# Consolidating and Dispatching Truck Shipments of Mobil Heavy Petroleum Products

DAN O. BAUSCH
*INSIGHT, Inc.*
*19820 Village Office Court*
*Bend, Oregon 97702*

GERALD G. BROWN
*Operations Research (OR/BW)*
*Naval Postgraduate School*
*Monterey, California 93943*

DAVID RONEN
*School of Business Administration*
*University of Missouri—St. Louis*
*8001 Natural Bridge Road*
*St. Louis, Missouri 63121*

Mobil Oil Corporation consolidates and dispatches truck shipments of heavy petroleum products—lubricants in packages and in bulk—from 10 lubricant plants nationwide. It dispatches hundreds of orders daily either individually or as consolidated truckloads, using a very nonhomogeneous fleet of Mobil-controlled and contract vehicles and common carriers. Shipment schedules may span several days and include stops to pick up returned drums or entire trailers. Shipping costs depend upon the vehicle used, the shipment size, the locations of the stops, and the route distance and time. Candidate consolidations are generated automatically or with dispatcher assistance. Then, the dispatcher uses optimization to select a minimal-cost set of schedules. Mobil has been using this system for three years, reducing annual transportation costs by about $1 million (US).

Stand not upon the order of your going,
But go at once.—Shakespeare, *Macbeth*

A computer-assisted system consolidates and dispatches truck shipments of packaged and bulk lubricant (lube) products at Mobil Oil Corporation. The system consolidates deliveries automatically or with dispatcher assistance, creating a large set of promising candidate schedules costed for each available truck type. The candidate schedules may include

Copyright © 1995, Institute for Operations Research and the Management Sciences
0091-2102/95/2502/0001$01.25
This paper was refereed.

INDUSTRIES—PETROLEUM/NATURAL GAS
TRANSPORTATION—ROUTE SELECTION

many truckload and less-than-truckload (LTL) shipment alternatives for any given order, may offer to ship orders earlier than absolutely necessary if this looks appealing, and may include stops to pick up returned drums or entire trailers. The system then solves a set-partitioning model to select a least-cost portfolio of schedules, which it then converts into truck routes. Ronen [forthcoming] gives a review of previous literature on dispatching all varieties of petroleum products.

**Mobil's Lube Operations**

Mobil Oil Corporation markets over 1,000 heavy petroleum products—lube oils, greases, and waxes—in over 2,000 product-package combinations. A small number of these products are fast movers, but the majority are slow movers and special products. The most familiar products are Mobil 1 motor oil, other Mobil motor oils, and industrial greases. Mobil markets these products to service stations, chain stores, and industrial customers, such as steel mills or automobile plants.

Mobil operates 10 lube plants in the continental US (Figure 1). These plants receive their base oil stocks in bulk from refineries and receive additives, mostly packaged, from refineries and outside vendors. Some lube plants blend products primarily to stock, while others blend primarily to order. Each plant blends products common to all plants, and each blends specialty products that it ships to the other plants for distribution. Each lube plant has a collocated mixing warehouse that supplies products to customers, distributors, packagers, and other warehouses.

Most of these products have high weight per unit volume: Truck weight limits are usually reached before truck volume limits. However, some products are light (for example, flaked wax), and some have bulky packaging (for example, drums).

There are about 10,000 customer accounts, each assigned to a primary supply source from one of the 10 lube plants. Mobil may ship from a secondary source in case of product shortage. Mobil reviews primary source assignments periodically. Dynamic sourcing is impractical because of the limited availability of special products.

A customer order usually consists of multiple products. Mobil fills as much of the order as possible, and back-orders the remainder. Infrequently, it may reassign entire orders prior to dispatch to a secondary source. Most plants and distribution warehouses have a local delivery region in which they assign customers to regular delivery days during the week. Customers not in a local delivery region may receive shipments on any day of the week.

Mobil usually pays for shipping heavy products. Some customers, however, prefer to pick up their own orders at the lube plant. Customer orders that are smaller than a minimum size are shipped freight collect or picked up by the customer.

In addition to transporting products outbound, trucks transport inbound bulk and packaged additives and packaging materials from suppliers, empty drums and trailers returned from customers, and occasional product returns from customers. They haul the empty drums to reconditioning facilities, which are not located at the warehouses.

**Centralized Dispatch**

Mobil has centralized its order taking and dispatching for heavy products at

**Figure 1: Mobil lubricant plants are located nationwide. Each order is sourced at one of these sites and either shipped individually or consolidated into truckloads with other orders.**

Malvern, Pennsylvania. It takes over 10,000 orders a month for packaged and bulk products, mainly by telephone. Mobil delivers about three-quarters of these orders, and customers pick up the rest.

Mobil dispatches each lube plant delivery area separately. It may ship orders earlier than required by customer service guidelines if the products are in stock and if it is economical to do so.

Mobil operates its own fleet of tractors, tank (bulk) trailers which may have multiple compartments, package trailers (box vans), and straight trucks. It uses this fleet, located at the lube plants, primarily for local truckload and LTL shipments, some long-haul truckload, and pickups. The Mobil fleet is operated for a single daily shift (with overtime) five days a week. The fleet is not uniform: Trucks may have different lengths, compartments, and equipment (Figure 2). Mobil also uses dedicated carrier trucks to supplement its own fleet for local delivery and long haul, and uses contract and common carriers for truckload and LTL long-haul shipments.

Mobil commits two separate dispatches each day for each lube plant, one for bulk orders with its fleet of tank trucks and the other for packaged products. The overlap between these two dispatches is minimal; however, they have some issues in common:

— What modes of transport should be used?
— What mix of owned and hired vehicles is best?
— Which carriers are most attractive?
— Which orders can be consolidated?
— Can pool or stop-off deliveries save

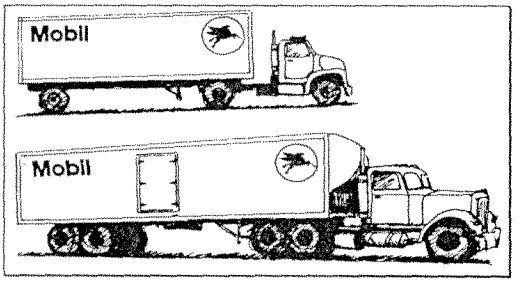


**Figure 2: Mobil has a mixed truck fleet. The 27-foot and 45-foot package trailers shown may be equipped with special delivery equipment. There are many other truck types.**

money?
— Should any future orders be shipped early?

Bulk tank trucks may have compartments so that they can carry multiple products. Most bulk shipments are tank truckloads. Very few bulk loads require multiple-stop delivery routes, and the tank trucks usually make multiple trips per day. The major problem in bulk dispatching is to coordinate production planning with dispatching, that is, to make sure that the bulk product will be available when the truck shows up for loading.

A packaged order may weigh from several pounds up to a truckload. Mobil may consolidate and ship orders by any of the following modes:

— Local delivery by owned trucks,
— Local delivery by hired trucks,
— Overnight delivery by owned trucks,
— Packaged good carriers (common carrier, United Parcel Service, and so forth),
— Truckload by contract carrier,
— Truckload by contract carrier, with stop-offs,
— Truckload to pool point (and a trailer switch), with local delivery beyond, and
— Outbound deliveries interspersed with order pickups or a trailer back haul.

Each of these shipment modes has its own costs and operating rules. These policies and restrictions include

— Truck and trailer capacity limits on weight, volume, and number of drums;
— Carrier limits on number of stop-offs and the extra driving distance these stop-offs require;
— Interstate and intrastate restrictions;
— Limits on the availability of special equipment;
— Limits on operating radius;
— Requirements for minimum empty truck capacity before pickup is permitted; and
— Time windows.

Mobil-owned and dedicated carrier trucks—the controlled fleet—must return to their source at the end of a schedule, and there are a number of ways to plan for this (Figure 3). Contract carriers may accept a route with intermediate stops as a truckload at a reduced truckload rate as long as the stops are not too far off-route (Figure 4). We can ship individual orders LTL via common carrier, but this costs more for good reason (Figure 5).

**System Development History**

Mobil began exploring the automation of heavy products dispatch following its success with centralized computer-assisted dispatch of light products in the early 1980s, reported later by Brown, Ellis, Graves, and Ronen [1987] and still in use today. In 1984, Mobil commissioned us to help formulate the requirements for a heavy-products dispatching system and to determine the feasibility of acquiring or

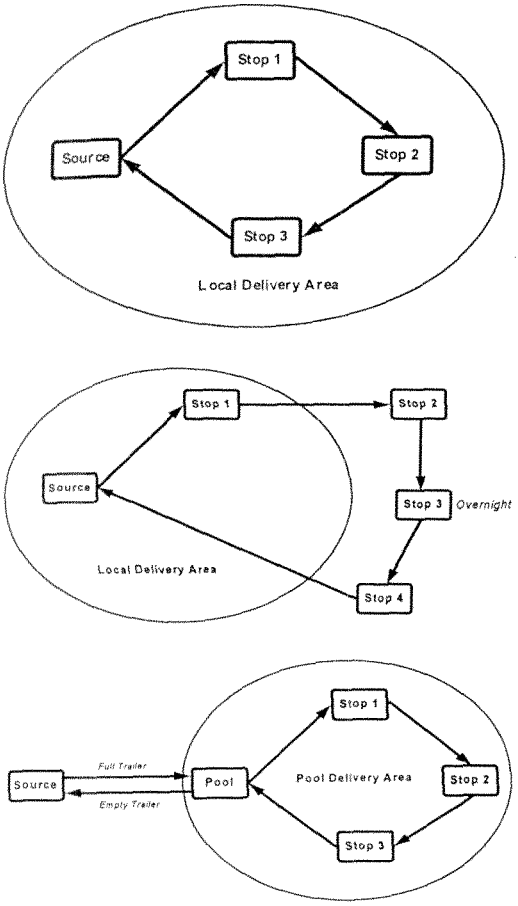


**Figure 3: Proprietary trucks and dedicated carriers—the controlled fleet—generally return to their sources. Shown are three ways to form truck schedules that do this. A local delivery area (shown at the top) may be defined within some geographic region or simply a radius from the source. Overnight deliveries outside the local delivery area are possible (center). A controlled truck can operate as a dedicated remote carrier and be supplied via linehaul to a pool location (bottom).**

developing such a system. We identified significant potential savings, but no off-the-shelf software package was available to do the job. Following some additional economic analysis, Mobil asked us to build a proof prototype of the most innovative technological component, an optimizing heavy-products dispatch module.

Although in 1984 Mobil centralized order taking on a computer in Woodfield, Illinois, dispatching was manual and decentralized at each lube plant. The company assigned each customer location to a geographical region and within that region to a master delivery route. On a given day of the week, each lube plant made deliveries only to specified regions and within those regions delivered customer orders via the master delivery routes.

The objectives of our prototypic effort were

— To give Mobil a better grasp of potential cost savings and their sources,
— To provide a clearer picture of the support required for a dispatching system, and
— To evaluate the feasibility of using the order-taking computers already in place for dispatching.

The prototype model clustered orders into truckloads and assigned truckloads to trucks using a generalized assignment model. However, the variety of trucks and

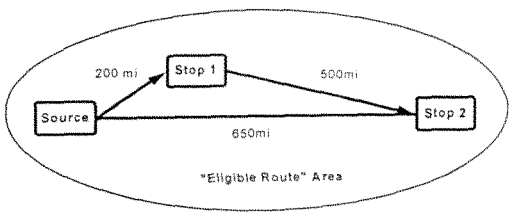


**Figure 4: Contract carriers will accept truckload routes with intermediate stops, as long as the stops add no more than some percentage (say 10 percent) to the direct distance from source to last stop. In this example the direct distance is 650 miles, permitting a route distance of 715 miles. The carrier will accept this 700-mile stop-off route as a truckload.**

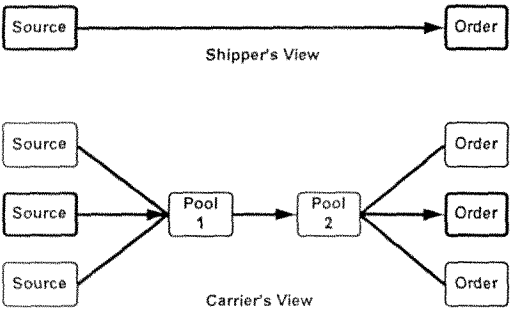


**Figure 5: Shipping LTL via common carrier looks simple enough. The carrier's view reveals why LTL costs more: The LTL carrier has to find its own opportunities for truckload consolidation but eventually deliver our order.**

transportation alternatives, combined with the complexity of the dispatching environment, necessitated refining the solutions using a dozen heuristic steps. The prototype was run with one month of actual data from one representative lube plant and compared with the actual, independent manual dispatches. Potential transportation cost savings were in the range of seven to 10 percent. The prototype was also tested for fleet size and mix decisions (for example, what is the best size and mix of Mobil's controlled fleet on a given day?).

Although the prototype was successful, the next step posed significant technical problems. The order-taking computers turned out to have no working scientific computer language, and connecting them to microcomputers (for data download, computation, and upload) was infeasible. In addition, the existing computers were already becoming overloaded and were due for replacement at any time. Developing a dispatching system was put on hold, pending upgrade of the host computers. In the meantime, as a result of the prototypic effort, Mobil appended new geographic data to the orders to facilitate better manual dispatches.

In 1989, Mobil centralized dispatching of heavy products at the same location as light products in Valley Forge, Pennsylvania and moved order taking to a new mainframe computer. These changes resurrected the potential for computerized dispatch. The new mainframe provided enough computing power to solve more aggressive mathematical models of the dispatch. Mobil decided to pursue a dispatching module hosted as a background transaction in its operating system.

In 1990, Mobil started development of user and database interfaces and asked INSIGHT to develop the dispatch module. We put the system into production on schedule in June 1991, and it has been operational since. Mobil has since asked for only one minor revision of a report in the dispatching module.

**System Design and Operation**

The heavy-products computer-assisted dispatch (HPCAD) system operates with Mobil's heavy-products system (HPS), where the orders reside. All user and corporate database and systems interfaces, as well as the additional data files needed for this new application, have been developed by Mobil information system personnel in close cooperation with the dispatching staff at Valley Forge.

The centerpiece of HPCAD is the INSIGHT dispatching module, which consists of a dispatch data importer, a schedule generator, a rater, an optimizer, and a dispatch solution exporter. The dispatch data (Table 1) are checked for completeness and consistency—this is important

**Source**
- Dispatch date
- Bulk or package dispatch
- Source name
- Location
- Speed-by-distance table
- Dispatch version
- Local dispatch policy parameters

**Orders**
- Customer and order identification
- Delivery/pickup indicator
- Scheduled data
- {trip, standard distance, and time}
- Location, including state
- Weight, volume, drums
- Preferred truckload common carrier
- Truckload cost
- Preferred LTL common carrier
- LTL cost
- Alternate source location (pool point trailer switch, transfer point)
- {smaller package (for example, UPS) cost}
- {time window}
- {special equipment}
- {preassigned requirement for any of the following: carrier, truck, trip, stop}
- Optimizer suggestion for date, carrier, truck, trip, stop

**Trucks**
- Identification
- Start location, if not source location
- {return to start location indicator}
- Package or bulk truck
- Operating radius around start location
- Contract type: proprietary, dedicated, common carrier truckload, or common carrier LTL
- {can do pickups}
- Intrastate, interstate, or both
- Weight, volume, drum capacity (package), or volume-by compartment (bulk)
- Special equipment
- Costs per day (minimum), hour, distance, stop, overnight, under-, overtime hour
- Shift hours under-, overtime, minimum, maximum
- Distance minimum, maximum
- Maximum number of stops per trip

**Georeference**
- Map for determining location-to-location distances and times

**Future workdays**
- Calendar for proximate future orders

**Table 1: Dispatch data are integrative and voluminous. Some optional fields are listed in curly braces, some other fields need not always be defined, HPCAD may edit some fields, and optimizer suggestions are outputs.**

because HPCAD integrates a variety of data types and sources that would not normally reside together. The system diagnoses and treats minor data problems by applying default rules with appropriate prescriptive messages. Major data problems result in error messages and must be externally corrected before the dispatch can resume.

The schedule generator aggregates trucks into truck types; trucks within a truck type are essentially identical for the work at hand. For each truck type, it selects all compatible orders (deliveries and pickups) and from these generates all candidate feasible work schedules (routes). Although Mobil prefers to include all alternate schedules, their number can be limited by enforcing quotas on the number of candidate schedules for each truck type, or for each truck category, or for each order, and so forth. This has been necessary in practice only for local deliveries from the largest lube plant.

A work schedule may consist of several consecutive trips, each starting and ending at the origin. When the schedule generator considers work schedules with multiple-stop trips, it generates the trips by a sweep heuristic [Gillett and Miller 1974]. Each sweep collects work by defining a ray from the origin to a starting seed order and then rotating the ray about the origin clockwise (or counter-clockwise) until the truck is full. A separate sweep is performed for each compatible seed order. Each sweep usually results in more than one feasible combination because as the sweep progresses, subsets of orders also form feasible combinations. For example, if a truck type is limited to three stops per trip, and we number orders encountered in a sweep 1, 2, and 3, and these orders will fit on the truck, the following feasible combinations

will be generated: 1; 1, 2; 1, 2, 3. The generator generates and retains only feasible combinations: for instance, interstate orders cannot be added to a combination on an intrastate truck type.

Unfortunately, trip costs for common carrier trucks defy approximation as a simple function of distance, weight, time, or number of stops—they are neither concave nor convex functions of these (for examples, review Figures 4 and 5). The schedule generator sequences feasible order combinations into stops within each trip to minimize distance or time either by full enumeration (two or fewer stops) or by a quadratic assignment heuristic. These sequences are constrained by dispatch policy. For instance, the distance between any successive pair of stops must not exceed a maximum by truck type. Also, sequenced truckload trips for commercial carriers must have a total route distance to the last stop that is no greater than a given percentage of the direct distance from departure to the last stop (Figure 4). Policies such as these help create face-valid routes that will be accepted and driven. Given capacities by truck type, stops for deliveries and pickups must also be feasible in the sequence driven. For controlled truck types that return to known locations (that is, proprietary and dedicated trucks), straight driving distance, time, and number of stops suffice for cost calculations. For other truck types, the sequence is ended at an anchor order that has that truck type as its preferred carrier and thus contributes a preferred rate to use for trip cost calculation.

Each work schedule generated is sent to a rater which calculates costs for that particular truck type; the rated schedules are then sent along with elastic penalty costs to the X-System [INSIGHT 1990] in the form of an elastic set-partitioning model (appendix). The dispatch solution exporter breaks truck types back up into individual trucks and suggests date, carrier, trip, and stop for each order, giving the detailed schedule of stops, distances, hours, and cost. For each truck type and each truck category (Mobil, dedicated, common carrier truckload, and common carrier LTL), it displays a summary of the number of orders dispatched, the total cost, and the average cost per order. Overall gauges of dispatch performance include cost per distance, cost per weight, cost per weight-distance, cost per volume, and penalty costs.

Figure 6 shows the dispatching process. Dispatching begins after dispatchable orders—those passing credit checks and for which inventory is available—are extracted

## The fleet is not uniform.

from HPS open orders. These are passed through a system called RATETRAC, which appends to each order its preferred carriers and rates.

Because order and truck data change daily, a preview of order workload and trucks available is always called for, with possible adjustment of dispatch policy parameters and correction of data blemishes. Minimal manual data entry is required, but dispatcher assistance is always welcome. For instance, the dispatcher may add, change, and rerate orders after the batch process through RATETRAC. Orders can also be manually preassigned to a carrier,

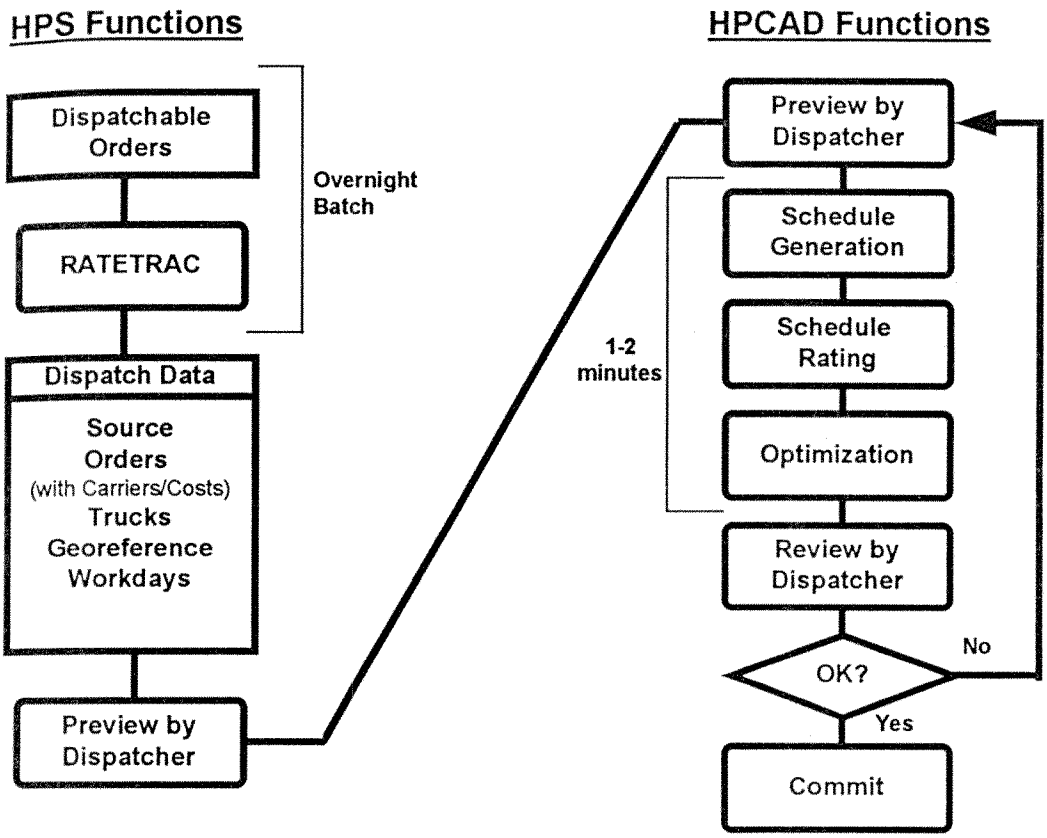


Figure 6: Heavy products computer-assisted dispatch (HPCAD). Every day, for every lube plant nationwide, for package and bulk products, Mobil's heavy products system (HPS) provides orders, and RATETRAC suggests for each the best common carrier truckload, LTL carrier, and cost. The entire dispatch data scenario is assembled for dispatcher preview. Schedule generation, schedule rating, and optimization take about 15 compute seconds, yielding a response time of just over a minute or so. Most important, the dispatcher can easily review suggested solutions, compare solutions, save, restore, or navigate among any tentative solutions, and manually preassign anything. Dispatch commitment prints shipping documents at the supply site.

or to a particular truck, or trip, or stop, or any combination of these.

A shift is then submitted for dispatch. Within a minute or two, the dispatcher receives an optimized result. A great deal of effort has been invested in HPCAD to make it easy and intuitive for the dispatcher to accept a dispatch, reject it and start over, or amend part or all of a prior dispatch and try again. The dispatcher can pursue and compare multiple parallel notions about what will work best that day.

**Development Challenges**

Several collateral topics have required effort. Among these are mechanisms to provide driving distances and times, carrier rates and delivery costs, consideration of future orders for early delivery, and appropriate solution technology.

Truck routes begin at sources and go to

customer locations but may continue with multiple additional stops: Trip standards are available that give driving distances and times between sources and customer locations, but it is impractical to keep records of distances and times between every pair of customers. We designed a georeference system with the help of the dispatchers that divides the country into dispatch zones. Each zone is a small contiguous area, perhaps several miles across, with a center location and one or more neighbor zones directly connected to it for purposes of truck transport. Every customer location is coded with a zone. For trips between a pair of locations in the same zone or neighboring zones, the truck travels directly from location to location. Otherwise, we construct a shortest-time or shortest-distance driving path from location to neighboring zone center and hence via legs between successive neighboring zone centers until the truck can drive from a zone center neighboring the destination zone directly to the destination location (Figure 7). We estimate driving times based on a table giving speed as a function of leg distance (longer distances have higher overall average speeds). Through extensive testing we have validated that these routes are realis-

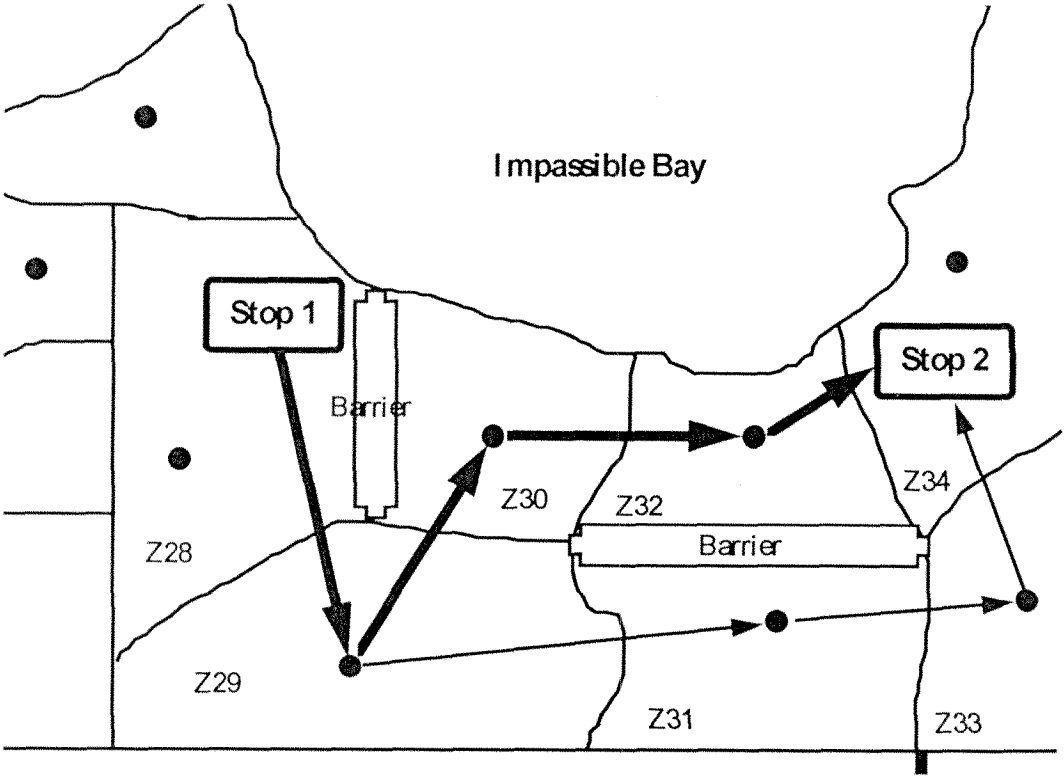


**Figure 7: "Geo-zones" are designed by dispatchers. Dots show zone centers. Trucks can drive from each zone directly to neighbor zones. The barriers show contiguous zones that are not neighbors. Driving routes begin at a location, then pass through successive neighbor zone centers until they reach a zone center neighboring the destination location, from which they proceed directly to the destination.**

tic and that the route distances and times are reasonably accurate, which makes sense in light of the dispersion of orders over a relatively large geographic area.

Common and contract carriers are an option for virtually all orders. Carrier rates are numerous, complex, change constantly, and depend on a variety of factors, including the carrier, product, shipment size, origin, and destination. HPCAD probably could not accommodate such a large database and still operate reliably and responsively. Manual extraction of rates for dispatching is impossibly tedious. Fortunately, when this issue arose Mobil was putting the finishing touches on RATETRAC, a system that maintains current carrier rates for other applications. RATETRAC manipulates a large and rather unwieldy database and can only process one order per invocation. Mobil provides an HPS procedure that extracts sourced, dispatchable orders and then submits each order to RATETRAC in an overnight process. It recommends a best truckload and LTL carrier and rate for each order.

Manual dispatching is staggeringly complex.

Future orders may be scheduled early if the products are available and if it makes sense to move the work up (that is, if a truck is not full but is scheduled to pass in the vicinity of a compatible future delivery). But shipping too many orders early may starve the controlled fleet in the future, leaving too little work to keep precommitted trucks and drivers busy. Accordingly, when the sweep encounters a good future-order candidate, we may include it with a current-order combination. However, future orders do not have to be moved up: Each future order is given a penalty for not shipping, which is its common carrier cost divided by the number of workdays left until the order must be shipped (that is, the number of remaining daily opportunities to ship the order). This heuristic always permits the future order to be ignored at a modest penalty, unless it appears in a schedule with other orders at an even more attractive cost.

Heuristic methods can be helpful in speeding construction of compatible order combinations, in sequencing these into face-valid candidate schedules, and in limiting the exponential numbers of candidate schedules when this is necessary. However, heuristics have not proven reliable for selecting which particular schedules to dispatch. For this we have had to resort to optimization: In particular, we solve set partition integer linear programs, and these optimizations require care (appendix).

**Evaluation**

Mobil has been using HPCAD for three years, using it at least once each day for packaged products and separately for bulk at each lube plant. Dispatchers usually experiment several times before committing a dispatch, especially with packaged products, which are harder to dispatch. Individual dispatches involve as many as 40 trucks and 250 orders.

In 1992, Mobil performed a thorough audit of HPCAD, including independent manual and HPCAD-assisted dispatching of the same workload. Originally, we forecast that HPCAD would offer potential savings of about $700 thousand annually.

The audit revealed annual savings of about $1 million.

About 77 percent of dispatches with HPCAD differ significantly from manual dispatches. HPCAD loads more weight and makes more stops per trip on company-controlled trucks. Manual dispatchers tend to schedule straight-line trips, making all deliveries on the way out to the last stop, or on the way back; trips from HPCAD include cloverleaf patterns and some total surprises. For example, Figure 8 shows that, counter to intuition and common practice, it may make perfect sense to build truckloads with large early deliveries and small deliveries later. HPCAD trips may look funny at first, but they are driveable and they save money. Other HPCAD routes arouse initial curiosity until it becomes clear that HPCAD can afford to customize each trip for each truck type, rather than laboriously build a general-purpose trip by hand first and then find a truck to put it on. For example, HPCAD suggests trips that only company-controlled trucks can drive, because they cross boundaries that commercial carriers will not. Rather than use LTL, HPCAD includes more orders as stop-offs on truckload shipments: This difference alone accounts for savings in the range of $40-to-$300 for each dispatch. HPCAD makes early delivery of future orders, which for one order saved about $600.

Additional benefits include

— Fewer dispatchers needed (especially to

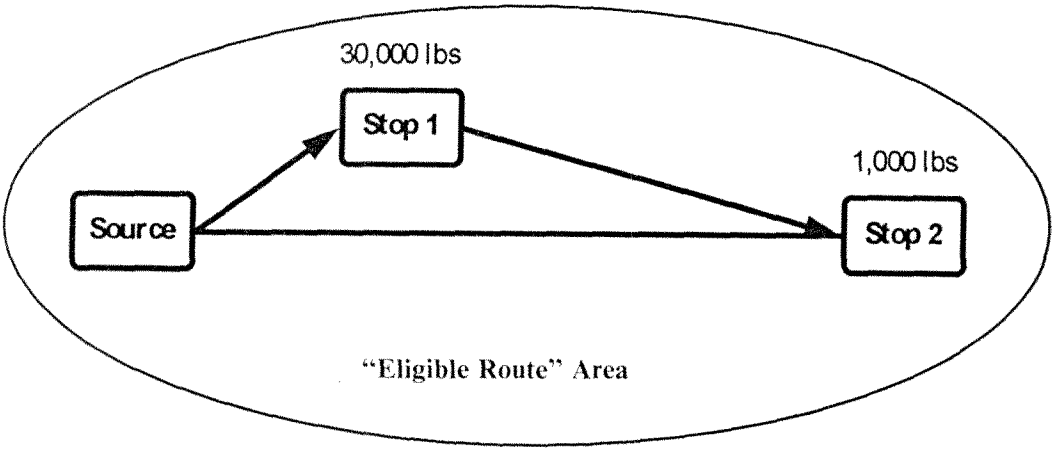


| Two individual shipments | |
|---|---|
| Stop 1 TL | = $450.00 |
| Stop 2 LTL | = $150.00 |
| Total | = $600.00 |

| One shipment with stop-off | |
|---|---|
| Stop 1 Drop | = $ 50.00 |
| Stop 2 TL | = $500.00 |
| Total | = $550.00 |

**Figure 8: Cost-based routes may not be intuitive. This route violates the conventional rule of thumb that the heaviest deliveries should be last on the route. Here, paying a full truckload rate for 1,000 pounds and dropping off 30,000 pounds for a small stop-off charge is cheaper than paying for two individual shipments.**

sequence routes),
- More consistent service (orders are assigned to compatible trucks),
- Reduced lead time for customer orders,
- Faster training for new dispatchers, and
- Answers to questions about the size and mix of the controlled truck fleet.

**Summary and Insights**

How do you save money in an intensely competitive industry in which truck operators are efficient, dispatchers are smart and good at what they do, and everybody already works hard? Two of the biggest problems are lack of information and little time to act on information: HPCAD transforms these problems into opportunities.

Manual dispatching is staggeringly complex and is always performed under severe time pressure. Perhaps the most innovative contribution of our work is cost visibility. Dispatchers can now grasp the costs of their decisions—this cost visibility has never been practical before. When dispatchers start talking about scheduling in terms of cost, rather than feasibility, managers are pleased.

Traffic folks, the ones who negotiate rates, benefit from this approach. There are advantages to looking at the big business picture, rather than focusing exclusively lane-by-lane on points per hundredweight mile. It's as important to know what shipping capabilities to negotiate for as it is to negotiate the rates.

Nonoptimizing heuristic methods are easily explained and can be fast, but in our prototypic experience, their performance and reliability depended upon assumptions we cannot make about our costs, orders, or trucks. Even if we build good schedules with heuristics, we have little confidence that heuristics will reliably, globally exploit the opportunities these schedules offer.

Optimization is key for success. Sometimes common sense, experience, and heuristic methods suffice, sometimes not. The important point is that there is no way to know for sure. If the optimum is not known, sooner or later someone will discover a better solution. Once that happens, and only once will suffice, credibility is permanently compromised.

It's hard to save money when operations are already fine-tuned and efficient. Without optimization, it is difficult and risky to try to assess subtle trade-offs and distinguish between close alternatives.

## HPCAD is intended to help dispatchers, not replace them.

Shipping goods by truck is an essential activity in a developed economy and one that most people take for granted—as long as shipments continue uninterrupted. In 1991, 1.2 million tractors and 3.6 million full or semi trailers conveyed about 2.7 billion tons of goods between US cities at a cost of $167 billion (that is, about three percent of the gross domestic product [GDP]). Local truck deliveries cost another $110 billion (two percent of GDP). Petroleum products accounted for 27 billion ton-miles [*Transportation in America* 1993]. Clearly, improving the efficiency of truck shipping has significant economic impact.

Mobil has been using HPCAD for several years to dispatch orders for heavy products in bulk and packages from lube plants to customers. HPCAD has not only

saved Mobil considerable amounts of money, it has helped Mobil to improve its customer service. HPCAD quickly generates large numbers of attractive, face-valid alternate shipment schedules—sets of orders sequenced as feasible routes for particular trucks in a heterogeneous fleet; dispatchers select a global, minimum-cost dispatch from these by quickly solving a set-partitioning problem. HPCAD is intended to help dispatchers, not replace them. Dispatchers can contribute their own ideas and compare alternatives; they control all aspects of a dispatch and have the final authority to commit it.

**Acknowledgments**

We thank Hal Schwartz of Mobil Oil Corporation, whose foresight and persistence brought about this application; Richard Powers of INSIGHT Inc., who facilitated work on this project; Rick Rosenthal of the Naval Postgraduate School, who was deeply involved in the early prototype work; Rob Dell, Tom Halwachs, and Kevin Wood of the Naval Postgraduate School for scrupulous manuscript reviews; and Ron Fiano, Andy Gibbons, Terry Hamlin, John Makhoul, and the rest of the Mobil project team who made things happen.

**APPENDIX**

The elastic set-partitioning model (ESP) selects a least-cost portfolio of delivery schedules for each truck type.

Indices:

$w \in W$ orders (deliveries, pick-ups);
$t \in T$ truck types; and
$s \in S$ alternate schedules.

Induced index sets:

$S_w$ schedules including order $w$; and
$S_t$ schedules requiring truck type $t$.

Data:

$COST_s$ cost of schedule $s$;
$\underline{WP}_w$, $\overline{WP}_w$ penalties for under-, overperforming order $w$;
$TRUCKS_t$ Number of trucks of truck type $t$; and
$\underline{TP}_t$, $\overline{TP}_t$ penalties for under-, overusing truck type $t$.

Decision variables:

$y_s$ binary decision variable equals one when using schedule $s$,
$\underline{\omega}_w$, $\bar{\omega}_w$ elastic constraint violation variables for under-, overperformance of order $w$; and
$\underline{\tau}_t$, $\bar{\tau}_t$ elastic constraint violation variables for under-, overutilization of truck type $t$.

ESP Model formulation:

subject to

$$\sum_{s \in S_w} y_s + \underline{\omega}_w - \bar{\omega}_w = 1 \quad \forall w, \tag{1}$$

$$\sum_{s \in S_t} y_s + \underline{\tau}_t - \bar{\tau}_t = TRUCKS_t \quad \forall t, \tag{2}$$

$$y_s \in \{0, 1\} \quad \forall s, \tag{3}$$

$$\underline{\omega}_w, \bar{\omega}_w \geq 0 \quad \forall w,$$

$$\underline{\tau}_t, \bar{\tau}_t \geq 0 \quad \forall t, \tag{4}$$

minimize

$$\sum_s COST_s y_s + \sum_w (\underline{WP}_w \underline{\omega}_w + \overline{WP}_w \bar{\omega}_w) + \sum_t (\underline{TP}_t \underline{\tau}_t + \overline{TP}_t \bar{\tau}_t). \tag{5}$$

Each partitioning constraint (1) seeks to schedule an order; a lower violation represents a common carrier shipment (where permissible, at the appropriate LTL or UPS cost), while an upper violation results in a high disruption penalty (which is worth analyzing if it happens). Each cardinality constraint (2) seeks one schedule per truck of truck type $t$; where a lower violation represents idleness of such trucks (at a penalty representing the appropriate idleness cost), and an upper violation incurs a schedule disruption penalty (more trucks may be needed). Stipulations (3) and (4) respectively specify discrete schedule selection and nonnegative elastic logical

variables. The objective function (5) includes the transportation cost of selected schedules as well as penalties.

Each alternate schedule $s$ is associated with a truck type $t$ and conveys a set of orders that constitutes one or more trips or routes of sequenced stops. The derivation of $COST_S$ includes finding for each feasible truckload combination of orders a minimal-distance feasible sequence that can be driven as a trip route, with stop-by-stop deliveries and pickups. If truck type $t$ returns to its starting location, multiple trips may be concatenated into one shift. Each route is costed event by event, with $COST_S$ the total cost of all trips and stops in the shift schedule.

A typical problem may have 250 orders and 40 trucks, which translate into about 275 constraints and several thousand binary selection variables. These problems are reliably solved in a minute or so, using an integrality tolerance (best incumbent objective cost less lower bound on this cost, expressed as a percentage of objective cost) of 0.1 percent using the X-System [INSIGHT 1990].

The X-System employs several nontraditional solution methods to produce these results, including elastic constraints, row factorization, cascaded-problem solution, prereduction, constraint-branching enumeration, and a fast dual simplex procedure.

Elastic constraints may be violated at a given linear penalty cost per unit of violation. Every constraint here is elastic. The X-System represents the elastic variables and penalties logically and exploits elasticity during optimization, concentrating on the active, or taut, constraints. Setting these elastic penalties warrants some thought: Here, some of the penalties play a direct role in the dispatch (for example, $\underline{WP}_w$ is the LTL or UPS cost for order $w$). Even when the penalties represent a high model disruption cost, one wants penalties that are meaningful when they are necessary and neither too low (soft) nor too high (hard). Moderation is a virtue. Fast, good-quality solutions are the reward.

Row factorization identifies and exploits sets of constraints that share a common special structure. Brown and Olson [1994] give set-partition examples along with a number of other applications to demonstrate the value of this approach in comparison to the traditional methods used by well-known commercial optimizers. Constraints (2) have at most one unit coefficient associated with each variable and thus qualify as generalized upper bounds. Exploiting this factorization reduces computation time dramatically.

Even the linear programming relaxations of ESP (that is, where stipulations (3) for binary variable values are relaxed to continuous unit bounds) can be difficult to solve directly. Cascaded problem solutions permit a particularly difficult model to be solved incrementally: we solve a sequence of submodels, analyze subsolutions and maintain records for the role played by each constraint and each variable, maintaining variables that would otherwise not be part of a submodel at their last-known values. Eventually, we can use recorded variable values as an advanced starting point for solving the entire model. Here, we follow the advice of Bausch [1982] and define subproblems by labeling constraints and variables as follows: Sort the constraints (1) by cardinality of $S_w$ and then break these sorted constraints up into, say, five roughly equal-sized components and label each constraint with its component number (that is, separate orders with fewer schedules from those with more). Label each variable $y_s$ with the minimum incident component number (that is, label each schedule $s$ with the component number of the order which is included in the fewest other schedules). Label all constraints (2) "0." Next, solve the following sequence of subproblems, where each of these is identified by "(min-label, max-

label)": (0, 1), (2, 2), (0, 2), . . . , (0, 5).

Prereduction of model instances prior to optimization, that is, seeking and removing structurally redundant features, can reveal hidden problem structure and speed up computation. For example, suppose some order $w$ is present only in a single schedule $s$, then the binary variable $y_s$ can be fixed to one and the constraints (1) for order $w$ and all other orders accompanying $w$ on schedule $s$ can be relaxed and ignored. Removing such isolated orders may isolate more orders, permitting further reductions. A number of well-known reductions like this one are easy to apply. We prefer that the problem generator be smart enough to detect and unambiguously diagnose such redundancies while creating the model. After all, the generator knows a lot more about the data and the problem than the optimizer does. For instance, the example above may actually be due to an order that requires special equipment available only on one truck type. Another reduction is trivial for the generator to explain but requires duality for a mathematical justification: A schedule will never be selected if its cost exceeds the sum of the costs to ship its orders as individual packages via LTL or UPS. We use the X-System prereduce function to tell us in a fraction of a second whether the problem generator is generating "good" models. Our goal is models that cannot be prereduced at all.

Constraint branching is a variation of branch-and-bound integer enumeration that selects a branch variable on the basis of its direct influence and the indirect effects of the values it will induce for other structurally dependent variables. For instance, constraints (1) dictate that if a schedule $s$ for order $w$ is fixed to one, then all other schedules carrying that order (those in the set $S_w \backslash s$) may be set to zero. Constraint branching speeds up integer enumeration. Branch variables are selected for restriction based on an estimate of the full elastic cost consequences of such restriction (that is, a beneficial interaction exists between elasticity and constraint branching).

Finally, this model has far fewer constraints than variables, and our dual simplex procedure solves it about 10 times as fast as the primal. (This also turns out to hold for the partitioning examples solved with primal by Brown and Olson [1994].)

Overall, these special techniques perform reliably on Mobil's heavy product dispatches hundreds of times daily. In fact, a personal computer is more than adequate for solving such problems in a minute or two.

**References**

Bausch, D. O. 1982, "Computational advances in the solution of large-scale set covering and set partitioning problems," MS thesis, Naval Postgraduate School, Monterey, California, October.

Brown, G. G.; Ellis, C. J.; Graves, G. W.; and Ronen, D. 1987, "Real-time, wide area dispatch of Mobil tank trucks," *Interfaces*, Vol. 17, No. 1, pp. 107–120.

Brown, G. G. and Olson, M. P. 1994, "Dynamic factorization in large-scale optimization models," *Mathematical Programming*, Vol. 64, Series A, No. 1, pp. 17–51.

Gillett, B. E. and Miller, L. R. 1974, "A heuristic algorithm for the vehicle-dispatch problem," *Operations Research*, Vol. 22, No. 2, pp. 340–349.

INSIGHT, Inc. 1990, "XS(F) mathematical programming system," Alexandria, Virginia.

Ronen, D. forthcoming, "Dispatching petroleum products," *Operations Research*.

*Transportation in America* 1993, Eno Transportation Foundation, Landsdowne, Virginia.

---

W. R. Baulkwill II, Supervisor, Transportation Programs, USS&L Logistics Support, Room 6A-909, Mobil Oil Corporation, 3225 Gallow Road, Fairfax, Virginia 22037-0001, writes, "We have reviewed the article submitted to you and have found it to be accurate in its description of our business. We installed the HPCAD op-

timizer (HPCAD stands for Heavy Products Computer Assisted Dispatch) in 1991 and it is working out very well for us. This optimizer has allowed us to reduce headcount in the dispatching function and improve our distribution economics. This algorithm has provided a very satisfactory return on our initial investment and since implementation, we have had to do very few enhancements."