enough space capacity to hold inventory (space constraints)?* and *Do we have enough resources to ship to our stores (shipping constraints)?* Every time the answer is no, the retailer has a capacity problem to resolve; or, if you prefer, it has a *constraint*.

But before we launch into a massive software search for expensive and narrowly focused optimization algorithms to solve these constraints, it might make sense to consider another alternative. Most constraint issues could be solved by extending capacity planning logic to the retail stores and retail distribution centers (RDCs). Capacity planning is a process that involves predicting future constraints based on sales forecasts, projected inventory, and replenishment requirements as calculated by a time-phased planning system. The concept of time-phased planning and the resulting capacity plans enable users to identify future constraints with enough possible time to pursue alternatives. If you can see the future, 90 percent of the time you can solve and eliminate the possible constraint before it becomes a problem.

Time-phased planning systems that generate capacity plans are not a new concept. They have been used very successfully for years in manufacturing. Most retailers, however, do not have a time-phased planning system in place. Or, if they do, it is only implemented at the retail distribution center level. Why the RDCs but not the stores? Product planning volumes for retail stores are so large that, until recently, there was no computer big enough and cheap enough to make it cost effective for them to run a time-phased planning system at the store level.

Without these time-phased planning systems, retailers do not have visibility into their own supply chain to see possible future constraints; instead, they can only address constraints as they show up. The retail store, then, is the last frontier for time-phased planning. Once time-phased planning systems are implemented there, the true benefits of the capacity plans they generate will be realized throughout the supply chain.

The Origins of Capacity Planning

Before we look at how capacity planning can be applied to the retail industry, it is important to understand its origins. The first attempts at capacity planning can be traced back to the mid-1970s with the introduction of materials requirements planning (MRP) inside the four walls of the factory. MRP reduced the constraint of not having enough material for production by ordering material at the right time and in the right quantity. Although this step was beneficial, material often sat on the factory floor waiting for people and production equipment to be free. At this point, we began to recognize that ordering material had to be balanced with the constraint of available production capacity. This realization led to the development of closed-loop MRP and ultimately MRP II (manufacturing resource planning), where all of a factory's resources could be planned and managed inside the computer. Today these systems are typically referred to as enterprise resource planning (ERP) systems.

These systems provided major breakthroughs in the way we could manage the flow of materials and products in and out of factories. Despite these advances, manufacturing bottlenecks and constraints still occurred in certain instances. People discovered that production could not proceed faster than its slowest point or machine. Thus the Theory of Constraints (TOC) and constraint programming (CP) techniques were created to help manufacturers sidestep production bottlenecks.

Deployment, Not Production Constraints, the Real Issue

For the most part, efforts such as MRP and MRP II have focused primarily on those constraints that affect production. They have largely overlooked constraints that affect the deployment and delivery of goods from the factory to the store. This has proven to be a major oversight. My experience with the consumer packaged goods (CPG) industry has taught me that the majority of constraint issues affecting the CPG supply chain occur during deployment and not during production.

Many of the constraints the CPG market faces are due to the promotional nature of the industry. CPG manufacturers sell as much as 40 to 70 percent of their products to retailers and wholesalers during promotional periods. In certain product categories, many wholesalers only buy from deal to deal. On average, most manufacturers promote their products three to four times per year. Typically, their promotions are no more than three to four weeks long. And store level promotions rarely exceed one week. In the case of the grocery industry, as an example, that means roughly $140 billion to $250 billion worth of product is manufactured and distributed to retail DCs over a 12-week period.

As part of the annual budgeting process, manufacturers' promotional plans are incorporated into financial and operating plans. This allows manufacturing people to know ahead of time when promotions are scheduled and to plan their production to accommodate any capacity constraints. For example, let's assume that the maximum amount of product a manufacturer can produce is 1,000 units per day. Let's also assume that the manufacturer knows that a promotion for the product will hit in three months and that the forecast predicts it will create a demand of 10,000 to 12,000 per day.

André Martin has been a practitioner, consultant, educator, and author in the supply chain arena for the past 35 years and specializes in supply chain integration.

This means that the manufacturer has no choice but to start producing product and holding it as inventory several weeks before the promotion starts.

On the average, CPG manufacturers start making product 10 to 12 weeks ahead of a promotion and deploy the product to their distribution centers. Thus, they are producing very large amounts of inventories ahead of the stores' needs and are deploying this inventory to their own DCs in anticipation of retail promotional orders. Manufacturers must do this because they have little flexibility to significantly add or reduce production capacity without creating a major financial impact. Instead, CPG manufacturers must purposefully build up their inventory before a promotion.

In other words, 40 to 70 percent of the time, manufacturers already have produced 100 percent of their product and deployed it into their own DCs six to eight weeks before that product will be sold in retail stores. Therefore, in the CPG industry, manufacturers do not have production constraints 40 to 70 percent of the time. The problem is not with production constraints; instead the real issue is one of proper deployment from the manufacturer's DCs to the retail DCs and stores.

Managing Constraints in Retail

Managing constraints in the retail supply chain, then, encompasses more than just production. It reaches outside the factory's four walls all the way to the final consumer at the retail store level. To manage constraints across a retail supply chain, companies must match the supply and demand of capital, people, equipment, and space across every node in the supply chain from supplier to retail store. Exhibit 1 shows a retail supply chain with the two flows of supply (product) and demand (actionable information). Managing constraints across such a supply chain would involve matching these two variables of supply and demand at each node in the exhibit.

EXHIBIT 1

The Retail Supply Chain

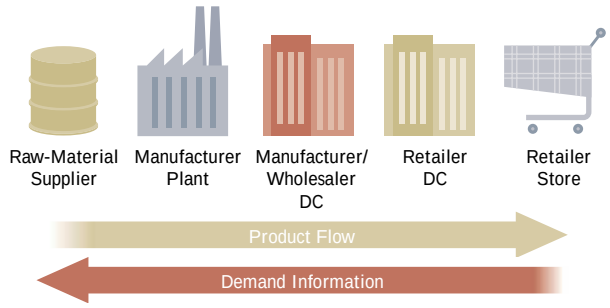

The first and most important of these variables is consumer demand, which starts at the retail store level. Consumer demand is pure demand; it represents what the consumer truly wants. It is free of all the noise that can distort the true nature of real consumer pull: promotions, forward buys, end-of-quarter or end-of-year inventory pushes, price wars, and new product introductions, among others. These distortions are part of the real world. We must deal with them, while, at the same time, we try to get as close to visibility of pure consumer demand as possible.

The second variable is supply. Supply constraints revolve around questions such as: *Do I have enough manufacturing capacity to meet customer demand today, tomorrow, and the next day? Do I have enough space to put away the inventory of frozen food I am receiving today, tomorrow, and the next day? Do I have enough trucks to deliver all my store orders today, tomorrow, and the next day? Do I have enough people, equipment, and receiving doors to receive all the supplier orders scheduled to arrive today, tomorrow, and the next day? Do I have enough capital authorization in my open-to-buy budget to purchase the products my purchasing system is telling me to order today, tomorrow, and the next day?*

The trick is to match demand and supply at every level of the retail supply chain inside the replenishment release horizon within a given supplier replenishment schedule. This matching can be accomplished through a time-phased planning system. Such a system provides you with the ability to simulate events in the future in an attempt to predict future constraints. Time-phased planning systems at the RDC level are commercially available today by software vendors such as Manugistics, i2, and Logility. (The sidebar on page 65 explains a time-phased planning system in greater detail.)

Although only recently introduced to retail, time-phased planning systems promise to revolutionize the way retail supply chains are managed. To fully demonstrate this potential, a brief comparison between current retail replenishment systems and time-phased planning systems is in order.

The replenishment systems currently being used in retail have a number of limitations. Specifically, they:

- Focus on execution rather than on both execution and planning. For example, most replenishment systems focus on the next order but do not look beyond this order.
- Tend to be departmental. For example, there are inventory systems, transportation systems, financial systems, buying systems, and so on.
- Tend to be different at each level of the supply chain rather than providing a single approach that spans the entire supply chain. Currently, there are store-level computer-assisted ordering (CAO) systems, distribution center CAO systems, and manufacturing CAO systems, but none of these are connected or integrated.

Even though they are derived from the MRP II and ERP systems that have been adopted nearly universally in manufacturing, time-phased planning systems represent a fundamental departure from the retail replenishment systems currently in place. The new technology:

- Allows you to do planning and execution at the store and RDC in a single system. In essence, execution can be thought of as the near-term horizon in a planning system. For example, the planning system calculates projected shipments from a distribution center to a store or from a supplier to a distribution center. Ordering is just a change in status from planned to released.

What Is Time-Phased Planning?

By Darryl Landvater

Time-phased planning is a simple calculation of the type we have all done from time to time. For example, it is the kind of calculation you would do with your checkbook. Assume you are paid on the first of the month, the mortgage is due on the 15th, and you want to put a down payment on a boat. Before you make the down payment, you would do a simple calculation to see if you will have enough money in the checking account on the 15th to pay the mortgage.

Similarly, if you are planning for the supply of products to a store, one of the first questions you would ask yourself is, "What do I have on hand?" Let us assume the answer is 29. The next question you might ask is, "What do I expect to sell?" Let us assume the answer is seven per day. Another question you might ask is, "Do I have any shipments in transit?" Let us assume the answer is no.

The figure below shows the first answers to these questions in the form of a calculation.

Store 1:
On hand = 29; Safety stock = 14; Order quantity = 12

	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7
Forecast	7	7	7	7	7	7	7
In transit							
Projected on-hand	22	15	8	1	-6	-13	-20
Schedule: receipts							
Schedule: shipments							

The projected on-hand line is a projection of what the on-hand balance in the store would be over time. Notice in day three that the projected on-hand balance goes below the safety stock of 14, and in day five that the store would be out of stock if the actual sales happened as forecast (seven per day).

Because the order quantity is 12, a shipment of 12 should arrive on day three and have been shipped a day earlier on day two (since the leadtime is one day). This is shown in the figure below.

Store 1:
On hand = 29; Safety stock = 14; Order quantity = 12

	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7
Forecast	7	7	7	7	7	7	7
In transit							
Projected on-hand	22	15	20	13	6	-1	-8
Schedule: receipts			12				
Schedule: shipments		12					

As a result of this planned shipment to the store, the projected on-hand balance in day three is no longer below the safety stock of 14. However, the next day, the projected on-hand does go below the safety stock again. The time-phased planning logic will continue this calculation, creating planned shipments in the future as far out as desired. In this case, we are only planning seven days into the future; in most actual implementations of these systems, the planning typically extends 20 to 52 weeks into the future.

This is shown in the figure below.

Store 1:
On hand = 29; Safety stock = 14; Order quantity = 12

	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7
Forecast	7	7	7	7	7	7	7
In transit							
Projected on-hand	22	15	20	25	18	23	16
Schedule: receipts			12	12		12	
Schedule: shipments		12	12		12		

The supplier replenishment schedule is the last line. It shows the projection of planned shipments from the supplier (a retail distribution center in this case) to the store. The same type of planning is done for the remaining stores. The sum of the planned shipments to the stores creates what is called dependent demand for the supplier (or the retail distribution center). It illustrates a fundamental principle in time-phased planning systems first articulated by Joe Orlicky, who wrote the first book on MRP in the late 1960s: "Never forecast what you can calculate."

The best and most accurate projection of demand for the distribution center is the sum of the planned shipments (or supplier schedules) for the stores. This is far more accurate than statistically forecasting the demand at the distribution center based on past shipments to the stores (a typical method in many systems). For example, if many of the stores were low on inventory last year, the history of shipments from the distribution center to the stores would show a spike in demand. However, if this year the stores were not low on inventory, this spike in demand would not materialize. However, a statistical forecasting system looking at last year's shipments would predict a spike, resulting in excess inventory. Conversely, if the stores were low on inventory this year but had adequate inventory last year, then a spike in demand would occur this year, even though the forecasting system would not have predicted it. As a result, the distribution center would go out of stock.

The figure below shows the time-phased plan for a retail distribution center supporting Store 1 and a second store.

	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7
Store 1 shipments	0	12	12	0	12	0	0
Store 2 shipments	0	0	24	0	24	0	0

DC 1:
On hand = 60; Safety stock = 34; Order quantity = 48

	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7
Dependent demand	0	12	36	0	36	0	0
In transit							
Projected on-hand	60	48	60	60	72	72	72
Schedule: receipts			48		48		
Schedule: shipments		48		48			

In this case, the last line is the supplier replenishment schedule for the distribution center. It represents the best estimate of what the retail distribution center needs from its supplier over time. This supplier replenishment schedule is used to make adjustments to the manufacturer's production plans to more closely match the output of the plant to the customer's needs. If the manufacturer receives a more accurate picture of the future that is kept up to date on a daily or weekly basis, customer service goes up, inventory goes down, and manufacturing productivity improves.

Darryl Landvater has been André Martin's consulting associate for the past 25 years.

- Provides access to a single set of numbers, which allows everyone to "sing from the same hymnbook." In other words, time-phased planning provides *one system* that manages inventory both at the retail store and at the distribution center. No longer does each have to have its own entirely separate system. As Exhibit 2 shows, the product plans, transportation plans, capacity plans, and financial plans for all the stores and RDCs are developed from the same set of numbers based on the planned orders table and the forecast table. As a result, the effects of decisions in one department are visible to the others. For example, no longer will a buyer unknowingly purchase more material than can be stored and create a space constraint for distribution. Instead, the distribution staff and the buyer can see the impact ahead of time and develop a plan that either handles the volume or changes the buy.

Time-phased planning systems can help companies accomplish the goal of matching supply and demand. More often than not, the constraint is on the supply side of the equation. In retail this is normal, because the demand side of the supply chain generates the requirements for capital, people, equipment, and space. Often a capacity problem can be resolved quickly by making more capital, people, equipment, and space available. For example, a retailer may decide to have its employees work two hours of overtime to receive today's orders. Or, it may decide to hire additional trucks to deliver tomorrow's store orders on time.

Although we can sometimes solve these capacity problems quickly, optimizing the retail supply chain implies being able to do so in the most economical manner possible. Is it financially viable to work two hours of overtime to receive today's orders? Is it profitable to hire additional trucks to deliver tomorrow's store orders? Time-phased planning systems provide the information to answer these questions.

Conventional wisdom has focused our attention on the supply side of managing constraints and optimizing the supply chain. We are left with the impression that demand is fixed and cannot be changed. Quite the opposite is true. On the retail side of the supply chain, demand can be changed to match supply in the majority of situations. This does not mean changing pure consumer demand. Rather, it means changing the store's distorted demand to the distribution center or the DC's distorted demand to the manufacturer. Think about it: Unless there is an out-of-stock situation at the retail store shelf, chances are we are simply replenishing inventory most of the time. Time-phased planning systems with one set of numbers give companies greater visibility and understanding of demand. These systems allow you to see whether the product is truly out of stock or whether the store or retail DC has just reached a minimum inventory level.

If we have a time-phased planning system that starts at the store level and runs across the retail supply chain, we have visibility into both current and future customer demand. With better visibility of demand, we can better resolve any constraint that comes up. This type of visibility is a prerequisite to optimizing the retail supply chain.

EXHIBIT 2

Working With One Set of Numbers

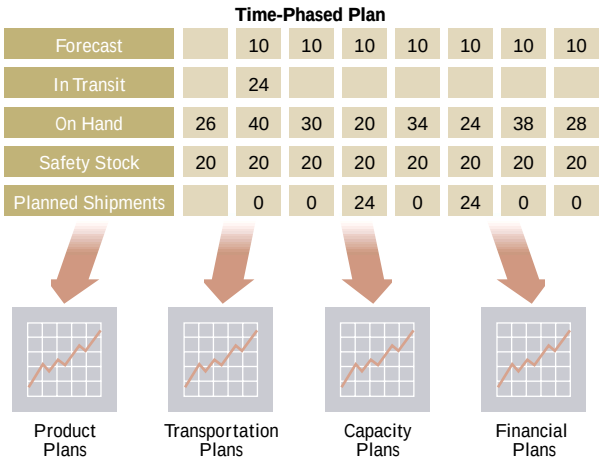

Visibility into the supply chain can help retail suppliers optimize their production schedules and yield better results. This is particularly true when an emergency arises. For example, suppose a key account has an order for delivery tomorrow, and we know it is a back order. If our production capacity is constrained, we can look at today's production schedule and swap the key account's back order with another order that can be delayed by a day or so because it is only replenishing inventory. Similarly, if the optimizer recommends working overtime, it would make more sense to work overtime to satisfy a true customer back order than to satisfy an order to replenish inventory.

Time-phased planning systems provide this type of visibility. As the examples suggest, this visibility does more than enable superior decision-making. It also provides retailers that use these systems with the ability to plan capital, people, equipment, and space capacity requirements into the future.

Time-Phased Planning in Action

Time-phased planning can be used to alleviate one of the retail industry's biggest constraints—delivery to retail distribution centers. Have you ever asked a retailer why it requires suppliers to make appointments for delivery to its RDCs? The answer might well be, "We have an appointment system so as to avoid traffic jams at our RDCs." In other words, if an RDC has the capacity to receive and unload 50 trucks a day, it certainly does not want 60 or more trucks to arrive on the same day. This would cause traffic jams, delivery problems, and delays. Consequently, an appointment system that limits daily deliveries to available capacity makes sense. This is why the appointment system is prevalent in the world of retail. It is necessary because traditional ordering systems are disconnected from the distribution organization that receives the supplier orders. Today in retail, one system orders and another system receives.

With a time-phased planning system, we have visibility of future deliveries into the RDC as far out as it makes sense.

Exhibit 3 shows a simple example that represents a receiv-

ing capacity requirement plan for the next eight days for an RDC. This plan was calculated by the time-phased planning system. It compares the required receiving capacity for all the orders scheduled to arrive at the RDC against the available receiving capacity. This information is visible to anyone in the company who needs to know. Looking at Exhibit 3, we can see immediately that we will have a receiving capacity problem on days three and seven. Thanks to this visibility, the distribution people can work with the buyers inside a single integrated system to resolve these future potential problems. Put another way, a supply constraint can be avoided.

The beauty of this information is that it identifies issues for resolution and allows people to work together to solve them. In this environment, people become proactive because the system gives them extra time to work out a solution. Experience has shown that people are very creative in finding solutions to problems. Unfortunately, they too often do not have the right information on a timely enough basis to use their judgment and expertise. Therefore, we often find ourselves in a firefighting mode as opposed to a fire-prevention mode.

With a time-phased planning system operating at both the store and RDC level, we have a true simulation of a retail business inside the computer. The information we are working with reflects reality. The time-phased information generated by the system represents future activities in buying, receiving, warehousing, and inventory investment in the RDCs. It also represents future shipping, receiving, and inventory investment in the stores. This information is then translated into the various "languages" of the company's different functional areas. To the buyers and financial people, the information represents future purchases and projected inventory investment in units and dollars. To the distribution people, the information represents future supplier receipts, warehouse space requirements, and future shipping activities to the stores. To the store people, it represents future receiving activity and how much inventory they will have in their stores.

The retailer that has this type of time-phased planning system can do away with the traditional appointment system. In fact, retailers that use the system effectively can do capacity planning for their receiving operations at the same time that they place orders with suppliers. In this way, receiving capacity problems are identified and resolved quickly, which results in suppliers' receiving time-phased delivery schedules with valid shipping dates. This approach eliminates the need to have an appointment system after the ordering is done. In addition, suppliers no longer have to call for appointments before shipping orders, as the delivery dates that appear on their delivery schedules are valid.

EXHIBIT 3

Available Capacity vs. Required Receiving Capacity

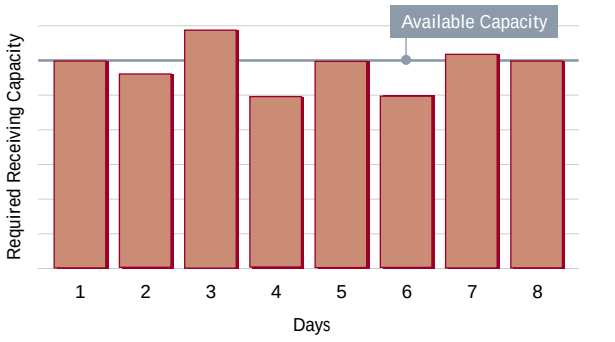

The example we just covered is but one illustration of capacity planning in the world of retail. As stated earlier, constraints across the retail supply chain come in different flavors. With a time-phased planning system, we can plan for a capacity requirement—whether it is capital, people, equipment, or space within a retail organization. The good news is, once we can plan for different types of capacity requirements, we can identify future potential constraints and resolve more than 90 percent of them before they ever become real constraints.

Until recently, retailers did not have systems that could plan ahead and calculate future requirements for products and capacity. Consequently, they had no choice but to implement different systems such as the appointment system. Once you are able to generate a valid simulation of a retail business inside the computer and keep it up to date, everything changes in the way you manage and control the flow of products across the retail organization. Suddenly things are simplified, solutions to problems are less complicated, and the obvious attraction to sophisticated software solutions changes dramatically. Though the advantages are clear, this approach is a new way of thinking that is foreign to most in the retail world.

A Straightforward Solution

Oliver Wight, a pioneer of MRP II, once said, "Solutions to a problem we don't understand always appear to be complex and very sophisticated, but once we understand the problem, the solution is always simple." This comes back to the question I posed at the beginning, Do we need optimization algorithms to solve retail supply chain constraints? In the great majority of cases, the answer is no.

The day you can plan your capacity requirements for capital, people, equipment, and space is the day you can identify and resolve 90 percent of your future constraints. The constraints you are left with are the here-and-now type of constraints inside your replenishment release horizon. Most of these constraints can only be resolved by spending more money than you normally would want to: for example, shipping the next order immediately by a faster and more expensive mode of transport to avoid a stockout.

As the old saying goes, "You can solve any problem; it's only a matter of time and money." Time-phased planning systems give you more time to solve your constraint problems. Those few constraints that are left can only be solved with money. In the end, the bottom line is this: It all depends on what is at stake and who the customer is. And those types of issues are a management call, not a computer call.