Interfaces
Vol. 38, No. 1, January–February 2008, pp. 40–50
ISSN 0092-2102 | EISSN 1526-551X | 08 | 3801 | 0040
DOI 10.1287/inte.1070.0331
© 2008 INFORMS

THE FRANZ EDELMAN AWARD
Achievement in Operations Research

Coca-Cola Enterprises Optimizes Vehicle Routes for Efficient Product Delivery

Goos Kant
ORTEC, 2800 AL Gouda, The Netherlands and Faculty of Economics and Business Administration, Tilburg University, 5000 LE Tilburg, The Netherlands, g.kant@uvt.nl

Michael Jacks
Coca-Cola Enterprises Inc., Atlanta, Georgia, mjacks@na.cokecce.com

Corné Aantjes
ORTEC USA, Six Concourse Parkway, Atlanta, Georgia 30328, caantjes@ortec.com

In 2004 and 2005, Coca-Cola Enterprises (CCE)—the world's largest bottler and distributor of Coca-Cola products—implemented ORTEC's vehicle-routing software. Today, over 300 CCE dispatchers use this software daily to plan the routes of approximately 10,000 trucks. In addition to handling nonstandard constraints, the implementation is notable for its progressive transition from the prior business practice. CCE has realized an annual cost saving of $45 million and major improvements in customer service. This approach has been so successful that Coca-Cola has extended it beyond CCE to other Coca-Cola bottling companies and beer distributors.

Key words: transportation scheduling; vehicle routing; distribution optimization.

Coca-Cola Enterprises (CCE) is the world's largest marketer, producer, and distributor of Coca-Cola Company products. These products extend beyond traditional carbonated soft drinks to beverages, e.g., still and sparkling waters, juices, isotonics, teas, and energy, milk-based, and coffee-based drinks. CCE distributes Coca-Cola brands, e.g., Coke, Dasani, Sprite, Barq's, Fresca, Hi-C, Nestea, Powerade, and Minute Maid, and also beverage brands of several other companies. In 2005, CCE distributed two billion physical cases (containing 42 billion bottles and cans), representing 20 percent of the Coca-Cola Company's worldwide volume. While CCE is a publicly traded company, the Coca-Cola Company owns 36 percent of its stock.

Coca-Cola has outsourced its production and distribution to its bottling and distribution companies, of which CCE is the largest. CCE distributes syrup from the Coca-Cola plants to 64 bottling plants; it distributes bottled and canned beverages from the bottling plants to the distribution centers, and from the distribution centers (depots) to the final retail outlets (i.e., stores and vending machines) where customers buy the products. The operations research (OR) application we discuss plans the distribution of products from over 430 distribution centers to 2.4 million final retail outlets.

Figure 1 shows the current CCE territory in North America; it also operates in parts of Europe. CCE franchise territories encompass a population of 400 million people. This represents 80 percent of the population in the United States and Canada and all of the populations of Belgium, continental France, Great Britain, Luxembourg, Monaco, and the Netherlands. It employs approximately 74,000 people, 54,000 vehicles, and 2.4 million vending machines, beverage dispensers, and coolers.

The CCE fleet is the second largest in the United States after that of the US Postal Service. It has grown from 13,000 vehicles in 1986 to 54,000 today. Because

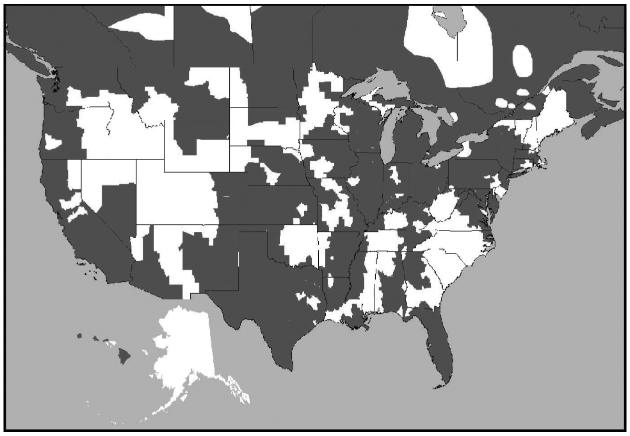

Figure 1: CCE operates in 46 states in the United States and all 10 provinces in Canada.

of this enormous growth and the competitive nature of the beverage industry, the need to optimize product delivery became apparent. Therefore, CCE's objectives were to:

- Provide world-class customer service;
- Optimize its labor and assets;
- Reduce natural resource consumption; and
- Provide its employees with a productive but realistic working day.

It has achieved these goals by implementing a route-optimization model through a cooperative arrangement with the ORTEC software company and Tilburg University. ORTEC provided the software and handled the implementation; Tilburg University developed the optimization algorithms and adjusted the model to CCE's specific needs. The project began in 2004; by the end of 2005, more than 300 CCE dispatchers were using the route-optimization model. CCE met the objectives of the project including annual cost savings of $45 million and improved customer service.

Product Delivery at CCE

CCE's challenge is the daily construction of optimal routes of orders from each distribution center (depot) to the retail outlets. In the literature, this is known as the vehicle-routing problem (VRP). CCE must deal with standard constraints, e.g., vehicle capacity and maximum route duration, and also with some nonstandard constraints that we list below:

- Requirements for a specific *vehicle type* and *equipment* per order and/or outlet (e.g., lift gate, military base, qualifications, certifications, technical equipment, or license). Because serving some locations requires a truck with specific equipment, the number of truck types has increased from 2 to 15; these truck types vary considerably in size, capacity, base location, cost structure, and available equipment. Because of these restrictions, some vehicle types are "scarce," i.e., the number of vehicle types that can service a specific delivery order may be limited. As an example, centers in large cities require the use of small truck types.
- Retail outlet *time windows* and *appointments* (i.e., tight time windows). Retail outlets continuously compress time windows, and many stores require CCE to deliver products either before they open or after they close.
- Driver's *working hours* and *start time* that must conform to the Department of Transportation (DOT) hours of service rules. Drivers are also required to carry a commercial driver's license that identifies the type of vehicle that DOT allows them to drive and includes any endorsements (e.g., permission to haul hazardous materials, drive vehicles with air brakes, or drive an articulated tractor-trailer vehicle).
- *Traffic patterns* to avoid certain areas, such as city centers during rush hours, at specific times of day.

In addition, CCE must address these "beverage-specific" constraints:

- Very specific and accurate *unloading and merchandizing times*. These vary by outlet types and vehicle types, e.g., modified sideload (a delivery vehicle consisting of multiple customer-specific pallets), normal sideload (a traditional beverage-body vehicle with individual pull-up side doors), and order fulfillment system (OFS—a type of delivery vehicle that uses a racking system to hold individual customer orders organized in delivery sequence).
- Specific *loading rules* for certain vehicles and vehicle types, such as pocket loading (used to organize customer-specific orders onto one or multiple pallets), OFS, containers for chilled products, and multiple trailers.

• Minimum *order-size requirement f*or certain vehicle types.

The overall objective is to assign all the delivery orders in the correct trip sequence, such that they are carried out by available vehicles and at a minimum cost, while respecting all constraints and providing excellent customer service. To some customers, excellent customer service means, "give me the same driver." While this is not always possible to accommodate, we try to limit the number of drivers a particular customer sees to three or four. The units of measurements of the cost function in this objective are expressed in various forms: per day, per route, per stop, per mile, per hour, per overtime hour, per cube, and per case.

CCE organized a team of stakeholders and consultants to evaluate route-planning software packages. It sought a solution that could fulfill its needs and a supplier who could act as a partner. In 2003, after a thorough selection process, it selected ORTEC and contracted with the company to implement a comprehensive route-planning system based on the following criteria: (1) proven relationship with the enterprise software company, SAP, (2) an 80 percent match with the business requirements, (3) commercially available solutions for some of the additional requirements, (4) proven development capability (as illustrated in other projects), (5) demonstrated willingness to modify its software, (6) competitive value proposition and pricing, and (7) strong relationships with universities for delivering state-of-the-art algorithms for optimization purposes.

CCE Optimization Algorithms

Since Dantzig and Ramser (1959) first studied VRP, researchers have spent much time and effort developing solutions because it plays a central role in distribution management. We refer the reader to two websites (http://www.sintef.no/static/am/opti/projects/top/vrp and http://neo.lcc.uma.es/radi-aeb/WebVRP) that provide excellent VRP overviews and include many variants and algorithms. In its simplest form, VRP entails constructing planned routes for vehicles that service retail outlets with known demand, starting from a central depot.

The objective is to assign all retail outlets to trips, including their sequence of visiting, and to assign all trips to the existing truck fleet and available drivers, such that we minimize the overall cost and satisfy all the constraints mentioned above. The cost function consists of a fixed cost per truck, per day, and a variable cost per mile, per type of truck. Each driver type has a fixed cost per hour and a fixed cost per hour of overtime. Theoretically, we can express this problem as a mixed-linear integer programming problem. However, as the literature shows, a MIP-formulation is intractable, even for a small number of locations. Therefore, the VRP algorithms in the ORTEC software (SHORTREC) are heuristic. They consist of the following:

(1) Construction algorithms to construct initial routes;

(2) Local-search improvement algorithms to improve routes;

(3) Clustering functions to optimize the visual attractiveness of a plan;

(4) Assignment of drivers to routes to optimize the driver's region-based knowledge.

Construction Algorithm

We began the construction algorithm by considering the s*avings-based construction algorithm* (Clarke and Wright 1964). The algorithm begins by considering each retail outlet as a separate route and then tries to find improvements on this solution ("savings") by adding the outlets of one route to another route without changing the order in which drivers visit the outlets. We adapted the algorithm to address the CCE-specific constraints, e.g., truck types, loading profiles, and strict time windows. Poot et al. (2002) provide details.

Although our results were promising, we were not satisfied. Therefore, we considered the *sequential insertion algorithm* (Solomon 1987). We start with a set of routes and a set of nonserved outlets. For a particular route, we then add nonserved retail outlets to the current plan by inserting them at the "best" position. When it becomes impossible to insert a retail outlet into the current route, the algorithm begins a new route. It constructs one route at a time as follows:

Step 1: *Select a vehicle*. Select the largest vehicle not assigned to a route and not tried in an earlier iteration.

Step 2: *Select a first retail outlet* (*i.e., seed retail outlet*) *for this vehicle*. Select the retail outlet that is the most difficult and which is feasible using the vehicle selected. A *difficult outlet* is one that can only be served by a very limited number of truck types, and requires a long distance to its depot and a strict time window. If no retail outlet meets these criteria, go to Step 1; otherwise assign this retail outlet to the current vehicle.

Step 3: *Add retail outlets to the new route*. Build a list of candidates to insert into the route. A candidate is a nonserved retail outlet that is feasible based on vehicle type, capacity, and region constraints, and is not too distant from the current seed retail outlet. If fixed and variable unload and merchandising time is allowed, insert the outlet at the best feasible sequence position, i.e., the position that satisfies all constraints and has the lowest additional insertion cost. This is the extra cost that is caused by the extra distance, driving, unloading, and waiting time needed to serve this retail outlet. If no retail outlets can be inserted, go to Step 4.

Step 4: *Move the route to a smaller vehicle*. When the route is full with respect to the total working time, check if the route can be assigned to a smaller (and cheaper) truck type. If the route is at capacity after a truck has been assigned, check if another route is necessary.

Local Search Improvement Algorithm

The methods for improving the routes consist of the following improvement algorithms:

• Improve the sequence of outlets in a route by attempting to move the outlets on this route. If the number of outlets is very limited, we evaluate all possibilities.

• Exchange or move outlets between two routes. We consider whether moving a set of outlets from one route to the other would help, and if exchanging outlets between routes would be an improvement.

• Exchange or move outlets between more than two routes. (This is a generalization of the exchange between two routes mentioned above.)

The following criteria apply:

(1) The set of outlets to be exchanged should be logical, e.g., the outlets in a village may comprise a logical group that should not be divided into separate routes.

(2) The truck types of the two routes should be such that exchanging the outlets is feasible, and that the capacity and other overall constraints remain satisfied.

(3) If we consider the change of driving times and mileages to these sets of exchanged outlets, then driving times and mileages should not increase by much; ideally, they should improve. An increase in driving time should be compensated by lower cost (because the truck is cheaper) or an equal decrease in waiting time.

(4) Similarly, this holds for the change from the sets of outlets. Based on total change in driving times and stop times to and from these sets of outlets, it is easy to check if the new working times are allowed.

We generalized the concepts to address multiple route improvements, but also to accept steps that are not improvements, e.g., using tabu-search techniques. Applying this technology and filtering out unpromising possibilities, the local-search improvement method becomes extremely fast. Overall, we achieved a calculation time of approximately 8 to 10 minutes of computation time on a typical run.

Clustering Function for Visual Attractiveness

CCE dispatchers prefer visually attractive routes, i.e., only a few overlapping routes (Solomon 1987), retail outlets that are clustered together on a route, and routes that are compatible with the results of the previous business practice and dispatch system. "Visual attractiveness" is important in deciding to accept the plan results even though it can increase the cost of a plan. To avoid visually nonattractive results, we added a "clustering function" into the algorithm and defined it as follows:

(1) If a route consists of n stops, then the $n/2$ stop is the "center stop," which we call C. For each stop in this route, we compute the distance (or driving time) to C. We total these distances (or times) over all stops on the route and multiply by a "clustering penalty parameter" (CP) to create a total penalty cost (Figure 2).

(2) Tuning CP is necessary to achieve the correct results. A route consisting of only one stop has a CP of zero. If we set CP too high, then the algorithm would prefer a larger number of routes and most routes would consist of only one stop. A high CP

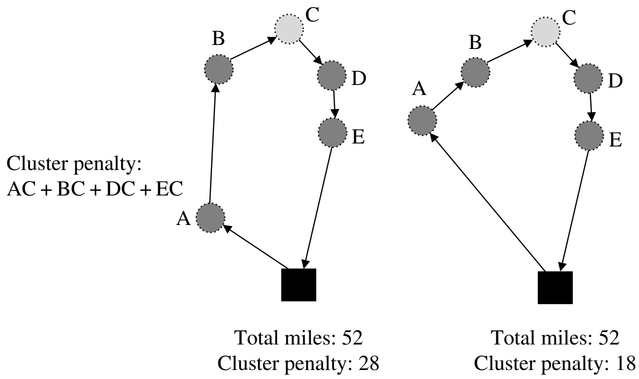

Figure 2: We incorporated this clustering function into the overall cost function and used it to evaluate the improvement algorithms.

means that the cluster cost from other stops to the center stop in a route is very high—so high that it might become cheaper to move these stops into a separate route, thereby increasing the number of routes. A route with only one stop (the center stop) has no other stops and thus no clustering cost.

(3) To mitigate the effects explained in Step 2, we added additional constraints, exceptions, and special rules for minimizing the number of routes while optimizing the costs into the algorithm.

Figure 2 shows an example of the effect of the clustering function.

Dispatchers can tune the CP parameter; they can run several scenarios with different CPs, compare the differences in mileage and working hours, and estimate the driver acceptance. In practice, the dispatchers set a high CP initially, possibly leading to more routes than necessary and nonoptimal transportation costs. Thus, the created routes are reasonably close to the routes in the prior practice. Once the dispatchers and drivers become accustomed to the new routes, we reduce the CP; as a result, we lower the costs. In this way, we accommodate change at a tolerable place, while still optimizing the costs eventually.

Assignment of Drivers to Routes

Drivers have regional knowledge and relationships with the employees at the retail outlets. Therefore, the final OR element is ensuring that, as much as possible, drivers visit the outlets with which they are familiar. Thus, we defined an *anchor point*—the center of the working area of each driver. After calculating the optimal routes, we compute the cost of each route from each anchor point (with its corresponding driver) provided the driver is qualified to drive on this route (otherwise, the cost is infinity). The cost of a route from a driver's anchor point is the sum of the "distances" (driving times) from all outlets in the route to this anchor point. We then apply a linear assignment method, i.e., assign each route to the driver who is allowed to execute this route, including the assigned truck (not all drivers are allowed to drive all vehicles) to minimize the total anchor point cost. As Figure 3 shows, the driver with anchor point B will be assigned to the solid black line route. In this way, every driver serves his or her own regions as much as possible, without incurring additional operating cost because this does not change the optimization of the routes.

When assigning the driver to a trip, we check several constraints: vehicle availability, capacity, and vehicle types (Figure 3). Tilburg University proposed these algorithmic changes and ORTEC developed them in SHORTREC. ORTEC has incorporated most of the changes in its beverage-industry solution to make them applicable to other users. It developed the interfaces specifically for CCE. Figure 4 shows an overview of SHORTREC, including the

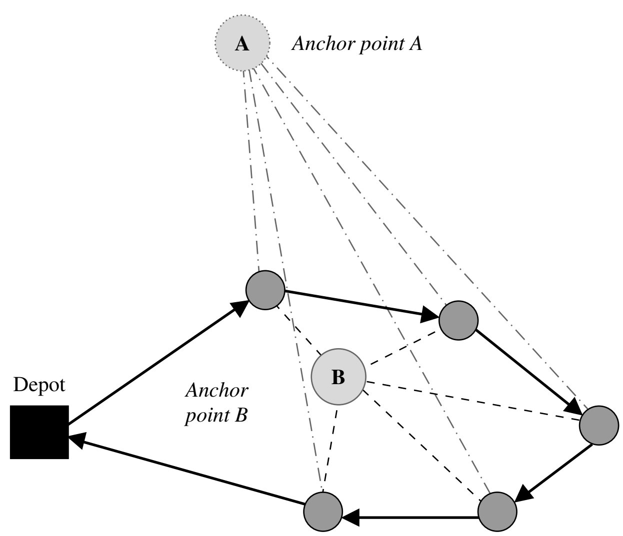

Figure 3: We assign routes to anchor points based on the sum of the dotted lines; anchor point B is preferable to anchor point A.

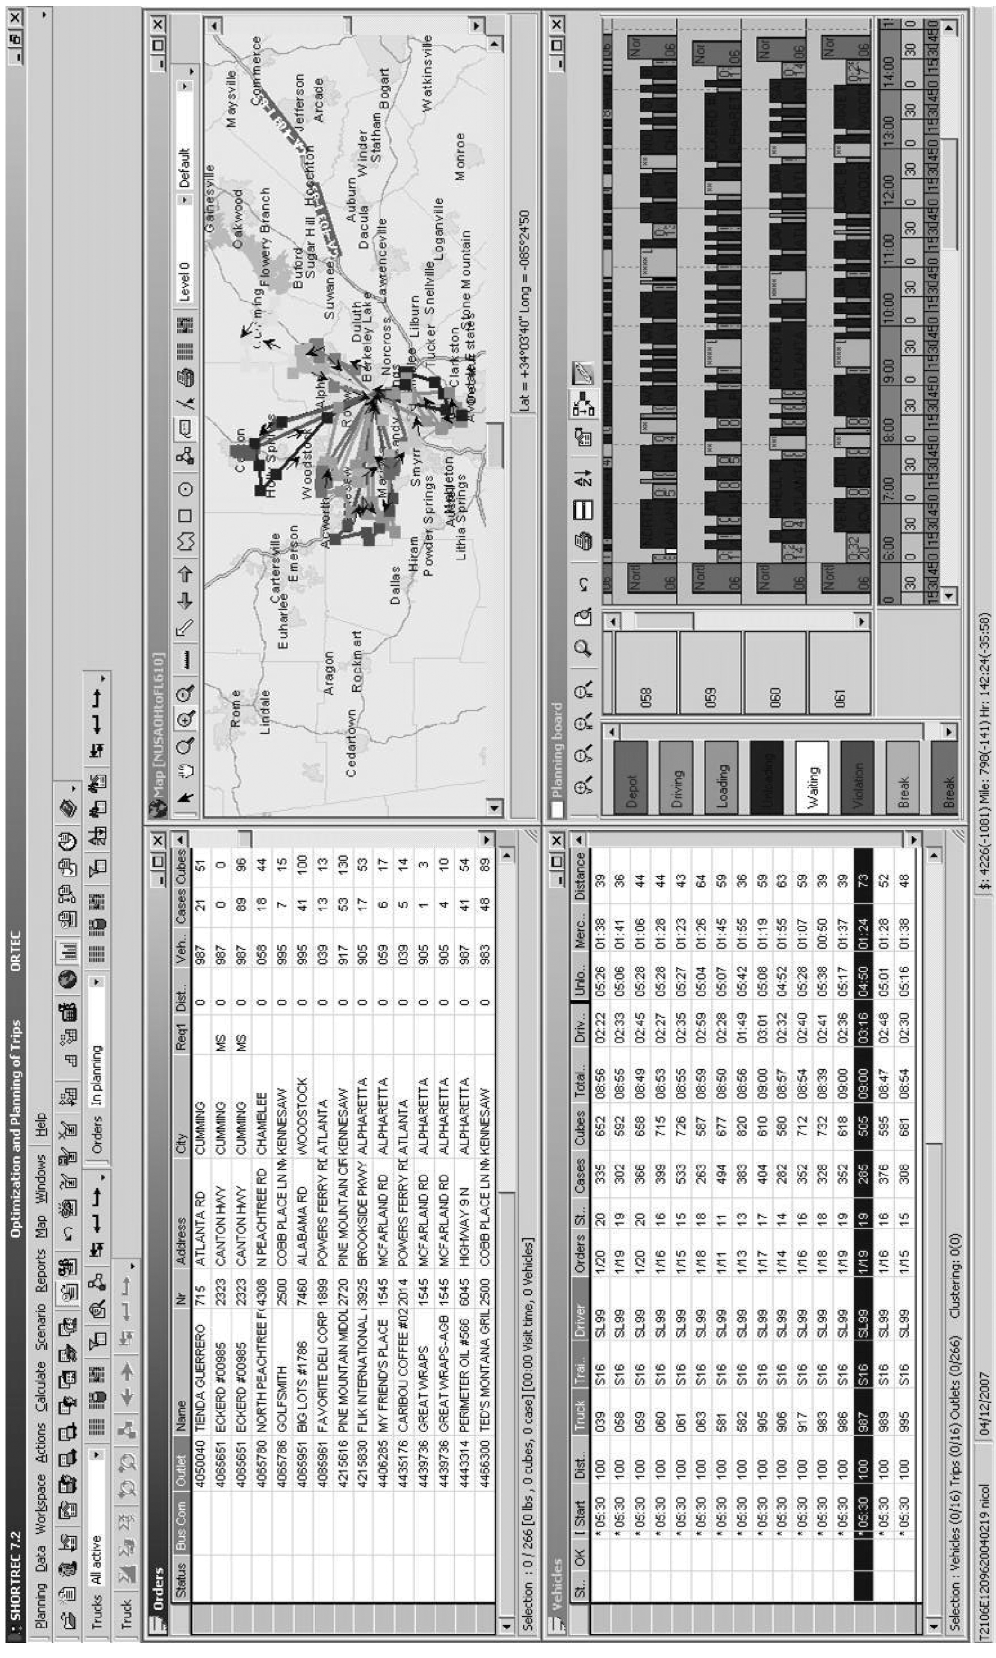

Figure 4: The screen shows a total review for the dispatcher including vehicles, orders, map, and Gantt-chart.

standard graphical user interface (GUI) and components used.

The SHORTREC-Enabled Dispatch Process

Each day, each dispatcher receives the orders that must be delivered on the following day, reviews them for accuracy and completeness, and transfers the orders into SHORTREC for route optimization. During this process, the dispatcher creates a new SHORTREC session for each depot or central warehouse. A depot may include between a dozen and 50 routes to schedule daily, and a dispatcher may be responsible for dispatching more than one depot. If so, the dispatcher repeats these steps for each depot. The route-optimization process involves assigning orders to the proper vehicle. The dispatcher typically runs several what-if scenarios to find the "best" solution. For example, the dispatcher might experiment with changes to time windows and discuss the possible monetary benefit with the customer, and might also run scenarios on driver requirements or entertain special customer requests. After examining the what-if analysis to determine if any inputs should change, the dispatcher reviews the solution from a vehicle perspective.

If any input data is erroneous, the dispatcher may not realize that there is a problem until he/she sees the output. This typically means that the relevant input must be corrected and the optimizer rerun. The dispatcher has the ability to split the screen to review all the impacted routes. This view shows the orders, the trucks used, and a map with a Gantt chart for the task associated with the routes. The dispatcher can view a single route or multiple routes simultaneously (Figure 4).

When the routes are satisfactory, the dispatcher finalizes them and initiates the process of placing beverages on pallets and onto the trucks. When the drivers arrive each morning, the trucks are loaded and ready for them; a handheld device that includes that driver's delivery schedule for the day is also available.

Implementation at CCE

CCE is organized into divisions; a division covers a European country or one or more states in the United States. In early 2004, CCE and ORTEC developed a plan to implement SHORTREC for product delivery across all divisions and dispatchers throughout the CCE organization within an 18-month time frame. It required at least 17 weeks to implement SHORTREC in a division with two or three dozen depots; the implementation plan involved four divisions in parallel. We developed a three-phased approach.

- *Phase* 1: *Division preparation.* The central team began working with each division two months before coming on-site to provide basic training and preparation guidance (e.g., data cleaning and information gathering necessary for implementing SHORTREC). This review helped to identify potential gaps or opportunities to minimize risk. For example, we would discuss the differences between the old approach of using the same static routes each day and the new approach of determining new routes each day. We would also discuss their ranking for data readiness. This ranking reflected data accuracy on the key inputs that SHORTREC needed for optimization. It also gave us a snapshot of which divisions would need more time during the pre-implementation phase.
- *Phase* 2: *On-site implementation.* The implementation at each division began with a workshop to train the dispatchers. They had the opportunity to practice what-if scenarios to see if they could come up with that were superior to those of the other dispatchers. This focused their attention on competing with each other rather than complaining about the new method and system. Furthermore, it encouraged the culture of working toward continual improvement and lower costs, while still creating practical schedules. After eight weeks, we succeeded in getting approximately half the depots to use the OR software for planning.
- *Phase* 3: *Self-implementation.* By this point, the division had enough skill and knowledge to deploy the remaining depots on their own during this third phase. We defined the following key performance indicators for reporting during this phase: actual costs, mileage, hours, overtime hours, violations, and number of routes.

While we planned to continuously implement four divisions in parallel including Europe and Canada, Figure 5 shows only the US roll out.

We made some observations based on this experience. First, maintaining the master data (outlet delivery types, delivery time windows, geography codes,

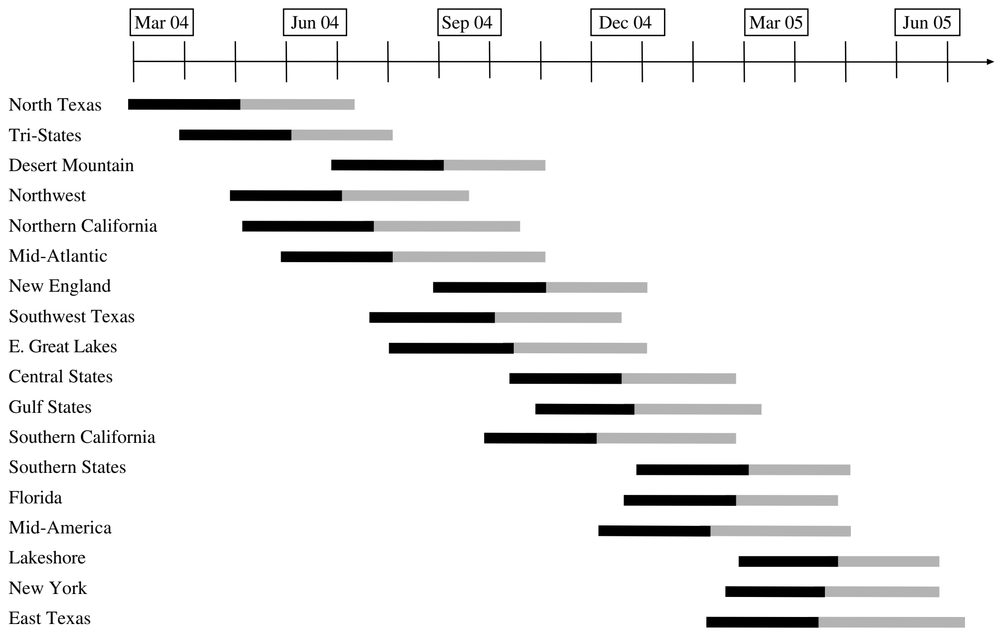

Figure 5: By the end of 2005, we had implemented SHORTREC in all the divisions that the slide shows. Some European countries have yet to deploy.

vehicle files, templates, merchandising times, etc.) properly is imperative. It is necessary to continuously monitor data accuracy and compare across divisions on a continual basis. The use of SHORTREC should enable dispatchers who have the correct attitude and skills to manage more locations, possibly across divisions. It provides them with the tools to make quick and educated decisions in less time and has enabled them to increase their responsibilities; they work more closely with sales teams to maintain master-data accuracy. In addition, we found we could train new dispatchers in less than a week.

Second, monetary savings from the implementation vary across divisions and sites. This is not surprising because it relates to several factors; these include how well the division or site did strategic route design and sales planning, how well it dispatched locations prior to using SHORTREC, how open sales centers were to accepting change, and local business practices. It is important to realize that dynamic dispatching involves change and change management.

Business Benefits

Overall, larger locations and those with growing business volumes show greater monetary savings. Because savings are, at least partially, a function of the number of drivers, vehicles, and dispatchers, it is easier to achieve savings when volume is growing rather than when it is declining. While we have many examples, we provide two as illustrations. The first example is among the 10 largest facilities in North America; it delivers approximately 20 million cases per year. Implementation began with the normal truck on August 31, 2004; it was completed with a sideload in the second week of December 2004. Its June 2005 year-to-date (YTD) business volume increased 18 percent when compared to the same six-month period in 2004; labor expenses increased by only 1 percent. Therefore, the June 2005 YTD delivery cost per case improved by 15 percent over the prior year. Although we cannot attribute all the savings solely to the implementation, there are other indicators that give evidence of its success. Missed deliveries (i.e., scheduled deliveries that have not been serviced) declined from 6.3 percent in 2004 to 2.4 percent in 2005. This reduction led to improved customer satisfaction (from the retail outlet perspective) and fewer lost sales (because more Coca-Cola products were present on the shelves). There are other indications of productivity improvement. Returns, which reflect a product that is taken to a retail outlet but not delivered (e.g., the product arrived outside of open hours, the order was inaccurate, or the driver could not find the store), dropped from 4.5 percent in 2004 to 3.7 percent in 2005. In the case of a return, the product must be delivered again and service to retail outlet management and customers is unsatisfactory. Cases delivered per hour improved from 86.3 percent in 2004 to 98.3 percent in 2005. Outlets serviced per day remained essentially flat. DOT hours-of-service rules limit the amount of time that a driver can spend actively delivering products. Optimizing delivery enabled drivers to complete more deliveries within the same time parameters and thus be more productive.

The second example is a typical mid-sized facility that delivers approximately 3.4 million cases per year using different delivery modes. This facility has been using SHORTREC for its dispatching process since December 14, 2004. Its June 2005 YTD volume increased only slightly (1 percent) over its June 2004 YTD volume. However, delivery labor costs decreased by 15 percent. This led to a 16 percent improvement in delivery cost per case. We conclude that this location is delivering approximately the same volume but using fewer resources. Favorable results are also apparent in other indicators for June 2005 YTD: missed deliveries dropped from 5.1 percent to 3.8 percent; returns dropped from 4.9 percent to 3.6 percent; cases per hour improved from 76 to 81.6, while outlets serviced per day stayed the same as in the prior year.

In two other divisions, the volume decreased slightly over YTD June 2004. However, both divisions reduced delivery labor costs (a 6 percent improvement over the prior year for North Texas and a 5 percent improvement for Central States). This led to a net savings in cost per case for both divisions (5 percent savings for North Texas and 2.6 percent for Central States). For June 2005 YTD, both divisions also made improvements in cases per hour, returns, and unproductive outlets over the prior year.

CCE Belgium also realized savings one year after the implementation. Despite inflationary pressures (e.g., a 3 percent driver wage increase from 2004 to 2005), it achieved overall yearly savings of 3 to 5 percent.

Across the CCE fleet, we estimate that we achieved a $0.03 improvement in delivery cost per case. This stems from the reduction in labor hours, fuel consumption, and vehicle usage (i.e., number of vehicles and wear and tear on existing vehicles). It represents $45 million in annual delivery cost savings, based on the 1.5 billion cases that dispatches planned using SHORTREC.

Service has also improved. Fewer missed deliveries have resulted in fewer lost sales and increased customer satisfaction (from retail outlet managers and consumers). The application of OR has led to greater delivery predictability. We have been able to realize time windows that are tighter and are based on the preferences of the retail outlet management.

We have also experienced qualitative improvements. Because our drivers drove fewer miles, we reduced pollution and our consumption of natural resources, such as fossil fuels. Fewer unplanned events reduced stress on our drivers and on the retail outlet employees. OR has enabled us to centralize CCE operations more completely. In 2004 and 2005, we installed SHORTREC at 25 divisions that had approximately 400 dispatchers. Because of the larger span of control and the increased overview that the plan provides, the dispatching process now takes place at 10 business units and requires only about 300 dispatchers.

An ORTEC press release (http://us.ortec.com/company/clients/case-coca-cola-enterprises/) reported that the SHORTREC success at CCE has been significant. In particular, the implementation methodology was important for this success. It is the basis for Goos Kant's (2006) inaugural speech at Tilburg University.

Portability

The approach was so successful that Coca Cola extended it beyond CCE to other Coca-Cola bottling companies (e.g., Coca-Cola HBC, which serves 25 countries that are mainly in Eastern Europe) and beer distributors (e.g., Carlsberg, Heineken, and Inbev). For example, we implemented SHORTREC at Inbev in France and Belgium, and realized a 100 percent return on investment within one year. For these clients, we generalized the method in two ways: first, we computed the optimal depot from which an outlet would be delivered; second, we computed the optimal frequency and days to deliver an outlet to achieve an overall balanced and clustered product delivery.

Future Development Opportunities

SHORTREC functionality includes several additional opportunities that can bring all CCE divisions to a higher level of efficiency. Examples include the following:

Multidepot Routing

SHORTREC includes a capability to optimize across multiple depots simultaneously. This means that:

(1) The algorithm can decide the depot from which to deliver to a specific retail outlet; and

(2) If a truck delivers along multiple routes, different depots could make the delivery.

While this would alter some existing sales center boundaries within the CCE organization, it would also provide more savings and a higher level of customer service. By proposing the most efficient delivery location for the outlet each day, SHORTREC make determinations about the sharing of workload and drivers across delivery locations. We should not underestimate this capability from a change management perspective; however, it would require a fundamental modification to the current information system and might impact upstream supply chain processes, such as forecasting.

Early in the software evaluation phase, we evaluated this capability to route multiple depots at once (i.e., in one session and with one responsible dispatcher), and determined that it was beneficial. The Belgium division, which has put this into practice in one of its locations, has found it to be advantageous.

Strategic Route Design

The delivery of products at CCE (and all other distribution companies of Coca-Cola Company products) takes place after visits by sales representatives (i.e., merchandising). These sales visits are intended to increase distribution-volume turnover (revenues); these volumes must also be distributed to the retail outlet. Thus, when a sales representative visits a retail outlet on Monday, the distribution volume will be delivered on, for example, the following Wednesday. Depending on the expected distribution volumes, retail outlets require one, two, or four sales visits (i.e., referred to as frequency 1, 2, or 4) in four weeks. A retail outlet with a high expected-distribution volume is classified as an important outlet. Therefore, the sales representative visits this outlet more often.

Retail outlets have different time windows and frequencies and their expected distribution volumes vary. These large differences in distribution volumes make it difficult to balance the volume per day. Delivery does not take place on weekends. We must also take into account that retail outlets are always served on the same day of the week, and retail outlets with frequency 2 are visited either in weeks 1 and 3 or in weeks 2 and 4. The sales canal is leading, i.e., we create routes only for the sales visits. However, during the planning and optimization process, we consider both the number of sales visits per day and the expected distribution volumes two days later. We must balance both measures as much as possible.

Hoendervoogt (2006) developed an algorithm for this based on the approach of Tan and Beasley (1984). Based on the "center of working areas for each day," the algorithm computes the cost for each delivery schedule (possible delivery days + volume) for each retail outlet. In the next step, it selects this delivery schedule for each retail outlet such that, for each day, the assigned volume and total working time is equal, while minimizing the total costs (i.e., the deviation from and to each center of working area per day). The results are promising and CCE has already begun pilots.

In addition, when the sales order for a retail outlet is booked, an optimal loading algorithm computes the required load (in volume, space, etc.) by optimizing the load on a pallet; as a next step, it optimizes the load of a trailer or a container by considering the length, width, and height of a pallet, a truck, and of each of the product types, as well as loading restrictions (i.e., overhang, stackability, and possible loading configurations). Because it is computing the required space based on the order in the sales and merchandising system, it can compute the optimal number of pallets at that moment, thus optimizing the truck load. If, for example, the calculated number of pallets is 1.8, then the next step is to propose to sell a default item to fill it up to two pallets: these pallets would be filled up with a (fast-moving) default item (or items). Therefore, both truck utilization and pallet utilization can be improved. Experiments at other Coca-Cola bottling companies are showing promising results.

References

Clarke, G., J. W. Wright. 1964. Scheduling of vehicles from a central depot to a number of delivery points. *Oper. Res.* **12** 568–581.

Dantzig, G. B., J. H. Ramser. 1959. The truck dispatching problem. *Management Sci.* **6** 80–91.

Hoendervoogt, A. 2006. Period scheduler: An algorithm for strategic planning. Master's thesis, Tilburg University, Tilburg, The Netherlands.

Kant, G. 2006. Ruim Baan! OR en IT in Transport en Logistiek. Speech, Tilburg University, Tilburg, The Netherlands.

Poot, A., G. Kant, A. P. M. Wagelmans. 2002. A savings based method for real-life vehicle routing problems. *J. Oper. Res. Soc.* **53** 57–68.

Solomon, M. M. 1987. Algorithms for the vehicle routing and scheduling problems with time window constraints. *Oper. Res.* **35** 254–265.

Tan, C. C. R., J. E. Beasley. 1984. A heuristic algorithm for the period vehicle routing problem. *Omega* **12** 497–504.