INTERACTIVE OPTIMIZATION FOR MOTOR CARRIER LOAD PLANNING

By

Warren B. Powell
Princeton University

and

Yosef Sheffi
Massachusetts Institute of Technology

HEADNOTE

In the summer of 1982, I.U. International Management Corporation, who owns several motor carriers, contracted the Massachusetts Institute of Technology and Princeton University to develop a computerized decision aid for performing the continuous operations planning of its LTL motor carriers. The software, which was completed after three years of research and development, is an interactive optimization code. It has been in continuous use at PIE (the forth largest national motor carrier) since June, 1984 as its main load planning tool. In addition it is used there for strategic planning and as an aid in marketing. This article describes the load planning problem and some of the main features of the package developed. It discusses some implementation issues and draws conclusions regarding the general applicability and merits of the approach used, that of interactive optimization, to other transportation and logistics problems.

INTRODUCTION

This article describes a computerized system for load planning developed for and installed at a large LTL firm--Pacific International Express (PIE). It presents a unique modelling approach to the simultaneous optimization of service lanes, the freight movement plan, and the distribution of empties. The model is also used for strategic planning to optimize terminal locations and test economic scenarios as well as for marketing decisions.

The decision support system (named **APOLLO** for Advanced Planner of LTL Loads) is based on both network algorithms and interactive optimization. It combines these elements to solve an extremely large integer programming problem involving millions of decision variables. The article describes the system developed including its use by the abovementioned carrier, and brings out some of the modelling principles employed and learned in the process of developing this decision support system.

THE LOAD PLANNING PROBLEM

LTL carriers typically operate a set of end-of-line (EOL) terminals and breakbulk terminals ("Breaks"). Each end-of-line terminal controls the pick-up and delivery of freight within a specified service area. For the purposes of the discussion in this article, however, each EOL terminal can be viewed as the point of origin and termination of freight. The breakbulk terminals are used for sorting and consolidating shipments, serving as freight transshipments points. Typically, each EOL terminal is associated with one "primary" break, which controls many of the operations of that EOL. (Occasionally an EOL is tied to more than one primary BB.)

To operate the consolidation network efficiently, each terminal has to follow a freight movement plan (FMP). This plan specifies what is the next downstream terminal that a shipment should be sent to, given its final destination. Terminal managers usually follow the plan except in cases in which real-time fluctuations in the freight characteristics necessitate a deviation from it. (Typically, company policy would include guidelines for allowed deviations.)

The freight movement itself is conducted over a network of established direct services, such as that shown in Figure 1. To ensure a high level of service, each direct service (shown as a link in Figure 1) is operated at a minimum daily or weekly frequency. The important decision variables in the load planning process (the process of developing a FMP) are then: (i) where to establish direct services, and (ii) how to route the freight over the network of direct services. Additional decision variables involve the distribution of empty trailers resulting from flow imbalances and the routing of TL freight (which is typically used for backhaul only).

The objective of load planning is to derive a FMP that minimizes the total system costs and maintains high service levels, subject to many constraints stemming from the nature of the operation. This problem is not amenable to standard optimization techniques due to the following reasons: (i) the network is

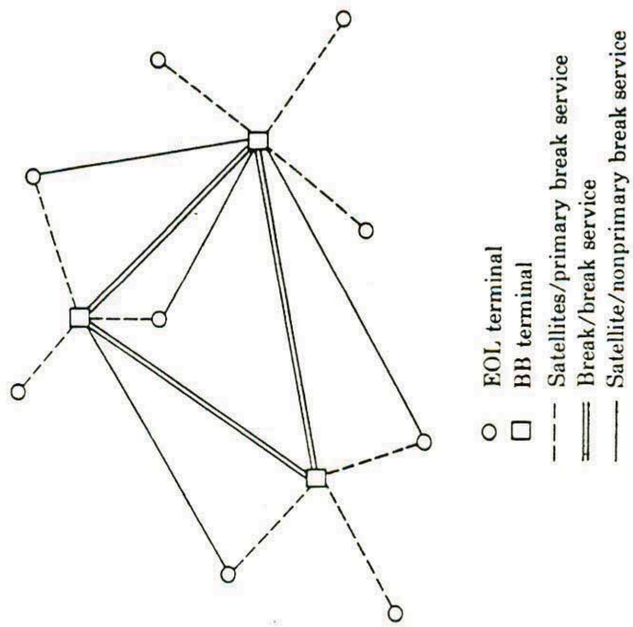

Figure 1

SCHEMATIC LTL NETWORK

very large; (ii) the analysis involves several (sometimes competing) objectives; (iii) the mathematical formulation of the problem involves nonconvex functions, discontinuities, and other complicating factors; (iv) the problem includes a spatial recognition dimension in which computers are particularly weak; and (v) the problem includes nonquantifiable parameters and constraints. These complications are not unique to load planning; they are apparent in many realistic transportation and logistics problems. For load planning problems these challenges are explained in the following paragraphs.

The size of the problem results from the fact that major LTL carriers operate 300-400 terminals (with a few operating many more). A carrier with, e.g., 300 EOL terminals and 30 breaks operates a network which includes over 100,000 origin-destination pairs and over 20,000 design links in the load planning network. In addition, there are over 3,000,000 routing variables (all of which are integer). The multiobjective aspect of the problem stems from the basic need to balance costs and level of service. In addition, there may be regional considerations and special marketing considerations which have to be accounted for.

The nontrivial nature of the cost functions can be exemplified by looking at one of the cost components--the linehaul cost over a given link in the load planning network. Due to the indivisibility of the vehicles and the need to maintain levels of service, the costs are a function of the flow of freight over that link. Level of service is maintained by establishing a minimum frequency of trailers which is operated regardless of freight levels. If freight levels are higher, trailers are dispatched whenever they fill up. (Typical minima are in the range of one trailer per day.) This means that the average cost function on each link is similar to that shown in Figure 2, which is the linehaul cost function used by APOLLO.

The unquantifiable aspects of the problems stem, in part, from the differences between the written plan and the execution of this plan in the field. Local conditions and real time freight level realizations cause deviations from the FMP. Some plans are better than others in providing options for the field personnel to perform local real time adjustments and may be preferable for this reason. Others are more amenable to actually being followed in the field and may be preferred for that reason. In addition, there are many factors that are

Figure 2

LTL COST FUNCTION WITH MINIMUM SERVICE

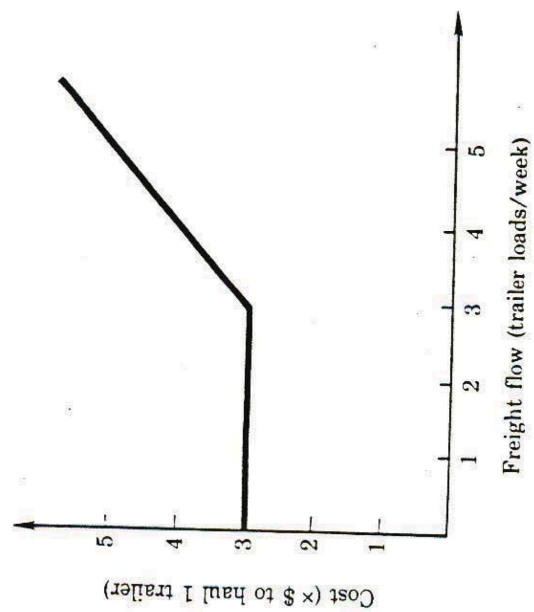

simply outside of the model domain, such as the analysts' and management's intuition regarding next month's freight flow patterns.

Mathematical programming formulations of the problem were suggested by several researchers including Powell and Sheffi[1] and Lamar and Sheffi.[2] The first abovementioned reference describes a heuristic with no known optimum-seeking properties, while the second can solve only very small problems (without including all cost items) to near-optimality. Other approaches to the problem are all based on deterministic simulation in which the load plan is specified and then the model determines the flow pattern or cost consequences of a given configuration. Examples include the linehaul model of Temple, Barker and Sloane,[3] a similar model used by Multisystems,[4] and the model developed by a team at ANR.[5]

LOAD PLANNING WITH APOLLO

APOLLO is designed as a tool for incrementally improving an existing load plan. It is used by PIE every month to plan the service and movement of freight for the next month. A typical session starts with reading in last month's freight movement plan (or the one that the analyst is currently working on) and a freight flows file (which is typically last month's freight modified by some forecast to reflect the estimate of next month's flows). All inputs to APOLLO are generated automatically from the company's standard MIS files, and thus no manual input and preparation are necessary before the analyst can work with the model.

The user typically starts the analysis by reviewing the current FMP (with next month's freight flow estimates), testing some field suggestions for changes and displaying various aspects of the system so he or she can develop a "feel" for potential improvements. Next, the analyst goes through a semistructured process (developed by the users) in which changes are suggested, tested, implemented, and documented (electronically). The monthly analysis process takes typically a few days and involves one or two analysts. Once the new plan is derived, it is reviewed by management (using APOLLO's reports and interactive capabilities) and sent electronically to the field, where the new freight routing guide is printed automatically on every bill of lading.

As mentioned before, one of the unique aspects of APOLLO is its interactive optimization nature. The interactive approach means that the responsibility for finding an improved load plan is divided between the analyst and the computer. Basically, the analyst has to determine the values of the network design variables (i.e., to decide where direct services should be established), while the computer is principally responsible for determining the values of the flow routing variables (i.e., for assigning the freight moving from each origin to each destination over the network of direct services). APOLLO, however, supports the analyst extensively in the quest for better network design decisions.

The help APOLLO provides the analyst in the abovementioned context includes a review of all the components of a given load plan; comparison between load plans; demonstration of the consequences of adding, deleting or enhancing flow on a given direct service; calculation of the impacts of rerouting any given flows; and suggestions as to which direct service should be tested.

Some decisions, however, are made automatically by the computer and not by the analyst. These include the routing of freight from origin terminal to destination terminal (subject to analyst's overrides), the distribution of empties, and the routing of TL freight. The flow routing (over the network of direct services) is performed extremely efficiently using a tailored set of shortest path algorithms. Speed is of the essence here since "what if?" questions have to be answered immediately to keep the "flow of analysis" going. The same logic is used for routing TL freight flows. The empties are distributed with an efficient transportation algorithm, subject to certain restrictions on drivers' return to their domicile.

Note that all these algorithms have to be executed over the entire network in order to respond to each "what if?" request regarding adding, deleting, or enhancing a direct service. Thus their efficiency is of prime importance. Efficiency is of even greater concern when APOLLO is used to suggest service changes (i.e., changes in the network of direct services). To come up with an estimate of the cost impacts of each suggested action APOLLO has to solve many times the problem of adding, deleting, or enhancing a single direct service.

THE INTERACTIVE FEATURES

Apollo offers an array of options as shown in Figure 3 which depicts its main working menu. This section discusses a small sample of these options in some detail in order to demonstrate their function and bring out both the complexity of the problem and the detailed level at which it was modeled.

Six of APOLLO's features are discussed here: (i) the total cost report; (ii) the path flow report; (iii) the add/drop direct service report; (iv) the suggestions editor; (v) flow overrides; and (vi) the grahical reports.

FIGURE 3

THE MAIN MENU

A P O L L O

L	- LOAD EXISTING NETWORK	M	- RUN STRATEGIC MODEL
S	- SAVE CURRENT NETWORK	I	- INCREMENTAL IMPROVEMENTS
B	- BEGIN NEW NETWORK	R	- REVIEW CURRENT SOLUTION
E	- READ EXTERNAL FILES	N	- RESTORE NETWORK
		X	- EXIT PROGRAM
ED	- EDIT DEFAULT PARAMETERS	EF	- EDIT FREIGHT FLOWS
EG	- EDIT GENERAL PARAMETERS	LO	- EDIT LINE OPERATION LINKS
ET	- EDIT TERMINALS	LP	- EDIT LOAD PLANNING LINKS
FL	- FLEET ALLOCATION MODEL	CC	- CHECK CONNECTIVITY
PR	- TURN PRINTER ON	WR	- WRITE COMMENT TO PRINTER

==>

The Total Cost Report

The total cost report (which is available to the user through several of the menus) is one of the ingredients in APOLLO's capability to give the analyst a "score card" as the analysis session progresses. This report lets the analyst check his/her progress throughout the analysis session by depicting several cost categories, as shown in Figure 4, including the costs of operating the current FMP vs. the costs of operating a base plan.

FIGURE 4

THE "SCORE CARD" TOTAL COST REPORT

CHANGE IN SYSTEM TOTALS RELATIVE TO BASE VALUES

	CURRENT	BASE	DIFFERENCE
LINEHAUL	35899	35778	121
LINEHAUL VARIANCE	11003	12733	-1730
B/B HANDLING COST	131665	167602	-35937
EOL HANDLING COST	28252	28252	0
EMPTY MOVEMENT COST	12584	13750	-1166
TRUCKLOAD COST	34952	34952	0
TOTAL SYSTEM COST	281355	320067	-38712

DO YOU WISH TO STORE THESE AS BASE VALUES? (Y/N)

To understand the cost items shown in the report, note that APOLLO artificially divides the transportation cost on every network link between the cost of carrying the freight and the cost of carrying the unused capacity on partially loaded trailers. The first item is referred to as "linehaul cost" and the second as "linehaul variance." Thus the first two lines in the cost report taken together show the total linehaul transportation cost systemwide.

The third line in the cost report depicts the handling cost at the breaks, and the fourth line the handling cost at EOL terminals. Naturally, breakbulk handling cost are a function of the amount of freight processed and they are therefore determined by the load plan. EOL costs, however, also vary with the load plan (even though the amount of freight handled by each EOL terminal does not), since these costs are a function of the number of direct services set out of each EOL.

In addition to the total cost report, the "score card" capabilities of APOLLO consist of several other reports available through other screens and options of the main menu. These include operating statistics such as load factors, freight rehandling rates, total miles driven, etc., all broken down by vehicle type, activity type and type of lane. In addition, APOLLO gives many level of service summaries (comparing the projected performance to published standards). Again, these reports are given at various levels of detail and broken down by several categories.

The Path Flow Screen

APOLLO provides many system reports to help the analyst answer "what if?" questions (in other words review a given plan). A menu of some of the available reports is shown in Figure 5 (this menu is accessible through option "R" of the main menu). The path flow report (option "P" on the above-shown screen)

FIGURE 5

THE REPORT MENU

REVIEW OF CURRENT SOLUTION

-------------------- SCREEN -------	------- REPORTS -------------------
C - SYSTEM COST SUMMARY	B - BREAKDOWN OF FLOW
D - DIRECT LOADS REPORT	G - GRAPHICAL REVIEW
O - LINE OPERATIONS REPORT	CT - TERMINALS DROPPED/ADDED
F - REVIEW DIRECT FLOW	CD - DIRECTS DROPPED/ADDED
T - TERMINAL SUMMARY	LS - LEVEL OF SERVICE ANALYSIS
P - PATH BETWEEN TERMINALS	LP - TERMINAL LOAD PLAN
DW - RUN DAY OF WEEK SIMUL	SD - SERVICE DAYS BY DAY OF WK
------------------- PRINTED -------	------- REPORTS -------------------
CP - COMPARE LOAD PLANS	TR - TERMINAL REPORTS
CA - CADE TAPE	DL - IN/OUTBOUND DIRECTS

==>

allows the user to track the freight movement plan between any two terminals. Other "what if?" reports include the list of direct services into and out of any given terminal, terminals' activity reports analyzing the flow on each direct service.

Figure 6 depicts the path of freight from terminal 004 (Portland, Maine) to terminal 215 (Miami, Florida) according to the current load plan. The report shows where the freight is handled, giving statistics on each of the direct movements in the path including distance, cost, transit time, frequency, and trailer type. (The last column in the report is reserved for an indication of a flow routing override mentioned later in this section.)

FIGURE 6

PATH BETWEEN TERMINALS REPORT

PATH BETWEEN TERMINALS

FROM: PORTLAND TO: MIAMI FL

	DIST MILES	COST $/PUP	TIME HOURS	SCHED TRTR	VEH TYPE	OVR
HANDLED AT: 004		95	6			
DIRECT FROM: 004 TO 092	440	210	13	12.30	PUPS	
HANDLED AT: 092		195	16			
DIRECT FROM: 092 TO 213	390	429	18	34.20	PUPS	
LEG 1: 092 TO 188	210	231				
LEG 2: 188 TO 213	180	198				
HANDLED AT: 213		155	11			
DIRECT FROM: 213 TO 215	350	385	9	8.70	PUPS	
SUMMARY 213 TO 215	1180	1469	5 DAYS (STD = 4)			

. . .

In addition, this report shows the actual path that a trailer would follow along a "line operations network" between the carrier's terminals. The freight on this trailer will not be handled when a trailer is relayed but the driver and possibly the tractor will change.

The Add/Drop Service Report

The interactive optimization framework allows the analyst to test the possible consequences of several actions. The strategic part of APOLLO (which is not described in this article) can be used to test new terminal locations and various economic scenarios. The tactical part of the model allows the analyst to test addition, deletion, or flow enhancement of direct services; changes in satellite to primary break connections; different service frequencies; alternative freight routing; various cost structures; and certain fleet compositions. All tactical network changes are made through a special "incremental improvements" menu which is shown in Figure 7.

FIGURE 7

THE INCREMENTAL IMPROVEMENT MENU

```
                I N C R E M E N T A L    I M P R O V E M E N T S
-------------------------------------------------------------------------------
C - TOTAL SYSTEM COSTS                  V - VIEW STORED CHANGES
R - REVIEW CURRENT SOLUTION             G - GRAPHICAL REVIEW
                                        H - ITERATE BREAKBULK HANDLING

------ MODIFY DIRECT SERVICES ---------- -------- MODIFY OVERRIDES ------------
D - ADD/DROP DIRECT SERVICE             I - INCREASE FLOW ON DIRECT
S - SYSTEMATIC NETWORK IMPROVEMENTS     U - SUGGESTIONS FOR INCREASING FLOW
M - ADD/DROP PRIMARY BREAK              L - LIST OVERRIDES
E - CLEAN UP OVERRIDES                  O - ADD/DROP OVERRIDES
-------------------------------------------------------------------------------
==>
```

As an example of a "what if?" question consider option "D - ADD/DROP DIRECT SERVICE (P)" of the incremental improvement menu, which can be used to test the impact of such action. (The paranthetical "P" means that the screen reports associated with this option will be printed if the printer is on.) To select the ADD/DROP SERVICE option, the user has to enter "D" and specify the upstream and downstream terminal identification of the direct service under study. If this direct is in the network, it is automatically tested for dropping, while if it is not, APOLLO automatically tests it for addition. The report associated with dropping a link from terminal 012 to terminal 213 is shown in Figure 8. It includes the change (in $/week) in all the cost items affected (negative numbers indicate cost reduction) if the direct under consideration is dropped. In addition, is shown the level of service impact of such a change. (Level of service is expressed as a "factor" which is the percent of transit time over standard.) The origin to destination transit time calculated by APOLLO includes linehaul time (calculated at a given average speed), terminal delays (for sorting and loading shipments), and schedule delays (waiting for the next outbound truck at every terminal on route). The report in Figure 8 shows the "old" (before the change) and "new" (after the change) service factors.

FIGURE 8

REPORT ON IMPACT OF ADDING A DIRECT SERVICE

ADD DIRECT FROM 012 TO 213

CHANGE IN COSTS		FOR REROUTED FLOW ONLY:	
LINEHAUL COSTS	-7	OLD SERVICE FACTOR	115
LINEHAUL VARIANCE	150	NEW SERVICE FACTOR	121
HANDLING COST AT BB	-331	CHANGE	6
HANDLING COST AT EOL	108		
HANDLING COST AT DEST	0	FLOW AFFECTED	3.90
EMPTY BALANCING COSTS	-143	LOCAL FREIGHT	1.50
TOTAL CHANGE IN COSTS	-223		

==> DO YOU WISH TO IMPLEMENT CHANGE? (ENTER Y/N/D/I)

The user can then choose the following: (i) to implement the change (in which case the new service is dropped from the load plan in core); (ii) to not implement the change; (iii) to obtain a detailed analysis of the change; or (iv) to obtain suggested routings that will increase the flow on the direct service under study and reduce cost.

The detailed analysis option lets the user examine closely changes in linehaul flows, changes in empty trailer movements, changes in the flows through breakbulk terminals, and changes in level of service. Both standard and graphical reports are available for all these functions. The detailed analysis option provides the analyst with an"audit trail" that can be used to validate the model calculations and build confidence in its logic. The option can also be used to study the full and detailed impact of a suggested change, including an assessment of who wins and who pays for each change.

The Suggestions Editor

As mentioned above, APOLLO helps the analyst by providing suggestions. It can perform user-directed searches and suggest service "add," "drop," and "increase flow" actions that may lead to better overall load plan. These searches are all included in options "S", and "U", of the incremental improvements menu in Figure 7.

Internally, these suggestions are generated by first ranking all the direct services or potential direct services in a specified set according to some easy-to-calculate criterion. (For example, if the user asks for suggestions for dropping services out of a given terminal, all outbound services from that terminal will be ranked by least flow.) APOLLO then performs a complete "what if?" optimization on the first n members of the ranked list (where n is user-specified), calculating the consequences of the actions request on that member (for example, the total cost impact of dropping a direct service). The suggestions are then stored in a special editor that can be accessed through option "V - VIEW STORED CHANGES" of the incremental improvements menu. An example of this editor is shown in Figure 9. The editor allows the user to review all the computer-suggested service changes which are ranked by impact on total cost.

The screen shown in Figure 9 depicts ten of the suggested improvements stored in the suggestion file. Each suggestion is shown with the associated action ("ADD," "DROP," "INC"--for increase flow, and "A/I"--for add and

FIGURE 9

THE SUGGESTIONS EDITOR

						CHANGES IN COSTS				SERVICE		
NO.	FROM	TO	FLOW	P-F	LINE	LVAR	H-BB	H-EOL	TOTAL	OLD	NEW	CHG
# 1 DROP	293	042	-0.4	0.0	0	-770	62	0	-1433	0	0	0
# 2 DROP	293	020	-0.5	0.0	1	-502	73	0	-1330	0	0	0
# 3 DROP	214	092	-3.1	0.0	-34	-827	446	-126	-1321	0	0	0
# 4 DROP	217	092	-3.1	0.0	-34	-775	446	-126	-909	0	0	0
# 5 DROP	216	092	-3.1	0.0	-34	-715	446	-126	-866	0	0	0
# 6 DROP	158	132	-0.9	0.0	109	-597	123	0	-824	0	0	0
# 7 ADD	293	101	3.1	0.0	-481	-134	-3	0	-801	0	0	0
# 8 INC	291	012	0.8	2.6	0	0	0	0	-585	0	0	0
# 9 INC	291	004	0.8	2.6	0	0	0	0	-585	0	0	0
#10 INC	291	021	0.8	2.6	0	0	0	0	-585	0	0	0

ENTER C (CONTINUE) T (TOP) I (IMPLEMENT) F (FORCE) D (DET. ANALYSIS) P (PRIMARY BREAK) E (ERASE) B (BEST N) R (REFRESH TOP N) increase flow). The next two columns identify the direct service for which the recommendation is made ("FROM" and "TO") and the fifth column depicts the amount of flow affected. The rest of the columns show the expected systemwide impact of each suggestion on cost and level of service.

The menu on the bottom of the screen shown in the figure allows the user to take certain actions with regard to the stored changes. These include various forms of implementation of any or all the suggested changes, detailed analysis and screen editing. This editor enables the analyst to work on a set of potentially good network improvements, making changes and getting feedback continuously.

The set of options that create the suggestions file, in conjunction with the editor shown above, make up APOLLO's ability to direct the analyst into a better solution by explicit recommendations.

The Treatment of Flow Overrides

As explained above, the "division of responsibilities" between APOLLO and the analyst is that the analyst decides on the network design while APOLLO routes the freight. By using "flow overrides," however, APOLLO allows the analyst to control even the detailed routing of freight. Thus the user can specify that the freight going from a given origin to a given destination should follow a certain route regardless of the minimum cost routing. Given this specification, APOLLO optimizes the routing of the unconstrained freight.

Overrides are used in APOLLO for three purposes: (i) to provide ultimate user control over the solution; (ii) to match current routing; (iii) for further optimization. These uses are explained below.

The ultimate user control is needed since the analysts have to be able to perform last minute fine-tuning before any solution is sent (electronically) to the field. Actual load plans include consideration outside APOLLO's scope that the user can implement by utilizing the override capabilities. These considerations also mean that the plan used to initialize the analysis process includes routes that do not follow APOLLO's routing logic. To create a base plan, APOLLO uses a set of overrides that are automatically generated as the data (e.g., from last month) is read in.

Lastly, overrides can be used to take advantage of unused capacity, thereby reducing the overall cost. To understand this, consider for example three terminals: 001, 002, and 003. If there is room in trucks going from 001 to 002 and from 002 to 003, it may be advantageous to send a shipment originating at 001 and destined to 003 to terminal 002 if sending it directly would require extra equipment on the direct 001-003 service. (This, naturally has to be balanced with the increased costs of handling the shipments in terminal 002.) Such an adjustment can be accomplished in APOLLO's environment by selectively inserting overrides. In fact, option "I - INCREASE FLOW ON A DIRECT" is a subprogram that finds beneficial overrides that should be inserted in order to enhance the flow on a given direct service.

Note that APOLLO offers full "what if" capabilities for override manipulation (adding, dropping, or switching), similar to those available for adding or deleting a direct service. In addition, APOLLO displays automatically the cost impact of dropping each override when it is listed.

The Graphical Reports

Computer graphics are of extreme importance in any interactive environment since they speed up the flow of information from the computer to the analyst. In addition, this information is conveyed in a form that people are used to when thinking about spatial problems. Consequently, computer graphics help the user recognize patterns that can enhance the load plan.

APOLLO is set to work on a regular CRT (IBM 3270 GA) arranged in tandem with a graphics terminal (TEKTRONIX 618). Thus the analyst gets graphical and numerical reports simultaneously. Furthermore, an analysis session can be performed continuously on the character CRT, while the graphics screen is used to display related information and to hold it as long as necessary.

APOLLO's graphics menu is shown in Figure 10. It allows the analyst to plot flows, connections, paths, and other network information. For example, Figure 11 depicts all flows into the Boston terminal that are greater than three trailer-loads per week. (The width of the rectangle on each link is proportional to the amount of flow on that link.) And Figure 12 shows all the flow impacts associated with changing a given override. (The empty rectangles in the figure designate links losing flow while the full ones designate links that gain flow.)

FIGURE 10

THE GRAPHIC REVIEW MENU

G R A P H I C A L   R E V I E W

M - REDRAW MAP	B - PLOT BREAKBULK LABELS
LO - LINE OPERATIONS NETWORK	E - PLOT END OF LINE LABELS
PO - PATHS OVER LINE OP NET	R - PLOT RELAY LABELS
FS - SYSTEM FLOWS ON DIRECTS	LD - LOAD PLAN INTO DEST
FP - FLOWS EOL TO PRIMARY BRK	FO - FLOWS OUT OF ORIGIN
FN - FLOWS EOL TO NONPRIM BRK	FD - FLOWS INTO DESTINATION
EM - FLOWS OF EMPTIES	SV - SERVICE VIOLATIONS
	TF - TRUCKLOAD FLOWS
WI - WINDOW IN	C - CHANGES IN FLOWS (I/D)
WB - BACKUP ONE WINDOW	G - SHOW LEGEND
WR - RESET WINDOW	EP - EDIT PLOTTING PARAMETERS

==>

The interactive features reviewed in this section are just a few of the ones included in the system. They exemplify the nature of the information flow between the analyst and the model. More details regarding the features shown and an explanation of all the others are contained in the program's user manual.[6]

IMPLEMENTATION ISSUES

Load planning was performed routinely by the industrial engineering staff at PIE long before APOLLO was developed. The process used to be performed every 3-6 months aided by a batch program that utilized lane-by-lane flow-based rules to route the freight. This program did not end up with an actual solution, since it could not be iterated to solve internal inconsistencies. In addition, it required a-priori specification of parts of the load plan and other extensive input

Figure 11
FLOW BETWEEN TERMINALS

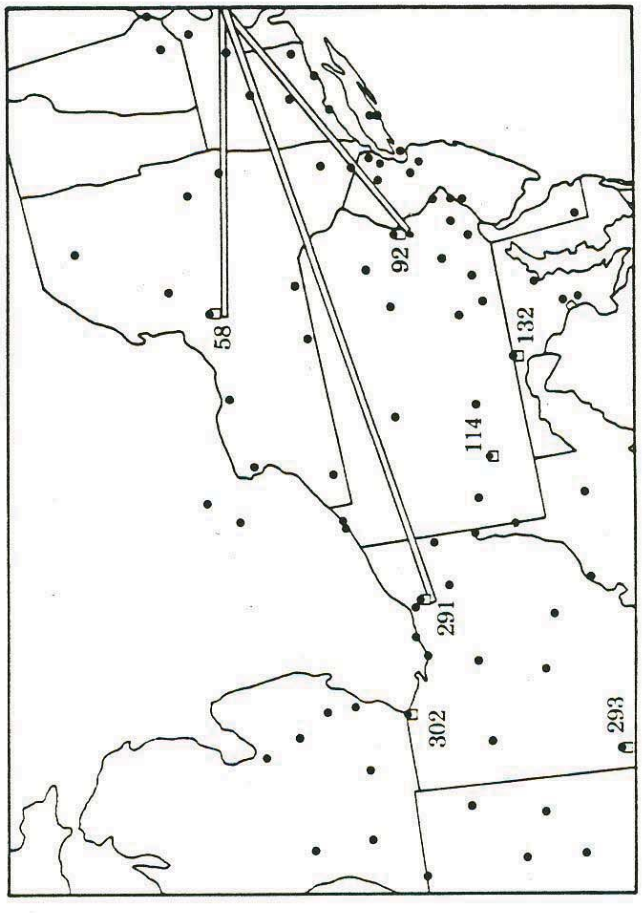

Figure 12

THE EFFECT OF CHANGING AN OVERRIDE
(full rectangle indicates significant greater flow—empty indicates significant losing flow)

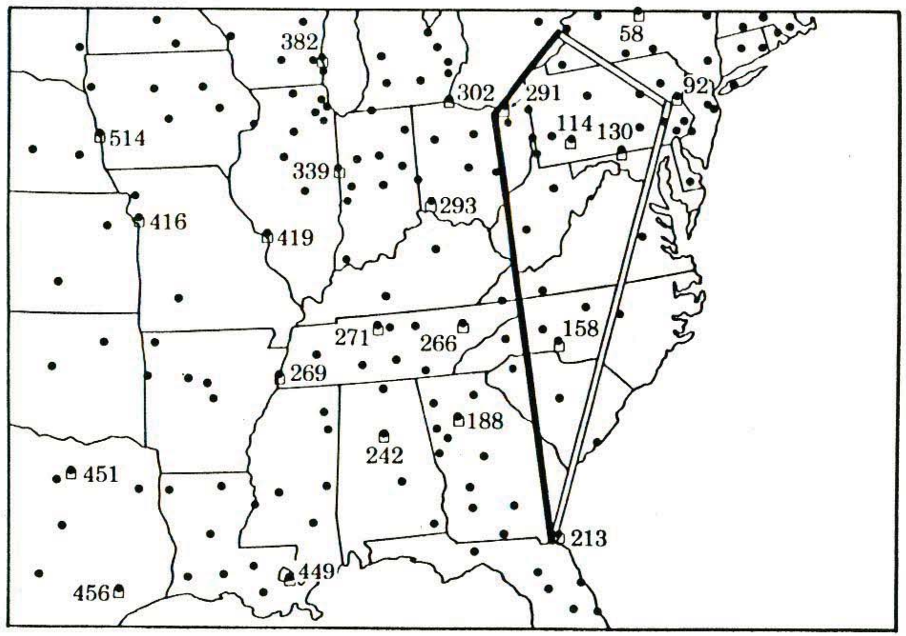

before it could be utilized. The output of the program (which typically included several thousand pages) was then reviewed manually and recommendations for changes in the current load plan were developed by the IE department and approved by a management committee. Lastly, a long list of load plan modifications would be entered to the computer to be then transmitted to the field.

The manual load planning process as described above suffered from several obvious shortcomings: (i) it was long, taking 4-6 months from start to implementation; (ii) the analysts could not account for any system effects while making local design decisions; (iii) decisions were not explicitly driven by cost and level of service considerations; and (iv) the manual review process and the data entry were prone to errors. It should be mentioned, though, that this process was typical for all large LTL motor carriers.

In addition to load planning the IE group at PIE has also been responsible for performing studies of system configurations based on management concepts, field suggestions, or internal initiatives. These include strategic questions (such as terminal location) and tactical questions (such as the assignment of EOL terminals to primary breaks, or the consequences of adding service for marketing purposes). These studies were performed manually, and were even more strongly subject to all the abovementioned shortcomings of the continuous load planning process.

The introduction of APOLLO changed the situation dramatically. It gave the IE department direct access to their own data (bypassing the MIS department), letting the planners use the model as frequently as needed and whenever desired. Another fundamental change that followed the installation of APOLLO is that the process of revising a load plan was trimmed from several months to several days. Similarly, the time needed to perform small load planning "projects" in response to management queries or field suggestions was cut from several days to several minutes. The new responsiveness of the IE department in this regard generated a much higher volume of suggestions and projects as well as the testing of many more options and ideas on each project, all leading to better load plans.

Finally, the manual data entry system was abandoned altogether in favor of full electronic transmission of new load plans from APOLLO to the MIS. The electronic transmission helped both to speed the process and to free it from human errors.

The amount of money saved by better load planning is difficult to estimate, especially since PIE has been in some turmoil during this period (in areas not related to load planning). An upper bound on the savings in the day to day operations can be estimated by using APOLLO to prepare a load plan for a past period and compare the costs of operating with the APOLLO's plan to those accrued using the load plan in effect then. Such studies have shown savings of 30-40 million dollars annually. In addition, the use of APOLLO to redo an existing breakbulk location study, has shown that the location chosen with the APOLLO was better than the location chosen with a manual approach by over 1 million dollars per year in operating costs.

Getting to the point where the model is the cornerstone of PIE's operations planning process took approximately three years of development. This process can be marked by four phases:

1. Initial problem conceptualization (7-8 months);
2. The development of the interactive model (8-10 months);
3. Major enhancements following users' reactions (6-9 months);
4. Minor enhancements following installation (12-14 months).

Of particular interest is the third phase which began when the IE department started testing the software. It became apparent that major new options were needed, the most notable one was the override capability. In order to use the model to generate a load plan electronically, the analyst had to be able to control all the model functions with no exceptions. Thus detailed routing override capability became a necessity. A full description of the override logic and its intricacies is given by Powell and Kostidis.[7]

Another major option that was developed at this phase was the "audit trail" or the detailed analysis capabilities. The problem here was that APOLLO, by accounting for systemwide effects, arrived at numbers that were not always intuitive (when solving subproblems in response to "what if?" questions). Furthermore, in some instances these numbers contradicted the numbers derived by taking only "local" effects into consideration. To develop confidence in the model's results, the users had to be able to trace the origin of every number.

Later on, the analysts kept using the audit trail in order to be informed of all the system effects of various changes.

Other major modifications had to take place due to external reasons. The first one included the merger of Ryder Truck Lines and Pacific Intermountain Express into Ryder/PIE Nationwide and later the merger of the combined company with Helmes-Barnes Truckline. (The combined company is now called PIE.) These mergers involved new data base structures due to some changes in operations philosophies (and, of course, an increase in problem size). The second set of modifications resulted from Federal regulations that went into effect in April 1983, allowing the use of double trailer combinations on all interstate and other designated highways. Following these regulations all LTL motor carriers started using heavily double trailer combinations and the model had to be adapted to the use of this type of vehicle in addition to the traditional semi-trailer units.

The next section looks at the general approach used in the development of APOLLO and draws general conclusions on guidelines for the development of interactive optimization logic for supporting transportation and logistics decisions.

CONCLUSIONS

The focus of this article is on the development of a decision support system for LTL load planning. The interactive optimization approach taken here, however, offers a viable modelling style in tackling many large and complicated problems (such as most transportation and logistics issues) in which operations research models cannot be expected to provide an optimal solution. With this class of problems, optimization models should be viewed as decision aids, providing insight and sometimes recommendations, to be used by experienced analysts.

This view of the role of models in the managerial decision-making process leads to the use of software that lets the decision maker control, direct and modify the solution process. Interestingly, several researchers have taken such approaches in developing computerized decision aids for routing and scheduling problems. (See for example, Bodin, et al.,[8] Cullen et al.,[9] and Krolak, et al.[10])

These problems parallel load planning in that they are often large, mathematically ill-behaved, include multiple objectives, and involve spatial pattern recognition.

A key ingredient in developing interactive optimization systems is the separation between the systematic, well-defined subproblems that can be solved by an algorithm, and the "fuzzy" nonquantifiable parts for which the analyst is responsible. The major challenge in building such models then is the design of the flow of information from the computer to the analyst and from the analyst to the computer.

Interestingly, "traditional" challenges of algorithmic efficiency and solution speed take on extra importance in the context of interactive optimization. The reason is that each subproblem has to be solved very fast in order to create a continuous analysis session.

Issues that are unique or particularly important in the design of interactive optimization tools can be classified into three categories: (i) the flow of information from the computer to the user; (ii) the flow of information from the user to the machine; and (iii) special features that enhance the analysis session. These categories are discussed below, summarizing our experience with APOLLO.

The flow of information from the computer to the user should be focused, and at varying level of detail. The following features were found helpful to this end: (i) On-screen summary reports (which should, for the most part, be generated automatically when the algorithm solving a subproblem terminates); (ii) use of terminology familiar to the user; (iii) highlighting (e.g., colors, screen intensity, reverse video, underlining, etc.); and (iv) computer graphics.

To facilitate the flow of information from the analyst to the computer the following attributes of APOLLO were found beneficial (i) requiring the user to supply only limited amount of information (all inputs should have default values); (ii) use of familiar terminology for input (as with the information flow from the machine to the user); (iii) easy interaction with minimum number of key strokes; (iv) lateral menu structure (so a small number of screens are necessary to accomplish any task); and (v) command stacking (to enable the experienced user to issue several commands at once).

Special features of APOLLO that were found useful in creating a productive analysis environment include the following: (i) a score card (letting the analyst know how he/she is doing during the analysis session); (ii) a search focus (where the computer helps the analyst focus his/her attention on areas where potentially good solutions exist); (iii) a road map (where the computer provides review capabilities so that the analyst knows all the features of a current solution); and (iv) an audit trail (where the analyst can verify all numbers obtained from the various optimization algorithms).

Similar features are likely to be of help in building decision support software for many other transportation and logistics problems. These are also features that model users should request from developers.

In general, one of the major advantages of interactive models is that they tend to be utilized a lot more readily than their batch counterparts (or "black box" models). The main reason for this is that by integrating the user's experience, judgment and intuition into the process, the solutions are usually better and more implementable. Also, the higher degree of user control over the model's recommendations helps create the right atmosphere for the model's use, i.e., an atmosphere in which the model does not replace the analysts, but instead enhances their capability to understand their own operation better and come up with good solutions. This turned out to be the case with APOLLO which is currently in continuous use for the routine tactical planning, for strategic planning and for special projects at PIE.

Several limitations of interactive optimization approaches should be considered, however, before such systems are developed or specified. These include the following: (i) long analysis periods (due to the interaction between the analyst and the machine); (ii) subjectivity (different analysts may come up with different results using the same model); and (iii) more programming effort is required (due to the high volume of input/output programming and the strong need for a user-friendly environment).

These limitations should be viewed in light of the advantages offered by interactive optimization models. This article does not try to argue that such models should always be used. Instead, it looks at one problem area for which interactive optimization was very appropriate and has proven to be very successful. It seems, however, that many transportation and logistics problems share the ingredients of the load planning problem and therefore may be tackled with similar methods.

ENDNOTES

1. Powell, W. and Y. Sheffi (1983), "The Load Planning Problem of Motor Carriers: Problem Description and Proposed Solution Approach." Transportation Research A 17(6), 471-480.

2. Lamar, B. and Y. Sheffi (1985), "Optimal Network Design for LTL Trucking." Working Paper CTS 85-17, Center for Transportation Studies, MIT, Cambridge, MA.

3. Temple, Barker, and Sloane (1978), Motor Carrier Linehaul and Breakbulk Planning Model, Version III, Program's User Manual, Temple, Barker, and Sloane, 33 Hayden Avenue, Lexington, MA 02173.

4. Multisystems (1979), Description of Motor Carrier Breakbulk Planning Model, Program's user manual, MULTISYSTEMS, 1050 Massachusetts Avenue, Cambridge, MA 02138.

5. Barber, H.H., F.M. Sharon and D.K. Sen (1981), "From Freight Flow and Cost Patterns to Greater Profitability and Better Service for a Motor Carrier." Interfaces II, No. 6, 4-20.

6. Transoft, Inc. (1985), APOLLO: Advanced Planner of LTL Loading, Version 4.1, User Manual, TRANSOFT, Inc., 15 Shuman Circle, Newton, MA 02159.

7. Powell, W. and I. Koskisidis (1985), "Shipment Routing Algorithms for Less-than-Truckload Trucking Networks." Working paper EES 85-8, Civil Engineering Department, Princeton University, Princeton, NJ.

8. Bodin, L., B. Golden, A. Schuster and W. Romig (1980), "A Model for the Blocking of Trains." Transportation Research 14B, 115-120.

9. Cullen, F., J. Jarvis, and H. Ratliff (1984), "Set Partitioning Based Heuristic for Interactive Routing." Networks 11(2), 125-144.

10. Krolak, P., W. Felts, and J. Nelson (1971), "A Man-Machine Approach Towards Solving the Generalized Truck Dispatching Problem." Transportation Science 6, 149-170.

ACKNOWLEDGEMENT

The authors thank I.U. International Management Corporation for their continued financial support for the authors' university research at M.I.T. and Princeton. In addition, our thanks go to the group of able and dedicated industrial engineers at PIE with whom we shared many long hours during the system development.

ABOUT THE AUTHORS

Warren B. Powell is Assistant Professor of Civil Engineering at Princeton University. He completed his undergraduate work at Princeton and earned the Ph.D. from the Massachusetts Institute of Technology. His research interests include the application of operations research methods to transportation problems with emphasis on trucking issues. Dr. Powell is Associate Editor of Transportation Science, and has published heavily in Transportation Research.

Yosef Sheffi is Head, Transportation Systems Division, Department of Civil Engineering at the Massachusetts Institute of Technology. He is also a faculty member at the MIT Center for Transportation Studies and the MIT Operations Research Center. He earned the Ph.D. in 1978 from MIT. Dr. Sheffi specializes in network optimization and other operations research methods and their applications to carrier operations planning and control. He has published over 50 papers in professional journals on transportation system analysis. He is author of Urban Transportation Network: Equilibrium Analysis with Mathematical Programming Methods, and is a frequent speaker in logistics industry forums as well as scientific symposia.