

Routing Optimization for Waste Management

Surya Sahoo, Seongbae Kim, Byung-In Kim

Institute of Information Technology, Inc., 2204 Timberloch, Suite 225, The Woodlands, Texas 77380
{surya@e-iit.com, skim@e-iit.com, bkim@e-iit.com}

Bob Kraas, Alexander Popov Jr.

Waste Management, Inc., 1001 Fannin Street, Houston, Texas 77002
{bkraas@wm.com, apopov@wm.com}

Waste Management (WM) obtains one third of its revenue from landfill disposals and two-thirds from waste-collection services. As most of the revenue comes from collecting trash, improving efficiency in operating the fleet improves the bottom line. After a flurry of acquisitions and a merger with USA Waste, WM found itself with a large fleet of vehicles whose routing, dispatching, maintenance, and management were decentralized. WM recognized that it could reduce operating costs by improving its use of assets. It contracted with the Institute of Information Technology to develop WasteRoute, a comprehensive route-management system that took into account WM's specific routing concerns and provided broad benefits. Initially, the target audience of the system was the dispatchers and indirectly the drivers. Sales and customer service also benefited because WasteRoute integrated the sales, customer service, and operations departments. The system reduced operating costs, provided better customer service, and determined appropriate prices. WM deployed WasteRoute across the nation beginning in March 2003. By the end of 2003, WM had 984 fewer routes, saving \$18 million. It estimated that its savings for 2004 due to the reduction will be \$44 million. As it extends the system to additional areas, it expects additional route reductions.

Key words: industries: transportation, shipping; transportation: vehicle routing.

Most US residents consider garbage collection a required service. As a result, there are many waste collection companies in the United States that compete based on price. Waste Management, Inc. (WM) is unusual in focusing on improving customer service to distinguish itself in the market.

WM made an enterprise-level investment in improving service to address several operational inefficiencies. First, after a string of acquisitions in some areas of competing companies, WM had overlapping routes with two trucks serving customers on the same street. It was obvious it needed to improve routing in these areas. Second, route planners or drivers sequenced the stops on the routes. When they had enough experience, this method worked fairly well, but it led to a host of problems for Sales and Customer Service. For example, a salesman wanting to offer service to a prospect might have difficulty determining the optimal route for the customer to ensure the best service at the least cost. Drivers were not required to inform the route planners if they changed

the sequence of stops, which made serving their customers nearly impossible when they were out sick or on vacation. Finally, to reduce information technology (IT) costs, WM sought a Web-based platform to handle routing.

Waste Management

With headquarters in Houston, Texas, WM is the leading provider of comprehensive waste-management services in North America. Its network of operations includes 293 active landfill disposal sites, 16 waste-to-energy plants, 72 landfill gas-to-energy facilities, 146 recycling plants, 346 transfer stations, and 435 collection operations (depots). Combined, these resources enable WM to offer a full range of environmental services to nearly 20 million residential customers and 2 million commercial customers throughout the US and Canada.

WM provides solid-waste collection services for residential, industrial, and commercial customers in

48 states, the District of Columbia, Canada, and Puerto Rico. With nearly 26,000 collection and transfer vehicles, it operates the largest trucking fleet in the waste industry and collects over 80 million tons of garbage a year. The company provides a wide range of services, from picking up household trash at single-subscription residences to providing comprehensive waste programs for large national companies with hundreds of locations. For these companies, WM's national accounts department develops customized programs, sometimes sending a representative to work on site to help them manage specialized and diverse environmental needs.

WM has improved its routing, maintenance, and utilization of labor and equipment in its 435 collection operations across the US and Canada. One of the most fundamental changes was the restructuring of the organization around geographic market areas. Before, WM operated as a collection of 1,200 separate business entities; now, these business units have been consolidated into 66 market areas. This simpler, more streamlined structure aligned the organization with the business strategy. Integrating all the assets in a market area (services, people, resources, and equipment) enables WM to manage business more effectively and at a lower cost. As a result, WM now has a market-based, customer-driven structure resulting in a stronger enterprise and a more competitive force in the market place.

Based on 2002 gross revenues, WM's waste-collection business contributed 68 percent of its total revenue. Within the collection business, commercial collection, residential collection, and industrial collection contributed 38, 32, and 29 percent to gross revenues, respectively (Waste Management, Inc. 2003). WM's achievements in IT were recognized by *CIO* magazine, which awarded it the CIO100 Award for 2001, 2002, and 2003.

The Problem of Routes

In changing from a decentralized to a centralized organization, WM recognized that managing the activities of its 19,600 daily routes was not trivial. WM typically operates vehicles six days a week, with a few running on Sundays. With an annual operating cost of nearly \$120,000 per vehicle, WM wanted to make every route as profitable and efficient as possible.

WM managers intended to reduce the overall operating expenses, whose key components are fixed vehicle costs, variable vehicle costs, and labor expenses. They decided that the best way to evaluate the results of a program to cut costs was to measure the reduction in vehicles. Eliminating one vehicle may eliminate five or even six route days and certainly contributes to reducing assets.

The collection business is in three major areas: commercial, residential, and industrial. Although they are very different, they all produce municipal solid waste and materials for recycling. Commercial customers include strip malls, restaurants, and small office buildings. A commercial route may include 60 to 400 customers and two or three disposal trips to a landfill each day. Depending upon the customer base, a driver may visit the same customer multiple times in one week. The weekly service schedule is fairly static, as most customers seldom change the frequency of service. Commercial routes usually contain about 20 percent volatility, or unplanned activity, throughout the course of the day. Blocked dumpsters or containers, extra services, and new customer stops alter planned routes. The difficulty in managing the efficiency and productivity of drivers lies in the details. On one route, a driver may pick up garbage for only 60 customers a day, while on another route in the same city, the driver may serve 300 customers in one day. However, the first driver may drive nearly 400 miles while the second drives only 150 miles. Measuring and comparing driver productivity was nearly impossible with WM's existing operational systems.

Residential customers are generally people living in private homes. In subscription service areas, where customers call WM individually for service, routes tend to be less dense and costs higher per customer than in contract areas where WM bids for and obtains the right to serve all homes in an area. In subscription areas, sales and operations staffs work to increase customer density to reduce cost per customer. Municipalities and homeowners' associations commonly contract with WM to limit the number of waste vehicles driving through the neighborhoods, thereby increasing the overall safety and appeal of the community. WM had difficulty measuring and comparing productivity within residential routes. Depending on




Figure 1: Because vehicles are permitted to travel and make stops only on the right sides of streets, residential collection requires arc-routing systems.

density, a residential route may serve between 150 and 1,300 homes a day. Compensation plans, geography, unions, and local experience levels all contribute to the productivity of a route. The frequency of service per week will vary based on the climate, geography, competition, and price of service. In northern states, customers commonly expect service once per week. In southern states, they typically expect service twice per week. Once WM determines the weekly frequency for a set of routes, it repeats this schedule every week.

The single largest factor distinguishing residential routes from commercial routes is the mandatory adherence to driving on one side of the street. Unlike drivers on commercial routes, those on residential routes are permitted to serve only customers on the right side of the street (Figure 1). Very few exceptions are granted for alleys and one-way streets. This means that solving the residential routing problem is different from solving the commercial routing problem. Point-to-point solutions are acceptable for commercial routes but residential routes require arc-routing solutions.

Industrial routes introduce a different routing problem. The differentiator between commercial routes and industrial routes is the size of the containers (dumpsters). A typical commercial container is eight loose yards, while an industrial container may range from 20 to 40 loose yards. Trucks can handle only one industrial container at a time and they commonly haul these large containers to the landfill, empty them, and have them back to the original customers' locations. To complicate the problem, WM uses many different approaches to improve the efficiency of this operation. For example, a truck may first deliver an empty container at the customer location, and then pick up the full container, travel to a disposal facility

and dispose of the contents. It might serve another customer with the same size container. The difficulty arises when a driver is scheduled to perform different types of services throughout the day for customers with different container sizes and different service requirements. WM must consider more factors, including driver-experience level, vehicle types, container types, material types, and security clearance, when creating industrial routes.

The Residential and Commercial Routing Problem Characteristics

Through WasteRoute, we mainly solved the residential routing problem and the commercial routing problem. Both residential and commercial collection problems can be classified as variations of vehicle-routing problems with time windows (VRPTW) but with additional constraints. In the literature, a typical vehicle routing problem (VRP) comprises a set of vehicles, stops, and a depot. Each vehicle starts from the depot, visits a number of stops, and ends at the depot. Depending on the nature of the application, VRPs may have different characteristics, including types of vehicles (homogeneous or heterogeneous) and number of depots (single or multiple). A VRPTW is defined as a VRP extended by additional time constraints associated with each stop. The depot also has a time window. Each vehicle has a single capacity constraint, such as maximum volume or maximum travel time. The objective function(s) generally minimizes total costs or total travel time. The VRPTW is NP-hard, and finding a feasible solution with a fixed fleet size is an NP-complete problem (Cordeau et al. 2002), and one frequently resorts to approximate or heuristic procedures in solving the problem.

As with a typical VRPTW, at WM, each operation site has a single depot per collection operation, a finite number of homogeneous vehicles, and a set of stops. In addition to those standard components of VRPTW, the routing problems at WM have landfill facilities, the feature unique to the waste-collection business. The constraints specific to landfills are a major component of the routing problem for waste-collection companies. When a vehicle is full, it needs to go to the closest available disposal facility. Each vehicle can and typically does make multiple disposal trips per day.

WM considers two main capacity constraints when creating a route: vehicle capacity and route capacity. Vehicle capacity is the maximum volume and weight that each vehicle can hold. Route capacity is the daily capacity for each driver: maximum number of stops, maximum number of lifts, maximum volume and weight that a driver can handle per day, and so forth. WM business rules set the route capacity constraints and route time. Vehicle capacity dictates when drivers should make disposal trips. If the vehicle capacity is larger than the route's volume capacity, drivers will make only one disposal trip and it will be the last stop before returning to the depot. If, however, the route's volume capacity is more than the vehicle capacity, the driver will make multiple disposal trips. Because multiple disposal facilities are available, we must carefully select the best. We assume that each vehicle starts at a depot and ends up at the depot with zero volume. Each vehicle has a one-hour lunch break between 11:00 AM and 1:00 PM. WM's typical problems range between 25,000 and 120,000 residential homes or 10,000 and 54,000 commercial stops per collection operation.

While authors of papers on the VRPTW consider minimizing the number of vehicles and total travel time to be the major objective, WM also considers the visual attractiveness or route compactness of solutions very important. We borrowed the term "visual attractiveness" from Poot et al. (2002); it depends on how stops are grouped into routes. A solution in which many routes cross over each other is less visually attractive than one in which no routes overlap. Balancing work among vehicles is also important in implementing a solution in the field. We summarize our problem as follows:

Objectives

- Minimize the number of vehicles,
- Minimize travel time,
- Maximize visual attractiveness,
- Balance workload among the vehicles.

Constraints

- Time windows of stops and the depot,
- Vehicle capacity (volume, weight),
- Route capacity (maximum number of homes (residential) or lifts (commercial), volume, and weight a vehicle can handle per day),
- Routing time limit per vehicle,

- Disposal trips (when a vehicle is full, it must go to a disposal facility),

- Driver's lunch break.

Before WM adopted WasteRoute, local route planners designed routes manually and distributed them to the drivers. Although some route planners had very limited visualization tools, the efficiency of the routes they planned depended greatly upon their experience and knowledge of the area, as well as the timely availability of the information they needed to plan routes to accommodate construction areas, weather, vehicle availability, and driver availability. Route planners had difficulty constantly taking all of this information into account when planning routes to ensure efficiency. To address disruptions on routes, drivers or dispatchers simply changed routes without reoptimizing the routes, resulting in routes crossing one another, less-than-efficient results, or poor customer service.

WM Selected IIT as a Solution Provider

WM organized a team of stakeholders and consultants to evaluate route-management software packages. WM was looking for an application that could be integrated into its existing IT infrastructure, preferably a Web-based solution.

WM decided that the best way to evaluate the initial list of 19 vendors was to put them head-to-head in a competitive situation. It held a competition in Houston during August 2001. It made its final evaluation comparing the quality of the routing algorithm engines in February 2002. WM evaluated and compared the algorithms' functional fit with its business rules. It evaluated the results of the algorithm tests based on meeting three objectives: (1) route reduction (cost savings), (2) workload balancing (number of routes) across days of the week, and (3) adherence to business constraints. The solution provided by the Institute of Information Technology, Inc. (IIT) met these requirements. IIT's solution also took into consideration turn restrictions or penalties, speed limits, and adherence to side of street for residential routes. In addition, given a four-hour processing-time restriction, the IIT software solved WM's problem in 35 minutes.

WM contracted with IIT to develop WasteRoute, a comprehensive route-management system that took

into account WM's specific routing concerns and provided broad benefits. In March 2003, three months after a team of people from both organizations began initial development, WM rolled out the first version of WasteRoute.

GIS and Optimization

A geographical information system (GIS) tightly integrates spatial information to managed data. To illustrate the contrast, consider a customer in a standard relational database. In a standard relational database, the attributes for a customer would typically include name, billing information, and address. The firm typically enters the address information manually in a system for managing customer relationships and uses it for mailings. With such a system, the firm could not determine customers' proximity to other customers or routes to support decisions about routing.

Based on a customer's address, a GIS would provide spatial information about the customer, including the customer's specific x , y , and z location in a coordinate system and the geometry of the customer's pickup site, usually a point but sometimes a line or a polygon. When it brings this spatial information together with other information having spatial components, such as streets, it can display truly unique aspects of the data. For example, it can display customers' locations on a map with the surrounding street network, various landmarks, other customers, or the collection facility.

Although planners can use the GIS component of WasteRoute to view customer locations on a map and manipulate information, they take advantage of the true power of the GIS when they construct a key component of the optimization engine, an origin-destination (OD) matrix. This matrix captures the distance and time traversed on a street network between any two points, taking into account such constraints as speed limits, directional attributes of streets (such as one-way directions or turn penalties and restrictions), and accurate segment distances.

The WasteRoute System

WasteRoute is a Web-based Java application, integrated with WM's other systems (Figure 2). The user interface is managed by applets, is launched from a

Web page, and runs within the Java virtual machine found in nearly all Web browsers. The applet consists of both tabular data (in forms and reports) and a map displaying customer locations, landfills, and other facilities against a background of streets, landmarks, and other geographical boundaries. The coordination of the navigation tree, customer data, and an interactive map makes the application simple to use. Ease of use is very important to WM, because local route managers, whose computer skills vary widely, use the application.

A given customer can have many stops on many days, each of which represents a location where the vehicle stops. A stop may have one or more orders, each of which represents a trash container or a residence. This representation models the business to support customer service at the container level but intelligently collapses orders from the same customer into a single stop as appropriate. For example, the optimization engine can view three containers at a strip mall, even though they are at different coordinates in the real world along a driveway, as one stop for routing purposes. The model is also sophisticated enough to account for the time taken to handle additional containers at the same location with configuration parameters. This database is automatically synchronized with the existing AS/400 infrastructure.

Deployment

From the beginning, WM understood that deploying an application like WasteRoute would fundamentally change the way it served its customers. Getting people to change for any reason is difficult, and failing to orchestrate the change risks failure. A team of two project managers (IT and business) and a group of four organizational-change-management personnel developed a plan to roll out WasteRoute to the organization. They designed the initial deployment plan with two tiers: tier one has 36 market areas consisting of larger collection operations, and tier two, 30 market areas consisting of smaller, more rural collection operations. They set the completion time line for tier one at the end of 2003 and for tier two at the end of 2004. They further decided to break tier one into three waves of 12 market areas for systematic rollout of WasteRoute. We conducted the first classroom training class in March 2003 for wave-one participants.

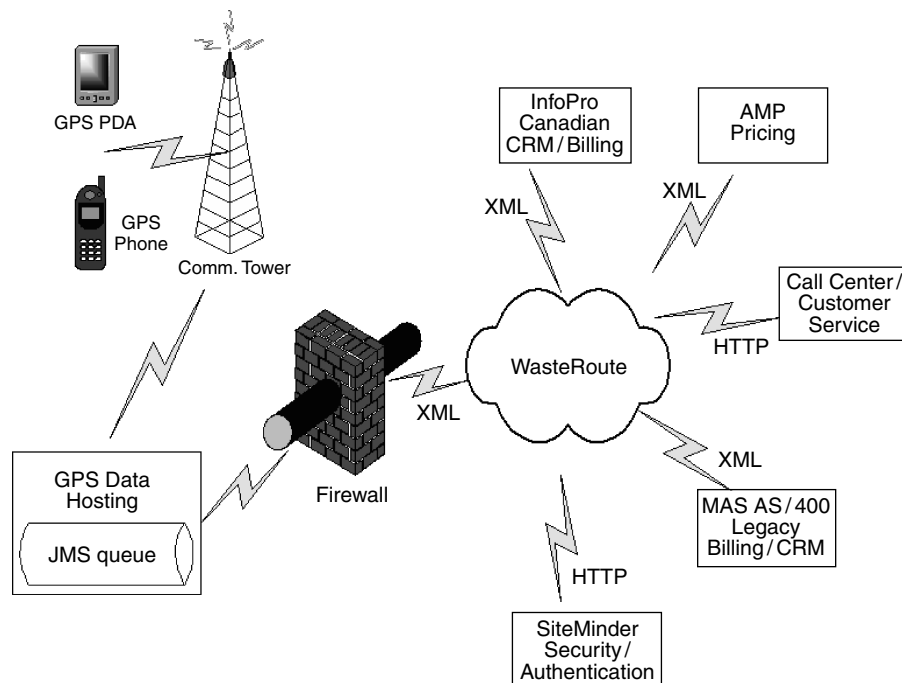



Figure 2: WasteRoute is a Java application which is integrated with various other enterprise systems at WM through Web services/HTTP protocol. The various applications provide data to WasteRoute and consume the services it generates. Customer and billing data for sites in the US reside in an AS400 system called MAS and for sites in Canada, reside in InfoPro. Contract residential customer data collected through GPS systems is also integrated. WasteRoute uses SiteMinder as a directory server for user authentication and authorization in accordance with WM's corporate standards. WasteRoute is also integrated with other systems, such as analysis and modeling for profitability (AMP) and customer-service applications.

To help us with training and expand our knowledge base, we had some of the route planners from the wave-two operation sites participate in the deployment of wave-one operations. Our objective was to train many people simultaneously while they gained on-the-job experience. This procedure worked very well. The training class consisted of two components: standard operating procedures and application training. After taking the initial class, the planners also took a series of self-paced Web-based courses that covered each of the two topics individually. While the exact savings in the various areas vary, Elgin, Illinois, for example, went from 10 routes to nine routes using WasteRoute. It improved the average productivity of the routes from 57.06 cubic yards of garbage per hour to 63.41 cubic yards per hour (Figure 3).


When companies adopt software, such as WasteRoute, the employees who must use it tend to be defensive and resist changing the way they do

their jobs. To maintain morale and build trust between employees and managers, WM drew up a proactive communication plan and included the employees at every step during the rollout. WM reduced the number of routes with the full support of the drivers because it did not lay off any drivers; it reduced staff through normal attrition. In Baltimore, Fleet Optimization Manager Renee Huff says, "Driver involvement has been a crucial success factor in implementing WasteRoute at our site." Huff projected the new routes that WasteRoute created on a screen and received valuable feedback from drivers on route density and capacity limitations. Giving the drivers ownership made them willing to adopt the new system and help to improve their routes.

We broke the deployment for each operation site into four steps: service-boundary analysis (SBA), weekly balancing, clustering stops, and optimizing routes. Market area managers and route planners use



(a) Before



(b) After

Figure 3: In the Elgin, Illinois area, WM reduced routes by one, a reduction of 10 percent. Before WasteRoute, that area required 10 nine-hour routes, and its productivity was 57.06 yards per hour. With the solution of WasteRoute, that area now requires nine nine-hour routes, and its productivity is 63.40 yards per hour. Each polygon shows an individual route.

WasteRoute extensively during these steps. During SBA, area staff could remove any overlaps between neighboring WM operations. In conducting weekly balancing, route planners try to move customers to different days to equalize workloads on the five days of the week. This step includes a review of customers' container sizes and frequencies. For a customer with a small container and service three or more times a week, WM can substitute a larger container and reduce the frequency of service, thus reducing WM's cost of serving that customer. This step can add time

to the deployment because such changes require personal phone calls by WM representatives to confirm them. In the clustering step, users look at the geographies of the routes without creating individual route sequences. During the optimization step, route planners sequence the daily routes using WasteRoute. After WasteRoute creates the routes, drivers and managers review them. Once the drivers approved of the new routes, the managers set a go-live date for the operation site, and WM notified the customers as necessary by phone, mail, or door hangers (notices hung on door knobs).

Impact and Business Benefits

After four years of emphasizing growth through acquisitions, WM needed to optimize its routes. WM strives for operational excellence. Corporate leaders identified WasteRoute as one of the building blocks it needed to achieve operational excellence. To sell this program to stakeholders and key personnel, we designated four key objectives: profitability, customer service, route optimization, and safety.

Profitability

In February 2003, CEO Maurice Myers remarked, "Over the past year, in pilots at 27 different hauling locations, we've reduced drivers and trucks by up to 17 percent and on average 10 percent. The company operates 15,000 commercial and residential routes. Our goal is to achieve the average pilot reduction of 10 percent. Each route is estimated to cost \$120,000 annually, so our potential cost savings are meaningful to our bottom line."

With a graduating deployment schedule, financial benefits increased on a nonlinear scale. WM began to realize benefits in the second quarter of 2003. Deployment to 36 market areas in 2003 realized savings at approximately \$18 million and costs at roughly \$10 million. WM estimated that its saving for the year of 2004 due to the route reduction of 2003 will be \$44 million. WM anticipates the savings to come in the long term. As each market area's employees become more familiar with the software and the software becomes institutionalized across WM, it will yield additional savings. The firm expects to make additional route reductions and thus cost savings as it extends WasteRoute to 30 market areas throughout

North America in 2004. WM expects its cash flow to increase by \$648 million over a five-year period. Over the same period, it estimates its savings in operational expenses will be \$498 million. Early in the deployment, WM realized that WasteRoute had impacts beyond saving costs. As routes increased in capacity, they could increase in revenue. Managers able to understand daily vehicle activities and sales could focus on targeted areas to increase route density.

Customer Service

As WM stabilized its routes, the dispatchers, customer representatives, and drivers communicated more effectively and can now produce more consistent and reliable service. Sales agents able to see the location of facilities or the area covered by a route can judge which customers to target. They can also close deals by offering those customers appropriate pricing options.

Route Optimization

Tom Derieg, vice present of fleet services and logistics and executive sponsor of WasteRoute, states, "This past year was one of patience and persistence for our team members and the many employees around the company who helped us analyze the efficiency of our routes. I am very pleased to report to you that we exceeded the goal Maury Myers set out for us in December 2002. We were asked to reduce 750 routes, and our 2003 cumulative accomplishment was a reduction of 984 routes (goals exceeded quarter over quarter in 2003). It is tough to admit that we had an efficiency issue, but based on our results, I am very proud to see how much this work has strengthened the company already." The most significant improvements were made by the Oregon market area, which reduced 46 routes against a goal of reducing 12, and their optimization efforts continue in 2004.

In a short span of four years, WM has transformed itself from a fragmented organization headed in many directions to a centralized corporation with a unified goal of operational excellence. The success of WasteRoute has brought WM close to realizing the benefits of its size and scale. This Fortune 100 company is now delivering on its promise of operational excellence and differentiated service as an example to the waste industry. Applying optimization to reduce

routes has made WM receptive to increasing the role for operations research techniques in other operational areas, such as locating facilities and delineating service boundaries.

Safety

Reducing routes continues to improve WM's operational efficiency and cost savings, and it has a positive impact on the environment and employees. Fewer trucks on the road noticeably reduce emissions and noise. Reducing travel during busy times of the day and traffic in the communities in which WM is a member are also noticeable benefits.

"Safety and WasteRoute are inextricably linked," explained Jim Schultz, vice president of safety. "The WasteRoute team recognized from the outset their responsibility to employees, customers, and people in the communities we serve to ensure safety is the foundation," he added. To do this, WM has made operational safety a cornerstone element in the WasteRoute implementation.

Conclusion

In 1999, Waste Management acquired businesses at the rate of one per day, so it was imperative that it searched for a way to manage this growth simply to ensure survival. During that same year, WM set the goal of achieving some business benefits by optimizing routes, and, with the help of the Institute of Information Technology, it realized this goal using WasteRoute. This GIS-based route-management application reduced WM's operational costs by (1) organizing routes to minimize overlap and thereby reduce the number of vehicles WM needed to serve its customers, and (2) sequencing the stops along a route to make the best use of fuel, driver schedules, and disposal trips. Integrating WM's existing systems with WasteRoute allows customer-service personnel to resolve service issues quickly because they have access to information about the customers. Marketing and sales personnel also use WasteRoute to offer appropriate pricing to prospects and to select areas to serve based on existing routes and facility locations.

The financial results from the initial deployment of WasteRoute are very positive. With the rollout of WasteRoute in 2003 and 2004, WM expects to reduce the number of collection routes by 10 percent. WM began

deploying WasteRoute in March 2003 with a net effect of 984 fewer routes at the end of the year, exceeding its goal of dropping 750 routes, and achieving savings of \$18 million. WM estimated that its savings for the year of 2004 due to the route reduction of 2003 will be \$44 million. The firm expects to make additional route reductions and thus cost savings as it extends WasteRoute throughout North America in 2004. WM expects its cash flow to increase by \$648 million over a five-year period. Over the same period, it estimates its savings in operational expenses will be \$498 million.

Appendix

The Optimization Problem

To explain the problem clearly, we present a mathematical programming model for a simplified version of our problem here. We simplified the problem by minimizing only travel time. In the real-world problem, we also try to minimize the number of vehicles, maximize the visual attractiveness, and balance the workload. This simplified model also considers single-vehicle capacity alone. In the simplified problem, we assume that the number of vehicles is given and consider a single objective function: minimize the travel time. We adopted the notation and basic VRPTW model of Cordeau et al. (2002) and changed them to incorporate the disposal operations and the drivers' lunch breaks.

The simplified waste VRPTW is defined on the network $G = (V, A)$, where the regular stops are represented by nodes 1 to N , the landfills are represented by nodes $N + 1$ to $N + M$, the lunch break is represented by $N + M + 1$, and the depot is represented by 0 in the vertex set $V = \{0, 1, 2, \dots, N, N + 1, \dots, N + M + 1\}$, and A is the arc set. For modeling purposes, the landfills are duplicated such that a vehicle visits a landfill at most once in a route. For example, if a vehicle visits a specific landfill three times in a route, we duplicate the landfill three times so that the vehicle virtually visits a landfill at most once. The maximum number of disposal trips for a vehicle can determine the number of duplications in a day because, in extreme cases, the vehicle visits the same landfill every time. Each node in V has a time window $[a_i, b_i]$ representing the earliest possible

service starting time and the latest possible starting time. The depot's time window $[E, L]$ has the special meaning of the earliest possible departure from the depot and the latest possible arrival at the depot. Each node has a service time s_i and demand d_i . The lunch break, the depot, and the landfills have zero demand. The depot has zero service time. A nonnegative travel time, t_{ij} , is associated with each arc $(i, j) \in A$. Travel time between the lunch break stop, $N + M + 1$, and any stop is set to zero.

Let K be the available vehicle set, N_k be the actual number of disposal trips for vehicle k , and C be the vehicle capacity. Let $x_{i,j,k}$, where $(i, j) \in A, k \in K$, equal 1 if arc (i, j) is used by vehicle k and 0 otherwise. Let time variable $w_{i,k}$, where $i \in V, k \in K$, be the start service time at node i when it is serviced by vehicle k . Let also $D_{i,k}$, where $i \in V, k \in K$, be the cumulative demand at node i for vehicle k . A route represents the travel path of a vehicle, and a subroute represents a sub-travel path that starts from the depot or a disposal facility and ends at a disposal location or the depot.

Simplified Waste-VRPTW Model

$$\min \sum_{k \in K} \sum_{(i,j) \in A} t_{i,j} x_{i,j,k}$$

subject to

$$\sum_{k \in K} \sum_{(i,j) \in A} x_{i,j,k} = 1 \quad \forall i \in 1, 2, \dots, N, \quad (1)$$

$$\sum_{j=1}^N x_{0,j,k} = 1 \quad \forall k \in K, \quad (2)$$

$$D_{i,k} \leq C \quad \forall k \in K, i \in 1, 2, \dots, N, N + M + 1, \quad (3)$$

$$D_{m,k} = 0 \quad \forall k \in K, m \in 0, N + 1, N + 2, \dots, N + M, \quad (4)$$

$$D_{i,k} + d_j - D_{j,k} \leq (1 - x_{i,j,k}) \text{Big}M \quad \forall k \in K, \\ i \in 0, 1, 2, \dots, N + M + 1, j \in 0, 1, 2, \dots, N, N + M + 1, \quad (5)$$

$$\sum_{j=1}^N d_j \sum_{i=1}^{N+M+1} x_{i,j,k} \leq C^* N_k \quad \forall k \in K, \quad (6)$$

$$\sum_{i \in 1, \dots, N, N + M + 1} \sum_{m=N+1}^{N+M} x_{i,m,k} = N_k \quad \forall k \in K, \quad (7)$$

$$\sum_{m=N+1}^{N+M} \sum_{j \in 1, \dots, N, N + M + 1} x_{m,j,k} = N_k - 1 \quad \forall k \in K, \quad (8)$$

$$\sum_{m=N+1}^{N+M} x_{m,0,k} = 1 \quad \forall k \in K, \quad (9)$$

$$\sum_{i=1}^{N+M} x_{i,N+M+1,k} = 1 \quad \forall k \in K, \quad (10)$$

$$\sum_{j=0}^{N+M} x_{N+M+1,j,k} = 1 \quad \forall k \in K, \quad (11)$$

$$\sum_{i=0}^{N+M+1} x_{i,j,k} = \sum_{i=0}^{N+M+1} x_{j,i,k} \quad \forall k \in K, j \in 0, 1, \dots, N+M+1, \quad (12)$$

$$w_{i,k} + s_i + t_{i,j} - w_{j,k} \leq (1 - x_{i,j,k}) \text{Big } M \quad \forall k \in K, (i, j) \in A, \quad (13)$$

$$w_{i,k} + s_i + s_{N+M+1} + t_{i,j} - w_{j,k} \leq (2 - x_{i,N+M+1,k} - x_{N+M+1,j,k}) \text{Big } M, \quad \forall k \in K, (i, j) \in A \quad (14)$$

$$a_i \sum_{j=1}^{N+M+1} x_{i,j,k} \leq w_{i,k} \leq b_i \sum_{j=1}^{N+M+1} x_{i,j,k} \quad \forall k \in K, i \in 0, 1, \dots, N+M+1, \quad (15)$$

$$E \leq w_{i,k} \leq L \quad \forall k \in K, i \in 0, 1, \dots, N+M+1, \quad (16)$$

$$x_{i,j,k} \in \{0, 1\} \quad \forall k \in K, (i, j) \in A, \quad (17)$$

$$N_k \geq 0, \text{ Integer} \quad \forall k \in K, \quad (18)$$

$$D_{i,k} \geq 0 \quad \forall k \in K, i \in 0, 1, \dots, N+M+1. \quad (19)$$

Constraints (1) impose the rule that each regular stop should be served by exactly one vehicle. Constraints (2) ensure that each route starts from the depot. Constraints (3) keep the amount of garbage collected at each stop within the vehicle capacity constraint. Constraints (4) ensure that the collected garbage volume is reset to zero once the vehicle has visited a landfill, as well as in the particular instance when the vehicle leaves the depot. Constraints (5) are in place to make sure that the collected garbage reflects the correct incremental volume of a particular stop when a truck visits that stop. By constraints (6), the actual number of disposal trips is calculated. Constraints (7) impose the condition that the number of trips from the regular stops to disposal facilities for a vehicle is equal to the number of actual disposal trips. Constraints (8) impose the condition that the number of trips from the disposal facilities to the regular stops should be one less than the number of disposal trips. The last trip from a disposal facility in

that route should be to the depot, and this will be achieved by constraints (9). Constraints (10) and (11) are introduced to add the lunch break for each route. Constraints (12) ensure that if the vehicle arrives at a stop, it must leave the stop. Constraints (13) through (16) make sure that the time constraints on both the route and vehicle are satisfied. Constraints (14) make sure that the travel time between before-lunch-break stop and after-lunch-break stop is included. Finally, constraints (17) through (19) impose binary conditions, nonnegative integer constraints, and nonnegative conditions on the variable set. As observed in the simplified waste-VRPTW model, the disposal trips and the lunch break make the problem much more complicated.

The Solution Approach

We developed and implemented an iterative two-phase algorithm in which the first phase generates an initial solution and the second improves it. First, we estimate the number of routes for a given site by considering the total workload and the maximum volume, yardage, number of lifts (commercial) or homes (residential), and allowable routing time of routes. Then, we initially (roughly) route the stops based on their locations and time-window constraints in the first step, which we call Phase 1. For the first initial route-building phase, we apply a hybrid method of balanced clustering and a greedy-insertion algorithm to enhance visual attractiveness. In this stage, we also balance the routes in terms of the workload so that a single vehicle can serve each route. The second phase of the algorithm accomplishes optimized route building using an extended traveling-salesman problem (TSP) and extended-insertion VRPTW and by applying metaheuristic algorithms. In this step, if any stops are not routed in a given route, it assigns those stops to their closest available routes. If after all the routes are optimized, there are stops remaining that have not been routed, the algorithm increases the number of routes and repeats the two steps from the beginning. We discuss the main algorithms of the procedure in the next sections.

Initial Route Construction Using Balanced Clustering Algorithms with Greedy Insertion

The first phase of the solution approach utilizes a k -means-variant-balanced-clustering algorithm (Kim

et al. 2004). Like the k -means algorithm, it selects an initial centroid seed stop for each route randomly as the first step and assigns the remaining stops to their nearest route. Then, it calculates a new centroid for each route. It clusters the stops based on the distances between the stops and the centroids. In other words, it assigns a stop to the route whose centroid is the closest to the stop. Each time it assigns a stop to a route, it considers the capacity of the route. The capacity of the route is defined by the maximum number of stops that the vehicle is allowed to visit, the volume or weight that the vehicle is allowed to handle,

and the allowable routing time. The algorithm estimates the routing time for each route each time it adds a stop to the route by using a fast TSP heuristic algorithm, which is a greedy insertion. When a route whose centroid is the nearest to the stop being processed has already reached its capacity, it assigns the stop to the next closest route whose capacity is not full.

We also developed and applied a simple improvement algorithm for enhancing visual attractiveness. The basic idea of the improvement is as follows: if a stop's current route is not the same as the stop's

Problem set	Best-known solution		IIT algorithm		Problem set	Best-known solution		IIT algorithm	
	Vehicle #	Distance	Vehicle #	Distance		Vehicle #	Distance	Vehicle #	Distance
C101	10	828.94	10	828.94	C201	3	591.56	3	591.56
C102	10	828.94	10	828.94	C202	3	591.56	3	591.56
C103	10	828.06	10	828.94	C203	3	591.17	3	603.37
C104	10	824.78	10	827.95	C204	3	590.60	3	601.11
C105	10	828.94	10	828.94	C205	3	588.88	3	588.88
C106	10	828.94	10	828.94	C206	3	588.49	3	588.49
C107	10	828.94	10	828.94	C207	3	588.29	3	588.32
C108	10	828.94	10	828.94	C208	3	588.32	3	588.32
C109	10	828.94	10	828.94	R201	4	1,252.37	4	1,279.98
R101	19	1,645.79	19	1,703.69	R202	3	1,191.70	4	1,105.65
R102	17	1,486.12	18	1,479.37	R203	3	939.54	3	1,000.33
R103	13	1,292.68	14	1,226.05	R204	2	825.52	3	774.41
R104	9	1,007.24	10	1,020.62	R205	3	994.42	3	1,099.99
R105	14	1,377.11	14	1,446.35	R206	3	906.14	3	971.85
R106	12	1,251.98	12	1,310.08	R207	2	893.33	3	854.01
R107	10	1,104.66	11	1,132.39	R208	2	726.75	3	743.27
R108	9	960.88	10	980.55	R209	3	909.16	3	1,019.73
R109	11	1,194.73	12	1,196.61	R210	3	939.34	3	1,000.26
R110	10	1,118.59	11	1,106.65	R211	2	892.71	3	848.92
R111	10	1,096.72	11	1,146.41	RC201	4	1,406.91	4	1,461.31
R112	9	982.14	10	994.16	RC202	3	1,367.09	4	1,224.31
RC101	14	1,696.94	16	1,705.72	RC203	3	1,049.62	3	1,134.20
RC102	12	1,554.75	13	1,537.24	RC204	3	798.41	3	853.45
RC103	11	1,261.67	11	1,324.17	RC205	4	1,297.19	4	1,362.71
RC104	10	1,135.48	10	1,173.91	RC206	3	1,146.32	4	1,135.19
RC105	13	1,629.44	14	1,577.90	RC207	3	1,061.14	3	1,204.67
RC106	11	1,424.73	12	1,450.65	RC208	3	828.14	3	897.29
RC107	11	1,230.48	11	1,248.44					
RC108	10	1,139.82	11	1,162.47					

Table 1: Because we could not find a publicly available benchmark problem set for waste collection, we used Solomon's (1987) well-known data set for VRPTW, which is much simpler than our problem, to evaluate our algorithms. The problem sets and their best-known solutions are available at <http://web.cba.neu.edu/~msolomon/>. The best-known solutions have been updated for the last 17 years after the problem sets were presented, and the current best-known solutions as of August 2004 came from 16 different works from all over the world. No single solution algorithm can generate the best-known solutions for all the problems. Most of the works are laboratory-level research experiments. Considering these facts, the results in this table, and the fact that ours has been successfully embedded in a real-world enterprise routing system and is being used, we can claim that our approach is highly competitive.

nearest route, move the stop to the nearest route when its move does not violate the capacity constraints of the routes and when its move improves the overall route quality.

Route Improvement by Extended-Insertion Algorithm and Metaheuristic Algorithm

We use an extended version of Solomon's (1987) insertion algorithm to improve a route generated in the first phase, and a combined simulated-annealing metaheuristic with the CROSS exchange local-search method of Taillard et al. (1997) for further improvement. While Solomon's insertion algorithm assumes that each vehicle departs from the depot, serves only one route, and returns to the depot, the extended-insertion algorithm considers multiple disposal-facility operations and the driver's lunch break for each route.


We treated the lunch break as a special stop with a time window between 11:00 AM and 1:00 PM and with a one-hour service time. We assume that the travel distance and time between the lunch break stop and a regular stop is zero. Each time a vehicle is full, it must visit a landfill facility to empty its load. To handle this, we applied an iterative process of splitting a route to Solomon's insertion algorithm. In the following description, a route represents the travel path of a vehicle and a subroute represents a sub-travel path that starts from the depot or a disposal facility and ends at a disposal location or the depot.

All stops are marked as *not routed* at the beginning of the algorithm. The algorithm creates an empty route for a vehicle and marks its lunch check as *not yet*. It initializes a subroute with a seed stop chosen as the farthest stop from the depot. While a route is initialized with a sequence (depot, the seed stop, depot) in Solomon's algorithm, a subroute is initialized with a sequence (depot, the seed stop, lunch break, the closest disposal to the depot, depot) in the extended-insertion algorithm. Because the route finally returns to the depot, the disposal facility closest to the depot is selected in an initial subroute.


After initialization, we insert repeatedly all stops that can be inserted without violating the time windows using Solomon's insertion method. Note that we do not consider vehicle capacity at this stage. Instead,

we consider only the time window constraints of the vehicle and stops. A simulated-annealing metaheuristic using the CROSS exchange local search method proposed by Taillard et al. (1997) improves the subroute constructed.

After the algorithm constructs a feasible subroute, it considers the vehicle capacity and splits the subroute. If the current subroute is within the vehicle capacity, it does not need to split the subroute and it adds the subroute directly to the current route. Otherwise, it adds a landfill facility in the sequence, splits the current subroute, and sets the rest of the stops after the splitting point to *not routed*. For the rest of the stops, it initializes a new subroute now with a sequence (landfill facility, new seed stop, the



(a) Best-known solution for RC205



(b) Our solution for RC205

Figure 4: We compare the visual attractiveness of our solution and a best-known solution. Although our solution for the problem set RC205 of Solomon's (1987) data set requires a longer travel distance than Bent and Van Hentenryck's (2001) solution, it has much less overlap between routes and so is much more visually attractive. The polygons represent the boundaries of routes, and the different shapes show the stops that belong to the different routes.

closest disposal to the depot, depot) and applies Solomon's insertion-heuristic algorithm to construct a subroute. It repeats this process until all the stops are *routed*.

Computational Experiments

Because we could not find a publicly available benchmark problem set for waste collection, we used Solomon's well-known data set for VRPTW to evaluate our algorithms. Note that Solomon's problem set does not require any disposal trips or a lunch break. The problem sets and their best-known solutions are available at <http://web.cba.neu.edu/~msolomon/> (Solomon 1987). Our method was able to generate competitive solutions compared with the best-known solutions (Table 1). Although our solutions are not as favorable when compared to the best-known solutions (in terms of the number of vehicles required and travel distance for some problem sets), our method usually generated more visually attractive solutions (Figure 4).

Acknowledgments

We could not have done the work reported in this paper without the committed hard work of many people from WM and IIT. We also appreciated invaluable help from

Anthony Brigandi, Sudhansu K. Baksi, Thomas Spencer III, Stephen C. Graves, and Mary F. Haight.

References

- Bent, Russell, Pascal Van Hentenryck. 2001. A two-stage hybrid local search for the vehicle routing problem with time windows. Technical report CS-01-06, Department of Computer Science, Brown University, Providence, RI. <http://citeseer.ist.psu.edu/454901.html>.
- Cordeau, Jean-Francois, Guy Desaulniers, Jacques Desrosiers, Marius M. Solomon, Francois Soumis. 2002. VRP with time windows. P. Toth, D. Vigo, eds. *The Vehicle Routing Problem*. SIAM, Philadelphia, PA, 157–193.
- Kim, Byung-In, Seongbae Kim, Surya Shao. 2004. Balanced clustering algorithms for vehicle routing problems. Institute of Industrial Engineers Annual Conference, May 15–19, Houston, Texas. Retrieved October 18, 2004 <http://www.xcd.com/iie/research/finalpapers/283.doc>.
- Poot, Alexander, Goos Kant, Albert P. M. Wagelmans. 2002. A savings based method for real-life vehicle routing problems. *J. Oper. Res. Soc.* 53(1) 57–68.
- Solomon, Marius M. 1987. Algorithm for the vehicle routing and scheduling problem with time window constraints. *Oper. Res.* 35(2) 254–265.
- Taillard, Eric, Philippe Badeau, Michel Gendreau, Francois Guertin, Jean-Yves Potvin. 1997. A tabu search heuristic for the vehicle routing problem with soft time windows. *Transportation Sci.* 31(1) 170–186.
- Waste Management, Inc. 2003. Investor Meetings, Boston, MA. <http://www.wm.com/newwm/investor/presentations/may20.ppt>.