

Safeway Designs Mixed-Product Pallets to Support Just-in-Time Deliveries

John Chao

Action Properties, 688 W. Dana St., Mountain View, California 94041, john@actionproperties.com

Matty Chen

Accenture, Inc., 1661 Page Mill Road, Palo Alto, California 94304, matty.chen@gmail.com

Aileen Deng

Rui Feng/Ultrarich Trading, Ltd., Cosco Tower, 183 Queen's Road, Central, Hong Kong, aileen_d@hotmail.com

Harold Miao

Standard & Poor's, 2000 University Circle, Suite 600, East Palo Alto, California 94303, harold_miao@standardandpoors.com

Alexandra Newman

Division of Economics and Business, Colorado School of Mines, Golden, Colorado 80401, newman@mines.edu

Sophia Tseng

Oracle Corporation, 100 Oracle Parkway, Redwood Shores, California 94065, sophia.tseng@gmail.com

Candace Arai Yano

Department of Industrial Engineering and Operations Research and the Haas School of Business, University of California, Berkeley, Berkeley, California 94720, yano@ieor.berkeley.edu

In 1998, Safeway milk plants began daily deliveries to most stores, doubling their delivery frequency. The new schedule increased the workloads for assembling orders at the milk plant by 20 percent because, under the new system, very few store orders called for full pallets of a single product, and store-specific mixed pallets were time consuming to assemble. To reduce the excess workload, we designed two types of standard mixed pallets that, in quantities of one to three, could satisfy part of the demands of many stores on most days, thereby reducing the need for store-specific mixed pallets. Our mathematical model optimizes the design of the standard mixed-product pallets. These mixed-product pallets can be assembled prior to the peak shipping periods and during other off-peak times to ease peak workloads and eliminate excessive overtime. Safeway implemented the solution without our assistance and continues to use it nearly five years after implementation with good results.

Key words: programming; integer, applications; industries: agriculture, food.

History: This paper was refereed.

Safeway's San Leandro, California, milk plant, built in 1976, is one the company's largest milk-processing facilities. Safeway is a major grocery chain with over 1,800 retail stores and over 30 food-processing facilities, including nine milk-processing plants, in the US and Canada. The San Leandro plant provides fresh milk and other beverages to over 250 stores in northern and central California and Hawaii.

The plant produces fluid milk, juice, and bottled water, accounting for 73 stockkeeping units that are

distinguished by liquid content and packaging. The plant is also a distribution center for 52 additional products, primarily milk derivatives, such as cheese and yogurt.

Within the facility's 125,000 square feet are receiving areas for pumping raw milk from delivery trucks into tanks, processing areas for pasteurizing milk and mixing juice drinks, storage areas for raw and finished products, order-picking and palletizing areas, and a staging area for loading outbound trucks that deliver to stores. Most of the 100 or so employees operate

processing equipment, conveyors, palletizers, or forklifts, or perform order-selection (picking) and packaging tasks. The facility operates one 10-hour shift, six days per week, and each hourly employee works four of the six days each week. The plant's processing capacity is approximately 20,000 gallons per hour.

In 1998, Safeway implemented a new delivery policy that roughly doubled the frequency of milk deliveries to stores. Instead of making deliveries two or three times a week to each store, it made daily delivery to most stores (except small or distant ones). This change had many benefits, which Safeway staff had already observed before we began the project: (1) fresher milk for consumers; (2) greater accuracy of forecasting (because of the shorter horizon) and better control of inventory at the stores, leading to fewer stockouts and less spoilage; and (3) lower labor requirements for material handling at the stores. When the firm delivered milk only two or three days a week, milk products consumed a lot of space in the stores' secondary storage areas (usually behind the refrigerated display case), and workers who restocked the shelves spent considerable effort simply searching for products. With daily delivery, these search costs declined dramatically. Safeway staff estimated that material-handling labor dropped by one to two hours per store per day. Although transportation costs increased, reductions in inventory holding and spoilage costs offset the additional costs.

The onset of daily deliveries also increased operating costs at the milk plant because workloads for picking and palletizing orders increased. For milk derivatives, such as yogurt and cheese, the number of orders has a greater effect on picking effort than does the quantity in each order, and the number of orders had doubled. For beverages, workloads increased because of the greater degree of customization required to assemble orders. Previously, many stores ordered single-product pallet loads of medium- to high-demand products that the plant could automatically assemble onto store-deliverable pallets at the end of the manufacturing process. With more frequent deliveries, however, a full pallet of a single product exceeded the typical store's order quantity, so the plant had to assemble more store-specific mixed pallets. As a consequence of the additional order picking and handling, what the plant had handled with

a 10-hour shift took at least 12 hours, and the employees were forced to work at a faster pace to fill all orders on time. Workers' morale declined, and they increasingly complained about excessive overtime.

Safeway management asked the student members of our team, for their senior project at the University of California at Berkeley, to find a way to make the process more efficient.

System Operations When Our Project Began

The plant processes three major categories of liquid beverages: milk, juice and juice drinks, and bottled water. Large tankers deliver raw milk throughout the day, which the plant first pumps into refrigerated storage tanks, then pasteurizes and homogenizes. The plant mixes concentrates for juice and juice drinks with filtered water to produce liquid juices, and it processes water for various bottled water products.

The plant's high-speed filling lines pour the products into plastic bottles or waxed cartons. These lines operate on a stable production schedule that considers line changeovers, limitations of the material-handling equipment, and products' shelf lives. Filling capacity is not regarded as a limiting factor.

Following the filling process, the automated processing system places all products in crates and then palletizes the crates prior to storing them in a refrigerated area immediately adjacent to the processing area. In the palletizing process, each of the 73 products proceeds through one of three sets of equipment depending on the product size and storage mode. Most products proceed through the same types of processing steps, but the equipment used depends on the product characteristics. A crew of 13 workers palletizes the products and loads trucks; each worker does a specific task that usually involves operating a piece of equipment (Figure 1).

The plant uses two types of pallets: nine-stack pallets, which hold nine columns, or stacks, each containing six crates, used primarily for storage, and lighter weight seven-stack pallets, which hold seven columns of six crates, used for delivery to retail stores.

One palletizing machine is dedicated to stacking crates of identical products as they come off the filling line onto seven-stack pallets that can be sent directly

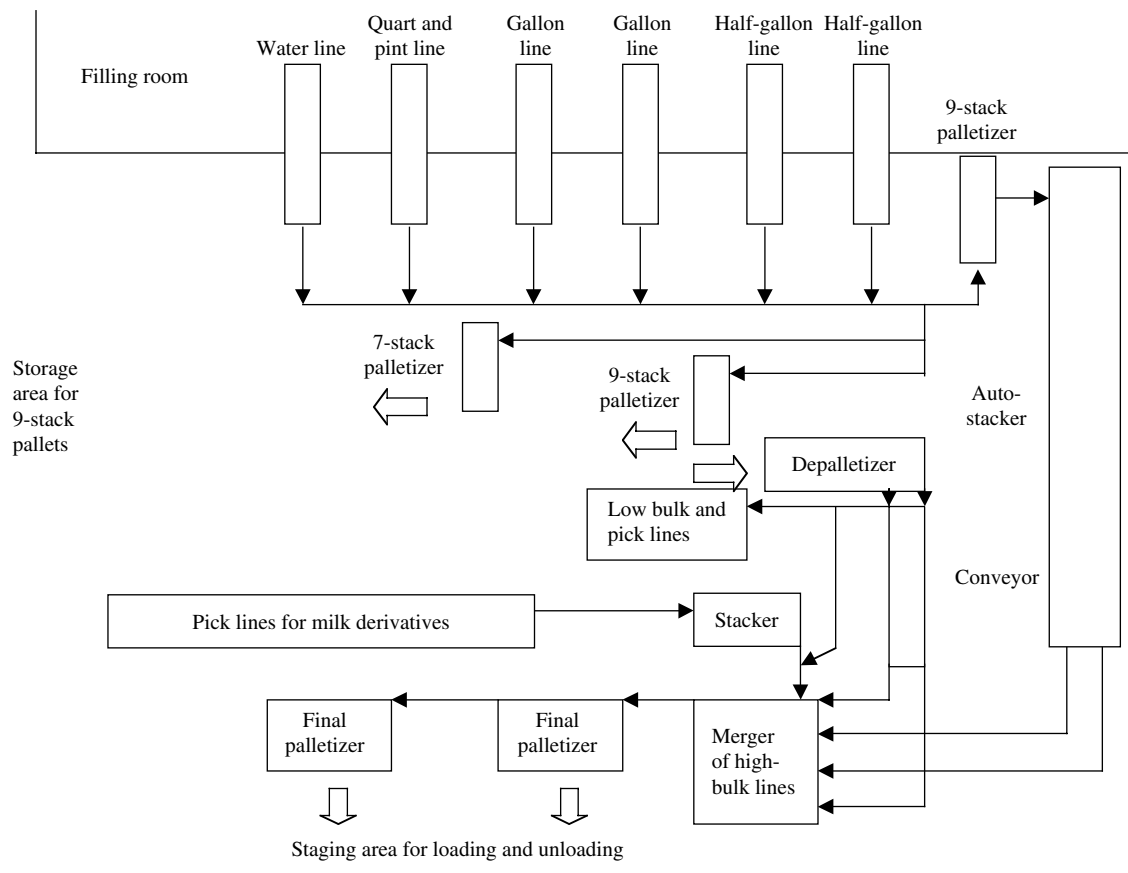


Figure 1: Milk products flow through the processing and palletizing systems. The filling equipment pours various types of milk into bottles and cartons of different sizes. The automated system places these filled containers into crates, and the crates continue to one of the palletizers for assembly onto pallets, which then proceed to the storage area via forklift or to the autostacker for storage in the form of stacked crates. Seven-stack, single-product pallets can be transported directly to stores, but the nine-stack pallets are designed for storage only. Forklift operators move the nine-stack pallets to the depalletizer for disassembly into constituent stacks. Final palletizing of customer orders requires combining products from the autostacker, depalletizer, and pick lines. The large, hollow arrows in the diagram show product movements requiring forklifts. All other arrows denote movement by automated or semiautomated conveyors or other material-handling devices.

to stores. Most of the palletizing machines, however, are designed for single-product, nine-stack pallets. Because the nine-stack pallets are quite sturdy, the plant can use them to move full pallets of heavy products around securely and can stack them on top of one another. The nine-stack pallets can be used only for storage at the milk plant, because they do not fit well in the truck trailers used for delivery and are too large and heavy to be handled by the forklifts at the stores. Thus, prior to shipping, the plant moves the nine-stack pallets by forklift from the storage

location to the depalletizing machine, which divides them into their constituent stacks. The plant combines these stacks with other products to make store-specific pallets.

The plant stores two of the four high-demand milk products (*high bulk* in Safeway vernacular) and all of the lower-demand milk products (*low bulk* in Safeway vernacular) on nine-stack pallets in the warehouse. (The low bulk products have substantial demand, but they are not the products with the highest demands.) Examples of lower-demand products are one-gallon

bottles of nonfat milk and half-gallon cartons of one-percent-fat milk. An example of a high-demand product is the half-gallon carton of two-percent-fat milk.

The remaining products fall into two categories. The first category consists of the other two high-demand products, which are stored in and handled by an autostacking machine, which is connected via conveyors to the final palletizing equipment (where store-bound mixed-product pallets are assembled). These items require no intermediate handling and can be automatically palletized at the touch of a button. (The operator indicates how many stacks of each product the autostacking machine should place on each pallet.)

The second category consists of products made from milk or milk derivatives (such as cottage cheese or yogurt) that the plant distributes. Safeway sells most of these items in pint or quart containers, and the stores order them from the plant as needed. Workers use a pic-to-light system to pick these items manually for each store order and place them in crates. In a pic-to-light system, lights show the locations of the items to be picked, reducing the times needed to find them. When a worker fills a crate or finishes picking an order, he or she sends the crates down a conveyor, and another worker stacks them on top of one another. These stacks then proceed on a conveyor to a final palletizer. Typically, two employees staff the pick line for the entire 10-hour shift.

One of two final palletizers combines the stacks of high-demand, lower-demand and picked products onto seven-stack pallets. To the extent necessary, workers add empty crates for stability and ease of handling during delivery, then wrap the pallets before loading them onto trucks. Because the final palletizer cannot finish the assembly of pallets until stacks arrive from three different sources (high-demand items from the auto-stacker, lower- and high-demand items from the depalletizer, and picked items), idle time often occurs on one or more of these feeder operations.

Order-Processing, Assembly, and Loading Timeline

Store orders arrive each day in waves, corresponding to truck-dispatch times. The first wave of store orders arrives at 8:00 A.M. Converting the orders into a truck-loading plan takes about two hours: computer programs at Safeway and at Summit Logistics, Safeway's

third-party logistics provider, determine the composition, weight, and volume of the pallets comprising each store's order and the vehicle routes. The truck-routing algorithm minimizes the number of trucks used each day as a proxy for minimizing total transportation cost. The routing decisions must consider the time required for unloading at the stores on the route, the distance between the stores and the plant (the furthest store is over 200 miles away), and the hours during which each store accepts deliveries. The processing takes two hours because the information systems are incompatible, preventing efficient transfer of data between the two systems, and because Safeway and Summit Logistics have different data needs. (For example, Safeway needs detailed, item-by-item order-quantity data, whereas Summit is concerned only with pallet weights and volumes.) Even if the data-transfer delays were eliminated, some time is needed to solve the truck-routing problem.

The truck-loading plan for the first wave of orders is available at 10:00 AM when the warehouse staff reports to work. The workers must pick the orders, palletize the goods, and load them onto the trucks within one hour to ensure that trucks can make the last delivery on each route while the stores still have unloading personnel available.

The plant repeats the process for two additional waves of orders. The second wave of orders arrives at 10:00 AM, and again, two hours elapse before the truck-loading plan is available, so the workers begin loading around noon. The second wave consists of shuttle orders: the plant loads the trucks with milk; they go to Safeway's distribution facility in Tracy (about 40 miles east of San Leandro), pick up groceries, and continue on to their destinations. The shuttle orders must leave San Leandro by 4:00 PM, giving the crew about four hours to assemble the orders and load the trucks. Consequently, workers are under less time pressure during the second wave. They process the third wave of orders throughout the rest of the day.

Problem Statement

After observing plant operations and discussing several potential remedies for the excessive overtime with the director of supply operations and the plant

manager, we decided to focus on decisions largely controlled by the San Leandro plant, not Summit Logistics, Safeway retail stores, or the Safeway corporate office. Consequently, we did not consider changing the vehicle-routing decisions or the store-ordering procedures. Instead, we sought to mitigate the adverse effects of daily deliveries on the plant's operations.

By instituting just-in-time delivery, Safeway created two types of problems for the San Leandro plant: a severe workload peak for final palletizing of the first wave of orders between 10:00 and 11:00 AM and a sustained increase in workload in the same area during the second and third waves of orders. The final palletizers were operating continuously throughout the first wave and most of the second, and even a short breakdown of one of them caused major disruptions. Space constraints, caused in part by the move to just-in-time deliveries (which shifted some inventory from the stores to the milk plant), prohibited the addition of another final palletizer. Thus, we needed a solution that reduced the workload on the final palletizers or enabled the plant to do some of the palletizing at off-peak times.

The alternative of reducing the workload on the final palletizers by moving some of the workload to other palletizers proved to be unworkable because the rather complex conveyor system was not designed to permit other palletizers to handle mixed-product pallets easily. Because demand for single-product seven-stack pallets had declined (because of just-in-time deliveries), the palletizer dedicated to creating those pallets had idle time but neither the conveyor connections nor a location suitable for handling mixed pallets. Consequently, we sought ways for the plant to do some final palletizing at off-peak times when the two final palletizers had idle capacity.

After studying the flow of products and information carefully, we hypothesized that by assembling in advance carefully configured seven-stack mixed pallets of high- to medium-volume items for direct loading onto trucks, we could shift some of the workload to nonpeak times, namely, early morning hours (before 10:00 AM), late morning hours when Summit Logistics was determining truck loads for the second wave of orders, and late afternoon hours when the final palletizers had unused capacity. By doing this,

we could speed order assembly during the critical morning peak.

By using preassembled (or prestaged) mixed-product pallets, we would also improve labor utilization in two ways. First, the palletizer operators could prepare the mixed-product pallets between the waves of orders, when they would otherwise be idle. Second, the mixed-product pallets would satisfy some of the demand that had imposed workloads on the autostacker and the depalletizer. The autostacker and depalletizer were secondary bottlenecks, so reducing their workloads would reduce the average cycle time per store order, reducing the idle labor in other parts of the assembly operation as well.

We decided to identify one or two mixed prestaged pallet (MPSP) configurations to maximize the number of such pallets that the plant could use, on average, each day. We had to consider a few practical issues. First, the plant could assemble efficiently only mixed pallets consisting of a combination of single-product stacks; mixes consisting of partial stacks would require too much manual handling. Second, conveyors were not configured to efficiently form prestaged pallets that combined high-demand products (on the autostacker) and lower-demand products (from the depalletizer). Such combinations would not improve labor efficiency nor reduce the workload on the bottlenecks, so we decided to avoid mixing the two product types. Third, the order-processing system was not set up to handle mixed pallets, much less multiple types of mixed pallets. If we assigned any product to more than one type of mixed pallet, we would have to select the best combination of the various mixed pallets for each store, that is, we would have to solve yet another decision problem. Thus, we decided to design only one mixed pallet for high-demand products and one for lower-demand products. Finally, it was impractical to change the configuration of the mixed pallets frequently because each change would necessitate reprogramming the order-processing system and could require a change in the material-handling equipment or pallet-assembly procedures. However, when configuring the mixed-product pallets, we had to account for daily demand variations to obtain an accurate assessment of how many mixed pallets the stores could use, which is the quantity we wish to maximize.

Demands exhibited only very mild seasonal variations, and consequently Safeway management considered day-of-week variations to be more important.

Given the set of products that might be included on the same mixed-product pallet, we initially formulated a nonlinear integer program in which the objective is to maximize the total number of mixed-product pallets that all stores could use across a set of representative days spanning several weeks. The primary constraints ensure integrality of the number of pallet positions assigned to each product and that their sum equals the total number of positions, that is, seven. Other complicated constraints collectively define how many mixed pallets a particular store can use on a given day for a given assignment of pallet positions. We rewrite the problem as a general integer program with a much larger number of decision variables and use the latter formulation in solving the problem.

Related Literature

We found no published literature on Safeway's problem. Although it may appear similar to an assortment problem, it is not, as the assortment of products is already specified and the question is how best to package the products to satisfy store demands. The most closely related class of problems is in the area of modular design. The oldest and most traditional version of the modular-design problem is to decide the composition of one, or a limited number of, distinct kits of components (Evans 1963, 1970; Shaftel 1971). With one or more standard finished products to be assembled using only collections of kits, the problem is to determine the composition of each kit so as to minimize the value of the excess (unneeded) components. Although this class of modular-design problems has parallels with our problem, it also has differences. In our problem, we have constraints on the number of components in the kit, and we are not allowed to have excess components. Our goal is not to satisfy a superset of the requirements at minimum overage cost, but to satisfy the largest possible subset of the requirements with constraints on the number of components in the kit. In addition, the traditional modular-design problem usually involves several to a dozen finished products. On the other hand, in our application, all of the finished products (customer

demands for a specific mix of products) are distinct, as two stores are unlikely to order exactly the same quantities of all products. Consequently, the number of finished products in our problem would be in the hundreds or thousands.

A more recent version of the modular-design problem (see Goldberg and Zhu 1989) concerns a family of related products, each of which requires a particular combination of capabilities or functionalities. These functionalities can be provided by components in a module or by other individual components in the product. The decision maker must determine which components to include in a common module or set of modules. Some substitution is allowed in that components with a higher level of capability, and usually higher unit costs, can satisfy the needs of a product with lower-level requirements. There is a fixed cost for designing the module, economies of scale in producing or procuring components, and savings in assembly from using a module. This modular-design problem has variants, but most involve similar cost trade-offs. We now explain the relationship between our problem and this version of the modular-design problem. Our mixed pallet is the module, and we can save assembly costs if we use a module. For the Safeway application, the fixed cost for designing the pallet is negligible, and practical (space) considerations, rather than module-design costs, made it impossible to consider using several types of modules. Stores may be able to use multiple modules of the same type, a situation that does not arise commonly in a modular-design setting. Substituting higher capability components is neither relevant nor allowed in the Safeway setting because the consumers choose the specific products (components) and do not generally desire an upgrade (for example, milk with a higher fat content).

Outside of the modular-design literature, we know of only one paper that concerns a problem related to ours. Swaminathan and Tayur (1998) developed an approach for designing "vanilla boxes," partially completed products that can be assembled prior to receiving customer orders for use in one or more finished products. After receiving customer orders, the plants assemble the finished products—desktop computers—from the vanilla boxes or from individual components using limited assembly capacity.

In Swaminathan and Tayur's model, the trade-off is between stocking less-complete and therefore more-generic vanilla boxes that can be used in a broad array of finished products and taking the risk of having too little assembly capacity to satisfy demand, versus stocking more-complete (less generic) vanilla boxes that can be used in a narrow range of products but require less assembly capacity to complete. Our problem differs in at least four ways: First, our vanilla boxes must contain a specific number of components, so we have limited flexibility with respect to how generic our vanilla boxes can be. Second, we can partially meet customer demands using multiple vanilla boxes, whereas the Swaminathan–Tayur model implicitly assumes that each finished product requires at most one vanilla box. Third, their model also implicitly assumes a limited mix of components in a vanilla box. For example, each computer contains one motherboard and one power supply. On Safeway's mixed pallets, however, the same product may occupy more than one position. Consequently, our related decisions are general integer variables, not binary variables. Fourth, we have no hard capacity constraint on assembly, because we could use overtime at a smaller cost than the cost of lost sales.

We can also view Safeway's problem as an unusual variant of a knapsack problem. The objective is to maximize the sum of the utilities of multiple evaluators (of the knapsack contents), one for each store-day combination. (We can restate the formulation with one evaluator for each store but cannot eliminate the multi-evaluator nature of the objective function.) Each evaluator's utility depends upon the combination of items in the knapsack (that is, items assigned to the mixed pallet) in a highly nonlinear way. We know of no research on knapsack problems in which there are both multiple evaluators and such complex utility functions.

Solution Strategy, Benefits, and Costs

Safeway's problem involves about 250 stores and 20 to 25 candidate products. Although the demands exhibit stable day-of-week patterns, demands on a specific day of the week (for example, Tuesday) do vary from week to week. Consequently, it is useful to consider demands over a horizon of several weeks and, if possible, during different seasons. (Consump-

tion of fluid milk exhibits very mild seasonal fluctuations in the region the plant serves.) The data we used in our analysis were limited, because Safeway purged order data from its system after it delivered orders. To obtain order data (for the 250 stores) we asked Safeway to save printouts for us for 18 days (six days per week for three consecutive weeks) during March 1999, a time of the year with average milk consumption. Had data been available from different times of the year, we could have included more data.

With a problem of this scale, we knew that state-of-the-art optimization software at the time of our project would not produce good solutions and might not even find feasible ones. We took a pragmatic approach and obtained solutions as follows: We carefully analyzed store order data to ascertain (1) which products had daily demands exceeding one stack for most stores and (2) the relative magnitudes of those demands. Based on this analysis, we identified the most promising candidates for inclusion in the mixed pallet. We then constructed a few initial candidate solutions. Starting with the best candidate solution, we applied variable reduction methods (appendix) to reduce the search space. We then enumerated the several remaining alternatives and chose the best one.

We applied our procedure to the demands that remained after eliminating the portion that could be satisfied by single-product pallets. We arrived at the following mixed prestage pallets:

The High-Demand MPSP:

Three stacks of Dairy Glen two-percent-fat milk (gallon),

Two stacks of Dairy Glen homogenized milk (gallon),

One stack of Lucerne two-percent-fat milk (half gallon),

One stack of Lucerne homogenized milk (half gallon).

The Lower-Demand MPSP:

One stack of Dairy Glen one-percent-fat milk (gallon),

One stack of Dairy Glen nonfat milk (gallon),

One stack of Lucerne whole milk (gallon),

One stack of Lucerne two-percent-fat milk (gallon),

One stack of Lucerne one-percent-fat milk (gallon),

One stack of Lucerne nonfat milk (gallon),

One stack of Lucerne one-percent-fat milk (half gallon).

Although we had data for only 18 days in one season, the plant managers considered the proposed MPSPs to be consistent with their knowledge of demand and store-order patterns. Using the available data, we estimated that, on average, the plant could use 40 high-demand and 60 lower-demand MPSPs each day. While performing this analysis, we learned that the number of single-product pallets shipped to stores declined from 80 per day to 30 per day when Safeway introduced daily delivery. This meant that the final palletizer workload had increased about 10 percent from 520 pallets per day to 570 per day (the plant shipped approximately 600 pallets a day before and after Safeway introduced daily deliveries). Taking into consideration the additional time required to handle smaller quantities, it is not surprising that the workday increased from 10 to 12 hours. However, by using the recommended MPSPs, the plant could reduce the workload on the final palletizers during its peak periods from 570 pallets per day to 470 pallets per day, well within the system's capability during a 10-hour shift. The plant could achieve a reduction of 100 pallets because of the demands satisfied by the 40 high-demand and 60 lower-demand MPSPs assembled off-peak.

To determine the time required to stage the extra 100 MPSPs, we performed time studies on the final palletizer. We found that its maximum capacity was about 70 pallets per hour. Thus, it would take about an hour and a half to prestage the 100 MPSPs the plant could use on an average day. Three workers could perform this process: one to dial in the combination of products, one forklift driver to feed nine-stack pallets into the depalletizer, and another forklift driver to transport the standard mixed pallets to the designated area for storage. However, because of timing considerations, the plant needed most of the MPSPs ready before it loaded the first wave of orders, and consequently the three workers had to perform the work on overtime prior to the start of the shift. We estimated that 70 of the 100 pallets had to be ready by 10:00 AM; the remainder could be prepared during the idle period between 11 AM and noon while Summit Logistics determined the truck-loading and routing plans for the second wave of orders, during idle periods in the afternoon, or even in the late afternoon

or early evening in anticipation of the next morning's first wave. By spreading the prestaging throughout the day, plant managers can adjust quantities in response to actual demands.

Our recommended solution entails scheduling just three workers to come in an hour earlier than usual to start prestaging the bulk of the MPSPs for that day. Thus, the net savings amounts to 23 person-hours of overtime each day—the difference between the two hours of overtime per day for all 13 members of the crew and one hour of overtime a day for three people. With overtime costs exceeding \$30 per hour, the estimated annual savings in overtime alone is over \$200,000.

The goods from three areas (autostacker (high demand), depalletizer (lower demand), and picked items) must be synchronized to form pallets on the final palletizer, and the workloads in these three areas are highly unbalanced (Figure 2). On average, without the MPSPs, the cycle time of the pick line is 20 seconds per order (40 seconds per order and two workers); however, the pickers must wait for the autostacker to finish assembling the high-demand products (in an average of 45 seconds) before beginning the next order, causing picker idle time of approximately 25 seconds per order. By prestaging

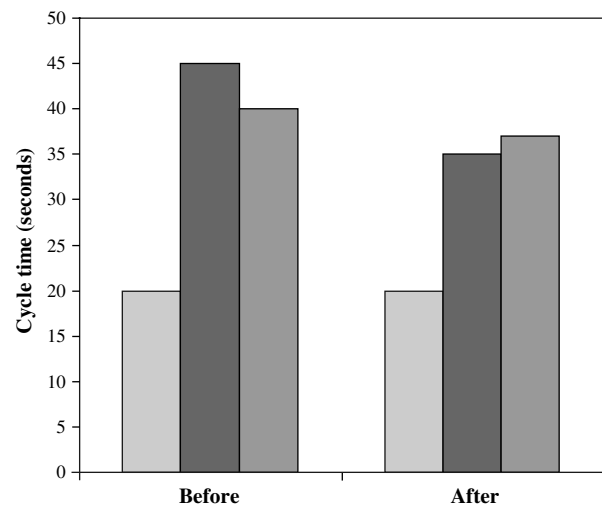


Figure 2: This bar chart shows the estimated average cycle time by store in the three assembly areas: pick line (light), depalletizer (dark), and autostacker (medium) before and after the introduction of standard mixed pallets.

MPSPs, we can reduce the average workloads on the final palletizers for the high- and lower-demand items by nearly 20 percent (a reduction from 570 pallets to 470 pallets for the same number of orders). That reduces the picker idle times and the average cycle time per order on the final palletizers. Because the contents and volumes of the orders vary and because they may consist of multiple pallets, it is difficult to estimate how much we can reduce the final palletizer cycle time in practice. However, even a five or 10 percent reduction could have a measurable impact on plant capacity and relieve stress during peak periods. Thus, introducing the mixed pallets should reduce the workload of the final palletizers and the cycle time for the remaining work.

Safeway needs to re-solve this problem only when demand patterns (the relative demands of the candidate products and the mix of store sizes) change so much that the benefit of reorganizing the infrastructure required to support new mixed pallets outweighs the cost.

Implementation

To alter such a long-standing system, we had to anticipate some change-management issues. We polled the workers in the palletizing, storage, and loading areas and found broad support for a change that would make their work less harried and reduce their overtime. Because many workers start their 10-hour shifts at about 10:00 AM, their normal shifts extend into the midevening hours, and two hours of overtime meant staying until 11:00 PM or later, which employees found very disruptive for their families. They were ready for change, even though it would reduce their overtime compensation.

When we finished our study in May 1999, Safeway was preparing to replace a 1970s vintage Y2K non-compliant computer system that performed part of the order-processing functions for the plant. At that time, Safeway did not expect to be able to implement our recommendation because it was unclear whether it could change the order-processing system to accommodate the mixed pallets. Five then-student authors of this paper graduated and scattered for new jobs and graduate school, and their project advisor (Candace Arai Yano) resumed her usual activities.

Unbeknownst to our team, as part of updating the computer system in late 1999 and early 2000, Safeway implemented a method to net out the prestaged mixed-product pallets and used the mixed-pallet solution that we suggested. As of spring 2005, nearly five years after implementation, the San Leandro plant continues to prepare 50 to 60 of these prestaged pallets each day prior to the 10:00 AM peak workload period.

The plant managers and employees decide how many mixed pallets to prepare based on their experience. This method is working well for the plant. One side benefit of using standard mixed-product pallets is that Safeway can take advantage of risk pooling, because high and low demands from the 250 stores tend to offset one another. This very high degree of risk pooling, combined with employees' experience during five years of implementation, enables the plant staff to decide quantities for mixed pallets and to phase their assembly during the workday.

If we had been asked to recommend an approach to determine the number of mixed pallets to prestage, we would have suggested using the newsvendor model. For the San Leandro plant, such an analysis would suggest that the number of mixed-product pallets should be set to a high fractile of the distribution of the aggregate demand for these pallets. A shortage of mixed-product pallets would cause additional work for the final palletizer during peak periods, for which the implicit penalty is high, while a controlled number of excess mixed pallets could easily be stored for the next wave of orders or until the next day, with the only penalty being a partial day's loss of freshness of the product.

Safeway has not performed a detailed analysis of the cost savings, and our team was not even aware of the implementation until late in 2002, so we also could not perform a systematic and timely comparison. Even if we had attempted such a comparison, it would have been difficult to measure the cost savings accurately because the business has changed (for example, the number of stores has grown and the milk-derivative products handled by the plant have changed). Nevertheless, Safeway managers report that the prestaged pallets have reduced the need for overtime and eased the stress of the morning peak. They are investigating advanced automation

that would provide more flexibility and efficiency in forming mixed-product pallets, which reflects a long-term commitment to the strategy of using mixed-product pallets and sufficient benefit to justify capital investments toward that end.

The mixed-pallet strategy has also freed up final palletizer capacity, which has made the system more robust to equipment failures and has allowed the plant to handle a growing number of stores. The plant was operating at about 120 percent of nominal capacity nearly every day (two hours of overtime in addition to a 10-hour work day), and because many stores are not open around the clock, the work had to be done within a specific time window. Without the standard mixed pallets, the plant probably would have had to change its material-handling system drastically to increase the rate at which it assembled orders. Safeway managers also reported improvements in on-time deliveries at stores, reducing the need for unloading staff to work overtime, smoothing the flow of work and reducing unwanted overtime, which has improved employee morale at the milk plant.

Discussion

Many manufacturers in the fashion-goods industry assemble cartons containing various styles, sizes, and colors for initial shipments at the beginning of a season. They commonly ship dozens or hundreds of pieces to each store. Quite commonly, they use only one type of mixed carton for a product family, so they may send a retail store one or two such cartons, depending upon the store's sales volume. This practice often leads to inventory imbalances—shortages of some sizes or colors and overages of others—because one type of mixed carton cannot match the demands of diverse stores. These manufacturers could use small numbers of mixed cartons of the size currently in use or several mixed cartons of smaller size(s), each with different proportions of sizes or colors, to help them match inventory to anticipated demand without the need for expensive manual picking of small order quantities. Our formulation could be adapted to help them design the mixed cartons.

At the time of our project, it was possible to obtain near-optimal solutions only for very small problems that were not of much practical value. Now—nearly

five years after Safeway's implementation—it is possible to optimally solve problems that could be regarded as scaled-down versions of the true problem Safeway faced (appendix), because both optimization and computing technology have advanced. For the purpose of exploring the sizes of problems that can be solved with today's technology, we brought on another team member (Alexandra Newman). Variable reduction techniques and easy-to-compute heuristics (appendix) show some promise. In our problem, however, reduction of the number of variables does not lead to a reduction in the number of constraints because of the structure of the constraint set. Furthermore, we expect the linear-programming bounds to be loose because few constraints are binding, so branch-and-bound methods may not be efficient unless we can develop tighter bounds.

Further research is needed to solve very large problems, such as the fashion-goods problem, where the number of product types, the number of units of each product, and the number of stores are much larger than they are in the Safeway application. Theoretical advancements for problems of this type may offer the promise of solving a great variety of real-world problems.

Appendix

Mathematical Formulation

For each product type (high-demand and lower-demand), the decision problem is as follows:

Notation

i : product index, $i = 1, \dots, I$.

s : store index, $s = 1, \dots, S$.

t : time period index, $t = 1, \dots, T$.

D_{ist} : demand for product i at store s in time period t (expressed in stacks).

Y_i : number of stacks of product i on a mixed pallet.

N_{st} : number of mixed pallets that store s could use in time period t .

Formulation

$$\begin{aligned}
 \text{(NLP)} \quad & \text{Max} \sum_s \sum_t N_{st} \\
 & \text{subject to} \sum_i Y_i = 7,
 \end{aligned} \tag{1}$$

$$N_{st} \leq D_{ist}/Y_i \quad \text{for all } i, s, t, \quad (2)$$

$$Y_i \text{ nonnegative integer} \quad \text{for all } i, \quad (3)$$

$$N_{st} \text{ nonnegative integer} \quad \text{for all } s, t. \quad (4)$$

The objective is to maximize the number of standard mixed-product pallets that all stores could use across all days. The first constraint ensures that the mixed-product pallet contains exactly seven stacks of products. The second set of constraints defines the number of such pallets that each store can use each day. These constraints are equivalent to

$$N_{st} = \min_i \{ \lfloor D_{ist}/Y_i \rfloor \}.$$

The remaining constraints ensure integrality and non-negativity of the decision variables.

The formulation above is compact but nonlinear. We linearize it as follows:

Additional Notation

- z_{iq} : 1 if product i is assigned q stacks on the mixed pallet, $q = 0, \dots, 7$.
- n_{istq} : $\lfloor D_{ist}/q \rfloor$ = maximum number of mixed pallets that store s could use in time period t , considering product i alone, if there are q stacks of product i on the mixed pallet; for $q = 0$, set n_{istq} to an arbitrarily large finite integer.

Formulation

$$(IP) \quad \text{Max} \quad \sum_s \sum_t N_{st}$$

$$\text{subject to} \quad \sum_q z_{iq} = 1 \quad \text{for all } i, \quad (5)$$

$$\sum_i \sum_q q z_{iq} = 7, \quad (6)$$

$$N_{st} \leq \sum_q n_{istq} z_{iq} \quad \text{for all } i, s, t, \quad (7)$$

$$z_{iq} \text{ binary} \quad \text{for all } i \text{ and } q=0, \dots, 7. \quad (8)$$

The objective function is the same as in the nonlinear formulation. Constraints (5) and (6) ensure that we allocate some number of stacks (possibly zero) for each candidate product and that the total allocation for all products is equal to seven. In constraints (7), the right-hand side of the inequality represents the number of mixed pallets that store s could use on

day t , considering product i alone. Thus, for a given (s, t) pair, the constraint ensures that the number of pallets allocated to store s on day t does not exceed the minimum of these values across all products.

It is not possible to eliminate the $q = 0$ terms and convert constraints (5) to less-than-or-equal-to constraints because without a $q = 0$ term for each i, s , and t , the right-hand side of constraint (7) would be zero if $z_{iq} = 0$ for all $q \geq 1$, which would overly constrain the value of N_{st} . To make the formulation as tight as possible, for each i, s , and t , we set

$$n_{ist0} = \max_i \{ n_{ist1} \}. \quad (9)$$

Each element in braces is the number of mixed pallets that store s could use in period t if we allocate exactly one slot to product i on the pallet, and if we consider product i in isolation. Thus, each element also is an upper bound on the number of mixed pallets that store s could use in period t if we allocate any space to product i , because the values of n_{istq} decline as q increases. Although we do not know in advance whether a specific product will be allocated at least one slot on the pallet, we do know that at least one product must be allocated one or more slots on the pallet. Consequently, by taking the maximum over all i , we have a valid upper bound on the total number of mixed pallets for store s on day t . In the absence of other restrictions, (9) defines the smallest valid values of n_{ist0} .

A formulation of a problem of the size that Safeway faces contains many dozens of general integer variables or alternatively, a few hundred binary variables. It also includes over 100,000 constraints, almost all of which are contained in (7). Very few of these constraints are binding, however, because for any (i, s, t) triple, the corresponding constraint may be binding only if on day t , store s orders a quantity of product i that is an integer multiple of the number of cartons of product i selected for the standard mixed pallet. Furthermore, the constraint will be binding only if the integer multiple is the smallest such integer multiple among all i . Although the vast majority of these constraints are not binding for any feasible solution, we need to retain them because small changes in the decisions can cause them to become binding. These aspects of the problem make it difficult to solve.

Variable Reduction

An upper bound on the number of mixed pallets with $z_{iq} = 1$ is

$$\sum_s \sum_t n_{istq}, \tag{10}$$

which represents the maximum number of mixed pallets that the stores could use if we were to consider product i alone, without considering constraints imposed by other products j , $j \neq i$. Suppose that we have a feasible solution whose objective value is P . Then, if the expression in (10) is less than P , we can eliminate $z_{iq} = 1$ from consideration. Generally speaking, this allows us, for each i , to eliminate $z_{iq} = 1$ for large values of q (for example, greater than two). We also fix the value of z_{i0} to 1 and z_{iq} to 0 for $q \geq 1$ in cases where z_{i0} is the only uneliminated variable for product i .

Heuristics

We consider two different methods for generating an initial feasible solution. In the first heuristic (H1), we rank the products in decreasing order of total sales volume across all stores and all days, and for the top seven products on the list, we set $z_{i1} = 1$.

Products are ranked differently in the second heuristic (H2). We compute

$$\sum_s \sum_t \delta_{ist} \quad \text{for all } i, \tag{11}$$

where $\delta_{ist} = 1$ if store s requires at least one stack of product i on day t . We then rank products according to this metric and set $z_{i1} = 1$ for the top seven products on the list. The metric in (11) ignores daily store demands above and beyond one stack and focuses on the number of store-day combinations on which demand for product i is one stack or greater. We included this metric after noting that realistic demand patterns would only infrequently lead to the selection of two or more stacks of a given product, and for many of the small to medium stores, the question was whether their demands were large enough to require a full stack of product i . The expression in (11) prioritizes products with these factors in mind, but for simplicity (as in H1), it does not account for interactions among products.

Computational Results

We conducted numerical experiments to assess the computational tractability of the (linearized) integer-programming formulation. Our problem dimensions correspond to those that Safeway would face if it were to consider demands over several similar working days. The scenarios have a mix of 80 small, 80 medium, and 90 large stores. We generated two sets of problem instances and solved each problem instance in three different ways. The first method is an implementation of the CLPEX solver, Version 9.0, without an initial feasible solution. The second method is a two-phase solution procedure in which we allow CPLEX 1,500 seconds to find an initial solution whose objective function value serves as the basis for variable reduction. Then, we solve the reduced problem, again using CPLEX. The third method is similar to the second except that we use the better of the objective values of H1 and H2 as the basis for variable reduction.

The problems in our first set have characteristics that are roughly analogous to the problem Safeway faces in designing the mixed-product pallet for lower-demand products. We considered 20 product types, 250 stores, and five days, which corresponds to problem instances with 160 binary variables and 25,021 constraints. We distinguish demand distributions by store size and by product. For each of the 250 stores, we randomly generate demands (number of stacks) for each of the 20 products (six low-demand, seven medium-demand, and seven high-demand) according to the distributions shown in Table 1.

Without the advantages of an initial feasible solution and variable reduction, we found optimal solutions in less than 40 minutes, on average, for the first

Product demand	Store size		
	Small	Medium	Large
Low	0 w.p. 0.5; 1 w.p. 0.5	DU[2, 5]	DU[4, 8]
Medium	DU[1, 3]	DU[5, 8]	DU[7, 11]
High	DU[3, 5]	DU[7, 8]	DU[8, 12]

Table 1: The demand distribution varies by store size and product demand level for each product in the first set of problems. Low, medium, and high levels of product demand are relative levels and do not correspond to Safeway’s high- or low-demand product categories. DU stands for discrete uniform.

Scenario	CPU time without variable reduction (seconds)	CPU time with variable reduction using the objective from CPLEX (seconds)	Variables in the problem after reduction using the initial objective from CPLEX
1	2,202	1,555	28
2	2,524	1,551	28
3	1,942	1,635	31
4	2,407	1,500	14
5	2,901	1,544	28
6	2,617	1,551	28
7	2,089	1,547	28
8	1,744	1,544	28
9	1,790	1,561	28
10	2,619	1,542	28

Table 2: The computational results for the first set of problems demonstrate the effect of variable reduction on the number of decision variables (reduction from 160 in the original problem), which in turn reduces CPU times.

set of problems (Table 2). If we allow CPLEX 1,500 seconds to identify a solution and then solve the problem that remains after variable reduction, the number of variables declines dramatically and the CPU time falls to 26 minutes, on average. The reduced problems are smaller than they appear because in many cases, only two variables remain for a given product, z_{i0} and z_{i1} . For such a product, we can replace these two variables with a single binary variable. Finally, by using the better of the heuristic solutions from H1 and H2—both of which turn out to be optimal for these realistic problem instances—and

Store size		
Small	Medium	Large
0 w.p. 0.2; 1 w.p. 0.4; 2 w.p. 0.4	DU[2, 6]	DU[5, 8]

Table 3: This table shows the demand distribution by store size for each product in the second set of problems. DU stands for discrete uniform.

using variable reduction, we can verify optimal solutions virtually instantaneously.

In our second set of problems, we distinguish demand distributions only by store size. Because product demands are less differentiated than in the first set, these problems are more difficult to solve. We had to restrict the number of store-days to find solutions within reasonable CPU times. We consider 20 product types, 250 stores, and a single day, which leads to problem instances with 160 binary variables and 5,021 constraints. For each of the 250 stores, we randomly generate demands (number of stacks) for each of the 20 products according to the distributions shown in Table 3.

The results are similar to those for the first set of problems, but the CPU times are slightly longer (about 45 minutes, on average) and the heuristics no longer provide optimal solutions, although they produce surprisingly good results (Table 4). Both the solution obtained by CPLEX after 1,500 seconds of CPU time and the better of the two objective values from H1 and

Scenario	CPU time without variable reduction (seconds)	CPU time with variable reduction using the objective from CPLEX (seconds)	CPU time with variable reduction using the objective from H1 or H2 (seconds)	H1 objective/Optimal objective	H2 objective/Optimal objective
1	3,553	2,010	611	0.994	0.976
2	2,925	2,041	647	0.986	0.979
3	2,222	2,129	770	0.976	0.967
4	3,140	2,184	718	0.963	0.989
5	3,143	2,149	668	0.990	0.973
6	2,844	1,952	581	0.990	0.981
7	2,064	2,071	586	0.973	0.992
8	2,558	1,989	559	0.989	0.979
9	3,266	2,052	507	0.977	0.979
10	2,115	2,064	546	0.979	0.974

Table 4: The computational results for the second set of problems show that using the objective value from the better of the two heuristics dramatically decreases the CPU times required to find the optimal solution because the objective values from the heuristics are near optimal.

H2 lead to reduced problems with 40 variables after variable reduction. The initial solutions from CPLEX are slightly better than the heuristic solutions, but the investment of 1,500 seconds to obtain the initial solutions from CPLEX outweighs the incremental benefit.

Most organizations would like to design a mixed pallet that will work well for an extended period of time because they must make changes in the order processing (software, physical infrastructure, and operational process) to integrate standard mixed pallets with single-product pallets and store-specific mixed pallets in filling customers' orders. Consequently, real problems may have many more days of data, for example, demands for a few dozen days selected from different times of the year. For problems with 250 stores, 20 products, and dozens of days, the number of constraints and the overall data-management requirements surpass the current capability of CPLEX. Moreover, it is not clear that a retail chain, such as Safeway, would use a sophisticated integer-programming software package when it can use simple heuristics similar to H1 and H2 that (are likely to) produce good solutions.

Acknowledgments

The authors are grateful to the Safeway staff who assisted in this project, particularly Kevin Phillips, director of manufacturing services, and the employees at the San Leandro Milk Plant.

References

- Evans, David H. 1963. Modular design: A special case of nonlinear programming. *Oper. Res.* **11**(4) 637–647.
- Evans, David H. 1970. A note on "Modular design: A special case of nonlinear programming." *Oper. Res.* **18**(3) 562–564.
- Goldberg, Jeff, Ji Zhu. 1989. Module design with substitute parts and multiple vendors. *Eur. J. Oper. Res.* **41**(3) 335–346.
- Shaftel, Tim. 1971. An integer approach to modular design. *Oper. Res.* **19**(1) 130–134.
- Swaminathan, Jayashankar M., Sridhar R. Tayur. 1998. Managing broader product lines through delayed differentiation using vanilla boxes. *Management Sci.* **44**(12) S161–S172.

Kevin Phillips, Director, Manufacturing Services, Safeway Inc., 5918 Stoneridge Mall Road, Pleasanton,

California 94588, writes: "I am writing to endorse and acknowledge the paper submitted by Candace Yano and the UC Berkeley IEOR engineering group titled, "Safeway Designs Mixed-Product Pallets to Support Just-in-Time Deliveries." It is important to note that the principles outlined in the paper are as important today as when the project recommendations were presented in mid 1999.

"When the San Leandro Milk Plant began exploring manufacturing and supply chain opportunities with the IEOR project team, key priorities and concerns were identified, including:

- Completing daily order selection activities and meeting scheduled departure times had become increasingly difficult due to a narrow window of time when store orders were available to begin the load-out process.
- Plant loading crew overtime was excessive (averaging +2 hours per day) because of fewer full pallet orders, increased mixed pallets, and more frequent small orders from retail. This translated into more hand selection activity and lower utilization of automated systems and palletizing equipment.

"After critical review of store order trends and plant activities, the team's recommendations proved to be extremely effective and offered many benefits, including:

- Significant daily labor savings for the 13 person crew.
- Increased ability to meet outbound traffic deadlines.
- Higher employee morale and resulting productivity gains.
- Higher load-out capacity to accommodate peak demand periods and future volume increases.

"The insight and key deliverables from the analysis continue to add value and represent sustainable improvement for the San Leandro Plant. These contributions have helped in reducing product costs for Safeway and further our effort to achieve operational excellence."